Developing Bibliometric Indicators of Research Performance in Computer Science: An Exploratory Study

Henk F. Moed and Martijn S. Visser

Centre for Science and Technology Studies (*CWTS*)
Leiden University, the Netherlands

Research Report to the Council for Physical Sciences of the Netherlands Organisation for Scientific Research (NWO)

February 2007

CWTS Report 2007-01

For more information please contact:

Dr. Henk F. Moed

Centre for Science and Technology Studies (CWTS)
Leiden University
Wassenaarseweg 52
Postbus 9555
2300 RB Leiden
The Netherlands

Tel.: *++ 31 71 5273940*
Fax: *++ 31 71 5273911*
E-mail: *Moed@CWTS.Leidenuniv.nl*

Preface

The project presented in this report was commissioned and funded by the Council for Physical Sciences of the Netherlands Organisation for Scientific Research (NWO). The project started early 2005, and was finished in December 2006. The work was supported by an Advisory Committee, installed by the Council for Physical Sciences, and consisting of the following members:

Dr A.P. Meijler (NWO, Council for Physical Sciences, Chairman of the Committee);
Prof. Dr. S. A. van der Geer (Leiden University, until September 2005);
Prof. Dr. H.J. van den Herik (University of Maastricht);
Prof. Dr. P. Klint (Centre for Mathematics and Computer Science, as from Sept. 2005).
Prof. Dr. A.P.J.M. Siebes (Utrecht University);
Prof. Dr. P.M.A. Sloot (University of Amsterdam);
Prof. Dr. F.W. Vaandrager (Radboud University Nijmegen);
Dr. E.W. van Zwet (Leiden University, as from September 2005).

The authors are most grateful to the members of the Advisory Committee for their important and stimulating contributions to the project, and to all researchers who have participated in the ‘verification round’ carried out in the project for their valuable comments. It needs emphasising that the authors are solely responsible for the content of this report. In addition, the authors wish to thank Drs. B. Bauland at VSNU for his assistance in collecting data pertaining to the QANU assessment, Dr. A. Hofmann at Springer (Heidelberg, Germany) for giving his permission to download and use full text documents from LNCS available at Springerlink.com, and Dr. S.R. Smith and Dr. G. Poindexter at the IEEE Computer Society for preparing a special data file with full texts of conference proceedings included in Computer.org, and for giving their permission to use these data in the study.

Executive Summary

During the past decades, bibliometric methods were successfully developed and tested as tools in the assessment of research performance in basic scientific disciplines such as Chemistry, Physics and Biomedical Sciences, using data from the *Science Citation Index* (*SCI*) or *Web of Science* (*WoS*) published by *Thomson Scientific*. From this experience a demand emerged to extend the application of bibliometric methods to other fields of science, including Computer Science.

On the one hand, the academic authorities of many universities and Research Councils expressed the need to obtain insight into the research performance in *all* departments in *all* fields of scholarship. On the other hand, evidence from statements by experts in the field, and from independent bibliometric research suggested that, in Computer Science, proceedings volumes rather than scientific journals constitute the main channel of written communication, and that the *Thomson* Indexes may not cover the written communication in this field sufficiently well.

Therefore, a pilot study was carried out, funded by the Netherlands Organisation for Scientific Research (NWO), aimed to expand the *Thomson* Indexes with source articles from a number of refereed proceedings volumes of important international conferences and to carry out a citation analysis in this compound universe of research articles published by Netherlands researchers in the field Computer Science. This report presents the outcomes of this study. It needs emphasising that the authors are solely responsible for the content of this report.

A source expanded citation analysis of Netherlands research papers in academic Computer Science was carried out in the following way. A database created at *CWTS* from *Thomson*'s *Web of Science* data (denoted as *CWTS-WoS database*) was expanded with source metadata and cited references from articles in proceedings of international conferences published by:

- Springer, in its Lecture Notes in Computer Science (*LNCS*). It should be noted that especially prior to 2003 only a selection of the *LNCS* volumes were covered by the Thomson Indexes.
- Association for Computing Machinery (*ACM*). The proceedings of over 200 recurring conferences are made available as part of the *ACM* Digital Library.
- Computer Society of the Institute of Electrical and Electronics Engineers (*IEEE*).

The decision to include conference proceedings from the digital libraries of *ACM* and *IEEE*/CS and the *LNCS*-series was based on the following considerations:

- The analysis of conferences appearing most frequently in the publication and reference lists of Netherlands academic computer scientists shows that the proceedings of these conferences were often published by *ACM, IEEE*/CS or as a volume of *LNCS*.

- The fact that the proceedings published by *ACM*, *IEEE*/CS, and in *LNCS*, are taken from a broad and large collection of conferences reduces the risk that certain important subfields of Computer Science are missed.
- In order to expand the database with the papers that have a sufficiently high level of quality, it was considered appropriate to select proceedings that are seriously refereed. The *ACM*, *IEEE* and *LNCS* proceedings generally meet this criterion.
- Several members of the project's Advisory Committee underlined their importance to the research community in Computer Science, and also to Netherlands researchers in this field.
- Both Springer, the publisher of *LNCS*, and *IEEE* have granted us permission and offered their assistance with collecting the data we needed.

Technical processing involved the collection of relevant full text documents mostly in PDF-format, their conversion from PDF to text format, parsing these text documents and extracting their reference lists. The bibliographic details uncovered in this way were then integrated with the *CWTS-WoS* database. Citations from conference papers to journal articles and vice versa were identified as well as citation linkages within the conference literature. Around 160.000 source articles were added covering the years 1996–2004, and over 1.6 million cited references contained in them. In this way, the total number of Computer Science articles in the compound database articles increased by 66% compared to that already included in the *WoS*.

This report presents the outcomes of the study. This report is based on the notion that bibliometric indicators have proven to be useful tools in the assessment of research performance in many scientific disciplines, provided that such indicators have a sufficient level of sophistication, that their pitfalls are taken into account, and that they are used in combination with other, more qualitative information on the research groups analyzed. Such tools do *not* aim to *replace* peer judgments, but instead *supplement* them as evaluation *support tools*. The report focuses on the following six main questions.

1. Is the expansion of the WoS database with a number of important conference proceedings technically feasible? Which problems arise and hoe can these be overcome?

It is concluded that expanding the *CWTS-WoS* database with conference proceedings sources is technically feasible, provided that their meta-data (including cited reference lists) is available in electronic form. However, it involves a lot of elementary data processing. The amount of work depends upon the nature and quality of the relevant meta-data of articles from these sources. If the meta-data have to be extracted from PDF documents, the process of data collection can be qualified as cumbersome. In this process, a part of the relevant data is lost. The study estimated that from about 17 per cent of source articles in PDF format (i.e., in IEEE and LNCS proceedings *not* included in the *WoS*), the cited reference lists were only partly extracted or not extracted at all. The major part of these papers is published in the early years 1996-1997. Although this loss of data does not necessarily invalidate the outcomes of bibliometric analysis, it could be useful to collect in a follow up study even more expert knowledge on these processes of data extraction, and examine whether their recall can be enhanced.

The study identified and solved many technical problems as regards the process of citation matching. A crucial problem is how to deal with the phenomenon that different versions may exist of one and the same publication, – e.g., as technical report, in several proceedings volumes, in a journal, and even included in a book –, or at least that publications made by an author in a particular year may have very similar if not identical titles. On the other hand, it must be emphasized that, even if cited references are not matched to the intended target publication, they are nevertheless matched to the correct author or group of authors. In other words, the problem is how to distribute in a proper way cited references to an author's work among the various publication versions emerging from it. A partial solution to this problem is outlined in Section 2.2, according to which priority is given to 'versions' published in journals included in the *Expanded WoS* database.

It is concluded that more work needs to be done in order to tackle these problems and thus further increase the accuracy of citation counts. In principle there are two lines along which one could proceed. *A 'classical' approach* conceives the individual publication as the basic unit in the citation analysis and further enhances the process of matching citations to individual publications. Citations to publications that are not published in journals or in well formatted conference proceedings need to be collected or at least checked *manually*. The same is true for papers published in different versions (e.g., as technical report, proceedings article, book chapter and as a journal paper), as it is extremely difficult to allocate merely by means of computer programs citations to the intentionally cited version of such papers. In an alternative approach, the basic unit of citation analysis would be a *concept,* - i.e., an idea, key finding or methodology – embodied in a *series* of publications, rather than an *individual* publication. However, there is as of yet no 'standard' methodology available based on this approach.

2. What is in the field Computer Science the relative importance of proceedings articles compared to journal papers, as reflected in citation links?

The results obtained in the study underline the importance of conference proceedings, – particularly *ACM*, *LNCS* and *IEEE* proceedings volumes –, in terms of the citation impact these volumes generated both upon proceedings and upon journals. The *Source Citation Rate*, an 'impact factor-like' citation impact measure, is for the proceedings volumes analysed in this study on average almost as high as it is for annual journal volumes.

The proceedings volumes analysed in the study tend to show a somewhat higher variability in their citation impact rates than annual journal volumes do. Both in the top and the bottom of the distribution of citation impact among sources, proceedings volumes are somewhat overrepresented. In other words, there are relatively more highly cited *and also* more poorly cited proceedings volumes than there are annual journal volumes.

The citation links among proceedings volumes of recurring (e.g., annual) conferences tend to be as strong as those among annual volumes of the same journal. These proceedings series reveal citation patterns that are statistically similar to those shown by a journal's annual volumes, if not stronger. The Report presents a listing of the top 50 annual journal or proceedings volumes in terms of their Source Citation Rate, and a table of the pairs of publication volumes that have the strongest citation links.

These findings corroborate outcomes from earlier studies and claims made by computer scientists as regards the importance of conference proceedings as channels of written communication in their field. Using citation impact of a publication source as an indicator of its importance, it follows that in the *WoS* database a number of important conference proceedings volumes is missing. These volumes tend to be as important as the journals that are covered by this database. For proceedings of (bi-) annual conference series, successive volumes tend to as important for one another as successive annual volumes of journals are for one another. Their inclusion in the *Expanded WoS* enhances the coverage of the important channels of written communication, and therefore provides a more accurate and a more valid bibliometric assessment of research performance in Computer Science.

3. How well does the WoS database expanded with ACM, LNCS and IEEE proceedings cover the publication output of Netherlands computer scientists?

The study created a database of the publications of Netherlands academic computer scientists during the time period 1996–2001, based on the publication lists in the self evaluation reports that were prepared for the *QANU* research evaluation conducted in 2003 and for the separate evaluation of the Computer Science Departments of Leiden University and Delft University of Technology. The publication details for CWI were collected from its Annual Research Reports (time period 1997-2001).

Before the inclusion of additional conference proceedings, 25 per cent of Netherlands Computer Science papers were included in the *CWTS-WoS* database. The expansion of the database raised this percentage to 35. As the number of additional conference papers with which the database was expanded increased sharply over time, the coverage rate increased as well, from 22 per cent in 1996 to 41 per cent in 2001. Since the major part of the papers from ACM, LNCS and IEEE sources added to the *Expanded WoS* database are from later years, the coverage rate of Netherlands Computer Science papers can be expected to further increase during de time period 2002-2006.

In order to analyse differences in coverage of the *Expanded WoS* database among subfields, Netherlands groups were categorized in a simple and pragmatic way, following the structure of the Netherlands Research Schools in this field. Even if this structure has purely national features that cannot be found in other countries, it can be used to show differences in coverage among collections of Netherlands groups with different cognitive orientations. These differences were found to be substantial. *Computing and Imaging* shows the highest coverage by the *Expanded WoS* database, reaching in 2001 a value of 53 per cent. For *Programming Research and Algorithmics* the coverage percentage for this year amounts to 41, and for *Information and Knowledge Systems*, – the largest subfield in terms of numbers of papers covered – to 32 per cent. The smaller subfields *Logic* and *Telematics* show in 2001 coverage rates of 36 and 24 per cent, respectively. Finally, *Mathematics*, covering the more mathematically oriented groups, shows a coverage rate of 53 per cent.

Another way to assess adequacy of coverage is calculating '*database internal*' coverage percentages, based on the extent to which (source) articles included in a publication database cite other (source) articles in that database. For disciplines such as physics, chemistry and many parts of biological sciences the internal coverage of the *WoS* is above 80 per cent. For computer science, expansion of the WoS database with ACM, LNCS and IEEE proceedings raised this percentage from 38 to 51. Although the increase is substantial, a level of 51 per cent should be qualified as moderate.

Even though coverage percentages per subfield are most relevant and informative, they depend upon the degree of self selection researchers have imposed themselves when compiling their publication lists. Therefore, it is also relevant and informative to examine one by one the sources in which Netherlands computer scientists published their articles and that were *not* covered by the *Expanded WoS* database. An evaluation of journals or conference proceedings requires a detailed knowledge of the fields they cover, and their communication networks. Citation impact measures provide useful tools to carry out such an evaluation.

Therefore, a list was compiled per subfield of the conference series with the highest numbers of Netherlands papers and not covered by the *Expanded WoS* database, indicating also the average citation impact of the Netherlands papers (measured in the citation universe of the *Expanded WoS* database). As far as concentration could be observed, several conferences were *local* rather than *international*, while the Netherlands papers published in them had on average a very low citation impact. Typical examples are the *Proceedings of the Conference of the Advanced School for Computing and Imaging (ASCI)*, and the *Proceedings of the Belgium-Netherlands Conference on Artificial Intelligence (BNAIC)*. Such national conferences and their proceedings may play an important role in creating and maintaining national networks, but from the perspective of assessing the contribution at the international research front their role seems less important.

But in all subfields the lists also contain proceedings of international, recurring conferences. Moreover, for several sources the citation impact of Netherlands papers was found to be substantial, reaching a level that is similar to that of many sources that are covered by the *Expanded WoS* database. Typical examples are (for a complete list of important sources mentioned in the verification round the reader is referred to Table 8.1 in Chapter 8):

Conference of the IEEE Communications Society (INFOCOM)
Conference on Parallel Problem Solving from Nature (PPSN)
Eurographics incl. Various Workshops
European Conference on Artificial Intelligence (ECAI)

European Conference on Software Maintenance and Reengineering (CSMR)
Genetic and Evolutionary Computation Conference (GECCO)
IEEE International Conference on Robotics and Automation (ICRA)
International Joint Conference on Artificial Intelligence (IJCAI)
International Conference on Very Large Databases (VLDB)
Text Retrieval Conference (TREC)
Proceedings of Recherche d'Information Assistée par Ordinateur (RIAO).

4. What are the statistical relationships between the bibliometric outcomes for the various Netherlands Computer Science departments on the one hand, and peer ratings of these departments given by the QANU Review Committee for Computer Science in 2004 on the other (Chapter 6)?

The study compared the outcomes of the bibliometric analysis per group to the quality ratings given to a group by the *Review Committee on Computer Science* in 2003/2004. This Committee carried out an evaluation, organised by the *Quality Assurance Netherlands Universities* (*QANU*), of 42 Computer Science groups in a number of Netherlands Universities. Such a comparison can be made from two distinct points of view. The first is that of validation of bibliometric indicators, using peer ratings as a benchmark A second point of view critically examines peer ratings of particular evaluation panels, while bibliometric indicators are applied as a benchmark. This report further develops *both* viewpoints. It does not assume the primacy of one of the two methodologies to assess research performance – peer review or bibliometric analysis – above the other.

The study calculated rank correlation coefficients between peer ratings and bibliometric indicators. These coefficients should be interpreted as purely *descriptive* statistics. It was found that the bibliometric indicator showing the highest rank correlation with the quality peer ratings of the Netherlands academic Computer Science groups, is the number of articles in the *Expanded WoS* database. This can be interpreted as evidence that the extent to which groups published in refereed international journals and in important conference proceedings (*ACM*, *LNCS*, *IEEE*) has been an important criterion of research quality for the Review Committee. Following this interpretation, the methodology developed in this study is on the right track. Any attempt to identify research quality or excellence in Computer Science should discriminate between truly important publications, and less significant ones.

The rank correlations between the quality rating of Netherlands computer science groups and citation impact indicators of their papers, are positive, but weak (around 0.2). In order to discuss this outcome further, the ratings for Computer Science were compared to those assigned in 2002 by the Review Committee on Chemistry to academic Chemistry groups Review. Compared to the Chemistry groups, the quality ratings of the academic Computer Science groups show a higher level and less variability, whereas their citation impact tends to be lower, and, – taking the coefficient of variation (the ratio of standard deviation and mean) as a standard –, shows more variability. Variability in citation impact and lack of correlation with peer ratings were also found *within* subfields.

One can argue that it is remarkable that the quality ratings of the various academic Computer Science groups are so similar one to another, whereas their citation impact reveals so much variation, and focus on the peer review process in which the ratings were generated. The authors of this report do *not* claim that the Review Committee for Computer Science (or any other review committee) *should* have based their judgments merely upon citation analysis, or that there should be a perfect or even a strong correlation between its ratings and citation impact. From the apparent weak correlation between peer ratings and relative or normalised citation impact indicators, one can *not* conclude that the peer ratings are basically invalid. But one could at least raise the question how the Review Committee for Computer Science evaluated aspects specified in its evaluation protocol as 'international recognition and innovative potential', the extent to which the work of group is 'at the forefront internationally' and generates an 'important and substantial impact in the field'.

The *QANU Review Committee for Computer Science* did *not* use the outcomes of the citation analysis presented in this report. A further development and validation of the indicators computed in this study could take place in a future peer review of the field. Application – be it experimental – in a peer review context ensures that background knowledge about the groups to be evaluated and the subfields in which they are active is taken into account in the interpretation of bibliometric indicators, and in this way establishes necessary conditions for their proper use.

5. How does the citation impact of journal and conference papers by Netherlands computer scientists compare to international levels of citation impact in the field Computer Science (Chapter 7)?

This study applied a methodology that is similar to the one used in many other studies of research performance. The principal difference is that in the current study the *WoS* database was expanded with *ACM*, *LNCS* and *IEEE* proceedings. Assuming that the publication database and methodology provide a sufficiently valid reflection of the research performance of *Netherlands* academic Computer Science, the following tentative conclusions could be drawn.

The citation impact of the Netherlands academic Computer Science groups is significantly above world average. An *overall normalised citation impact* of 1.30 was found (a level of 1.0 represents the world average), increasing to a level of 1.4 and 1.6 for papers published in the last two years (2000 and 2001). This increase in normalised citation impact is also visible in an analysis of papers in *journals* covered by the Web of Science, and for articles from each major *subfield*. The outcomes suggest that there has been a genuine increase in the average citation impact of Netherlands Computer Science papers during 1997-2001.

In addition, they suggest that the citation impact of Netherlands academic Computer Science has the same level of that of Netherlands academic research activities in other disciplines in the exact sciences. It was also found that among the top 10 per cent most frequently cited articles published world-wide in Computer Science, the number of papers by Netherlands academic computer scientists is 50 per cent higher than expected on the basis of the total volume of Netherlands publication output in the field.

The above observations relate to the aggregate set of all Netherlands computer science papers in the Expanded WoS database. It is equally informative to analyse *differences* among Netherlands groups. The outcomes of the exploratory and preliminary analyses presented in this report indicate a rather strong variability in the citation impact of Netherlands computer science, both at the level of individual papers and that of research groups, a variability that is not reflected in the quality ratings of the QANU Peer Review Committee evaluating a large segment of Netherlands academic Computer Science groups in 2003/2004.

6. How did the leaders of groups involved in the study react in the verification round to the listings of publication and citation data collected for their groups (Chapter 8)?

During the course of the project it was decided to carry out a verification round among the leaders of the groups involved in the study, enabling them to assess the accuracy of the data, and particularly the adequacy of coverage of the *Expanded WoS* database. A web site was created, showing per group publication and citation data collected in the study. It indicated an email address to which replies could be sent. Through this address we received reactions of 27 researchers. The response rate was about 40 per cent. Their comments are summarized in Chapter 8 of this report.

A first crucial issue was the completeness, data quality and policy significance of the publication lists of Netherlands computer scientists collected in the study. It was noted that important publications were not included in the lists, and that some publications were listed twice. In addition, it was emphasised that for several groups the lists were not representative for the recent performance of the group, as some researchers had left a group, and new researchers entered it during the time period from 2002 up to date. Though significant, these comments were somewhat unexpected, as they criticised the publication data collected from the QANU review and CWI institutional annual reports. It follows that for a number of groups the relevance of the bibliometric outcomes in current policy is rather limited.

The second issue relates to publication sources that were not covered by the *Expanded WoS* database, and that respondents qualified as important and therefore missing. They did not only indicate important conference proceedings to be added to the database, but also a number of scientific-scholarly journals that are not (yet) covered by the Web of Science. It must be noted that the major part of these missing journals are fully covered in recent years.

As regards proceedings volumes, many missing sources respondents qualified as important are *not* included in the lists of sources most frequently used by Netherlands computer scientists. The standardization of source titles has not been perfect, and the publication lists related to the time period 1996-2001, while respondents may have indicated sources that were founded in recent years, or sources in which they started publishing after 2001. But even if this were so, there seem to be major discrepancies between the two lists of sources.

7. Recommendations

- In order to obtain a global publication database with an adequate coverage of the field computer science, both from a bibliometric perspective and in the perception of Netherlands researchers in this field, the existing *Expanded WoS* publication database should be *further* expanded (and updated) with a number of additional important conference proceedings volumes. Suggestions for sources to be included were given above, and especially in Table 8.1 in Chapter 8.

- Experts in the various subfields could compile a list of the most important journals and conference proceedings that should be added to the currently existing *Expanded WoS* database. The citation impact rates of sources and their importance as perceived by respondents in the verification round would constitute valuable sources of information. From a practical point of view, the availability of sources in electronic format should also be taken into account.

- It is unlikely that the expansion of the database with these additional proceedings volumes would raise the internal coverage to a level above 80 per cent. Assuming realistically that it will be below 80 per cent, one should carry out ***within a further Expanded WoS database*** a type of citation analysis, in which citations to articles published in journals or proceedings that are not covered by the expanded database are also counted, and contribute to the citation impact rates of a group's publication output. The application made available in the verification round through the internet only included citations to articles published in journals or proceedings included in the *Expanded WoS* database.

- A future bibliometric study of Netherlands computer science should be based on publication lists that are provided – or at least *verified* – by the evaluated researchers themselves, applying standard data collection and verification procedures developed at CWTS. The publications to be analysed should relate to the most recent ten-year period (e.g., a study carried out in 2008 should analyse publications published during 1998-2007).

- Evaluated researchers justly demand that citation counts of their publications are accurate, especially since citation distributions are highly skewed. It appears to be extremely difficult to determine in a fully automated way accurate citation counts to publications that are not published in journals or in well formatted conference proceedings. In addition, if papers are published in different versions (e.g., as technical report, proceedings article, book chapter and as a journal paper), it is most difficult to allocate merely with computer programs citations to the intentionally cited version of such papers. A citation analysis of these types of papers should therefore be carried out – or at least checked – *manually*.

- In a bibliometric assessment of Netherlands research groups one is not merely interested in comparisons among Netherlands groups, but also – if not primarily – in comparing the performance of Netherlands researchers with that of their international peers. In the type of citation analysis recommended above, in which citations to articles published in sources not covered by the expanded database are also counted, there is as of yet *no* ‘standard’ methodology available that allows for international comparisons. In a follow up study, further research into the development of such a methodology is necessary.

- The list of conference proceedings in which Netherlands computer scientists published frequently reveal a huge scattering of published articles among sources in all subfields. The community of computer scientists in the Netherlands could address the question whether this scattering is a basic or ‘natural’ characteristic of the field (or a specific subfield), or whether it perhaps at least partly indicates that many groups in Netherlands academic computer scientists have not (yet) developed a well-considered strategy as regards the participation at international conferences and the selection of publication sources.

- It is important for a research community to create and apply quality standards. One way to come to such standards could be to identify particular conferences and publication sources that one can normally only enter if the work presented is of high quality. In addition, more selectivity in the selection of conferences and publication sources could increase the visibility and international standing of Netherlands academic Computer Science.

1 Introduction

This report presents the main outcomes of a methodological study, aimed to develop – i.e. construct and test – bibliometric indicators of research performance in Computer Science, based upon insight into the system of written communication in that field. The study focuses on the development of citation impact indicators, measured by the number of times articles published by a research group are cited in other documents published in important publication sources. The object of study is formed by the research activities conducted by researchers in the field Computer Science at General and Technical Universities in the Netherlands, and at the Centre for Mathematics and Computer Science in Amsterdam (CWI). Throughout this report, these activities will be labelled as *Netherlands academic Computer Science*.

Evaluation of research performance is traditionally carried out in peer reviews. During the past decades, bibliometric methods were successfully developed and tested regarding research performance in basic scientific disciplines such as Chemistry, Physics and Biomedical Sciences (Narin, 1976; Garfield, 1979; Martin & Irvine, 1983; Schubert et al., 1989; Moed, 1995; Van Raan, 1996). Bibliometric indicators have proven to be useful tools in the assessment of research performance in many scientific disciplines, provided that such indicators have a sufficient level of sophistication, that their pitfalls are taken into account, and that they are used in combination with other, more qualitative information on the research groups analyzed. Such tools do *not* aim to *replace* peer judgments, but instead *supplement* them as evaluation *support tools*.

The Centre for Science and Technology Studies (*CWTS*) has conducted numerous bibliometric studies of research groups in basic science disciplines, commissioned by directors of individual research groups, departments, institutions and organizations, or by policy makers of governmental departments located both in the Netherlands and abroad (Van Raan, 1996). In these studies data from the Science Citation Index (SCI) and related Citation Indexes published by the Institute for Scientific Information (*ISI*, currently *Thomson Scientific*) played a prominent role. *CWTS* has created a bibliometric database of all articles in journals processed by *Thomson Scientific* for the *Web of Science* (*WoS*). The *WoS* covers some 8,000 leading international journals from all domains of scholarship.

A common methodology applied in such studies collects articles that were published by a unit to be assessed in journals processed for the *Web of Science (WoS)*, and analyzes their citation impact, by counting the number of times they are cited in other papers published in *WoS* source journals. Such a methodology is denoted below as a ‘*Pure WoS* analysis’.

In discussions with Netherlands computer scientists during the setup phase of the study, several participants qualified Computer Science as a relatively new, rapidly developing field with a strong multi-disciplinary orientation, embracing both basic and more applied research activities. It was claimed that, particularly in the youngest subfields, proceedings volumes rather than scientific journals constitute the main channel of written communication; that the *ISI* citation indexes do not adequately cover the journal literature in all subfields; that there may be no clear hierarchy of journals based on their perceived importance; and that in some subfields the scientific article may not even be the most important carrier of scientific output.

Several studies carried out in the past have demonstrated the frequent use of conference papers in the field Computer Science (e.g., Goodrum et al., 2001). However, the claim that conference papers are important was not only based on their frequent use, but also on the particular role conference papers play in the communication process. As many conference papers do not end up in journals, Drott (1995) put forward that a conference paper may very well be a final product in itself. Based on discussions with computer scientists, Goodrum et al. (2001) concluded that in Computer Science a research article in a proceedings volume may serve the same purpose for the author as a journal article. Therefore, researchers may not seek to have their conference papers published in journals.

These comments and studies indicate that the *Web of Science* and related Citation Indexes published by *Thomson Scientific* (formerly Institute for Scientific Information (*ISI*)) may not cover the written communication in this field sufficiently well, since this database covers mainly (though not exclusively) scientific journals.

Therefore, a pilot study was carried out, funded by *the Netherlands Organisation for Scientific Research (NWO),* aimed to expand the *Web of Science* with source articles from a number of peer-reviewed proceedings volumes of important international conferences, and to carry out in this expanded universe a citation analysis of research articles published by Netherlands academic researchers in Computer Science. A source expanded citation analysis of Netherlands academic research papers in Computer Science was carried out in the following way. A database created at *CWTS* from *Thomson*'s *Web of Science* data was expanded with source metadata and cited references from articles in proceedings of international conferences published by:

- Springer, in its Lecture Notes in Computer Science (*LNCS*, only partly covered by the *Thomson* Indexes);
- Association for Computing Machinery (*ACM*); and
- Computer Society of Institute of Electrical and Electronics Engineers (*IEEE*).

Is must be emphasized that this database contains *all* articles published in the selected sources, i.e., papers published world-wide and *not* merely papers published by Netherlands researchers.

Chapter 2 addresses the first research question of the study: *Is the expansion of the WoS database with a number of important conference proceedings technically feasible? Which problems arise and how can these be overcome?* Section 2.1 describes how an expanded version of the *WoS* database was created. A major technical task carried out in the study was the collection of relevant full text documents mostly in PDF-format, their conversion from PDF to text format, parsing these text documents, and extracting the relevant metadata including their reference lists. The bibliographic data uncovered in this way were then integrated with the *CWTS-WoS* database. The new database is denoted as *Expanded CWTS-WoS* database or briefly as *Expanded WoS* database or simply as *WoS+* throughout this report. Section 2.2 gives an outline of the methods applied in the process of *citation matching* in the new database. Finally, Section 2.3 describes the creation of a database of publications made by Netherlands computer scientists. It is labelled as the *NL-CS publication database* throughout this report, and covers publications published during the time period 1996-2001. Publications in this database were subjected to a first, exploratory citation analysis, presented in Chapters 5, 6 and 7.

Chapter 3 presents a number of *methodological* issues that play an important role in later chapters. Section 3.1 deals with citation analysis. It presents base principles, and specifies the various publication and citation universes that were explored in the study. Section 3.2 introduces the concept of internal coverage of the WoS database, and distinguishes four types of bibliometric studies. It provides a methodological background of the study. Section 3.3 briefly describes a series of bibliometric indicators that were calculated in later chapters. The principal indicator proposed is a measure that expresses the citation impact of a group's papers, relative to the world citation average in the (sub-)field(s) in which it is active.

Chapter 4 focuses on the role of conference proceedings in Computer Science. Analysing the *Expanded WoS* database described in Chapter 2, it addresses the following research question: *What is in Computer Science the relative importance of proceedings articles compared to journal papers, as reflected in citation links?* It presents citation 'impact factors' for the main journals and conference proceedings included in this study. For conference proceedings, impact factors can be calculated that are similar to journal impact factors published by the Institute for Scientific Information for its covered journals. To the extent that such impact factors are calculated in the proposed study, they will be used merely to assess the sources' importance in the written communication system. They will *not* be used as surrogates of actual citation impact of the papers published by the groups involved.

Chapter 5 discusses the *coverage* of the newly created *Expanded WoS* database. The research question addressed is: *How well does the Expanded WoS database described in Section 2.1 cover the publication output of Netherlands computer scientists included in the NL-CS database, described in Section 2.3 (the NL-CS database)?* An analysis by subfield calculated several indicators of adequacy of coverage. One indicator is the percentage share of NL-CS publications published in sources (journals or proceedings) included in the *Expanded WoS* database, denoted as *external* coverage indicator. Another is an *internal* coverage indicator, based on the extent to which articles included in the *Expanded WoS* database cite other articles included in this database.

As a first test, the bibliometric results for Netherlands research groups were compared with the ratings of peer evaluations of the field Computer Science conducted in 2003/early 2004 on behalf of the Quality Assurance Netherlands Universities (*QANU*).

Chapter 6 addresses the following research question: *what are the statistical relationships between the bibliometric outcomes for the various Netherlands Computer Science departments on the one hand, and peer ratings of these departments given by the QANU Review Committee for Computer Science in 2004 on the other?*

Chapter 7 presents the outcomes of a preliminary analysis of the publication output of Netherlands researchers included in the NL-CS database described in Section 2.3. The principal research question addressed is: *How does the citation impact of journal and conference papers by Netherlands computer scientists compare to international levels of citation impact of those types of papers in Computer Science?* This type of indicator has been calculated for many groups in many fields in the natural and life sciences. Its use enables one to make comparisons across disciplines, comparing, for instance, Netherlands academic Computer Science with Netherlands academic Chemistry or Physics.

In the last phase of the study presented in this report, a *verification round* was carried out among the leaders of the Netherlands academic groups in Computer Science included in the study. **Chapter 8** presents a summary of their reactions, especially those related to the coverage of the *Expanded WoS* database. This section provides a list of proceedings volumes that respondents qualified as important in the (sub-)field(s) but were *not* included in the *Expanded WoS* database. It also gives the replies of the authors of this report to the comments made in the verification round. Finally, **Chapter 9** presents a discussion and conclusions as regards the methodology developed in the study. In addition, it makes a number of recommendations that should be taken into account in the setup of a future bibliometric study of the research performance in Netherlands Computer Science.

2 Data Collection

2.1 The expansion of the WoS database with ACM, LNCS and IEEE conference proceedings

As mentioned in Chapter 1, *CWTS* has created a bibliometric database of all articles in journals processed by Thomson Scientific for the *Web of Science* (*WoS*). The *WoS* covers some 8,000 leading international journals from all domains of scholarship, including around 350 journals in the field Computer Science. The *CWTS-WoS* database currently covers the time period 1980–June 2006 and is regularly updated. The database now contains some 28 million source articles and 500 million cited references. It is a bibliometric rather than a bibliographic database, primarily aimed to enable statistical analysis of the data included.

The selection of proceedings volumes for inclusion in this *WoS* database constituted a complex and difficult issue. *Three* main approaches were applied. The *first* was based on an analysis of the conference proceedings sources that were actually used by Netherlands academic computer scientists. A database was created of publications they made during the time period 1996-2001. This database is further described in Section 2.3 below. An effort was made to standardise the names of the conferences they covered, and a frequency table was generated of (recurring) conference proceedings in which the Netherlands computer scientists had *published frequently*.

This list revealed a strong scattering of Netherlands publications among proceedings sources. Among proceedings from recurring conferences ranked on top, several related to local or national conferences. It was concluded that the proceedings most frequently used by Netherlands computer scientists were not necessarily the most important ones in the field. In addition, it was recognised that proceedings of important international conferences in a field were not necessarily used frequently as publication outlets by Netherlands computer scientists.

Therefore, a *second* approach was carried out, analysing reference lists in papers by Netherlands computer scientists that were published in sources covered by the Web of Science. The analysis aimed at creating a frequency table of the conference proceedings that were *most frequently cited* by Netherlands computer scientists. In order to achieve this, numerous conference names had to be standardised. This is a cumbersome task, as the information on the titles of the conference proceedings cited in *WoS* articles is limited and often incomplete. Actual capturing and standardization of citations to a particular conference proceedings is only feasible if one has a file with complete bibliographic meta data on all articles published in these proceedings. But since the *WoS* database was not yet expanded with additional proceedings volumes, such a file was not available.

In spite of the effort that was undertaken, the picture that emerged from this exercise showed little concentration of sources. Therefore, this particular analysis did not provide a sufficiently useful indication of which sources should be included. In addition, the only reference lists available for analysis were drawn from journal articles. This restriction to reference lists from journals carries a certain bias with it. In some subfields of Computer Science, communication networks using mainly conferences may exist, that are quite separated from those using journals. Such communication networks could have remained invisible in this approach.

Given the exploratory, methodological nature of the study, it was therefore decided to adopt a *third* approach, identifying existing electronic libraries containing peer-reviewed proceedings of important (recurring) conferences, and to expand the *WoS* database with source metadata and cited references from articles in proceedings of international conferences published by:

- Springer in its Lecture Notes in Computer Science (*LNCS*). It should be noted that a part of the *LNCS* volumes were already covered by the Web of Science. In the study described in this report articles from 460 missing *LNCS*-volumes published during 1996–2004 were added. The series Lecture Notes in Artificial Intelligence (LNAI) is a part of LNCS.
- Association for Computing Machinery (*ACM*). The proceedings of over 200 recurring conferences are made available as part of the *ACM* Digital Library.
- Computer Society of the Institute of Electrical and Electronics Engineers (*IEEE*). This Society granted us graciously permission to use data from over 400 recurring conferences that are accessible through the *Digital Library of the Computer Society of IEEE*, denoted as *IEEE/CS* or shortly as *IEEE* in this report.

The decision to include conference proceedings from the digital libraries of *ACM* and *IEEE*/CS and the *LNCS*-series was based on the following considerations.

- The analysis of conferences appearing most frequently in the publication and reference lists of Netherlands academic computer scientists shows that the proceedings of these conferences were often published by *ACM*, *IEEE*/CS, or as a volume of *LNCS*, although the absolute numbers were rather low.
- In view of this limited amount of concentration, it seems more appropriate to include a large collection of proceedings than to focus on the proceedings of a few particular conferences. The fact that the proceedings published by *ACM*, *IEEE*/CS, and in *LNCS*, are taken from a broad and large collection of conferences reduces the risk that certain important subfields of Computer Science are missed.
- It was also noted that many conference proceedings are not seriously refereed, or not refereed at all. In order to expand the database with papers that have a sufficiently high level of quality, it was considered appropriate to select proceedings that are seriously refereed. The *ACM*, *IEEE/CS* and *LNCS* proceedings generally meet this criterion.
- Several members of the project’s Advisory Committee who are experts in various subfields of Computer Science have indicated that the conference proceedings

published by *ACM* and *IEEE*/CS and those included in the *LNCS*-series indeed constitute important scientific communication channels for the research community in Computer Science, also for Netherlands researchers in this field.

- Finally, there are also pragmatic considerations. Both, Springer the publisher of *LNCS* and *IEEE/CS* have granted permission and offered their assistance with collecting the data we needed. *ACM* had already extracted the cited references from the proceedings articles they publish and makes them available to the general public. The fact that we only had three sources from which we received our data made the task of data processing somewhat easier.

Is must be emphasized that the expanded database is a *global* database. It contains articles published in the selected sources by authors from *all* countries, *not* merely by Netherlands researchers. The numbers and trends identified in this chapter are *global* and relate to the database as a whole, *not* to the collection of papers published by Netherlands authors. The next subsections present details on the processing of data from the three data libraries that were added to the Web of Science.

ACM

Data processing related to the conference proceedings contained in the *ACM Digital Library* was much simplified by the fact that complete reference lists were already extracted from the full text articles and could easily be downloaded from their publicly accessible web-pages. Other bibliographic information relating to the author names, titles and page numbers was also available in BibTeX format. *ACM* has used the extracted cited references to link them to other papers within their Digital Library. It is not entirely clear how reliable their data is, although *ACM* suggests that OCR errors may be found. These errors would also occur if one would extract reference lists from their PDF-files. For a small number of proceedings we found that no cited references had been extracted. Some of these proceedings were being published jointly with *IEEE*. For that reason we decided to use the data originating from *IEEE* for those proceedings that were available in both digital libraries. Out of a total of 982 volumes of 200 recurring conferences covered by the *ACM* Digital Library, 60 volumes did not include cited references.

LNCS

Figure 2.1 gives the evolution during 1996-2004 of the number of LNCS volumes published, the number processed by Thomson/ISI and included in the *WoS*, and the number added to the *Expanded WoS* database, respectively. This figure shows that the number of volumes in the *LNCS* series grows fast from 159 in 1996 to 420 in 2004. At the time of the data collection phase for this study not all volumes of the *Lecture Notes in Computer Science* were available in electronic form, as Springer was still processing back volumes. Especially, volumes from 1996 and 1997 were missing, while coverage from 1998 onwards was practically complete. We included in the *Expanded WoS* database only LNCS volumes that were electronically available. Therefore, for the years 1996 and 1997 the number of volumes added to our database is limited. For these years, there is a serious gap in our database.

As mentioned before, *Thomson Scientific* indexed some volumes of the *LNCS* series that were published during 1996 – 2002 and all volumes from 2003 onwards, although frequently with a substantial time lag. In the end, 460 volumes remained that were available online and not present in the *CWTS-WoS* database. The additional volumes were for the most part published in 2001 and 2002.

Figure 2.1: LNCS Volumes added to the Expanded WoS database

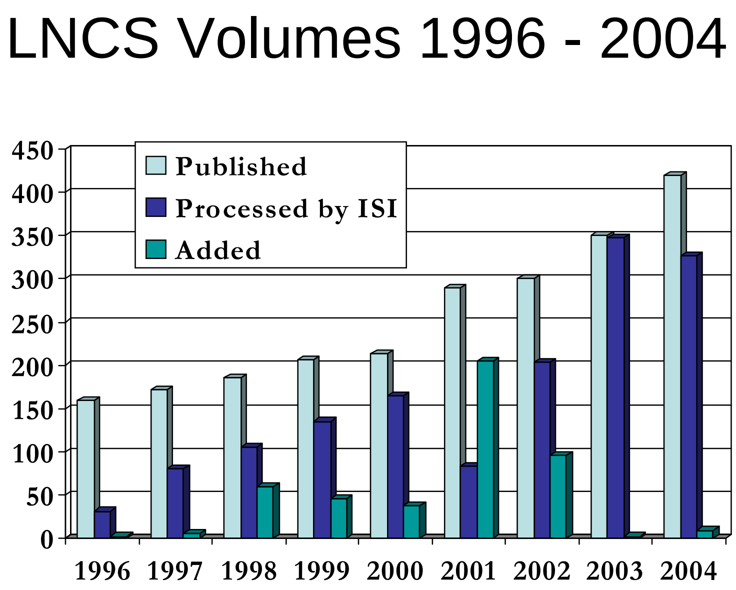

IEEE/CS

For this study we had the disposal of the complete *Digital Library of the Computer Society of the IEEE*, containing information about more than 100,000 items included in conference proceedings published by *IEEE*. Most bibliographic information was available in XML-files, except for the cited references, which had to be extracted from the attached PDF-files. The total number of conference proceedings from *IEEE* that were added to the *CWTS-WoS* database amounted to 1,379.

The extraction of cited references from PDF-files, i.e., from papers in LNCS volumes not included in the *WoS* and IEEE proceedings, proved to be a major, difficult task in this project. ***Table 2.1*** summarizes the main sources of error generated in this process, and their frequency of occurrence. As a first step the text contained in the PDF files had to be extracted. There are many software packages available to perform this task but none of them works perfectly. No text at all, scrambled text, distorted word order and OCR-errors were among most frequently encountered problems. Some packages could not properly handle text in two separate columns, as was the case for the PDF-files from *IEEE*. The most satisfactory results were obtained by using *pdftotext*, a tool that is part of the Xpdf library, an open source PDF viewer available under the GNU General Public Licence.

Although the performance of this software package was relatively good, loss of information caused by the conversion of PDF to text appeared inevitable.

Table 2.1: Estimated loss of information in extracting cited references from PDF-Documents (LNCS volumes not included in the WoS and IEEE proceedings only)

Source of error	***Nr. Documents***
Text was not or only partially extracted	5.5%
No reference list recognized	3.5%
No reference separator recognized	8%
Total	**17%**

Apart from a few files that were damaged or had been protected against text extraction, there were no errors reported by the software we used. Nevertheless, from 5.5 per cent of the PDF-documents text was not or only partially extracted, making them useless for further processing. In 9 per cent of the remaining PDF-documents no reference list was found. A sample of these were manually checked, and it was found that in most of them no cited references were mentioned, but that in almost 40 per cent of the cases the reference list was missed due to distorted or missing text, obtained from the conversion from PDF to text.

Typical reference separators such as numbers between brackets are particularly susceptible to errors related with the extraction of text from PDF-documents. These reference separators frequently cause a distortion of the word order and are also prone to OCR related errors. The software that was used to parse the resulting text files after conversion, could not handle cited references that were not separated from one another by a distinct reference separator. As a consequence, from 8 per cent of the PDF-documents not containing a recognizable separator between individual cited references, reference lists could not be processed. We estimate that around 17 per cent of all papers from the proceedings published by *IEEE* and in *LNCS* volumes not included in the *WoS* was affected by one of the three types of errors mentioned above. These papers are *not* lost as (potentially) cited articles, but as their reference list could not be extracted, they do not appear as citing articles. In addition, it must be noted that this problem particularly pertains to PDF files from conference proceedings published in the *earlier* years of the time period considered.

We also analysed the distribution across processed sources of the percentage of source papers for which no reference list was found in our data, and calculated the coefficient of variation. Its value was much lower for 2004 proceedings than it was for volumes published in 1996: 12.6 against 30.4. This suggests that the bias caused by this problem is decreasing over time. It must also be noted that in the citation analysis carried out in this study, particularly of Netherlands computer groups presented in Chapters 6 and 7, citing sources published in 2004 are far more important than those from 1996.

Figure 2.2: The origin of Computer Science papers included in the expanded *CWTS-WoS* database 1996 – 2004

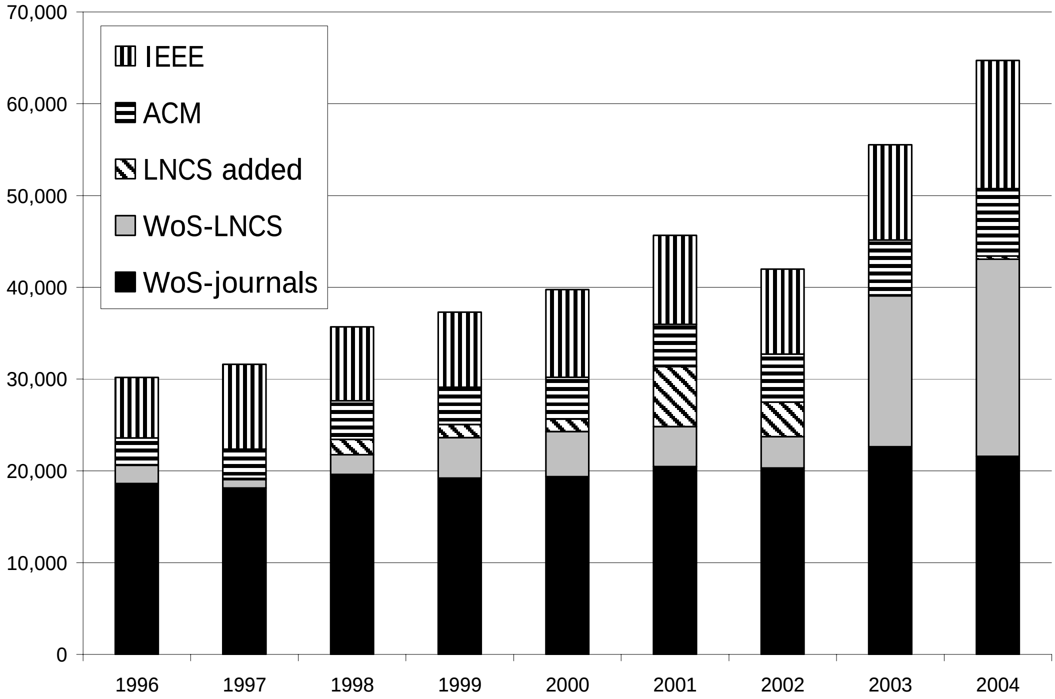

Figure 2.2 shows the origin of research papers in the field Computer Science included in the expanded *CWTS-WoS* database during 1996 – 2004. The solid black and grey areas jointly comprise the papers indexed by Thomson Scientific. The non-solid parts represent the conference papers that were specifically included for the current project. The size of the field more than doubled in these nine years. The growth came solely from the addition of conference papers, whereas the number of journal articles remains almost the same. Among the collections of conference proceedings, *LNCS* is by far the largest in 2004. Practically half of all Computer Science papers indexed by Thomson Scientific in 2004 was published in the Lecture Notes in Computer Science series.

All in all, about 160,000 proceedings articles were added to the database, jointly containing 1.6 million cited references. The number of articles in the expanded database covering the field Computer Science increased with 66 per cent compared to the number of papers in this field that were already included in the *WoS* database itself.

2.2 Citation matching

In a citation impact analysis one should distinguish between a *citing* or *source side* on the one hand, and a *cited* or *target side* on the other. Target articles are those that are subjected to a citation analysis. They are the (potentially) cited articles. Source articles are documents from which cited references are extracted. They (potentially) cite the target papers subjected to a citation analysis. The methodologies that can be applied in matching cited references to (intended) target articles strongly depend upon the amount of information that is available as regards the cited references themselves, – in other words, about the documents that are cited in the reference lists of source papers.

In the *WoS* database and related Citation Indexes produced by *Thomson Scientific*/*ISI*, the information on cited references is limited. *Thomson*/*ISI* extracts the name of the first author of the cited document, the year of its publication, its starting page number, its volume number (if applicable), and the name of the cited source. The length of a cited source string included in the database is at most 20 characters. On the other hand, *Thomson*/*ISI* puts an enormous effort in reformatting cited references from the source articles they process. The extracted information is normally sufficient to carry out accurate citation matching, particularly for target papers published in scientific/scholarly journals. With respect to citation matching in the *CWTS-WoS* database of cited references in *WoS* journals to target articles published in *WoS* journals, the reader is referred to Chapter 14 in Moed (2005). A series of important variations or discrepancies in cited author name, volume number, starting page, publication year and source title were taken into account.

In the conference proceedings sources added to the *WoS* database, the full cited reference string is available, including the title of the document, all authors, and a full description of the cited source. All this information can in principle be used to match cited references and targets. On the other hand, the cited reference data are not as properly structured and reformatted as those processed by *Thomson*/*ISI*. In many cases, conference proceedings volumes do not have volume numbers, and their titles appear in a very large number of variations, which makes it more difficult to identify them as targets in cited reference lists. In addition, their reference lists are generally less accurate than those found in *WoS* journals. Starting page numbers of target documents (be it journal articles or proceedings papers) are often omitted.

In the current study cited references contained in *the proceedings sources added to the WoS database* were matched to target documents on the basis of a match-key that includes parts of the name of the first author, the year of publication, and significant words from the title of a document. Several important variants or discrepancies in the author name and the publication year were taken into account. From the title of a (citing/source or cited/target) document the four longest words were selected. As regards matching cited references *in WoS journals* to the added proceedings volumes, a match-key was applied, containing parts of the name of the first author, the year of publication, and ‘significant’ words from the source (=proceedings volume) title. In this methodology, acronyms of conference titles play a key role.

In the current stage of the project, it is hardly possible to provide reliable estimates of the accuracy of the various types of citation matching processes. A first test suggests that the percentage of false–positives, i.e. the percentage of cited references erroneously matched to a target, is in the order of magnitude of several per cents. But it cannot be assumed that these mismatches are evenly distributed among targets.

An author – or a group of authors – may publish a piece of work in a series of publications in number of sources. For instance, it may be subsequently published as a Ph.D. thesis of the primary author, as a Technical Report, in one or more conference proceedings papers, and in one or more journal articles. Eventually, it may even be further expanded and published as a book. If the titles of all these publications are identical or very similar, it appears to be extremely difficult – if not impossible – to establish, in a fully automatic way, from a ‘raw’, unformatted reference list which version of the work is cited.

When the information provided in a cited reference is ambiguous or missing, it is simply impossible to identify which version is cited. But even if at least some information on the cited source is available, in bibliographic databases, publication lists of research groups, or scientists’ personal Curricula Vitae the source title may appear in so many variations, that a computerized approach based on an analysis of meta-data only (including cited references) may not in all cases be capable to establish the correct match between a cited reference an a collection of potential or ‘candidate’ targets.

A second reason why it is so difficult to match cited reference to the proper, ‘intended’ target is the following. The *Expanded WoS* database built in this study contains *ACM*, *LNCS* and *IEEE* proceedings volumes only. Publications in other journals or conference proceedings not covered by the *Expanded WoS* database are therefore not included. The same is true of other types of publications, such as monographs, book chapters, edited works and technical reports.

The set of potential targets of a particular cited reference contains only articles published in sources (journals or proceedings volumes) included in the Expand *WoS* database. For instance, if a particular author has published in a given year a monograph, this book is not in the database, and therefore, it is not included in a set of potential or ‘candidate’ targets. The database is ‘unaware’ of this publication, and therefore it cannot take into account in the citation matching process. It must be noted that in this study there is information on the publication output of Netherlands academic computer scientists only.

A first example of the difficulties in matching cited references to targets relates to a series of publications, all with the title “A scalable peer-to-peer lookup service for internet applications” by I. Stoica et al. (2001). It is assumed here that these publications represent different versions of one and the same ‘paper’. It is in the *Expanded WoS* database the paper with the highest number of citations from *ACM*, *LNCS* and *IEEE* conference proceedings. During the time period 2000–2004 it is cited 388 times from these proceedings volumes. It was published in at least the following versions:

- In Proceedings *ACM*/SIGCOMM '01 Conference, San Diego, August 2001;
- In the *ACM*/SIGCOMM Computer Communication Review, vol. 31, 2001;
- In *IEEE*/*ACM* Transactions on Networking, vol. 11, 2001 (according to University/ personal website);
- As Technical Report: TR–819, MIT.

The major part of the cited references indicates the Proceedings of *ACM*/SIGCOMM 01 as (cited) source. There are few references to the technical report version, while no references were found to the *ACM* Transactions on Networking. A number of references, particularly those in more recent papers cite the version in *ACM*/SIGCOMM Computer Communication Review. The citation matching methodology applied in this study assigned all cited references to this latter version. As a rule, if a paper is published both in a proceedings volume and in a journal, citations were assigned to the journal version.

A second example of the difficulties in matching cited references to targets relates to the cases where authors published a book, and published in the very same year a proceedings paper either with (almost) the same title, or with a title from which the principal words are included in the title of the conference itself (e.g., a book entitled 'Computational Geometry', presented at a 'Symposium on Computational Geometry'). In these cases, the citation matching methodology applied in this study may have assigned cited references intentionally given to the book to the proceedings paper. It should be noted again that books are not included in the *Expanded WoS* database.

The citation matching methodologies developed thus far in this study do not in all cases handle these problems in a proper way. The factor highlighted in the first example above has a negative effect upon the citation rates of proceedings articles relative to those of journal articles, whereas the factor outlined in the second example has a positive effect. On the other hand, it must be emphasized that, even if cited references are not matched to the intended target publication, they are nevertheless matched to the correct author or group of authors. In other words, the problem is how to distribute in a proper way cited references to an author's work among the various publication versions emerging from it.

It is much more difficult to estimate the percentage of true–negatives, i.e., the cited references that intentionally cite a particular target and therefore could or should be matched to that target, but that were not picked up in the matching process. In the analysis of cited references contained in added proceedings papers their share is expected to be lower than it is for cited references in *WoS* journal articles, since in the former several types of variations in publication title are taken into account, whereas in the latter the information on the cited source available in the database may be too limited. In a follow-up study, the citation matching process should be further developed and analysed, and more reliable estimates of its accuracy should be obtained.

We checked the 50 most frequently cited publications by source proceedings volumes (*ACM*, *LNCS* and *IEEE*), and the 50 most frequently cited publications by *WoS* source journals. In the first set we found 5 (journal or proceedings) papers to which the citation matching procedure assigned citations that were intentionally given to a book of the same first author, published in the same year and with an almost identical title. In the second set we found 3 cases in which authors who had published a paper in a particular proceedings volume, but who had also published a book in the same year with a title that showed a strong resemblance to the title of the conference. Excluding the books, the fraction of citations to other article versions than those identified in the citation matching process was in most cases very low (typically less than 5 per cent). In 4 cases it was around 20 per cent, and in three cases between 50 and 80 per cent.

We also checked the consequences of the cases mentioned in the previous paragraph for the outcomes of the bibliometric analyses presented in this report. It was found that, among the papers published by Netherlands academic Computer Scientists, there was one proceedings paper that had collected numerous citations to a book of the same authors and with the same title. This apparent error was corrected and does not affect the results presented in Chapters 5-7. It was also found that the ranking of journal and proceedings volumes presented in Chapter 4 was not affected by any of the cases.

2.3 Database with publications of Netherlands academic computer scientists

This study relates to all groups that were evaluated in the research evaluation of the field carried out by the Quality Assurance Netherlands Universities (*QANU*) in 2003/2004, supplemented with the groups subjected to a separate evaluation at Leiden University (UL) and Delft University of Technology (TUD). In addition, it includes all groups at the Centre for Mathematics and Computer Science (CWI). As a result the proposed study focuses on research activities in what can be termed as ‘Core’ academic Computer Science (‘kern–informatica’) at Netherlands universities.

The publication data of the *QANU* groups were obtained from the information provided by the groups within the framework of the *QANU*, and those of UL and TUD from the reviews carried out by these two universities themselves. The time period covered by these publications is 1996–2001. All universities involved formally agreed that these publication data were (re-)used in the current study. In this way, no additional collection of publication data was necessary.

According to the specifications in the QANU/VSNU protocol, publication lists of a group (or ‘programme’) included only articles published by researchers who were a member of that group at the date their articles were published (or at least at the date these were submitted for publication). As a consequence, the publication lists of a group do *not necessarily* contain the complete publication output of each group member made during the time period considered (1996-2001). For instance, if a researcher joined a group in

1999, his or her papers published prior to 1999 were *not* included in the publication list of that group.

CWI publications were obtained from the Annual Research Reports available at the CWI home page. Since the arrangement of research activities into groups and departments at CWI changed radically in 1997, it was decided to collect publications with respect to the time period 1997–2004.

A bibliometric database enables a user to conduct statistical analyses, and to link it – even on an article-by-article basis – with other, external literature databases. In the study described in this report a 'bibliometric' database was created of the complete publication output of all groups mentioned above. This database includes both articles published by these researchers in *WoS* covered journals, as well as their publications in non *WoS*-covered sources such as proceedings volumes, multi-authored books, and monographs. A bibliometric database of scientific literature reformats bibliographic data on scientific publications, splits such data into bibliometrically relevant data elements (particularly. each contributing author, the year of publication, name of source, and – whenever applicable – the source's volume number and starting page number), de-duplicates relevant data elements (e.g. journal titles), and identifies multiple occurrences of one and the same publication in the database.

The database with NL-CS publications was matched – on a paper by paper basis – to the database with articles from *WoS* journals and additional conference proceedings volumes. In addition, in the citation universe of the *Expanded CWTS-WoS* database, a citation analysis was carried out of all publications. The outcomes of these analyses are presented in Chapters 6 and 7.

3 Methodology

3.1 Target and citation universes

As outlined in Section 2.2, in a citation impact analysis one should distinguish a *citing* or *source side* on the one hand, and a *cited* or *target side* on the other. Target articles are those that are subjected to a citation analysis. They are the (potentially) cited articles. Source articles are documents from which cited references are extracted. They (potentially) cite the target papers subjected to a citation analysis.

In the analyses presented in Chapter 4 of the structure of the communication system, and of the importance of journals and conference proceedings papers therein, one can distinguish a number of sub-universes in terms of the types of documents included either as targets or as sources. They are schematically presented in ***Table 3.1***.

Both targets and sources were subdivided into three classes: journal articles, papers in *LNCS*, *ACM* and *IEEE* conference proceedings, and all other papers, e.g., articles in journals not processed for the *WoS*, articles in other conference proceedings than the three groups mentioned above, contributions to books, monographs, and Ph.D. theses. Technical reports were not included in the analysis, neither as targets, nor as sources. Table 3.1 illustrates that one may for instance study the impact that journal articles have made upon proceedings articles, and vice versa.

Table 3.1: Citation universes in analyses of importance of journals and proceedings presented in Chapter 4

Cited	*Citing*		
	Journals (*WoS* excl *LNCS*)	Proceedings (*ACM, LNCS, IEEE*)	Other
Journals (*WoS* excl *LNCS*)	Impact of journals upon journals	Impact of journals upon proceedings	Not available in this study
Proceedings (*ACM, LNCS, IEEE*)	Impact of proceedings upon journals	Impact of proceedings upon proceedings	
Other	Not available		

The types of analysis presented in later chapters of this report of the performance of Netherlands academic groups in Computer Science, are schematically presented in ***Table 3.2***. It also indicates the type of analysis that was carried out and the chapter in which the results are presented. The core analysis is labelled as the *Expanded WoS* analysis, in which both citing sources and cited target papers are included in the compound database of *WoS* articles and *ACM*, *LNCS* and *IEEE* proceedings papers. In this analysis a series of indicators is calculated, presented in Section 3.3 below, including an indicator comparing the citation rate of a group's articles to the world citation average in the fields in which it is active.

Table 3.2: Citation universes in analyses of NL papers presented in Chapter 5-7

Cited	*Citing*		
	WoS	*WoS+ACM+LNCS+ IEEE*	Other
WoS	*Analyses at the level of (anonymous) groups in Chapter 7*		Not available in this study
WoS+ACM+LNCS+ IEEE		*Analyses at the level of (anonymous) groups in Chapters 5, 6 and 7*	
Other (NL only)		*Analyses in Chapter 5 at the level of publication type and source*	

For methodological reasons, this report also presents indicators based on a 'pure' *WoS* analysis of the Netherlands groups included in this study, in order to examine differences and similarities between the outcomes of this analysis and those obtained in the *Expanded WoS* analysis. In this way one may assess the extent to which the expansion of the *WoS* database 'makes a difference'.

Finally, this report presents a citation analysis of all other articles and books published by Netherlands groups, i.e., all articles not included in the *Expanded WoS* analysis. Since in this analysis citation impact data are only available for Netherlands papers, the indicators are basically simple, absolute counts, and allow for rough comparisons among Netherlands groups only. No comparison can be made between Netherlands groups and foreign groups.

From the experiences collected in this study and described in Section 2.2 it is concluded in Section 9.3 that it is extremely difficult to determine in a fully automated way accurate citation counts to individual publications that are not published in journals or in well formatted and structured conference proceedings, or to publications appearing in different versions (e.g., as technical report, proceedings article, book chapter *and* as journal paper). It is therefore recommended that in a follow up study a citation analysis of these types of papers published by researchers under evaluation is carried out – or at least checked – *manually*.

The current study collected citations to these types of publications merely on the basis of computer algorithms. The outcomes were checked for a limited number of cases only. The outcomes were not considered to be sufficiently accurate to be presented at the level of *research groups* in this report. On the other hand, they can be used to analyse some more general patterns in the data. The main purpose is to give a rough estimate of the citation impact generated by the various types of target articles *not* included in the *Expanded WoS* database (e.g., proceedings papers, book chapters, monographs), and of the sources (journals or proceedings volumes) not covered by the *Expanded WoS* database, in which Netherlands researchers have published at least three papers.

3.2 WoS coverage and 4 types of bibliometric studies

A study recently conducted at *CWTS* analyzed reference behaviour of scientists/scholars in the various broad domains of scholarship (Moed, 2005). All source papers included in the *Thomson Scientific/ISI* Citation Indexes were categorized into 15 broad domains: Molecular Biology & Biochemistry, Biological Sciences related to Humans, Biological Sciences (primarily) related to Plants, Clinical Medicine, Physics, Chemistry, Engineering, Applied Physics & Chemistry, Mathematics, Geosciences, Psychology & Psychiatry, Economics, Social Sciences related to Medicine, Other Social Sciences, and Arts & Humanities.

Computer Science was not a separate domain in this study. Computer Science publications can be expected to spread among the 15 domains listed above, with high shares in Engineering, Mathematics and Physics. A basic assumption underlying this study holds that the documents cited in *WoS* source journals constitute the cognitive base upon which papers published in these journals are built, and that an analysis of cited references may provide insight into the structure of the scholarly communication system and the extent to which it is covered by *WoS* source journals.

A first attempt was made to operationalize the concept of adequacy of coverage of the *ISI* Citation Indexes in the following way. For all *WoS* source papers assigned to a broad domain the cited references were collected, i.e. the documents cited in the reference lists of these papers. A next step determined the percentage of cited references published in journals processed for the *WoS*. This percentage is conceived as an internal adequacy of coverage indicator. If the *WoS* would completely cover a broad domain, one would expect that all cited references would be published in journals processed for these Indexes. In this case, the adequacy of coverage indicator would amount to 100 per cent.

Table 3.3: Internal coverage percentages of the Thomson Scientific/ISI Citation Indexes

Internal Coverage Percentage			
80-100%	***60-80%***	***40-60%***	***<40%***
Biochem & Mol Biol	Appl Phys & Chem	Mathematics	Other Soc Sci
Biol Sci – Humans	Biol Sci – Anim & Plants	Economics	Humanities & Arts
Chemistry	Psychol & Psychiat	Engineering	
Clin Medicine	Geosciences		
Phys & Astron	Soc Sci ~ Medicine		

The outcomes are presented in ***Table 3.3***. It shows that that in the main field Engineering the role of journals is less prominent than it is in basic science fields such as Physics and Chemistry, and that the *WoS* source journals cover the former field less adequately than the latter two. For instance, in the broad domain Molecular Biology & Biochemistry the internal *WoS* coverage percentage exceeds 90 per cent, for Physics and Chemistry it is between 80 and 85 per cent, whereas for Engineering it was found to be near 50 per cent.

A tentative classification was built of four types of bibliometric studies in function of the adequacy of *WoS* coverage of the field of inquiry. It is presented in ***Table 3.4***. Essentially, this classification represents a central hypothesis underlying past and particularly future research and applied work in the field of research performance assessment at *CWTS*. This research includes further operationalizing the concept of adequacy of *WoS* coverage in the various scientific or scholarly research fields. In Table 3.4, the qualifications of *WoS* coverage are purely qualitative and need to be further quantified in future research. The current study aims at providing a substantial contribution in this direction.

Table 3.4: Four types of bibliometric studies

Type	*Label*	*Target / Cited*	*Source / Citing*	*WoS coverage*	*Internal coverage % (tentative)*
1	'Pure' *WoS*	*WoS*	*WoS*	Excellent	80-100%
2	Target expanded citation analysis	*WoS*+non *WoS*	*WoS*	Good	60-80%
3	Source expanded citation analysis	*WoS*+non *WoS*	*WoS*+non *WoS*	Moderate	40-60%
4	No citation analysis at all			Poor	<40%

In order to further explain this classification, some technical terms should be specified. In a citation impact analysis one should distinguish a citing or source side and a cited or target side. Target articles are those that are subjected to a citation analysis. Source articles are documents from which cited references are extracted. The total collection of cited references in source articles constitutes the universe within which citation counting takes place. Articles not included in this universe cannot contribute to citation counts of a target article, even if they actually cite that target, as the citation is not recorded in the database. On the other hand, target articles not published in source journals may be cited in other papers from the citing or source universe.

According to the tentative classification, in fields with an excellent *WoS* coverage, for which the internal coverage percentage is typically between 80 and 100, it is generally sufficient in a citation impact analysis to take into account as target articles only those that are published in *WoS* source journals, and to use the total collection of cited references in *WoS* source journals as citation universe. This type of analysis is labelled above as the 'Pure' *WoS* analysis.

If *WoS* coverage in a field is not excellent, but can nevertheless be qualified as good, with internal coverage percentages typically between 60 and 80, the scheme suggests to expand the collection of target articles analyzed in the 'pure' *WoS* analysis by including target articles that are not published in *WoS* source journals (a ***target expanded*** citation analysis, Type 2 in Table 3.4). In a recently finished pilot study on research activities in Computer Science at the Flemish Free University of Brussels (VUB) this approach was explored. In this approach it is assumed that, although the collection of source (citing) documents is incomplete, their cited references still may provide reliable citation impact estimates, to the extent that the *WoS* source articles constitute a representative sample of a wider population of citing sources. For instance, if in a field a limited number of important journals or proceedings volumes is not covered by *WoS*, whereas *WoS* coverage can still be qualified as good, it can be expected that papers published in those not-covered media are on average sufficiently frequently cited from *WoS* source journals in order to be able to assess their impact.

If *WoS* coverage of a field is moderate, with internal coverage percentages typically between 40 and 60, it becomes questionable whether such an assumption of representativeness of cited references in *WoS* source journals is still valid. For instance, if proceedings of annual international conferences play a dominant role in the communication system in a field, it can be expected that there is heavy citation traffic among proceedings of annual conferences from subsequent years. If such proceedings are not covered by the *WoS*, these citation relationships remain invisible, as citations are merely extracted from *WoS* source journals. In that case, it seems appropriate to expand the universe of citing sources with articles in proceedings volumes from a range of subsequent years. Such an approach is labelled as a ***source expanded*** citation analysis (Type 3 in Table 3.4).

Finally, if *WoS* coverage in a field is poor, showing internal coverage percentages below 40, it is questionable whether it is useful conducting a citation analysis based on *WoS* data, even if target or source universes are expanded. This is particularly true in fields that are fragmented into schools of thought and hampered by national or linguistic barriers. It is to be expected that in such fields alternative approaches, not based on citation data, are more fruitful than citation impact analyses.

The study presented in this report aims at a further development of the central hypothesis presented in schematic form in Table 3.4 above, and particularly to the development of a ***source expanded*** citation analysis, by expanding the *WoS* publication and citation universe with articles and their cited references published in a number of conference proceedings.

The web application created in the verification round (see Chapter 8) only included citations to articles published in journals or proceedings included in the Expanded WoS database. It thus presented the raw counts of a *Type I* citation analysis, but *within a citation universe of the expanded WoS database*. The analyses presented in Chapters 6 and 7 of this report are also based on this type of citation analysis.

As outlined above this report also presents outcomes of a citation analysis of all other articles and books published by Netherlands groups, i.e., all articles not included in the *Expanded WoS* analysis. This type of citation analysis can be denoted as a *Type 2* analysis, but carried out within the citation universe of the *Expanded WoS* database. The outcomes were *not* considered sufficiently accurate to be presented at the level of *research groups* in this report. They are merely used to analyse some more general patterns in the data, and provide rough estimates of the citation impact generated by the various *types* of target articles *not* included in the *Expanded WoS* database (e.g., proceedings papers, book chapters, monographs), and of the sources (journals or proceedings volumes) not covered by the Expanded WoS database in which Netherlands researchers have published frequently.

3.3 Bibliometric indicators

The primary level of aggregation applied in the study is that of a *research group*. It is assumed that the research group represents the natural unit of scientific activity in Computer Science, and therefore constitutes the appropriate unit of analysis. Although individual scientists play an intermediary role in the data collection process and the names of leading professors may be used as a label to indicate a research group ("the group headed by professor P"), the proposed study does not aim at producing rankings of individuals on the basis of their personal bibliometric scores. It should also be noted that Ph.D. supervisors are not always co-author of the papers published by their doctoral students. The groups included in this study, and the bibliometric data collected for them were further specified in Section 2.3 of this report.

Table 3.5 below presents a list of indicators calculated in the analysis of Netherlands academic Computer Science articles. These indicators are calculated within a particular publication and citation universe. As explained above, in the case of the *Expanded WoS* analysis, this universe consists of all *WoS* journal articles *plus* the papers in *ACM*, *LNCS* and *IEEE* conference proceedings. In the case of the *'Pure' WoS* analysis the universe coincides with the *WoS* database. In Table 3.5, the universe in which the analysis is carried out is labelled simply as the 'global publication database'. Journals or proceedings volumes will be denoted as 'sources'.

The most important bibliometric indicator calculated in this study is a so called *normalised citation impact indicator*. This indicator, denoted as ***CPP/FCSm*** (see Table 3.5), compares the average impact of a research group's articles to the world citation average in the subfields in which the group is active, by calculating the ratio of these two parameters. A ratio of 1.0 indicates that the average impact of a group's articles is equal to the world average, and a ratio above 1.0 that the group's citation impact is above world average. The normalisation takes into account the number of papers a group has published, the papers' age distribution, the type of paper, and citation characteristics of the subfield.

Differences in citation characteristics among research fields constitute an important 'disturbing' factor in citation analysis, and hamper the interpretation of its outcomes. Simple citation indicators based on absolute number of citations are seriously affected by such differences. This is not only true for indicators of the citation impact of research groups, but also for *journal impact factors*. For instance, in the fields Mathematics and Computer Science a journal impact factor of 1.0 is relatively high (compared to other journals in those fields), whereas in Physics it represents an average value, and in Biochemistry a relatively low value. Therefore one should not directly compare journals from different fields merely on the basis of the value of their journal impact factors as published by Thomson Scientific/ISI in its Journal Citation Reports.

Skewness is a crucial property of the distribution of citations amongst articles published by a research group. Normally a limited number of a group's papers is relatively highly cited, whereas the major share of its papers receives a low number of citations or no citations at all. It is because of this skewness that it is so important to collect complete lists of articles published by a group. If one misses a highly cited article, or if one erroneously assigns a highly cited article to a group, the outcomes of a citation analysis may be highly inaccurate, and the interpretation invalid.

The normalised citation impact indicator is based on average values of citation distributions that tend to be skewed. An alternative approach is to focus on the 'top' of these distributions, by identifying in a scientific field all *highly cited articles* – for instance, the 1, 5 or 10 per cent most frequently cited papers – and determining the number of a group's papers in this global set of 'top' articles in terms of citation impact. This study calculates indicators based on the number of papers published by a group that are among the *10 per cent* most frequently cited articles in a field.

Table 3.5 uses the term 'research unit'. This is a purely neutral term, indicating any aggregate of researchers, for instance, all members of a research group, all researchers in a university, or all scientists located in a particular country. The word unit refers to 'unit of analysis' rather than 'unit of measurement'.

Table 3.5: List of bibliometric indicators calculated in the study

Symbol	***Definition***	***What it measures/why it is useful***
Standard bibliometric indicators(Chapters 6 and 7)		
P	The number of articles (normal articles, letters, notes and reviews) published by a research unit in sources (journals, proceedings) covered by the global publication database	Article or publication output: reflects both the *size* of a unit (number of active researchers) and the *publication productivity* of its members (papers per researcher)
C	The number of citations recorded in the global publication database given to all articles subjected to the citation analysis.. Author self-citations are excluded.	'Raw' total citation impact (excl/incl author self citations): reflects both a unit's number of publications (P), their age (older papers may collect more citations than younger ones), type (e.g., reviews attract more citations than normal articles) and subfield citation characteristics (e.g., biochemistry shows higher citation levels than mathematics)
C+sc	The number of citations recorded in the global publication database to all articles subjected to an citation analysis, including author self-citations.	
CPP	The average number of citations per publication. Author self-citations are not included.	'Raw' citation impact per paper. Corrects for differences in number of publications among units, but is affected by the other factors indicated above
CPP/FCSm	The impact of a research unit's articles, compared to the world citation average in the subfields in which the research unit is active. A '+' ('-') symbol directly after the numerical value indicates that the impact of the research unit's articles is significantly above (below) world (subfield) average. Author self-citations are not included.	This is the 'crown' indicator of citation impact. It takes into account the number of papers a unit has published, the papers' age distribution, type of paper, and citation characteristics of the subfield. A value above one indicates that the average citation impact of a unit's papers is above world average.
CPP/JCSm	The impact of a research unit's articles, compared to the average citation rate of the research unit's set of sources in which it published. A '+' ('-') symbol immediately after the numerical value indicates that the impact of the research unit's articles is significantly above (below) the average citation rate of the source set. Author self-citations are not included	Indicates citation impact compared to journal/source average. It is especially useful in multi- or inter-disciplinary fields that cannot be defined in terms of a set of journals (journal categories)
JCSm/FCSm	The impact of the sources (e.g. journals, proceedings volumes) in which a research unit has published (the research unit's source set), compared to the world citation average in the subfields covered by these sources	Calculated for journals this indicator can be conceived as a *normalised* journal impact factor.
Pnc	The percentage of articles not cited during the time period considered, excluding author self-citations.	This indicator relates to the *bottom* of the distribution of citations amongst cited articles published by a unit, and is complementary to indicators based on the *mean* of that distribution.
% Selfcit	The percentage of author self-citations. An author self-citation is defined as a citation in which the citing and the cited paper have at least one author in common (either a first author or a secondary author).	Author self citations are normally excluded from the counts, since the impact analysis aims at measuring the citation impact upon research activities *outside* the evaluated unit itself

Symbol	*Definition*	*What it measures/why it is useful*
Indicators of 'top' or highly cited publications (Chapter 7)		
P top 10%	The actual, absolute number of papers published by a unit that are among the 10 % most frequently cited of similar (in terms of type, age and subfield) papers during the time period considered; citations are counted during the first 4 years after publication date (a 4-year citation window).	Focuses on the *top* of the world distribution of received citations amongst cited articles in a subfield, and is complementary to indicators based on the *mean* of that distribution, especially the crown indicator CPP/FCSm
E (P top 10%)	The expected number of papers amongst the top 10 %, based on the number of papers published by the research unit in the time period considered.	This is 10 % of that portion of a unit's publication output that can be followed during at least 4 years (the citation window)
A/E (P top 10%)	Indicates the relative contribution of a unit to the upper percentiles of the citation distribution in the time period considered (=P top 10% / E (P top 10%))	A ratio above one indicates that the number of a unit's papers among the world top 10 % is higher than expected
Coverage indicators (Chapter 5)		
External coverage	The percentage of articles published in sources included in the global publication database (in this study the Expanded *WoS* database) relative to the total number of publications made by the research unit.	It can be denoted as publication coverage indicator. A value of 100 % means: all papers are published in sources (journals, proceedings) covered by the global database
Internal coverage	Gives for all articles published in sources included in the global publication dataset, the percentage of their cited references that are published in sources in the global database	It can be denoted as cited reference coverage indicator. A value of 100 % means: all cited articles are themselves published in sources covered by the global database
C/Ctot	Gives the percentage of citations to articles published in sources included in the Expanded *WoS* database (C), relative to the total number of citations to all types of publications (Ctot)	It can be denoted as a citation coverage indicator.
Special indicators (Chapters 4 and 6)		
FTE	Full time Equivalents research time, not counting Ph.D. students	These data were only available for groups included in the VSNU/QANU review of NL Computer Science
P/FTE	Number of published articles (P) per FTE research time	Denoted as publication productivity
C/FTE	Number of citations (C) to published articles, per FTE research time	Denoted as citation productivity
SCR	Source Citation Rate, the average number of citations received during the first three years after publication date by articles published in a source (journal or proceedings volume)	This measure is similar to the journal impact factor. Calculated for a journal, it is about 2 times the 'official' Thomson/ISI impact factor.

4 Citation Patterns in Computer Science

4.1 The role of conference proceedings and journals

A first analysis is based on all articles in the field Computer Science and published during the time period 1996–2004. The field Computer Science comprises all journals assigned by *Thomson Scientific* to journal categories related to Computer Science, and all *ACM*, *LNCS* and *IEEE* conference proceedings volumes added in this study to the *CWTS-WoS* database. The total number of articles amounts to 377,371. Of these, 174,870 (46.3 per cent) were published in journals, and 202,501 (53.7 per cent) in proceedings.

In a next step, all citations to these articles were identified, made during the same time period 1996–2004 in the same set of Computer Science articles. In this way a sub-universe was created of Computer Science articles in the *Expanded WoS* database and their citation links. Both cited and citing sources were categorized into journals and proceedings sources. Citation patterns were analysed across source types (journals versus proceedings), for instance, the frequency at which proceedings papers cite journal articles and vice versa. The analysis is based on the time period 1996-2004, because only for this time period relevant data are available. However, there is no reason to assume that outcomes based on longer time periods would be significantly different from those presented below.

Table 4.1 focuses on the '*citing side*'. It presents per citing source type the average number of references per article to journal and proceedings papers, respectively. Table 4.1 shows that a journal paper published during 1996–2004 contains on average 1.46 references to other journal papers published during the same time period, and 0.79 references to proceedings articles. Expressed in percentages, 64.9 per cent of references in journals are to journals and 31.1 per cent to conference proceedings.

For references in proceedings articles the situation is reversed: 36.7 per cent of references are to journal articles, and 63.2 per cent to proceedings papers. It should be noted that, in the sub-universe, proceedings articles contain on average more references than journal articles do: 2.86 versus 2.25 references per article (see row 'Total').

But one should keep in mind that these outcomes relate to a sub-universe of articles in Computer Science journals and proceedings volumes published during 1996–2004. References given to articles published before 1996 are not included in this universe. It should also be noted that the values presented for citations and references are averages over a range of available years. It can be shown that journal articles contain on average a higher number of cited references than proceedings papers do. On the other hand, cited references in the latter tend to be published more recently than those in the former (Visser and Moed, 2005).

Table 4.1: References per article per type of source

Cited		Citing			
Source type	***Publ 96-04***	***Journals***		***Proceedings***	
		Refs/Article 96-04	*%*	*Refs/Article 96-04*	*%*
Journals	174,870	1.46	64.9 %	1.05	36.7 %
Proceedings	202,501	0.79	31.1 %	1.81	63.2 %
Total	377,371	2.25	100.0 %	2.86	100.0 %

Table 4.2 presents the '*cited side*' of the citation links within the sub-universe. This table shows that a journal article receives on average 2.68 citations (column 'Total'). The average citation impact of proceedings papers is 2.49, which is only 7 per cent lower than that for journal articles. Proceedings articles receive the major part of their impact (72.7 per cent) from other proceedings articles in the sub-universe, and only 27.3 per cent from journals. The impact of journal articles comes for 54.5% from other journals, and 45,5 percent form conference proceedings papers. In other words, the impact of journal articles upon proceedings papers is higher than that of proceedings papers upon journal articles.

Table 4.2: Citations per article per type of source

Cited		***Citing***					
Source type	***Publ 96-04***	***Journals***		***Proceedings***		***Total***	
		Cites/Article 96-04	*%*	*Cites/Article 96-04*	*%*	*Cites/Article 96-04*	*%*
Journals	174,870	1.46	54.5%	1.22	45.5%	2.68	100%
Proceedings	202,501	0.68	27.3%	1.81	72.7%	2.49	100%

Table 4.3 presents more details on the citation impact of conference proceedings and journals. In the *entire Expanded WoS* database – and *not* merely in the Computer Science sub-universe analysed in Tables 4.1 and 4.2 above – the citation impact of proceedings and journals was calculated, expressed in a measure that is similar to the journal impact factors calculated by *Thomson Scientific*/*ISI*.

The *journal impact factor* for a journal J in year T is defined as the ratio of the number of citations in year T to articles published in J during years T-1 and T-2, and the number of articles (more precise: 'citable documents') published in J in years T-1 and T-2. As a rule, journals publish papers each year. But a conference proceedings volume is a unique source that is published only once, even though a series of (bi-)annual conferences may produce a series of conference proceedings volumes, one for each conference.

Therefore it was decided to calculate an impact factor-like indicator based upon one single 'cited' year, denoted as *Source Citation Rate* (***SCR***). For the year T and for a source S (proceedings volume or journal) the ***SCR*** is defined as the number of citations received during the years T, T+1 and T+2 by articles published in S in year T, (i.e., received during the first three years after publication date, the year of publication included), divided by the number of articles published in S in year T. As a rule of thumb, for journals the SRC is about two times the journal impact factor calculated by *Thomson Scientific*/*ISI*.

Table 4.3 presents key parameters of the distribution of the number of articles published in a year, as well as the *Source Citation Rate*s, among annual journal volumes and proceedings volumes. Data relate to proceedings volumes or annual journal volumes published during 1996–2002.

Table 4.3: Distribution of three bibliometric measures among sources*

Source type	***N***	***Mean***	***P10***	***P25***	***P50***	***P75***	***P90***	***P95***
Number of articles								
Journals	2,303	57.0	14	23	39	69	123	167
Proceedings	2,783	45.4	13	21	31	51	88	117
Source Citation Rate								
Journals	2,303	1.92	0.36	0.74	1.34	2.46	4.22	5.51
Proceedings	2,783	1.98	0.31	0.67	1.35	2.47	4.15	5.94
% Uncited articles								
Journals	2,303	45.7	16.7	28.6	44.2	60.4	76.9	85.7
Proceedings	2,783	46.5	15.1	28.6	45.5	64.2	80.0	88.9

* P10, P25 ,P50, P75, P90, P95: The 10th, 25th, 50th (=median), 75th, 90th and 95th percentile of the distribution

Table 4.3 shows that annual journal volumes contain on average more articles than proceedings volumes do. The mean number of articles is 57.0 for the former and 45.5 for the latter, while the median values are 39 versus 31. The 75th percentile of the distribution for journals is 69 and for proceedings 51. This means that 25 per cent of journals with the highest number of articles have at least 69 in a year, whereas 25 per cent ‘thickest’ proceedings volumes contain at least 51 papers.

The citation impact of the proceedings volumes is similar to that of annual journal volumes. The mean *Source Citation Rate*s for journals and proceedings volumes are 1.92 and 1.98, respectively, and the median values 1.34 and 1.35. The 10th and 25th percentile of the SRC distribution for proceedings are slightly lower than those for journals, whereas the 95th percentile is slightly higher. The percentage of articles in a volume that are uncited during the first three years after publication date, is for journals on average 45.7 per cent and for proceedings volumes 46.5. The 10th percentile is for proceedings slightly lower than that for journals, and the 75th, 90th and 95th percentile somewhat higher.

In order to further compare journal and proceedings volumes, all volumes were ranked by increasing value of a variable, and their percentile ranks were established. For a number of percentile ranges, the share of journal and proceedings volumes was calculated. The outcomes are presented in ***Table 4.4***. Table 4.4 shows that in the set of 5,086 volumes analysed, 45.3 per cent are annual journal volumes, against 54,7 per cent proceedings volumes. Focusing on the number of articles published (second and third column in Table

4.4), this table shows that in the class of volumes with percentile ranks between 95 and 100, i.e., the top 5 per cent, 64.3 per cent are journals and 35.7 per cent conference proceedings volumes. It follows that in this class, journals are overrepresented: among the top 5 per cent of volumes with the highest number of papers, there are relatively more journals than there are in the total population.

Focusing on the *Source Citation Rate* (***SCR***), Table 4.4 reveals that both among the top 5 per cent and bottom 5 per cent of volumes (either journals or proceedings) in terms of *Source Citation Rate*, proceedings volumes are slightly overrepresented. In both segments the share of proceedings volumes amounts to 60 per cent, which is higher than the overall percentage of 54.7 obtained for proceedings volume in the total set. Considering the percentage of uncited articles during the time window considered, one observes the same phenomenon: proceedings volumes are overrepresented in the top and the bottom of the distribution, a finding that is consistent with that for the ***SCR*** indicator.

Table 4.4: Shares of journal and proceedings volumes per range of percentile ranks

Percentile	***No Articles***		***Source Citation Rate***		***% Uncited articles***	
	% Jrnls	***% Procs***	***% Jrnls***	***% Procs***	***% Jrnls***	***% Procs***
All	45.3	54.7	45.3	54.7	45.3	54.7
0–5 (=Bottom 5%)	46.9	53.1	39.8	60.2	34.6	65.4
5–10	36.2	63.8	44.9	55.1	49.6	50.4
10–25	39.8	60.2	42.5	57.5	46.3	53.7
25–50	36.9	63.1	48.7	51.3	48.0	52.0
50–75	48.3	51.7	44.8	55.2	47.2	52.8
75–90	51.6	48.4	44.6	55.4	42.2	57.8
90–95	57.9	42.1	51.2	48.8	40.6	59.4
95–100 (=Top 5%)	64.3	35.7	41.2	58.8	39.6	60.4

4.2 Rankings of publication sources

Table 4.5 presents a list of the 50 proceedings and annual journal volumes with the highest *Source Citation Rate*, calculated within the *Expanded WoS* publication and citation universe (including *ACM*, *LNCS* and *IEEE* proceedings). Table 4.5 shows that among the top 50 sources in terms of citation impact, 70 per cent are conference proceedings volumes, and 30 per cent annual journal volumes.

Table 4.5: Top 50 volumes with the highest Source Citation Rate

Rank	***Source type***	***Seg-ment***	***Source***	***Year***	***Arti cles***	***Cites***	***Source Citation Rate***	***% Uncited articles***
1	JRNL	*WoS*	*ACM* Computing Surveys	2002	15	489	32.6	0.0
2	PROC	*ACM*	SOSP '01: P 18th *ACM* Symp on Operating Systems Principles	2001	17	404	23.8	5.9
3	PROC	*ACM*	SIGGRAPH '01: P 28th Annual Conf on Computer Graphics and Interactive Techniques	2001	65	1234	19.0	1.5
4	PROC	*ACM*	SIGGRAPH '00: P 27th Annual Conf on Computer Graphics and Interactive Techniques	2000	59	1098	18.6	1.7
5	PROC	*ACM*	ISCA '00: P 27th Annual Int Symp on Computer Architecture	2000	28	506	18.1	0.0
6	PROC	*ACM*	WSNA '02: P 1st *ACM* Int Worksh on Wireless Sensor Networks and Applications	2002	15	263	17.5	0.0
7	PROC	*ACM*	ISCA '97: P 24th Annual Int Symp on Computer Architecture	1997	30	518	17.3	0.0
8	PROC	*LNCS*	Designing Privacy Enhancing Technologies, Int Worksh on Design Issues In Anonymity and Unobservability, Berkeley, CA, USA, July 25-26, 2000	2001	11	189	17.2	18.2
9	PROC	*ACM*	SIGCOMM '99: P Conf on Applications Technologies Architectures and Protocols For Computer Communication	1999	24	402	16.8	4.2
10	PROC	*IEEE*	29th Annual Int Symp on Computer Architecture (ISCA'02)	2002	27	445	16.5	3.7
11	PROC	*ACM*	SIGGRAPH '98: P 25th Annual Conf on Computer Graphics and Interactive Techniques	1998	45	722	16.0	2.2
12	PROC	*ACM*	POPL '02: P 29th *ACM* Sigplan-Sigact Symp on Principles of Programming Languages	2002	30	467	15.6	6.7
13	PROC	*IEEE*	28th Annual Int Symp on Computer Architecture (ISCA'01)	2001	24	367	15.3	0.0
14	PROC	*ACM*	MOBICOM '02: P 8th Annual Int Conf on Mobile Computing and Networking	2002	26	392	15.1	7.7
15	PROC	*ACM*	SIGCOMM '01: P 2001 Conf on Applications Technologies Architectures and Protocols for Computer Communications	2001	23	343	14.9	4.3
16	PROC	*ACM*	PODS '02: P Twenty-First *ACM* Sigmod-Sigact-Sigart Symp on Principles of Database Systems	2002	27	395	14.6	3.7
17	JRNL	*WoS*	Journal of Computational Biology	2002	50	719	14.4	12.0
18	JRNL	*WoS*	VLDB Journal	2002	15	214	14.3	13.3
19	PROC	*IEEE*	13th *IEEE* Visualization 2002 (VIS'02)	2002	75	1059	14.1	6.7
20	JRNL	*WoS*	Bioinformatics	2002	304	4276	14.1	6.9

Rank	***Source type***	***Seg-ment***	***Source***	***Year***	***Arti cles***	***Cites***	***Source Citation Rate***	***% Uncited articles***
21	JRNL	*WoS*	User Modeling and User-Adapted Interaction	2001	19	267	14.1	0.0
22	JRNL	*WoS*	*ACM* Transactions on Graphics	2002	86	1193	13.9	3.5
23	PROC	*ACM*	MOBICOM '01: P 7th Annual Int Conf on Mobile Computing and Networking	2001	29	399	13.8	0.0
24	PROC	*ACM*	MOBIHOC '02: P 3rd *ACM* Int Symp on Mobile Ad Hoc Networking & Computing	2002	21	287	13.7	0.0
25	PROC	*ACM*	MOBICOM '00: P 6th Annual Int Conf on Mobile Computing and Networking	2000	28	378	13.5	7.1
26	JRNL	*WoS*	*ACM* Computing Surveys	2001	16	214	13.4	6.3
27	JRNL	*WoS*	Journal of Molecular Graphics	1996	28	360	12.9	17.9
28	PROC	*ACM*	SOSP '97: P Sixteenth *ACM* Symp on Operating Systems Principles	1997	23	295	12.8	0.0
29	JRNL	*WoS*	Data Mining and Knowledge Discovery	1998	24	307	12.8	8.3
30	PROC	*ACM*	SIGGRAPH '99: P 26th Annual Conf on Computer Graphics and Interactive Techniques	1999	51	650	12.7	0.0
31	JRNL	*WoS*	Human-Computer Interaction	2002	23	288	12.5	8.7
32	JRNL	*WoS*	Bioinformatics	2001	213	2655	12.5	7.0
33	PROC	*IEEE*	26th Annual Int Symp on Computer Architecture (ISCA'99)	1999	26	320	12.3	7.7
34	JRNL	*WoS*	*ACM* Computing Surveys	1998	10	123	12.3	20.0
35	PROC	*ACM*	SIGCOMM '97: P *ACM* Sigcomm '97 Conf on Applications Technologies Architectures and Protocols for Comp. Communicat.	1997	24	288	12.0	4.2
36	PROC	*ACM*	POPL '98: P 25th *ACM* Sigplan-Sigact Symp on Principles of Programming Languages	1998	32	383	12.0	3.1
37	PROC	*ACM*	POPL '99: P 26th *ACM* Sigplan-Sigact Symp on Principles of Programming Languages	1999	25	299	12.0	4.0
38	JRNL	*WoS*	MIS Quarterly	2002	37	441	11.9	5.4
39	PROC	*LNCS/ WoS*	*LNCS*-Vol 1241	1998	22	262	11.9	4.5
40	PROC	*ACM*	JAVA '99: P *ACM* 1999 Conf on Java Grande	1999	20	233	11.7	0.0
41	PROC	*IEEE*	Second *IEEE* Worksh on Mobile Computer Systems and Applicat.	1999	12	139	11.6	8.3
42	PROC	*ACM*	ISCA '96: P 23rd Annual Int Symp on Computer Architecture	1996	28	324	11.6	3.6
43	PROC	*ACM*	PLDI '00: P *ACM* Sigplan 2000 Conf on Programming Language Design and Implementation	2000	30	337	11.2	3.3
44	PROC	*ACM*	SIGCOMM '98: P *ACM* Sigcomm '98 Conf on Applications Technologies Architectures and Protocols for Comp. Communicat	1998	26	292	11.2	0.0

Rank	***Source type***	***Seg-ment***	***Source***	***Year***	***Arti cles***	***Cites***	***Source Citation Rate***	***% Uncited articles***
45	JRNL	*WoS*	Journal of Molecular Graphics & Modelling	1997	35	388	11.1	28.6
46	JRNL	*WoS*	Computer Communication Review	2001	81	897	11.1	23.5
47	PROC	*IEEE*	25th Annual Int Symp on Computer Architecture (ISCA'98)	1998	33	365	11.1	3.0
48	PROC	*ACM*	SIGCOMM '02: P 2002 Conf on Applications Technologies Architectures and Protocols for Computer Communications	2002	25	272	10.9	4.0
49	PROC	*ACM*	MOBICOM '99: P 5th Annual *ACM*/*IEEE* Int Conf on Mobile Computing and Networking	1999	28	303	10.8	7.1
50	PROC	*LNCS*/ *WoS*	*LNCS*-Vol 1294	2001	37	395	10.7	2.7

A next analysis focuses on the citation links among sources. For each citing and cited volume the number of articles published in that volume was determined, as well as the number of citations given in the citing volume to the cited volume. Next, a simple measure of strength of the citation link among volumes was calculated. If ***Pcd*** and ***Pcg*** represent the number of articles in the citing and cited volume, respectively, ***CITES*** the number of citations from the citing to the cited volume, and SQRT the arithmetic Square Root Function, the variable ***STRENGTH*** is defined as follows:

STRENGTH = ***CITES*** / SQRT(***Pcg*** * ***Pcd***).

Table 4.6 presents for cited and citing volumes with at least 20 articles the 25 links with the highest value of the strength as defined above. It shows that there are relatively many links between conference proceedings volumes of (bi-)annual conferences. Apparently, there is a heavy citation traffic from one such volume to another. The citation strength is comparable to that among different annual volumes within the same journal.

Table 4.6: Citing and cited volumes with the strongest citation links

Cited Volume			*Citing Volume*			*Citation Links*	
Ty-pe	***Year***	***Source***	***Type***	***Year***	***Source***	***Ci-tes***	***Strength***
P	2000	SIGGRAPH '00: P 27th Annual Conference on Computer Graphics and Interactive Techniques	P	2002	SIGGRAPH '02: P 29th Annual Conference on Computer Graphics and Interactive Techniques	102	1.62
P	2001	SIGCOMM '01: P 2001 Conference on Applications Technologies Architectures and Protocols for Computer Communications	P	2004	*LNCS*-Vol 2429	42	1.60
P	2001	SIGGRAPH '01: P 28th Annual Conf on Computer Graphics and Interactive Techniques	P	2002	SIGGRAPH '02: P 29th Annual Conference on Computer Graphics and Interactive Techniques	105	1.59
P	2000	SIGGRAPH '00: P 27th Annual Conf on Computer Graphics and Interactive Techniques	J	2003	*ACM* Transactions on Graphics	121	1.57
P	2000	SIGGRAPH '00: P 27th Annual Conf on Computer Graphics and Interactive Techniques	J	2002	*ACM* Transactions on Graphics	108	1.52
P	2001	SIGGRAPH '01: P 28th Annual Conf on Computer Graphics and Interactive Techniques	J	2002	*ACM* Transactions on Graphics	106	1.42
P	2000	SIGGRAPH '00: P 27th Annual Conf on Computer Graphics and Interactive Techniques	P	2001	SIGGRAPH '01: P 28th Annual Conference on Computer Graphics and Interactive Techniques	85	1.37
J	2003	*ACM* Transactions on Graphics	J	2004	*ACM* Transactions on Graphics	135	1.32
P	1999	Java '99: P *ACM* 1999 Conf on Java Grande	P	2000	Java '00: P *ACM* 2000 Conference on Java Grande	26	1.30
P	2001	SIGGRAPH '01: P 28th Annual Conf on Computer Graphics and Interactive Techniques	J	2003	*ACM* Transactions on Graphics	99	1.22
P	2002	SIGCOMM '02: P 2002 Conf on Applications Technologies Architectures and Protocols for Computer Communications	J	2004	Computer Communication Review	64	1.21
P	1998	Spatial Cognition, An Interdisciplinary Approach To Representing and Process. SpatiaKnowledge	P	2000	*LNCS*-Vol 1849	30	1.21

Cited Volume			*Citing Volume*			*Citation Links*	
Ty-pe	***Year***	***Source***	***Type***	***Year***	***Source***	***Ci-tes***	***Strength***
P	1998	SIGGRAPH '98: P 25th Annual Conf on Computer Graphics and Interactive Techniques	P	2000	SIGGRAPH '00: P 27th Annual Conference on Computer Graphics and Interactive Techniques	62	1.20
J	2002	*ACM* Transactions on Graphics	P	2004	Sca '04: P 2004 *ACM* SIGGRAPH/Eurographics Symp on Computer Animation	67	1.19
P	1996	SIGGRAPH '96: P 23rd Annual Conf on Computer Graphics and Interactive Techniques	P	2000	SIGGRAPH '00: P 27th Annual Conference on Computer Graphics and Interactive Techniques	76	1.18
P	1998	SIGCOMM '98: P *ACM* SIGCOMM '98 Conf on Applications Technologies Architectures and Protocols for Computer Communication	P	1999	SIGCOMM '99: P Conference on Applications Technologies Architectures and Protocols for Computer Communication	29	1.16
J	2002	*ACM* Transactions on Graphics	P	2003	Sca '03: P 2003 *ACM* SIGGRAPH/Eurographics Symp on Computer Animation	64	1.15
P	1996	SIGGRAPH '96: P 23rd Annual Conf on Computer Graphics and Interactive Techniques	P	1998	SIGGRAPH '98: P 25th Annual Conference on Computer Graphics and Interactive Techniques	64	1.14
P	2003	*LNCS*-Vol 2079	P	2003	*LNCS*-Vol 2740	24	1.12
J	1999	MIS Quarterly	J	2004	MIS Quarterly	27	1.10
P	2001	Cross-Language Information Retrieval and Evaluation, Worksh of Cross-Language Evaluation forum, CLEF 2000	J	2004	Information Retrieval	27	1.10
J	1996	*IEEE* Transactions on Information Theory	J	1998	*IEEE* Transactions on Information Theory	172	1.09
P	2003	*LNCS*-Vol 2406	J	2004	Information Retrieval	31	1.08
P	2002	29th Annual International Symp on Computer Architecture (ISCA'02)	P	2003	36th Annual *IEEE*/*ACM* Int Symp on Microarchitecture (MICRO-36)	34	1.08

5 Coverage of the Expanded WoS Database

This chapter addresses the question as to which extent the expansion of the database with conference papers from *IEEE*, *ACM* and *LNCS* proceedings volumes improved coverage rates of Netherlands academic research output in Computer Science described in Section 2.3. A first analysis presented in Section 5.1 relates to the conference proceedings covered by the Expanded *WoS* database in which Netherlands academic computer scientists published at least three articles during the time period 1999-2001. The extent to which the *Expanded WoS* database covers the Netherlands output is further analysed from three different angles, presented in Sections 5.2 to 5.4. First we consider in Section 5.2 what share of *Netherlands publications* is included in the expanded *CWTS-WoS* database. Next, in Section 5.3 we analyse the *reference lists* in papers by Netherlands academic computer scientists, and give an impression of the importance of papers included in the *Expanded WoS* database for their ‘knowledge base’. Finally, in Section 5.4 we compare the *citation impact* of Netherlands papers included in the *Expanded WoS* database against the impact of Netherlands papers *not* included in this database.

5.1 Conference proceedings covered by the Expanded WoS database

As outlined in Section 2.3, the data collection of Netherlands academic Computer Science for this study is based on the publication details for the time period 1996 – 2001 reported in the self evaluation reports that were prepared for the *QANU* research evaluation conducted in 2003 and for the separate evaluation of the Computer Science Departments of Leiden University and Delft University of Technology. The publication details for CWI were collected from their Annual Research Reports.

In order to obtain an impression of the proceedings that are covered by the *Expanded WoS* database, ***Table 5.1*** gives per subfield a list of the proceedings in which Netherlands computer scientists have published at least three papers during the time period 1999-2001. For each proceedings the number of articles by Netherlands computer scientists is given as well as its share in the total number of all conference papers authored by Netherlands computer scientists that are included in the WoS Expanded database. This table also presents for each source a simple indicator of citation impact of the Netherlands papers published in it. This indicator is similar to the Source Citation Rate presented in Chapter 4. It gives the average number of citations the Netherlands papers in a source received during the first three years after publication date. A second impact indicator presented in Table 5.1 gives the Source Citation Rate for all papers published by researchers from all over the world.

Table 5.1: Proceedings *included* in the Expanded WoS database in which Netherlands computer scientists published at least three papers during 1999-2001

Subfield/proceedings	*Nr. Papers*	*% Papers*	*Average Citation Impact of NL papers*	*Average Citation Impact of all papers*
Computing and Imaging				
High-Performance Computing and Networking, International Conference and Exhibition	13	5.6%	0.38	0.33
International Conference on Pattern Recognition (ICPR)	9	3.8%	0.11	0.22
IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)	8	3.4%	0.50	1.96
SCG: Proceedings Of The Annual Symposium On Computational Geometry	8	3.4%	2.38	2.24
Euro-Par: Parallel Processing, International Euro-Par Conference	8	3.4%	0.00	0.41
International Workshop on Database and Expert Systems Applications	7	3.0%	0.43	0.36
Annual IEEE Conference on Computational Complexity (CoCo)	5	2.1%	1.20	1.24
Medical Image Computing and Computer-Assisted Intervention - MICCAI	5	2.1%	0.80	1.52
JGI: Proceedings Of The Joint ACM-Iscope Conference On Java Grande	4	1.7%	5.25	4.29
PPOPP: Proceedings Of The ACM Sigplan Symposium On Principles And Practices Of Parallel Programming	4	1.7%	6.25	2.79
STOC: Proceedings Of The Annual ACM Symposium On Theory Of Computing	4	1.7%	5.75	4.33
Automata, Languages and Programming, International Colloquium, ICALP	4	1.7%	1.00	1.79
Recent Advances in Parallel Virtual Machine and Message Passing Interface, European PVM/MPI Users' Group Meeting	4	1.7%	0.50	0.67
Visual Information and Information Systems, International Conference, VISUAL	4	1.7%	1.50	0.77
IEEE International Symposium on High Performance Distributed Computing (HPDC)	3	1.3%	0.00	2.43
IEEE International Conference on Multimedia and Expo (ICME)	3	1.3%	0.33	0.49
IEEE Symposium on Information Visualization	3	1.3%	2.33	1.30
MULTIMEDIA: Proceedings Of The ACM International Conference On Multimedia	3	1.3%	0.67	1.07
SMA: Proceedings Of The ACM Symposium On Solid Modeling And Applications	3	1.3%	0.00	1.27
VIS: Proceedings Of The Conference On Visualization	3	1.3%	2.00	2.85
Advances in Databases, British National Conference on Databases, BNCOD	3	1.3%	0.33	0.54
Principles of Data Mining and Knowledge Discovery, European Conference, PKDD	3	1.3%	2.33	0.72
Protocols for Multimedia Systems, International Conference, PROMS	3	1.3%	0.33	0.12
Scale Space Methods in Computer Vision, International Conference	3	1.3%	1.33	2.33
SOFSEM: Theory and Practice of Informatics, Conference on Current Trends in Theory and Practice of Informatics	3	1.3%	0.00	0.39
STACS, Annual Symp.on Theoretical Aspects of Computer Science	3	1.3%	1.67	1.50

Subfield/proceedings	*Nr. Papers*	*% Papers*	*Average Citation Impact of NL papers*	*Average Citation Impact of all papers*
Programming Research and Algorithmics				
Tools and Algorithms for the Construction and Analysis of Systems, International Conference, TACAS	9	3.2%	4.56	2.28
CONCUR - Concurrency Theory, International Conference	8	2.8%	2.38	2.22
Implementation of Functional Languages, International Workshop, IFL	6	2.1%	0.83	0.81
Model Checking Software, International SPIN Workshop	6	2.1%	1.67	3.21
Theorem Proving in Higher Order Logics, International Conference, TPHOLs	6	2.1%	0.67	1.16
SAC: Proceedings Of The ACM Symposium On Applied Computing	5	1.8%	0.60	0.49
Compiler Construction, International Conference, CC	5	1.8%	2.20	1.47
Object-Oriented Technology, ECOOP	5	1.8%	1.40	0.92
Euro-Par: Parallel Processing, International Euro-Par Conference	5	1.8%	0.00	0.41
Fundamental Approaches to Software Engineering, International Conference, FASE	5	1.8%	3.00	1.51
Automata, Languages and Programming, International Colloquium, ICALP	5	1.8%	1.20	1.79
Mathematical Foundations of Computer Science, International Symposium, MFCS	5	1.8%	0.20	0.82
Rewriting Techniques and Applications, International Conference, RTA	5	1.8%	2.60	1.25
IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)	4	1.4%	0.75	1.96
Computer Aided Verification, International Conference, CAV	4	1.4%	2.75	3.76
Perspectives of System Informatics, International Andrei Ershov Memorial Conference, PSI	4	1.4%	0.25	0.18
FME: Formal Methods for Increasing Software Productivity, International Symposium of Formal Methods Europe	4	1.4%	1.75	1.37
Foundations of Software Science and Computation Structures, International Conference, FOSSACS	4	1.4%	0.75	3.01
Formal Techniques in Real-Time and Fault-Tolerant Systems, International Symposium, FTRTFT	4	1.4%	1.75	0.91
Advances in Intelligent Data Analysis, International Conference, IDA	4	1.4%	0.00	0.37
IEEE International Conference on Automated Software Engineering (ASE)	3	1.1%	0.00	0.89
CHI: Proceedings Of The Sigchi Conference On Human Factors In Computing Systems	3	1.1%	4.33	1.72
Conference on Software Maintenance and Reengineering	3	1.1%	3.33	1.23
Annual IEEE International Conference and Workshop on the Engineering of Computer Based Systems (ECBS)	3	1.1%	0.00	0.31
International Conference on Software Engineering (ICSE)	3	1.1%	5.00	2.08
IEEE International Conference on Software Maintenance (ICSM)	3	1.1%	2.33	1.26
MULTIMEDIA: Proceedings Of The ACM International Conference On Multimedia	3	1.1%	0.67	1.07
VIS: Proceedings Of The Conference On Visualization	3	1.1%	1.67	2.85
Advanced Functional Programming, International School	3	1.1%	4.00	2.29
Formal Methods for Real-Time and Probabilistic Systems, International AMAST Workshop, ARTS	3	1.1%	3.00	1.05

Subfield/proceedings	*Nr. Papers*	*% Papers*	*Average Citation Impact of NL papers*	*Average Citation Impact of all papers*
Programming Research and Algorithmics (continued)				
Process Algebra and Probabilistic Methods, Performance Modeling and Verification: Joint International Workshop, PAPM-PROBMIV	3	1.1%	2.00	1.17
STACS, Annual Symposium on Theoretical Aspects of Computer Science	3	1.1%	3.33	1.50
Theory and Application of Graph Transformations, International Workshop, TAGT	3	1.1%	0.33	0.70
Information and Knowledge Systems				
Intelligent Agents. Agent Theories Architectures and Languages, International Workshop, ATAL	11	4.6%	3.09	2.48
International Workshop on Database and Expert Systems Applications	9	3.7%	0.33	0.36
Engineering of Intelligent Systems, International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, IEA/AIE	8	3.3%	0.25	0.13
Knowledge Acquisition, Modeling and Management, European Workshop, EKAW	7	2.9%	0.29	0.40
Formal Models of Agents, ESPRIT Project Model Age Workshop	7	2.9%	1.43	1.00
Text, Speech and Dialogue, International Conference, TSD	5	2.1%	0.00	0.13
HYPERTEXT: Proceedings Of The ACM On Hypertext And Hypermedia	4	1.7%	3.75	1.90
Issues in Agent Communication	4	1.7%	2.00	2.62
Advances in Databases, British National Conference on Databases, BNCOD	4	1.7%	0.25	0.54
Computers and Games, International Conference, CG	4	1.7%	0.75	0.51
Cooperative Information Agents, International Workshop, CIA	4	1.7%	1.75	1.67
Computer Science Logic, Annual Conference of the EACSL	4	1.7%	2.50	0.92
Genetic Programming, European Workshop	4	1.7%	0.75	1.20
CHI: Proceedings Of The Sigchi Conference On Human Factors In Computing Systems	3	1.2%	4.33	1.72
MULTIMEDIA: Proceedings Of The ACM International Conference On Multimedia	3	1.2%	1.00	1.07
SAC: Proceedings Of The ACM Symposium On Applied Computing	3	1.2%	1.33	0.49
International Symposium on Temporal Representation and Reasoning (TIME)	3	1.2%	0.67	0.45
Computational Logic - CL, International Conference	3	1.2%	0.33	0.47
Compositionality: The Significant Difference, International Symposium, COMPOS	3	1.2%	0.67	1.56
CONCUR - Concurrency Theory, International Conference	3	1.2%	1.00	2.22
Cooperative Information Systems, International Conference, CoopIS	3	1.2%	5.67	1.15
Principles of Data Mining and Knowledge Discovery, European Conference, PKDD	3	1.2%	2.33	0.72

Subfield/proceedings	*Nr. Papers*	*% Papers*	*Average Citation Impact of NL papers*	*Average Citation Impact of all papers*
Logic				
Annual IEEE Conference on Computational Complexity (CoCo)	8	10.1%	1.50	1.24
Knowledge Acquisition, Modeling and Management, European Workshop, EKAW	5	6.3%	0.40	0.40
Engineering of Intelligent Systems, International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, IEA/AIE	5	6.3%	0.40	0.13
STACS, Annual Symposium on Theoretical Aspects of Computer Science	5	6.3%	1.20	1.50
Intelligent Agents. Agent Theories Architectures and Languages, International Workshop, ATAL	3	3.8%	2.33	2.48
Cooperative Information Agents, International Workshop, CIA	3	3.8%	1.67	1.67
Computer Science Logic, Annual Conference of the EACSL	3	3.8%	3.33	0.92
Automata, Languages and Programming, International Colloquium, ICALP	3	3.8%	1.33	1.79
Formal Models of Agents, ESPRIT Project ModelAge Workshop	3	3.8%	3.00	1.00
Mathematics				
Principles and Practice of Constraint Programming - CP, International Conference	4	11.4%	2.75	1.57
High-Performance Computing and Networking, International Conference and Exhibition	4	11.4%	0.00	0.33
SAC: Proceedings Of The ACM Symposium On Applied Computing	3	8.6%	2.33	0.49
Computational Logic - CL, International Conference	3	8.6%	0.67	0.47

In order to analyse *differences among subfields* of Computer Science, a subfield classification system was applied that was inspired by, and partly based upon, the system of research schools. Groups that were *not* a member of any research school were assigned to one or more schools on the basis of a comparison of the titles of their research programmes with the descriptions of the research topics covered by the various schools. This assignment process also took into account the sources (journals, proceedings volumes) in which a group published its papers. In this way all groups were allocated to one or more subfields. The names of the subfields are derived from those of the related research schools. If a group was allocated to more than one subfield, all its papers were assigned to each subfield. The subfield *Telematics* consists of one group only. Since this report presents bibliometric indicators of individual groups only in an anonymous or aggregated manner, it was decided to delete this subfield from most tables below.

The authors of this report do not claim to have invented a valid, generally accepted classification of subfields in global computer science. Although there are sophisticated methodologies to develop such a classification, its development falls outside the scope of the current study. Instead, the authors have categorized Netherlands groups in a simple and pragmatic way, following the structure of the Research Schools in this field. Even if this structure has purely national features that cannot be found in other countries, it is useful in an attempt to show differences in publication practices (and corresponding coverage rates) among collections of Netherlands groups with different cognitive orientations.

Table 5.1 gives per subfield an overview of the conference proceedings covered by the *Expanded WoS* database in which Netherlands academic computer scientists published most frequently during 1999-2001. The next section focuses on the conference proceedings that are *not* covered by this database.

5.2 Coverage of Netherlands academic Computer Science papers by the Expanded CWTS-WoS database

The publication lists of Netherlands computer scientists relate to a wide variety of documents such as theses, journal articles, conference papers, monographs, book chapters, reports, and edited books. However, not all universities reported on all of these publication types. In order to produce results that can be compared with one another, we limited the total publication output in this section to those publication types that were available for *each* organisation, i.e. Ph.D. theses, conference papers, journal articles, books and chapters. Rather than applying a weighing scheme, we counted each publication as one and calculated the percentage included in the *CWTS-WoS* database before and after its expansion with the additional conference proceedings from *ACM*, *LNCS* and *IEEE*.

Table 5.2 presents the outcomes of a coverage analysis per year, and for the total time period 1996-2001. Before the inclusion of additional conference proceedings, one out of four papers (25%) was included in the *CWTS-WoS* database. The expansion of the database raised this ratio to slightly more than one out of three papers. The coverage rate of 35% is an aggregate figure for the total time period 1996–2001. As the number of additional conference papers with which the database was expanded increased sharply over time, the coverage rate increased as well, from 23% in 1996 to 41% in 2001. On the one hand this increase is at least partly due to the fact that within the three collections used for expansion, recent conference proceedings are more likely to be available online than older volumes. On the other hand, the higher coverage rate should also at least partly be explained by the increasing number of proceedings included in the three collections, particularly in the LNCS series.

Table 5.2: Expanded WoS coverage by year

Year	*Tot P96-01*	*in WoS*	*% in WoS*	*in Wos+*	*% in WoS+*	*+/-*
1996	832	133	16%	188	23%	7%
1997	1,012	218	22%	295	29%	8%
1998	1,096	298	27%	414	38%	11%
1999	1,111	323	29%	431	39%	10%
2000	1,193	327	27%	444	37%	10%
2001	***1,420***	***373***	***26%***	***578***	***41%***	***14%***
1996-2001	6,664	1,672	25%	2,350	35%	10%

In fact, the coverage percentages in Table 5.2 relate to the time period 1996-2001, whereas Figure 2.1 in Section 2.1 revealed that the major part of the papers from ACM, LNCS and IEEE sources added to the *Expanded WoS* database are from later years. Therefore, the coverage rates of Netherlands Computer Science papers can be expected to further increase during de time period 2002-2006. Since in this study no data were available on the publication output of Dutch computer scientists during that time period, the precise coverage rates after 2001 could not be determined.

Table 5.3 reveals differences in *Expanded WoS* coverage among subfields. *Computing and Imaging* shows the highest coverage by the *Expanded WoS* database. Its share of papers published during the entire period 1996-2001 and covered by the *Expanded WoS* database was found to be 40 per cent. It increased over the years and reached a maximum value of 53 per cent in 2001. For *Programming Research and Algorithmics* this coverage percentage for the year 2001 amounts to 41, and for *Information and Knowledge Systems*, – the largest subfield in terms of numbers of papers covered – to 32 per cent. The smaller subfields *Logic* and *Telematics* show in 2001 coverage rates of 36 and 24 per cent, respectively. Finally, *Mathematics*, covering the more mathematically oriented groups, shows the highest coverage rates. For this field the *WoS* itself provides already a coverage of 53% in 2001, while the expansion of this database with ACM and IEEE proceedings and LNCS volumes not covered by the *WoS* hardly raised the coverage.

To gain a better understanding of the coverage rates we further analysed the sources that were used for publication by Netherlands Computer Science groups, but that were *not* included in the CWTS/*WoS* database after expansion. In particular, our aim was to identify the conferences that appear most frequently in the publication lists of Netherlands Computer Scientists. As we are interested in identifying those conference series that are no part of the three large collections we used for expansion, we limited the analysis to the most recent three years *1999–2001*. Of the 2,212 *conference papers* published *during this time period* by the Netherlands research groups, 1,513 (68 per cent) were not included in the expanded CWTS/*WoS* database. The conference series names that relate to these proceedings were roughly standardized, in order to generate the overview of the most frequently appearing conference series.

Table 5.3: Expanded WoS coverage per subfield

Year	Tot P96-01	in WoS	% in WoS	in Wos+	% in WoS+	+/-
Computing and Imaging						
1996	233	31	13%	52	22%	9%
1997	279	52	19%	86	31%	12%
1998	323	81	25%	126	39%	14%
1999	317	84	26%	132	42%	16%
2000	307	84	27%	133	43%	16%
2001	361	117	32%	191	53%	19%
1996-2001	1,817	449	25%	720	40%	15%
Programming Research and Algorithmics						
1996	359	63	18%	95	26%	8%
1997	360	77	21%	104	29%	8%
1998	374	108	29%	158	42%	13%
1999	359	105	29%	144	40%	11%
2000	409	113	28%	158	39%	11%
2001	452	95	21%	187	41%	20%
1996-2001	2,305	561	24%	846	37%	13%
Information and Knowledge Systems						
1996	271	40	15%	48	18%	3%
1997	347	52	15%	75	22%	7%
1998	368	76	21%	115	31%	10%
1999	417	103	25%	128	31%	6%
2000	465	87	19%	121	26%	7%
2001	619	118	19%	196	32%	13%
1996-2001	2,474	476	19%	683	28%	9%
Logic						
1996	104	20	19%	26	25%	6%
1997	73	22	30%	26	36%	6%
1998	109	35	32%	50	46%	14%
1999	115	47	41%	56	49%	8%
2000	133	33	25%	45	34%	9%
2001	200	53	27%	72	36%	9%
1996-2001	733	210	29%	275	38%	9%
Telematics						
1996	25	4	16%	4	16%	0%
1997	29	4	14%	7	24%	10%
1998	30	3	10%	4	13%	3%
1999	22	3	14%	6	27%	13%
2000	29	2	7%	3	10%	3%
2001	33	5	15%	8	24%	9%
1996-2001	168	21	13%	32	19%	6%
Mathematics						
1996	1	0	0%	0	0%	.
1997	110	52	47%	55	50%	3%
1998	134	60	45%	67	50%	5%
1999	120	60	50%	67	56%	6%
2000	137	73	53%	80	58%	5%
2001	142	75	53%	77	54%	1%
1996-2001	638	320	50%	346	54%	4%

Table 5.4 shows per subfield the conference names in the proceedings in which Netherlands computer scientists published at least three papers during 1999-2001. In the standardization process we may have missed some name variants. Therefore, numbers of papers published by Netherlands computer scientists are indicative. In addition, this table gives for each source a simple indicator of citation impact of the Netherlands papers published in it. It gives the average number of citations the Netherlands papers in a source received during the first three years after publication date.

Three points need to be emphasized. *Firstly*, the indicator relates to the impact of papers by *Netherlands computer scientists only*. For articles published by researchers from other countries no citation data were available. *Secondly*, the lists of the proceedings are arranged by subfield. Groups from different subfields may publish in the same proceedings. In that case, the average citation rate of Netherlands papers in a proceedings volume may show differences among subfields. *Thirdly*, the impact indicator relates to citations made in papers *included in the Expanded WoS database*. Since the proceedings sources listed in Table 5.4 are *not* covered by the *Expanded WoS* database, the citations given in these sources to any Netherlands paper are *not* included in the counts. It is most difficult to give an accurate estimate of the number of citations from one volume to another, and of the effect these citations have upon the average citation impact of the Netherlands papers published in these volumes. A preliminary analysis suggests that for volumes of recurring conferences included in the *Expanded WoS* database, and for annual volumes of journals in that database, somewhere between 10 and 30 per cent of received citations are from the same conference or journal, respectively.

Per subfield proceedings with three or more papers account for 20 to 30 per cent of the total number of proceedings papers not covered by the *Expanded WoS* database and published by groups allocated to a subfield. In each subfield the tail of the distribution of published papers among proceedings is rather long. This outcome reveals the scattering of proceedings papers published by Netherlands academic Computer Science groups among conference proceedings.

As far as concentration can be observed, several conferences tend to be *local* rather than *international*. For instance, in the subfield *Computing and Imaging* 7.7% of the papers in proceedings not covered by the *Expanded WoS* database were published in *Proceedings of the Conference of the Advanced School for Computing and Imaging (ASCI)*. In *Information and Knowledge Systems* 6.6 % of the proceedings papers not covered by our expanded database were published in the *Proceedings of the Belgium-Netherlands Conference on Artificial Intelligence (BNAIC)*. But in all subfields the lists also contain proceedings of international, recurring conferences.

Table 5.4: Proceedings *not included* in the Expanded WoS database in which Netherlands computer scientists published at least 3 papers during 1999-2001

Subfield/proceedings	***Nr. Papers***	***% Papers***	***Average Citation Impact of NL papers***
Computing and Imaging			
Conference Of The Advanced School For Computing And Imaging (ASCI)	28	7.4%	0.18
Eurographics incl. Various Workshops	11	2.9%	0.73
Joint Eurographics IEEE TCVG Symposium On Visualization	11	2.9%	0.64
IEEE International Conference On Robotics And Automation (ICRA)	10	2.6%	1.20
Euromedia	6	1.6%	0.00
Belgium-Netherlands Conference On Artificial Intelligence (BNAIC)	5	1.3%	0.60
International Joint Conference On Artificial Intelligence (IJCAI)	5	1.3%	2.60
International Symposium On Mathematical Morphology (ISMM)	5	1.3%	0.60
IEEE Instrumentation And Measurement Technology Conference (IMTC)	4	1.1%	0.25
International Conference On Very Large Databases (VLDB)	4	1.1%	5.00
British Machine Vision Conference (BMVC)	3	0.8%	0.33
Not Unified	*288*	*75.8%*	*0.44*
Total	*380*	*100.0%*	*0.53*
Programming Research and Algorithmics			
Belgium-Netherlands Conference On Artificial Intelligence (BNAIC)	22	4.6%	0.05
Genetic And Evolutionary Computation Conference (GECCO)	13	2.7%	0.54
Procs World Multiconference On Systemics, Cybernetics And Informatics (SCI)	7	1.5%	0.00
International Joint Conference On Artificial Intelligence (IJCAI)	6	1.3%	0.33
Eurographics Incl Various Workshops	5	1.1%	0.40
Text Retrieval Conference (TREC)	5	1.1%	4.00
Workshop On Language Descriptions, Tools And Applications (LDTA)	5	1.1%	0.40
CMCS	4	0.8%	0.25
European Conference On Software Maintenance And Reengineering (CSMR)	4	0.8%	0.75
European Control Conference (ECC)	4	0.8%	0.25
European Symposium On Algorithm (ESA)	4	0.8%	0.00
Conference on Parallel Problem Solving from Nature (PPSN)	3	0.6%	0.67
Content Based Multimedia Indexing (CBMI)	3	0.6%	0.33
Dutch-Belgian Conference On Machine Learning (BENELEARN)	3	0.6%	0.00
EUROMICRO Conference	3	0.6%	0.00
Int. Conf. On Augmented, Virtual Environments And 3D Imaging (ICAV3D)	3	0.6%	0.33
Philips Workshop On Scheduling And Resource Management (SCHARM)	3	0.6%	0.00
Proceedings Of Recherche D'information Assistee Par Ordinateur (RIAO)	3	0.6%	0.67
Not Unified	*374*	*78.9%*	*0.34*
Total	*474*	*100.0%*	*0.36*

Subfield/proceedings	*Nr. Papers*	*% Papers*	*Average Citation Impact of NL papers*
Information and Knowledge Systems			
Belgium-Netherlands Conference On Artificial Intelligence (BNAIC)	49	6.5%	0.02
Genetic And Evolutionary Computation Conference (GECCO)	17	2.3%	0.88
Euromedia	13	1.7%	0.00
International Joint Conference On Artificial Intelligence (IJCAI)	12	1.6%	1.58
Text Retrieval Conference (TREC)	11	1.5%	2.18
International Conference On Enterprise Information Systems (ICEIS)	10	1.3%	0.40
Conference On Uncertainty In Artificial Intelligence (UAI)	9	1.2%	1.11
Summer Computer Simulation Conference (SCSC)	9	1.2%	0.11
Dutch-Belgian Conference On Machine Learning (BENELEARN)	8	1.1%	0.00
European Conference On Modelling And Simulation	8	1.1%	0.00
European Simulation Symposium (ESS)	8	1.1%	0.00
Information Theory In The Benelux	7	0.9%	0.29
International Conference On Cognitive Science (ICCS)	6	0.8%	0.17
Advances In Modal Logic (AIML)	5	0.7%	0.20
Eurographics Incl Various Workshops	5	0.7%	0.40
European Conference On Artificial Intelligence (ECAI)	5	0.7%	0.40
IEEE Vehicular Technology Conference (VTC)	5	0.7%	0.00
International Symposium On System Integration	5	0.7%	0.20
Proceedings Of Recherche D'information Assistee Par Ordinateur (RIAO)	5	0.7%	0.40
Conference On Parallel Problem Solving From Nature (PPSN)	4	0.5%	0.25
EUROSIM Congress	4	0.5%	0.00
International Conference On Acoustics, Speech, And Signal Processing (ICASSP)	4	0.5%	0.25
Int. Conf. on Augmented, Virtual Environments And 3D Imaging (ICAV3D)	4	0.5%	0.25
International Conference On Very Large Databases (VLDB)	4	0.5%	2.75
International Workshop On Deontic Logic In Computer Science (DEON)	4	0.5%	0.00
Int. Worksh. Evaluation of Modeling Methods In Systems An. & Design (EMMSAD)	4	0.5%	0.00
Procs World Multiconference On Systemics, Cybernetics And Informatics (SCI)	4	0.5%	0.00
CMG Computer Olympiad	3	0.4%	0.00
Conference On Logic And The Foundations Of Game And Decision Theory, Loft	3	0.4%	0.00
Content Based Multimedia Indexing (CBMI)	3	0.4%	0.33
IEEE International Conference On Systems, Man & Cybernetics (SMC)	3	0.4%	0.00
Int. Conf, On Parallel And Distributed Processing Techniques Applications (PDPTA)	3	0.4%	0.00
Not Unified	*511*	*67.7%*	*0.36*
Total	*755*	*100.0%*	*0.38*
Logic			
Advances In Modal Logic (AIML)	9	5.9%	0.22
Belgium-Netherlands Conference On Artificial Intelligence (BNAIC)	7	4.6%	0.00
Genetic And Evolutionary Computation Conference (GECCO)	7	4.6%	0.71
International Conference On Cognitive Science (ICCS)	4	2.6%	0.00
International Joint Conference On Artificial Intelligence (IJCAI)	4	2.6%	1.50

Subfield/proceedings	***Nr. Papers***	***% Papers***	***Average Citation Impact of NL papers***
Logic (continued)			
Conference On Logic And The Foundations Of Game And Decision Theory, Loft	3	2.0%	0.00
European Conference On Artificial Intelligence (ECAI)	3	2.0%	0.67
International Workshop On Deontic Logic In Computer Science (DEON)	3	2.0%	0.00
Not Unified	*112*	*73.7%*	*0.43*
Total	*152*	*100.0%*	*0.41*
Telematics			
Open European Summer School (Eunice)	4	7.3%	0.25
Annual Symposium Of The IEEE/LEOS Benelux Chapter	3	5.5%	0.00
Brazilian Symposium On Computer Networks (SBRC)	3	5.5%	0.00
Not Unified	*45*	*81.8%*	*0.31*
Total	*55*	*100.0%*	*0.27*
Mathematics			
Conference Of The IEEE Communications Society (INFOCOM)	8	9.2%	5.13
European Control Conference (ECC)	7	8.0%	0.00
IEEE Conference On Decision And Control (CDC)	4	4.6%	0.00
Not Unified	*3*	*3.4%*	*0.33*
Total	*65*	*74.7%*	*0.20*

Table 5.4 shows that the papers by Netherlands computer scientists in many sources have on average a low citation impact. For instance, for Netherlands papers published in the BNAIC, the citation impact per paper is only 0.05 (for *Programming Research and Algorithmics*) or 0.02 (for *Information and Knowledge Systems*), and for the ASCI proceedings 0.18. Even if the average citation rates of Netherlands papers in BNAIC and ASCI would be increased by 10-30 per cent, they would still be low compared to that of other proceedings volumes covered by the *Expanded WoS* database.

But certainly not in all conference proceedings listed in Table 5.4 Netherlands papers have such a low impact. For several sources the citation impact is substantial, and reaches a level that is similar to that of many sources in Table 5.1 that are covered by the *Expanded WoS* database. The following list contains sources in which Netherlands papers have an average impact above 0.66, i.e., that would probably have reached an impact of around 1.0 if citations from these sources would have been included in the database:

Computing and Imaging:

Eurographics incl. Various Workshops
IEEE International Conference on Robotics and Automation (ICRA)

International Conference on Very Large Databases (VLDB)

Programming Research and Algorithmics:

Text Retrieval Conference (TREC)
European Conference on Software Maintenance and Reengineering (CSMR)
Conference on Parallel Problem Solving from Nature (PPSN)
Proceedings of Recherche D'information Assistee par Ordinateur (RIAO)

Information and Knowledge Systems:

Genetic and Evolutionary Computation Conference (GECCO)
International Joint Conference on Artificial Intelligence (IJCAI)
Text Retrieval Conference (TREC)
International Conference on Very Large Databases (VLDB)

Logic:

Genetic And Evolutionary Computation Conference (GECCO)
International Joint Conference on Artificial Intelligence (IJCAI)
European Conference on Artificial Intelligence (ECAI)

Mathematics:

Conference of the IEEE Communications Society (INFOCOM)

This issue is further discussed in Chapter 9.

5.3 Coverage of reference lists by the expanded CWTS-WoS database (internal coverage)

To gain further insight into the extent to which the expanded *CWTS-WoS* database covers the research conducted at Netherlands academic Computer Science departments, we analysed the reference lists of their papers. It is important to note that this analysis considers a subset of the total publication output, since only those reference lists are available for examination that are contained in papers that are already included in the expanded *CWTS-WoS* database. For that reason it can be denoted as an ‘internal’ measure of the coverage rate. To what extent do these papers in the expanded *CWTS-WoS* database cite documents published in sources not processed for the expanded database? These documents may be published in types of sources that are not present at all in our database such as monographs, book chapters and reports, but they may also be published in journals and proceedings volumes not selected for inclusion in the database.

As the earliest year for which we included additional conference proceedings was 1996, we could only trace those references that are dated between 1996 and 2001. This is an important restriction in two respects. Firstly, it limits the analysis of cited references primarily to the more recent papers in our publication set, as papers from 1996 will have very few or no cited references that cite other documents from 1996. Secondly, this restriction affects the calculated coverage rates. This is due to the fact that some publications types, particularly reports, preprints and conference papers usually receive most of their citations shortly after publication, whereas it generally takes longer for books and journal articles to generate citation impact. Consequently, if one subdivides cited references into a class with recently published documents (typically 0–5 years old) and older publications, reports, preprints and conference papers tend to be overrepresented in the first class, and underrepresented in the second. Author self-citations were included, as we could not exclude all self-citations for non-source documents.

Table 5.5 presents an analysis of the field Computer Science in the entire database created in this study. It compares the internal coverage percentage in the WoS Expanded Database with that for the ‘Pure’ WoS database, in which papers from ACM, IEEE and additional LNCS proceedings are *not* included. This table shows that, compared to the ‘Pure’ WoS, the internal coverage of the field computer science in the Expanded WoS database increased from 38% to 51%. Outcomes for articles published in earlier years are very similar to those for the year 2004.

Table 5.5: Internal coverage percentages for Computer Science in WoS Expanded and ‘Pure’ WoS database

Database	***Nr articles in 2004***	***Nr Cited Refs to 0-5 year old articles***	***% Cited references published in database (Internal Coverage)***
‘Pure’ WoS (before expansion)	43,074	446,162	38%
Expanded WoS (after expansion)	64,717	558,709	51%

In order to analyse differences in internal coverage among subfields, ***Table 5.6*** presents the outcomes of a cited reference analysis per subfield. The second column in this table gives the number of papers published by Netherlands academic computer scientists in 2001 in sources (journals or proceedings volumes) covered by the *Expanded WoS* database. The third column gives the total number of cited references, contained in the papers counted in the second column, and published during the time period 1996-2001. Finally, the fourth column presents the percentage share of these references that were published in sources covered by the *Expanded WoS* database.

Table 5.6 shows that the highest internal coverage rates were found in *Mathematics* (60%), followed by *Computing and Imaging* (52%) and *Programming Research and Algorithmics* (46%). *Logic*, *Information and Knowledge Systems* and *Telematics* show coverage percentages of 40, 38 and 34 per cent, respectively. These outcomes are consistent with those obtained in the external coverage analysis presented in Table 5.3 in Section 5.1.

Table 5.6: Differences in internal coverage among subfields (based on reference lists in Netherlands papers only)

Subfield	*Nr papers in 2001 in sources covered by Expanded WoS database*	*Nr 0-5 year old cited references contained in these papers*	*% Cited references published in sources covered by Expanded WoS database.*
Computing and Imaging	191	2,431	52%
Programming Research and Algorithmics	187	2,050	46%
Information and Knowledge Systems	196	2,444	38%
Logic	72	763	40%
Telematics	8	79	34%
Mathematics	77	631	60%
Total	731	8,398	46%

To which extent do the internal coverage percentages calculated for Netherlands papers in the various subfields and shown in Table 5.6 differ from those for papers by colleagues outside the Netherlands who are active in the same subfields? In other words, are the outcomes for Netherlands papers in a subfield representative for the subfield as a whole? The categorisation of papers into subfields was carried out for Netherlands papers only. In order to obtain at least some indication, we analysed the reference lists in all papers published world-wide in the sources (journals and proceedings volumes) in which Netherlands groups allocated to a subfield had published at least one paper. The outcomes are presented in Table 5.7.

Table 5.7: Differences in internal coverage among subfields (based on reference lists in all papers published world-wide in sources in which Netherlands groups published at least one paper)

Subfield	*Nr papers in 2001 in sources covered by Expanded WoS database*	*Nr 0-5 year old cited references contained in these papers*	*% Cited references published in sources covered by Expanded WoS database.*
Computing and Imaging	19,525	186,105	60%
Programming Research and Algorithmics	15,634	153,871	55%
Information and Knowledge Systems	15,855	158,627	49%
Logic	11,955	117,473	61%
Telematics	3,750	28,151	39%
Mathematics	20,434	261,546	75%
Total			

Table 5.7 shows that the internal coverage percentages for all papers published ‘world-wide’ in sources in which Netherlands groups published at least one paper are generally somewhat higher than those percentages calculated for Netherlands papers only.

Table 5.8 presents statistics on the distribution of the coverage rate of both cited references and publications across *research groups*. Although the dispersion between research groups has been reduced by the expansion of the database, the coverage for a substantial number of groups remains moderate. In fact, Column 'Q1' in Table 5.8 shows that 25 per cent of groups have before expansion an internal coverage percentage of 21 per cent or lower, and after expansion at most 33 per cent. On the other hand, another 25 per cent of groups shows after expansion of the database an internal coverage of at least 51 per cent, and an external coverage of at least 46 percent (column 'Q3').

Table 5.8: Distribution of the coverage rate of cited references and publications among 67 Netherlands academic Computer Science groups 1996 – 2001*

		Mean	***Min***	***Q1***	***Median***	***Q3***	***Max***
Internal coverage (based on cited References)	Before expansion	31%	9%	21%	28%	40%	69%
	After expansion	42%	17%	33%	42%	51%	69%
External coverage (based on Publications)	Before expansion	26%	0%	16%	24%	35%	75%
	After expansion	38%	9%	28%	37%	46%	82%

* Min, Max: the minimum and maximum score in the distribution among groups. Q1, Q3: the first and third quartile of the distribution

5.4 Coverage of citations to Netherlands academic research in academic Computer Science by the expanded CWTS-WoS-database

Table 5.9 relates to *all types* of publications listed by the Netherlands academic Computer Science groups. It gives per type of publication the number of publications (***P***), the number of citations these publications have received during the time period 1996–2004, including (***C+sc***) and excluding (***C***) author self citations, and finally the citation per publication ratio (***CPP***) and the percentage of uncited publications (***%Pnc***), both with author self-citations included.

Table 5.9 shows that the total number of citations (***Ctot***) to all publications except those in the category 'other' amounts to 18,143, whereas the number of citations to articles in the *Expanded WoS* database (***C***) is 11,439. Consequently, 55 per cent of the total number of citations are taken into account of the *Expanded WoS* citation analysis presented below in Chapters 6 and 7. This percentage can be denoted as the citation coverage.

Table 5.9 also shows that articles published in sources included in the *Expanded WoS* database are on average much more frequently cited than the publications not included in this database: 4.86 against 1.55 citations per publication. Both journal and conference papers in the *Expanded WoS* database are cited on average more frequently than their counterparts not included in this database. Particularly for proceedings articles this difference is large: 3.57 versus 0.87, or a ratio between these two of 4.1.

In fact, Table 5.9 shows that about 27 per cent of proceedings articles are published in *ACM*, *LNCS* and *IEEE* proceedings, while these 27 per cent attract 60 per cent of all citations to conference proceedings articles. This reflects at least partly that the proceedings volumes included in the *Expanded WoS* database have a much higher impact than those that were not included. On the other hand, it must also be noted that the cited references in proceedings not included in the *Expanded WoS* database do *not* contribute to the citation counts, since they were not added to the database.

Table 5.9: Citations 1996 – 2004 in expanded *CWTS/WoS* database to the total publication output 1996 – 2001 of Dutch research groups in Computer Science broken down by type of publication

Type of Publication	P	C+sc	C	CPP	Pnc
Publications in Expanded WoS database					
WOS+ Journal Paper	1,309	10,204	7,718	5.90	26%
WOS+ Conference Paper	1,041	5,426	3,721	3.57	36%
Total WoS Expanded	***2,350***	***15,630***	***11,439***	***4.86***	
Publications not in Expanded WoS database					
Journal article	491	1,585	1,273	2.59	64%
Conference paper	2,858	3,570	2,486	0.87	76%
Chapter	473	925	754	1.59	73%
Theses	397	493	415	1.05	73%
Book	95	1,907	1,776	18.69	55%
Total not in WoS Expanded	***4,314***	***8,480***	***6,704***	***1.55***	
Grand total	***6,664***	***24,110***	***18,143***	***2.72***	
[‘Other’ publications	1,181	509	424	0.36	91%]

In the set of articles published in sources included in the *Expanded WoS* database, Table 5.9 reveals that proceedings articles have on average a lower citation rate than journal articles. This outcome differs from that obtained in the ‘world’ analysis of the field Computer Science presented in Chapter 4, where it was found that proceedings papers and journal articles tend to have on average the same citation impact.

Table 5.10: Citation impact ratio covered vs. non-covered proceedings papers per subfield

Subfield	*Impact ratio covered vs. non-covered proceedings papers*
Computing and Imaging	4.1
Programming Research and Algorithmics	4.2
Information and Knowledge Systems	4.3
Logic	4.0
Telematics	4.1
Mathematics	2.5
Total	4.1

An analysis by subfield revealed that the patterns observed in Table 5.9 can be found in each subfield. In other words, they are not determined by characteristics of one or two particular subfields. In all subfields, the citation impact of proceedings papers not included in the *Expanded WoS* database is much lower than that of papers in proceedings that are covered. ***Table 5.10*** shows that, except for Mathematics, in all subfields the ratio of the impact of papers in proceedings covered by the *Expanded WoS* versus that of papers in non-covered proceedings is around 4.

But it does not follow that all papers published in these non-covered sources have a low citation impact. Citation distributions are skewed, and even proceedings volumes of which the papers have on average a low impact, may contain articles that are relatively highly cited. A similar comment can be made for book chapters.

6 Comparison with *QANU* 2004 Peer Ratings

6.1 Introductory comments

The main objective of this analysis is to subject the bibliometric indicators to a first test, by comparing their values to peer ratings, and calculating simple correlation measures. The analysis does *not* aim at re-doing the QANU evaluation and it does not assume nor want to suggest that correlations should be perfect. The Committee took into account several dimensions of research quality, and did not only consider quantitative information, but also background information collected during site visits to the groups under evaluation. In addition, the Committee as a whole can be assumed to have a detailed knowledge of the various subfields in Computer Science. As outlined in Chapter 1, the authors of this report have the view that bibliometric indicators, when applied properly, provide useful *tools* in peer review processes. They cannot and should not *replace* peer judgments, but their use can make peer review processes more transparent.

In the sections below, rank correlation coefficients are calculated between peer ratings and bibliometric indicators, and among bibliometric indicators as well. These coefficients should be interpreted as purely *descriptive* statistics of a particular data sample.

6.2 The QANU Peer Review

In 2004 an international peer review committee evaluated the research performance of research groups in Computer Science at a number of Netherlands universities. The evaluation was organised by the Quality Assurance Netherlands Universities (*QANU*). In that year the evaluation of academic research at Netherlands universities, previously organised by the Netherlands Association of Universities (*VSNU*), was transferred to *QANU*. A Review Committee consisting of nine members evaluated the research groups in Computer Science at the following institutions: Eindhoven Technical University (TUE); Radboud University Nijmegen (KUN); University of Maastricht (UM); University of Amsterdam (UvA); Free University Amsterdam (VU); Utrecht University (UU); University of Groningen (RUG); and University of Twente (UT).

The Committee rated the research quality, productivity, relevance and viability of all groups on a five–point scale: 5 (excellent); 4 (good); 3 (satisfactory); 2 (unsatisfactory); and 1 (poor). These grades conform to the '*VSNU* Protocol 1998'. For more details the reader is referred to the Committee's final report (*QANU*, 2004), particularly to Appendix H. The grades given in this Appendix were used in the analysis presented in this Chapter. The data on publication output that each group provided to the Review Committee are the same as those that were analysed in the study presented in this report.

This chapter focuses on the two principal aspects taken into account by the Committee: quality and productivity. Quality was further specified as 'international recognition and innovative potential'. Five key publications provided by each research group, a group's research program description, and its lists of publications made during 1996–2001 constituted the main sources of information for the quality judgment. Excellent groups are 'at the forefront internationally', generate an 'important and substantial impact in the field' and are 'international leaders'.

Productivity was conceived as the ratio of publication output and research input. Publication output included Ph.D. theses and books (weight factor 2, edited works not included); academic publications; articles in conference proceedings; book chapters; and items of scientific software (all weight factor 1). Input was defined as the number of Full Time Equivalents research time in a group, not counting Ph.D. students.

6.3 Results

Table 6.1 presents Spearman rank correlation coefficients between 11 variables for 42 groups evaluated in the 2004 *QANU* assessment. It takes into account the Committee's grades for quality and productivity, the number of Full Time Equivalents Research Time (FTE) during the time period 1996–2001, and eight bibliometric indicators. Apart from the standard indicators outlined in Section 3.3, Table 6.1 includes two additional indicators:

- ***P/FTE***: the number of articles in the *Expanded WoS* database per FTE research time. This indicator can be termed as a publication productivity indicator, in the sense that it relates 'output' to 'input'.
- ***C/FTE***: the number of citations (received during 1996–2004 by a group's articles published during 1996–2001), per FTE research time. It can be denoted as a citation productivity indicator.

Table 6.2 compares the outcomes of the *QANU* quality assessment of Computer Science (2004) with those of the *VSNU* assessment of Netherlands Chemistry and Chemical Engineering, carried out in 2002 (*VSNU*, 2002). This table presents statistics on peer ratings, bibliometric indicators and their rank correlations.

Tables 6.1.and 6.2 reveal the following characteristics. The Chemistry Review assessed as many as 152 research groups, the Computer Science Review 42. Measured in terms of the number of FTE research time in the last year of the time period under evaluation (2001 and 2000, respectively), Computer Science groups were somewhat smaller than those active in Chemistry: 7.11 against 8.96 FTE.

The quality ratings made by the *QANU* Review Committee for Computer Science groups are generally higher then those given by the *VSNU* Chemistry Committee to Chemistry groups (4.30 versus 4.07), whereas the standard deviation of the *QANU* ratings is lower than that for the Chemistry groups (0.32 against 0.68). The mean citation impact (compared to the world citation average) of the Computer science groups is substantially lower than that of Chemistry groups (1.15 against 1.55), while the standard deviations are almost equal. The rank correlation coefficient between quality rating and citation impact is for academic Computer Science groups lower than it is for academic Chemistry groups: 0.22 against 0.38.

Table 6.1: Spearman rank correlation coefficients between 11 variables for 42 groups evaluated in the 2004 *QANU* assessment

.	*Pro-duc-tivity*	*P*	*C*	*CPP*	*%Pnc*	*CPP/ FCSm*	*JCSm/ FCSm*	*FTE 96-01*	*P/ FTE 96-01*	*C/ FTE 96-01*
Quality	*0.09*	*0.48*	*0.42*	*0.27*	*-0.25*	*0.22*	*0.13*	0.31	0.19	0.27
Productivity		*0.35*	*0.29*	*0.19*	*-0.22*	*0.17*	*0.09*	0.12	0.33	0.33
P			0.90	0.62	-0.50	0.60	0.46	0.71	0.42	0.59
C				0.88	-0.62	0.85	0.54	0.56	0.46	0.80
CPP					-0.72	0.97	0.46	0.29	0.45	0.90
%Pnc						-0.72	-0.11	-0.06	-0.46	-0.81
CPP/FCSm							0.42	0.30	0.44	0.88
JCSm/FCSm								0.34	0.13	0.31
FTE 96-01									-0.23	0.09
P/FTE										0.74

Table 6.2: Comparison *QANU* Computer Science (2004) versus *VSNU* Chemistry (2002) peer ratings and bibliometric outcomes.

Variables	*QANU Comp Sci 2004 (n=42)*	*VSNU Chemistry 2002 (n=152)*
	Mean ± STD	*Mean ± STD***
FTE in 2001	7.11 ± 3.94	8.96 ± 5.30
Quality Peer rating	4.30 ± 0.32	4.07 ± 0.68
Productivity Peer rating	4.10 ± 0.61	3.97 ± 0.82
Citation impact compared to world average (CPP/FCSm)	1.15 ± 0.78	1.55 ± 0.73
	Spearman's R	*Spearman's R*
Quality ~ Productivity	0.09	0.32
Quality ~ FTE	0.31	0.32
Quality ~ P	0.48	0.35
Quality ~ CPP/FCSm	0.22	0.38
Quality ~ C/FTE	0.27	0.27
Productivity ~ FTE	0.12	0.10
Productivity ~ P	0.35	0.51
Productivity ~ P/FTE	0.33	0.51
FTE ~ CPP/FCSm	0.30	0.20
P ~ CPP/FCSm	0.60	0.14
P/FTE ~ FTE	-0.22	-0.23
C/FTE ~ FTE	0.09	-0.09
CPP/FCSm ~ JCSm/FCSm	0.41	0.66

** STD: Standard Deviation. For an explanation of the indicator symbols: see Table 3.5 in Section 3.3.

It must be noted that Chapter 7 reports for Netherlands academic Computer Science as a whole a citation impact compared to the world average of 1.30, which is higher than the mean of 1.15 over all groups presented in Table 6.2. A first explanation of this discrepancy is that the analyses presented in this chapter include only those 42 groups that were evaluated in the *QANU* review or in the separate University reviews, whereas the outcomes presented in Chapter 5 are based on the total collection of 67 groups, including also those at CWI.

A second explanation is that there is a rather strong correlation (Spearman's R = 0.60) between the number of articles published by a group, and its citation impact. High impact groups tend to contribute a larger share of articles to the total Netherlands output than lower impact groups do. In the calculation of the citation impact of the total Netherlands publication output, a group's citation impact is weighted with the number of papers it has published, whereas in the computation of the mean value over groups presented in Table 6.2, a group's impact score is unweighted. Interestingly, in the collection of Chemistry groups the rank correlation between a group's citation impact and the number of articles it published is much lower (Spearman's R=0.14).

The bibliometric indicator that shows for academic Computer Science groups the highest rank correlation with the quality rating is the number of articles in the *WoS-Expanded* database (including *ACM*, *LNCS* and *IEEE* proceedings). The number of articles per FTE research time, an indicator that can be conceived as publication productivity, shows for groups in both fields a negative rank correlation of about -0.22 with the group size measured in FTE. The citation productivity, defined as the number of received citations per FTE research time, reveals low correlation coefficients of 0.09 and - 0.09, respectively. The citation impact compared to the world average has with group size a rank correlation of 0.20 and 0.30.

Using bibliometric 'productivity' indicators – published articles and citation impact per FTE research time – as performance measures, there is no evidence for an economy of scale. However, using Quality peer ratings as performance measures, both Computer science and Chemistry groups revealed a positive rank correlation with group size of around 0.30.

The productivity peer rating of academic Computer Science groups does hardly show a rank correlation with the quality rating (Spearman's R=0.09). Its correlation with the number of articles in the *WoS-Expanded* database (P) and the number of articles per FTE research time (P/FTE) are positive, though not significant, and somewhat lower than those obtained for Chemistry groups. The results presented in Tables 6.1 and 6.2 relate to the aggregate of research groups from various subfields. A secondary analysis revealed that variability in citation impact and lack of correlation with peer ratings were also found *within* subfields. The outcomes are further discussed in ***Section 9.4***.

7 Preliminary Results

This chapter presents the outcomes of a preliminary citation analysis of the Netherlands academic publication output in Computer Science. The database of Netherlands Computer science publications was described in Section 2.3. The methodology applied in the citation analysis was presented in Chapter 3 of this report. A short list of the indicators presented in this chapter is given in Table 3.5 in Section 3.3. The citation analysis presented below relates to the subset of articles of Netherlands academic Computer Science groups that were published in journals or proceedings volumes included in the *Expanded WoS* database, i.e. published in journals processed for the *WoS*, or in *ACM*, *LNCS* and *IEEE* proceedings volumes added to the *WoS* database. With the exception of Figure 7.3, other types of publications are not included in the Tables and Figures presented in this chapter.

7.1 Results for Netherlands academic Computer Science as a whole

Table 7.1 presents the main results at the overall level of all groups combined, i.e. for Netherlands Computer Science ('kern-informatica) as a whole. Citations are counted up to and including the year 2004, the last year for which the *CWTS-WoS* database was expanded. This means that for each publication year a different citation window length is applied. For example, a nine-year citation window is used for papers published in 1996 (time period 1996–2004), and a four year citation window was used for papers from 2001 (time period 2001–2004). In this manner, we counted for the overall level almost sixteen thousand citations (***C+sc***).

In a quarter of these citations, the citing publication and the cited publication had at least one author in common (***Self Cits***). These so called '*author self citations*' are excluded from all other indicators. The 2,350 papers (***P***) by Netherlands academic computer scientists received on average almost five citations (***CPP***). The skewness of the citation distribution becomes evident when we observe that 30% received no citations at all during the time period considered (***Pnc***).

The values of all indicators just quoted depend heavily on the timeframe in which they are measured. The ratio's ***CPP/JCSm*** and ***CPP/FCSm*** indicate how the average citation rate per publication compares to the average citation score of 'similar' papers, measured in the same timeframe. In other words, these indicators correct for differences among groups or organizations as regards the distribution of their publications among publication years.

In the case of ***CPP/JCSm***, 'similar' papers are defined as papers published in the same year and source, i.e. journal or conference proceedings. Thus, Table 7.1 shows that papers by Netherlands academic computer scientists receive 13% more citations on average than similar papers published in the same journals and conference proceedings. The '+'-sign behind the indicator shows that in this case *CPP* deviates significantly from

JCSm (at a confidence level of 95% according) to a statistical test developed by Schubert and Glänzel (1983). Because the present study does not use randomly sampled data, significance tests are not appropriate for inferential analysis. However, significance is reported here as an arbitrary criterion in deference to its widespread use in social science for exploratory analysis of non-random data.

In the case of ***CPP/FCSm*** the citation impact of Netherlands academic computer scientists is compared to the world citation average in the subfields in which they are active. Subfields are conceived of as particular sets of journals and/or conference proceedings belonging to the same scientific sub-discipline. For this project we defined one single subfield 'Computer Science' that encompassed all conference proceedings from *ACM*, *IEEE* and *LNCS* and all the journals that were classified into one of the Computer Science sub-disciplines discerned by *Thomson*/*ISI*.

Applying this method we find that 72% of the output of Netherlands academic computer scientists belongs to the subfield '*Computer Science*', followed by '*Applied Mathematics*' (7%), '*Mathematics*' (4%) and '*Electrical and Electronic Engineering*'.(3%). Comparing the citation impact of papers by Netherlands academic computer scientists to the average citation rates in this specific combination of subfields, it turns out that research articles by Netherlands academic computer scientists receive 30% more citations, the difference with the worldwide citation average being statistically significant. As the ratio of ***CPP/FCSm*** is even higher than ***CPP/JCSm***, it follows that Netherlands academic computer scientists chose to publish in the journals and conference proceedings that have a relatively high citation impact within the total collection of sources that is covered by the expanded *CWTS-WoS* database.

The last three rows of Table 7.1 show the results of an analysis focusing on 'highly cited' publications rather then average citation rates. They address the question as to how many papers that belong to the top ten percent most frequently cited papers in a field are (co-) authored by Netherlands academic computer scientists. A fixed citation window length of four years was applied. For each subfield in which Netherlands academic computer scientists were active, papers were ranked according to the number of citations they received during a 4-year period. Next, for each paper it was determined whether it belonged to the upper decile of this citation distribution or not. We only ranked conference papers, review and normal journal articles; letters were not included.

Table 7.1 shows that Netherlands academic computer scientists produced 344 such frequently cited publications in the time period 1996 – 2001 that belong to the upper decile of their citation distribution. This is 50 % more than one would expect based on the size of the Netherlands publication output and the shapes of the (discrete) citation distributions. This level of 50 per cent above expectation for highly cited papers is somewhat higher than the level of 30 per cent above world average that was obtained in the calculation of the normalised citation impact ratio ***CPP/FCSm***.

Table 7.1 Indicators for Netherlands Computer Science as a whole*

Symbol	*Definition*	*Score*	
P	The number of articles (normal articles, letters, notes and reviews) published during 1996-2001 by Netherlands academic computer scientists in sources (journals, proceedings) covered by the Expanded WoS database	2,350	
C	The number of citations recorded in the expanded WoS database given to all articles subjected to the citation analysis. Author self-citations are excluded.	11,439	
C+sc	The number of citations recorded in the Expanded WoS database to all articles subjected to an citation analysis, including author self-citations.	15,630	
CPP	The average number of citations per publication. Author self-citations are not included.	4.86	
CPP/FCSm	The impact of the Netherlands articles, compared to the world citation average in the subfields in which Netherlands computer scientists are active **	1.30	+
CPP/JCSm	The impact of the Netherlands articles, compared to the average citation rate of the set of sources (journals or proceedings) in which they were published **.	1.13	+
JCSm/FCSm	The impact of the sources (e.g. journals, proceedings volumes) in which the Netherlands articles were published , compared to the world citation average in the subfields covered by these sources	1.15	
Pnc	The percentage of articles not cited during the time period considered, excluding author self-citations.	30 %	
% Selfcit	The percentage of author self-citations. An author self-citation is defined as a citation in which the citing and the cited paper have at least one author in common (either a first author or a secondary author).	27 %	
P top 10%	The actual, absolute number of papers published by Netherlands academic computer scientists that are among the 10 % most frequently cited of similar (in terms of type, age and subfield) papers during the time period considered; citations are counted during the first 4 years after publication date (a 4-year citation window).	344	
E(P Top 10%)	The expected number of Netherlands computer science papers amongst the global top 10 %, based on the number of Netherlands papers published during the time period considered.	229	
A/E (P top 10%)	Indicates the relative contribution of Netherlands academic computer scientists to the global upper 10% of the citation distribution in the time period considered (=P top10% / E (P top10%).	1.5	

* Publication period: 1996-2001 (CWI: 1997-2001). Citation period: 1996-2004. Citation analysis relates to those publications in the NL-CS publication database (based on QANU/VSNU data and CWI Annual research Reports) that are published in journals or proceedings covered by the Expanded WoS (=WoS+ACM+LNCS+IEEE) database.
** A '+' ('-') symbol directly after the numerical value indicates that the impact of the Netherlands articles is significantly above (below) world average in the subfields covered (CPP/FCSm) or in the set of sources used (CPP/JCSm). Author self-citations are not included.

To examine the evolution of the citation impact of Netherlands academic Computer Science over time we conducted a trend analysis in which publications from each year were taken separately. A feature that can be observed is the marked increase in the number of publications from 188 in 1996 and 295 in 1997 to 578 in 2001. The relatively low number of papers in 1996 is mainly caused by the fact that no papers from CWI were included for this year (see paragraph 2.3), but the output of Netherlands academic computer scientists grows at a high rate even when 1996 is left out. This is partly due to the better coverage of the expanded *CWTS-WoS* in later years, but it also reflects the increase in the number of papers that were listed in the self evaluation reports.

To facilitate comparison of the citation impact between the different years of publication, we applied a fixed citation window of four years. This means that citations are counted up to four years after the publication date (the year of publication included). For publications from 1996 citations are counted during 1996 – 1999 (but *not* during 2000 – 2004), for publications from 1997 citations are counted during 1997 – 2000 et cetera. Naturally, the focus of this analysis is more on short term citation impact than in the overall analysis presented in Table 7.1. The number of citations received will be lower while the percentage of uncited publications will be higher. As we applied the same short term citation window when calculating the reference values ***FCSm***, the ratio ***CPP/FCSm*** gives an indication how the citation impact of publications by Netherlands academic computer scientists on the short term relates to the worldwide citation average.

Figure 7.1 depicts the behaviour of ***CPP/FCSm*** over time and reveals a sharp rise in citation impact of Netherlands academic Computer Science research, especially in the later years when the citation impact increased from around 25% to almost 60% above worldwide citation average. A secondary analysis revealed that the increasing trend displayed in Figure 7.1 is also visible in an analysis based on *journal* papers published in journals processed for the Web of Science – in other words, in the journal segment of our database. More importantly, each major subfield of Computer Science showed an increasing trend in the citation impact of the papers allocated to it, similar to that of the overall analysis presented in Figure 7.1. In other words, it is not true that only the research activities in one or two subfields are responsible for the positive trend this figure reveals.

The normalised citation impact indicator calculated in Figure 7.1 is a ratio of the actual and expected citation rate of Netherlands academic Computer Science papers. A secondary analysis revealed that the increase in this ratio during the time period considered is mainly due to an increase on the numerator of the ratio, the actual average citation impact of Netherlands papers, rather than a decline in the denominator (the world average). The interpretation of the trend is further discussed in ***Section 9.5***.

Figure 7.1: Trend in Impact per Publication Compared To World Subfield Average, 1996 – 2001/4

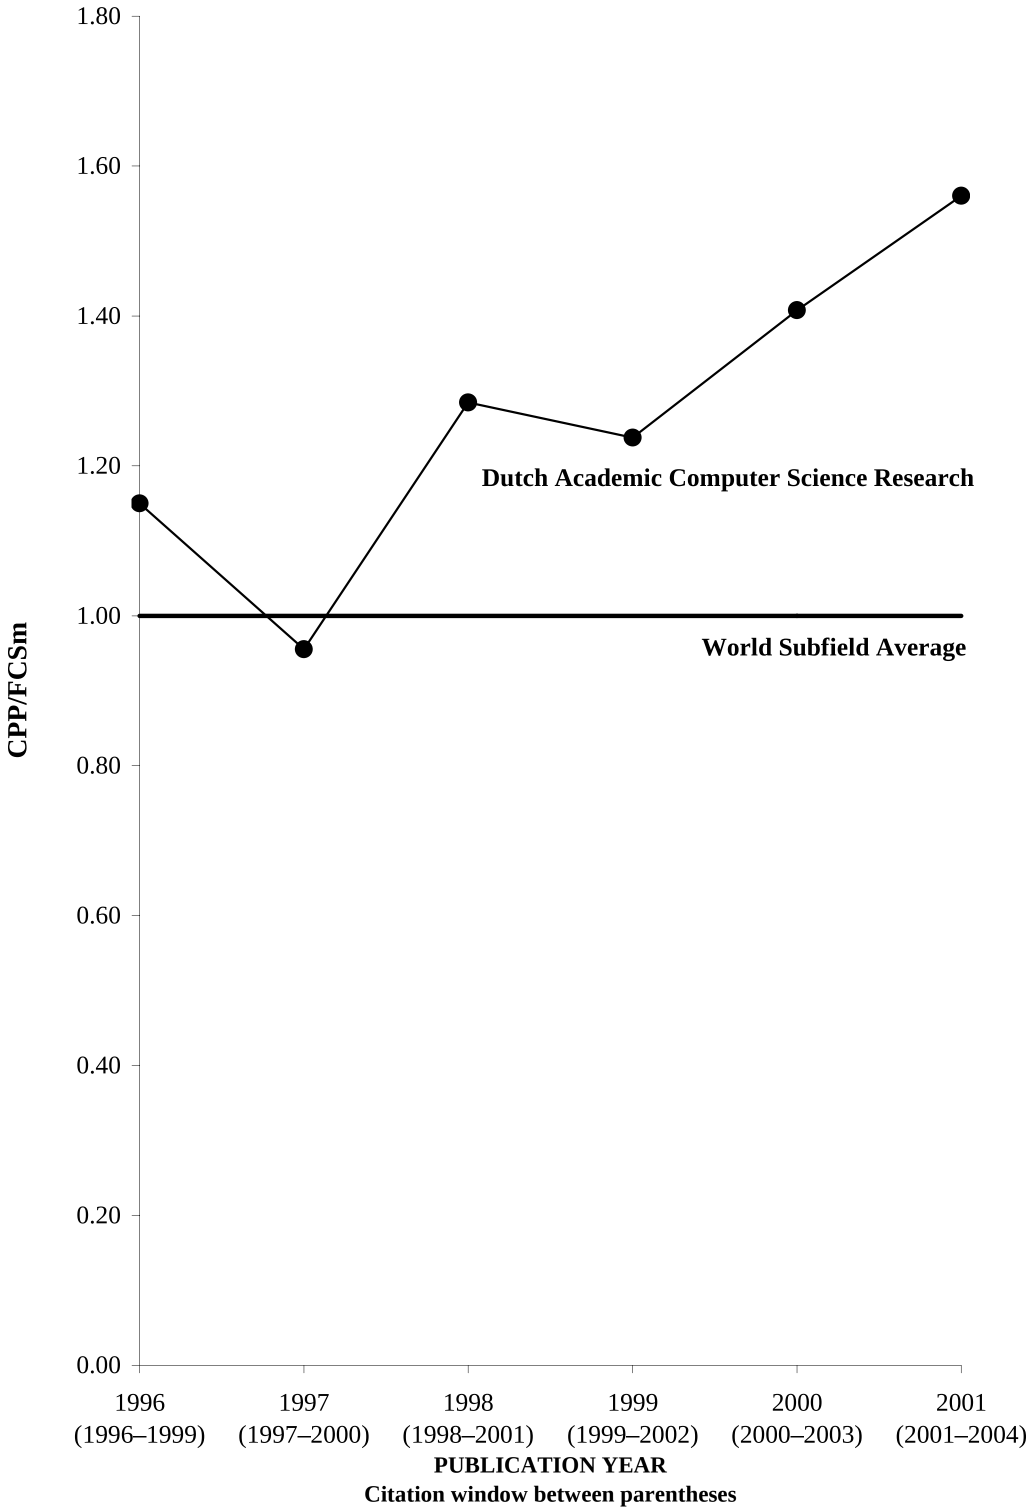

7.2 Results per (anonymous) research group

A next analysis also focuses on ***CPP/FCSm*** but relates to *research groups* instead of the total combined publication output. ***Figure 7.2*** is a scatter plot in which each circle or black coloured square represents an individual research group. The number of papers published during 1996 – 2001 (horizontal axis) is plotted against the relative citation impact measure in which the average citation rate of papers published by a group is compared to the world average citation rate in the subfield(s) in which the group is active (***CPP/FCSm***). The citation impact of research groups represented by black coloured squares is significantly different from the world average (at a 95% confidence level). Open squares indicate groups for which the citation impact does *not* differ significantly from this world average. Because the present study does not use randomly sampled data, significance tests are not appropriate for inferential analysis. However, significance is reported here as an arbitrary criterion in deference to its widespread use in social science for exploratory analysis of non-random data.

Figure 7.2 shows that the number of research groups with a citation impact above world average is almost equal to the number of groups below world average (35 above against 32 below). The fact that the citation impact of Netherlands academic computer scientists in general nonetheless compares favourably to world citation average should largely be attributed to differences in the number of publications across groups. Research groups with a relatively low citation impact tend to make a less than average contribution to the total publication output of Netherlands academic Computer Science.

The results presented in Figure 7.2 are based on the citation counts of papers in sources included in the *WoS Expanded* database. An interesting issue is how the computed ratios would change if the coverage of the Expanded WoS database would be further expanded with sources in which Netherlands researchers have published.

It must be noted that the Netherlands papers included in the Expanded WoS database can *not* be regarded as a *random* sample of the total population of papers published by Netherlands academic research groups of the Netherlands. In fact, the analyses presented in Chapter 5 indicated that among the sources the Netherlands computer scientists publish in, the ones included in the Expanded WoS database have on average a much higher impact than the sources that were left out. (cf. Table 5.8). Although this assertion may be correct as a general rule, it does not mean that all papers or even all sources that are not included in the Expanded WoS database have lower impact than the ones that are included. This raises the question to what extent the impact scores shown in Figure 7.2 are affected by the fact that some frequently cited publications have not been taken into account.

Figure 7.2: Impact in 1996 – 2004 of papers published in 1996 – 2001 by 67 Netherlands academic research groups in Computer Science compared to the World Subfield Average

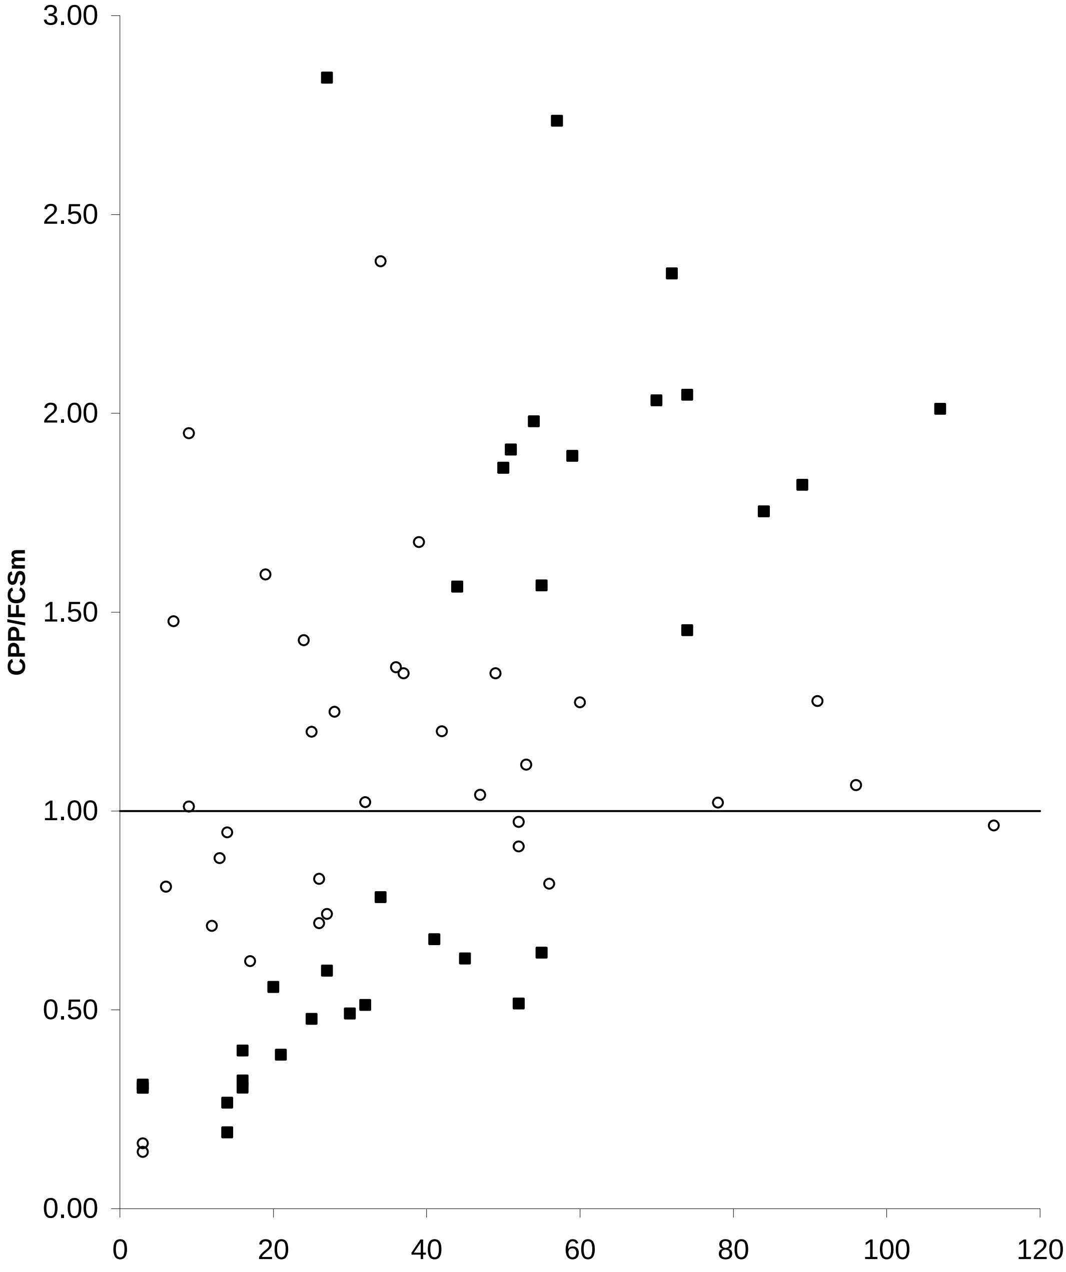

Black coloured squares above (below) the horizontal reference line represent groups for which the impact (CPP) is significantly above (below) world average (FCSm)

As argued in earlier chapters, the major advantage of restricting ourselves to the publications included in the WoS Expanded database universe is the ability to relate the mean impact score of research groups to an international impact score, the FCSm value. There is no such measure available for non-Expanded WoS papers. However, as the FCSm value is based on a very large collection of papers, we can safely assume that the FCSm value would probably not change very much if we would incorporate only a few additional papers. Based on this assumption, we were able to estimate the *maximum* value of the CPP/FCSm score if we had been able to include highly cited papers that have now been left out.

The estimation procedure was confined to conference papers, journal articles and chapters that were not incorporated in the Expanded WoS database. Books and theses have been left out because of their distinct role in the scientific communication process and the resulting deviant citation characteristics. Estimating FCSm values in the manner just described would not have been appropriate for books and theses.

As has been demonstrated in Chapter 4 there is heavy citation traffic within sources among journal volumes or conference proceedings volumes. For that reason, citation scores were increased by 33% for papers published in sources not included in the WoS Expanded database. In a first step the change in CPP/FCSm was estimated, for each individual non-Expanded WoS paper had it been included. After that, papers were included one by one, starting with the paper having the highest contribution to the CPP/FCSm until a maximum was reached.

Figure 7.3 shows the results of the estimation procedure described above. It shows that the estimated *maximum* of CPP/FCSm value is for most groups less than 15% above the original CPP/FCSm value, while for 75% of the groups the difference is less than 33%.

For the remaining 15 research groups - from which 9 have less than 20 Expanded WoS papers – a considerable difference can be observed between the original CPP/FCSm value and the Expanded WoS database had it also included their highly cited non-Expanded WoS papers. However, among these groups there is only one group for which their impact score increases from (just) below world average to above world average.

It needs emphasising that the deviations from the original values displayed in Figure 7.3 represent *a theoretical upper bound* of the citation impact ratio CPP/FCSm . Only highly cited papers are added. All other papers were ignored, even if they were published in the same sources as the added highly cited paper. But the figure provides at least some insight into the effect of limiting the citation analysis to (target) papers published in journals or proceedings covered by the Expanded WoS database. As such, it represents a first step in a more comprehensive study into this matter.

Figure 7.3: **Theoretical upper bound of the citation impact Netherlands academic research groups if their highly cited papers in non-Expanded WoS sources had been included**

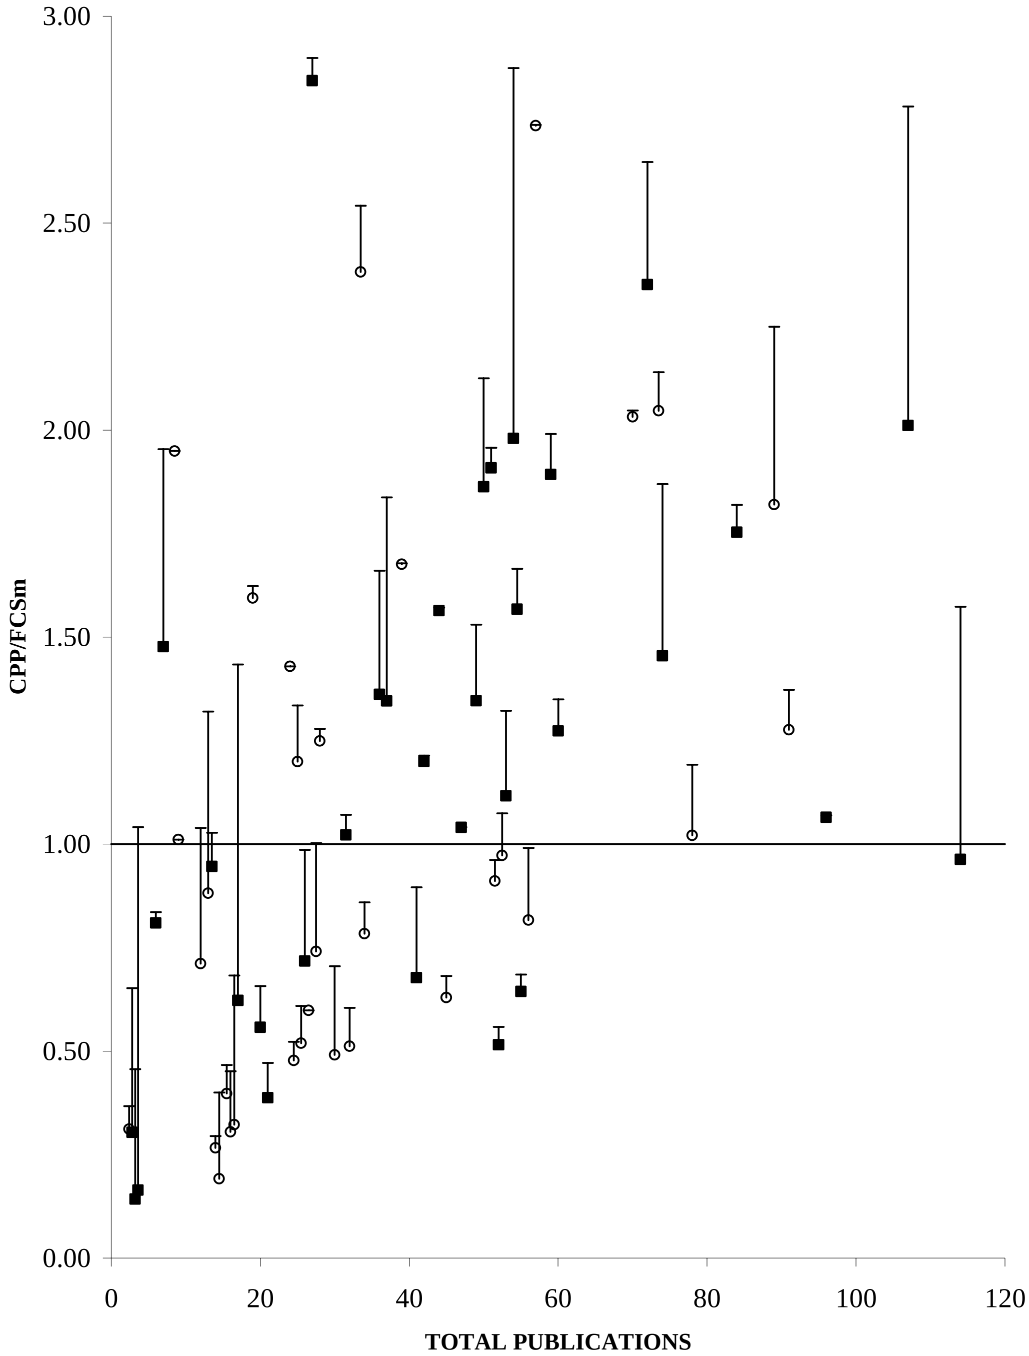

7.3 The effect of the database expansion upon citation impact indicators

Finally, we compare the outcomes from a '*Pure WoS*' analysis with those obtained in the *Expanded WoS* analysis. In the first, both the publication and the citation universe consists of all articles included in the *WoS* database. Papers – and their cited references – in proceedings volumes added to the database are *not* included. In the second, the publication and citation universe consists of all papers published in *WoS* journals or in *ACM*, *LNCS* and *IEEE* proceedings. ***Figure 7.4*** gives for each group on the horizontal axis the impact compared to the world citation average (***CPP/FCSm***) calculated within the *WoS*, and the vertical axis displays the value of the same measure, now computed in the *Expanded WoS* universe.

Figure 7.4: Comparison of *Pure WoS* and *Expanded WoS* citation impact of Netherlands academic computer science groups

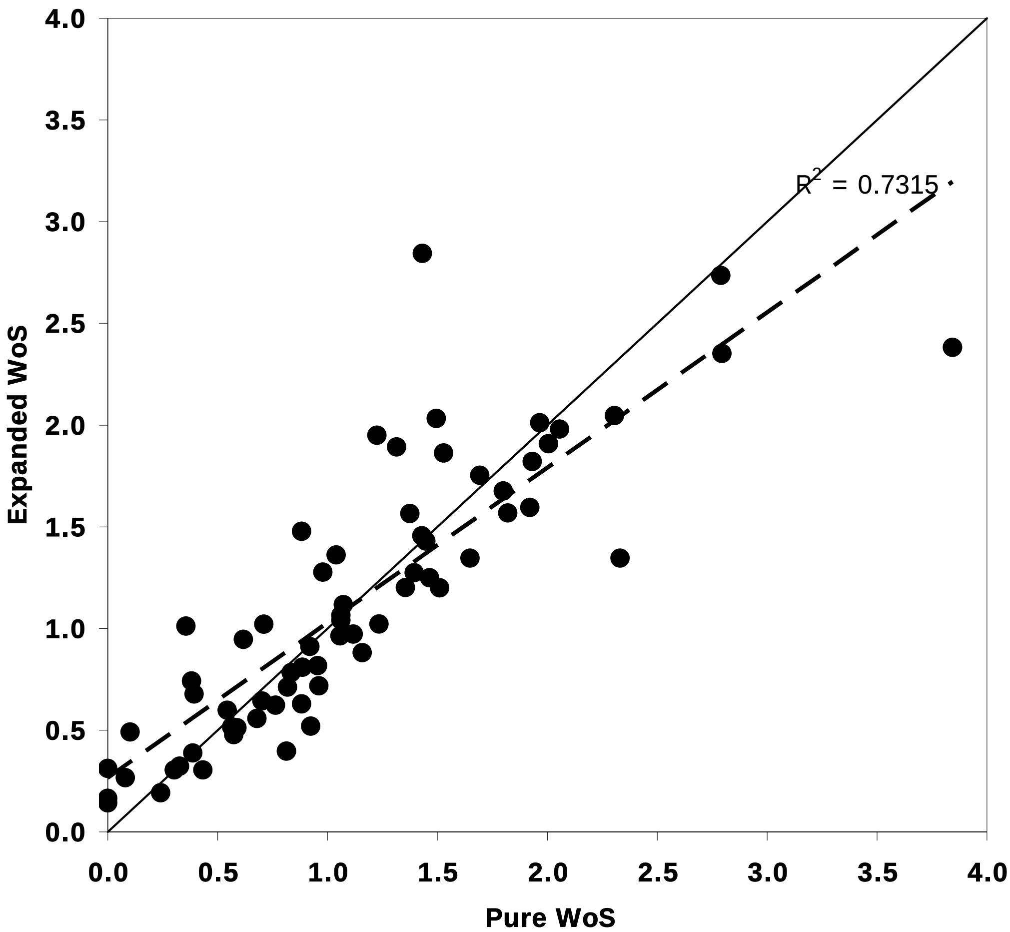

Figure 7.4 displays the linear regression line (dashed line), based on the *Pure WoS* scores as independent and the *Expanded WoS* rates as dependent variable. The explained variance amounts to 0.73, and the Pearson correlation coefficient to 0.82. Finding a positive correlation is not surprising at all. The articles and citations in the *Pure WoS* database are *also* included in the *Expanded* version.

The diagonal line represents points for which the values of the two measures are identical. The observed explained variance of 73 per cent indicates that a group's citation impact measured in the *Pure WoS* universe is statistically speaking in this particular case a good predictor of that computed in the *Expanded WoS* database. It must be noted that three apparent outliers have a strong effect upon the explained variance and the direction of the regression line. For some groups the *Pure WoS* and the *Expanded WoS* impact measures have similar values. These groups are represented as dots on or near the diagonal. Other groups are located above the diagonal. Interestingly, the major part of the groups is located somewhat below the diagonal. There are 19 groups that profit from the expansion of the *WoS* database in terms of their citation impact compared to the world average.

Table 7.2 presents the outcomes of a second analysis, in which all 66 groups were categorized into tertiles (upper, middle and bottom 33 per cent) on the basis of their normalised citation impact calculated within the *Pure WoS* universe, and also according to their citation impact obtained from the *Expanded WoS* universe. The diagonal in this table contains groups that did *not* move from one tertile to another when the *WoS* database was expanded. 18 Groups stayed in the bottom tertile, 13 in the middle and 17 in the upper tertile (i.e., the upper 33 per cent in terms of normalised citation impact). Thus, 72 per cent of groups (48 out of 66 groups) remained in the same tertile after expansion of the database. 28 Percent of groups moved either from or towards the middle tertile, while no group moved from a bottom to the upper tertile. 5 Groups in the middle tertile in the *Pure WoS* analysis moved to the upper tertile and 4 for the bottom tertile after expansion of the *WoS* with ACM and IEEE proceedings and with LNCS volumes not covered by the *WoS*.

Table 7.2: Comparison of Pure WoS and Expanded WoS impact categorisations

Expanded WoS *(=WoS+ACM+ LNCS+IEEE)*	***Pure WoS***		
	Bottom Tertile *(impact<0.81)*	Middle Tertile	Upper Tertile *(impact>1.37)*
Bottom Tertile *(impact<0.72)*	18	4	0
Middle Tertile	4	13	5
Upper Tertile *(impact>1.35)*	0	5	17

In needs emphasising that the two measures displayed in Figure 7.4 and Table 7.2 are largely dependent. The purpose of this figure and table is to show the differences in outcomes among the two universes that show a substantial overlap, and *not* to compare the citation impact of papers *added to* the *WoS* database to that of papers *already included* in the *WoS*. The latter type of analysis is presented in ***Figure 7.5***. It shows substantial differences in the average citation impact of the two sets of papers. The two measures do not show a significant linear or rank correlation at a 95 per cent confidence level.

Figure 7.5: Comparison citation impact of proceedings versus journal papers

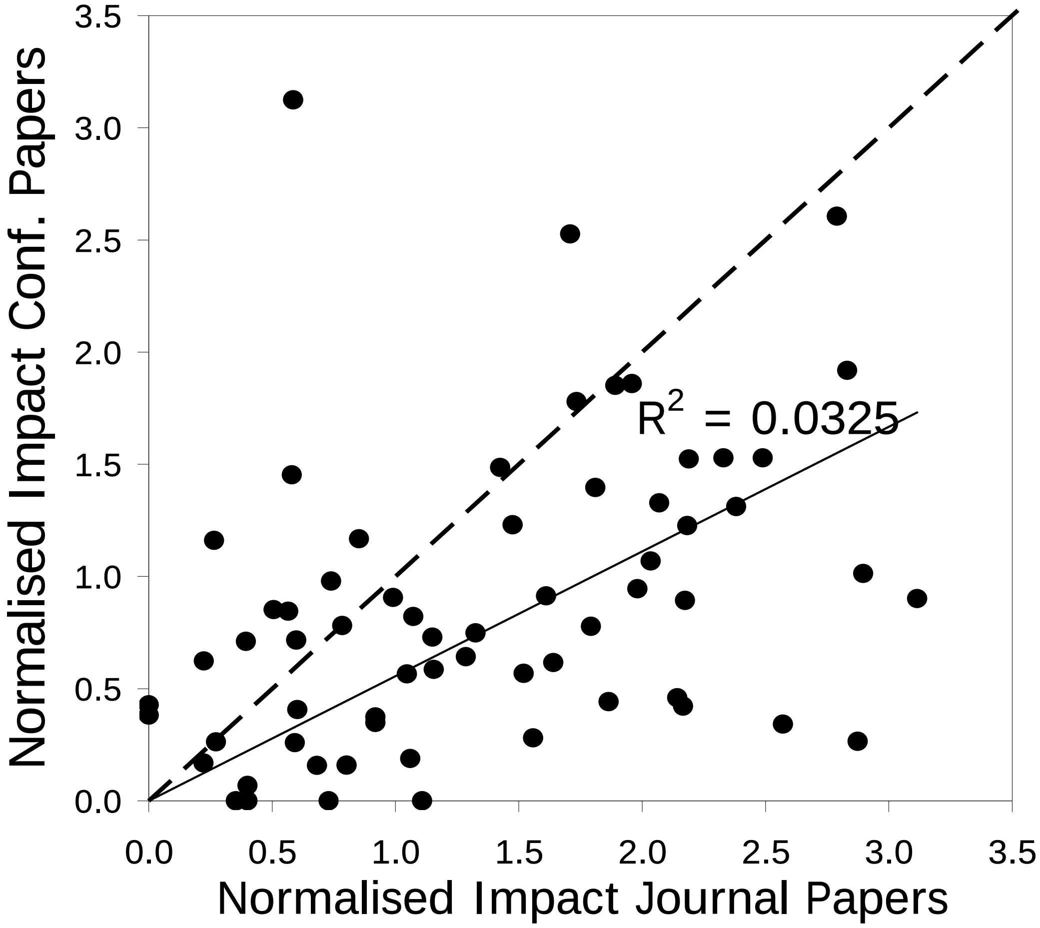

8 Outcomes of the Verification Round

8.1 Introduction

During the course of the project it was decided to carry out a verification round. The objectives of the verification process were twofold:
(i). To give the Netherlands computer scientists involved in the study an insight into the data on their publication output and citation impact that were collected in the study.
(ii). To enable the Netherlands computer scientists involved in the study to assess the accuracy of the data, and particularly the adequacy of coverage of the database that was used in a citation analysis of their publications, denoted as the *Expanded WoS* database

A web site was created, enabling a researcher to call up the list of publications from his or her group, displaying for each paper the following information:
a) Whether the paper was included in the selection of journals and conference proceedings covered by the *Expanded WoS* database.
b) For each paper included in this expanded database, the citation score accumulating in 2004, as well as a list of all articles in the *Expanded WoS* database that cited the paper.

An email letter was sent to 67 group leaders inviting them to participate. Four specific questions were addressed:
1. Are there striking, major inaccuracies in the publication and citation data (as far as one can see)?
2. How good is the coverage of the *Expanded WoS* for the subfield in which a group is active?
3. Which important sources (journals, conference proceedings volumes) are missing in this database, and why are they important?
4. Are there any frequently cited publications from your group published in other sources than the ones that are included in the *Expanded WoS* database, and from which source (e.g. Citeseer, Google Scholar) do these data originate?

The verification website indicated an email address to which replies could be sent. Through this address we received reactions of 27 researchers. Hence, the response rate was about 40 per cent. The respondents did not merely address more technical points related to accuracy of data collection and adequacy of coverage, but also issues of a more general nature. The comments are summarised below, grouped into 10 main issues. The summary focuses on more general aspects, and does not give details. The replies of the authors of this report are given in *italic*.

8.2 Summary of the comments

***1. Completeness, data quality and policy significance of publication data collected for a group**.*

As outlined in Section 2.3, the publication data of the university groups were obtained from the information provided by the groups themselves within the framework of the QANU and related reviews. The time period covered by these publications is 1996-2001. All universities involved formally agreed that these publication data were (re-)used in the current study. In this way, no additional collection of publication data was necessary. CWI publications were obtained from the CWI annual research reports. Since the arrangement of research activities into groups and departments at CWI changed radically in 1997, it was decided to collect publications with respect to the time period 1997-2004.

Many reactions addressed or criticised the completeness, data quality and policy significance of the publication lists collected in this way. It was noted that important publications were not included in the lists, and that some publications were listed twice. In addition, it was emphasised that for several groups the lists were not representative for the recent performance of the group, as some researchers had left a group, and new researchers entered it during the time period from 2002 up to date.

It is a significant, though somewhat unexpected outcome of the study that so many researchers criticised the publication data collected from the QANU review and institutional annual reports. Papers not listed in these sources could not be taken into account. As outlined in Section 2.3, it was decided to use these data sources in order to avoid that that the groups involved had to be requested to deliver their publication lists while not long before most of them had already provided this information for the QANU review. Given the methodological, experimental nature of the study, and in view of the foreseen use of the outcomes of the QANU review for validation purpose, the decision to use these data sources is in the view of the authors acceptable. It needs emphasising that in the calculation of indicators, papers listed twice were counted only once. It has become clear that for several groups the publication output during 1996-2001 does not give an adequate picture of the performance of the members currently active in the group. It follows that for those groups the relevance in current policy of the bibliometric outcomes obtained in this study is rather limited.

2. Partial coverage of sources included in the Expanded WoS database.

It was noted and criticised that several journals and proceedings series were only partly covered by the *Expanded WoS* database. A typical example is the *Journal of Functional Programming*. Not all papers published in this journal during 1996-2001 are included in the *Expanded WoS* database. The *WoS* only covers papers from a limited number of years, not from all years. Therefore, not all papers a group published in this journal may have been included in the *Expanded WoS* database.

This does not reflect errors in the process of matching NL-CS papers to articles in the WoS Expanded, but rather partial coverage of the database. A similar comment can be made for papers in LNCS proceedings that were not available in electronic form. This problem can in principle be overcome, but its solution involves a substantial amount of additional work.

3. Discrepancies between citation counts from the WoS Expanded database and those derived from the internet version of the WoS

Several researchers noted for specific papers discrepancies between the citation counts extracted from the *Expanded WoS* database on the one hand, and those obtained from a manual citation search in the internet version of the Web of Science, included in Thomson Scientific/ISI's Web of Knowledge. Generally, the latter scores were higher than the former.

The citation counts generated in this study relate to the time period 1996-2004, whereas to the best of our knowledge the manual searchers by respondents counted all citations up to date, thus including citations received in the years 2005 and 2006. This difference in the length of the time period covered in the citation analysis is an important source of discrepancies. It should be noted that the counts generated in this study include citations from ACM, LNCS and IEEE proceedings that were added to the WoS.

4. Limitation to research articles, database years

Several respondents noted and criticised that certain types of papers were not included in the citation analysis and qualified as 'out of scope'. They also criticised that monographs, book chapters and doctoral theses were not included.

The citation analysis carried out in this study considered as target articles only research articles, reporting outcomes of original research, either in journals or in proceedings volumes. Letters to the editor, editorial materials and bibliographies were not taken into account, even if they (particularly bibliographies) were cited rather frequently. Omitting such papers as targets from the analysis does not reflect an error in a strict sense, but is the consequence of an assumption made in the study. The study assumes that the contributions a group makes to the advancement of knowledge at the research front are primarily embodied in research articles. In a bibliometric assessment of these contributions, it is therefore appropriate to take into account research articles only. It does not follow that other types of papers, and especially books, are unimportant or of no value, but their function in the communication process tends to be different from that of research articles. It should be noted, however, that several respondents, especially from the subfields ***logic*** *and* ***information and knowledge systems*** *(particularly* ***artificial intelligence****) questioned the validity of this assumption in these subfields. They argued that especially books may represent important, original contributions to the research front in these subfields, and that work reported in doctoral theses is not always published later in journal or proceedings articles.*

Another issue of a more technical nature concerns papers that were not included due to a discrepancy between the publication year and database year. In the current study we collected publication data that were ***processed*** *for the WoS database during the time period 1996 – 2001. Therefore, papers that formally have been published during 1996 – 2001 may not have been included if they were processed in 2002 or later. This can be due to a time lag in processing journal issues for the Citation Index but may also occur if a journal issue appears after its cover date. Normally, this only relates to a few papers published at the end of the time period considered. However, in the current study it was found to be a more substantial matter, pertaining to 133 papers in the NL-CS database. This was mainly caused by a backlog in processing several LNCS volumes, possibly related with the policy change at Thomson Scientific regarding the coverage of the LNCS series. It would be advisable to take this finding into account in a possible follow up study.*

5. Inaccurate data handing at CWTS

Many researchers claimed to have found severe errors, resulting from inaccurate data collection and data handling at CWTS.

As indicated above, many types of 'errors' or omissions are consequences of theoretical assumptions underlying the study, or result from special properties or peculiarities in the data sources used. Among the comments made by the researchers in the verification round, we detected two cases that are (most probably) due to inaccurate data handling at CWTS. In the first case, a paper published by a particular group was linked to a different article in the Expanded WoS database with similar meta data. In fact the two papers were published in the same year, had the same volume number and starting page number, and the first characters of the names of the first authors were identical as well. The Expanded WoS paper to which the paper was linked erroneously, was cited rather frequently, and substantially raised the citation counts of the group to which it was allocated. A second case relates to the manual data collection of publication lists from annual research reports. For one group and one particular year all papers were missed. The authors of this report do apologize for the errors that were made. However, it needs emphasising that in the standard data collection and verification processes carried out by CWTS these errors would have been detected and corrected.

6. The problem of multiple positions in different institutions

Some researchers raised the issue how the data sources used in the study (from the QANU review and CWI annual research reports) dealt with the phenomenon that a researcher may have a position at more than one institution, especially at a university and at CWI. They had the impression that several university groups had listed papers emerging from CWI, in view of the double position of the group leader.

A secondary analysis revealed that for several university groups a substantial number of papers listed in the QANU review were also included in CWI annual research reports.

This outcome indicates that the issue of double positions and double counting of research articles deserves special attention in a follow up study (if there will be one).

7. Comparisons with citation data from Citeseer and Google Scholar

Several groups verified the publication and especially the citation data in a most rigorous way. They listed citations to their papers given in important sources not included in the Expanded WoS database, and indicated even the names of these citing sources. Other respondents compared the numbers of citations from the Expanded WoS database with citation counts obtained from Citeseer or Google Scholar, and often detected discrepancies, that in several cases were attributed to errors in the Expanded WoS citation counts.

The authors of this report are most grateful to those who have so thoroughly verified the citation data. Their comments and additions are currently analysed in detail, and are most useful. Differences in citation counts among Web of Science, Citeseer, Google Scholar and Elsevier's Scopus constitute an important topic of empirical research, and the respondents' comments provide more insight into this topic. The coverage of the four databases differs from one another, and these differences can be expected to explain at least partly observed discrepancies between citation counts extracted from them. On the one hand, large discrepancies provide a good reason for checking data more carefully. But a general problem is that it is not always clear which sources are covered by the various databases. This is especially, though not exclusively, true for Google Scholar. In our view it is in the measurement of the impact of publications at the research front appropriate to take into account only citations in research articles published in peer-reviewed sources and reporting original research findings. To the best of our knowledge it is unclear to what extent the sources covered by Google Scholar meet this criterion. For instance, it is not unlikely that Google Scholar covers substantial numbers of theses of master students. Although citations in such sources are significant, it is in our view questionable whether they should be taken into account in an analysis of citation impact at the international forefront in the field. It is appealing to search for citations in such a large universe of sources, and the fact that it is not precisely known which sources are covered makes citation searches even more exciting, as one can browse through the citing documents one by one. But if one merely counts citations, without evaluating the sources giving them, the problem arises: what does one count, or what do the counts reflect? A basic notion underlying the citation analysis undertaken in this study is that it should be carried out in a 'transparent' universe of sources containing contributions that meet the standards of originality, methodological soundness and significance of outcomes.

8. Important sources missing

Many respondents indicated sources (journals and particularly proceedings volumes) that were not included in the *Expanded WoS* database but were nevertheless important enough to be added to the database. They stated that both the publications in these sources and the cited references therein should be added to the database, in order to give more complete account of their group's performance, and particularly its citation impact.

Table 8.1 *gives per subfield a list of the sources (conference proceedings or journals) that one or more respondents qualified as important for their (sub-) field, and that were not covered or partially covered by the Expanded WoS database. Partial coverage means that only a part of a source's papers are included in the database, for instance, those published in specific years or volumes. Since the creation of this database and its adequacy of coverage constitute the core of the study, this table will be discussed in detail in* ***Section 9.3****.*

Table 8.1: Journals and proceedings qualified by respondents as important but not covered or partially covered by the Expanded WoS (=WoS+ACM+LNCS+IEEE) database

Subfield and type of source	*Coverage*		*Average citation impact of NL papers according to Table 5.4*
Computing and Imaging			
- Journals			
Computing Supplement	no		
Fundamenta Informaticae	partially	as from 2001	
- Proceedings			
Advances In Discrete And Computational Geometry, Contemporary Mathematics	no		
Eurographics	no		0.73
European Signal Processing Conference (EUSIPCO)	no		
Mathematical Morphology and its Applications to Image and Signal Processing (ISMM'96)	no	ISSM covered as from 1998	0.60
Proceedings Joint Eurographics-IEEE TCVG Symposium Visualisation (VisSym)	partially	2002, 2003	0.64
- Other			
CRC Handbook Of Discrete And Computational Geometry	no		
Programming Research and Algorithmics			
- Journals	***Coverage***		
ACM Transactions on Computational Logic	no		
Expositiones Mathematicae	partially	as from 2003	
Formal Aspects of Computing	partially	as from 2004	
Journal of Functional Programming	partially	as from 2001	
Logical Methods in Computer Science	no		
Mathematical Structures in Computer Science	partially	as from 2005	
Nordic Journal of Computing	no		
- Proceedings			
Electronic Workshops in Computing (eWiC)	no		
Haskell Workshop proceedings	partially	as from 2002	
International Colloquium on Automata, Languages and Programming (ICALP)	partially	1999	
International Conference on Concurrency Theory (CONCUR)	partially	2000-2002	
International Workshop on Evaluation of Modeling Methods in Systems Analysis and Design (EMMSAD)	no		
Logic Programming, Intl Conference	partially	as from 2001	
Operating Systems Design and Implementation (OSDI)	partially	up to 1999	
Proceedings Summer Computer Simulation Conference (SCSC)	no		
The logic programming paradigm	no		
USENIX Annual Technical Conference	no		

Subfield and type of source	*Coverage*		*Average citation impact of NL papers according to Table 5.4*
USENIX Symposium on Internet Technologies and Systems (USITS)	no		
- Other			
Electronic Notes in Theoretical Computer Science (ENTCS)	no		
Handbook Of Automated Reasoning	no		
Handbook on Graph grammars and computing by graph transformations	no		
Images of SMC research	no		
Kreiseliana, about and around Georg Kreisel	no		

Information and Knowledge Systems

- Journals	***Coverage***		
ACM Sigmod Record	partially	as from 2000	
ACM Transactions on Computational Logic	no		
ACM Transactions on Internet Technology	no		
Bioinformatics	partially	as from 1998	
Biostatistics	partially	as from 2002	
BMC Bioinformatics	partially	as from 2002	
BMC Biotechnology	partially	as from 2002	
BMC Genomics	partially	as from 2002	
Briefings in Bioinformatics	partially	as from 2004	
Computational Economics	no		
Formal Aspects of Computing	partially	as from 2004	
Genetic Programming And Evolvable Machines	no		
Genome Biology	partially	as from 2003	
IEEE Transactions on Computational Biology and Bioinformatics	no		
IEEE Transactions on Intelligent Transportation Systems	partially	as from 2001 (Vol.2)	
Journal of Autonomous Agents and Multi Agent Systems	partially	as from 2000	
Journal of Biomedical Informatics	partially	as from 2001	
Journal of Universal Computer Science	partially	as from 2001	
Logical Methods in Computer Science	no		
Mathematical Structures in Computer Science	partially	as from 2005	
Nordic Journal of Computing	no		
OMICS: A Journal of Integrative Biology	partially	as from 2004	
PLoS Computational Biology	no	started in 2005	
PLoS Genetics	no	started in 2005	
The New Review of Hypermedia and Multimedia	no		
World Wide Web Journal	no		
- Proceedings			
European Conference on Artificial Intelligence (ECAI)	no		0.40
European Conference on Computational Biology	no		
Intelligent Systems in Molecular Biology	no		

Subfield and type of source	*Coverage*		*Average citation impact of NL papers according to Table 5.4*
International Colloquium on Automata, Languages and Programming (ICALP)	partially	1999	
International Conference on Concurrency Theory (CONCUR)	partially	2000-2002	
International conference on REsearch in COmputational Molecular Biology	partially	missing: 2001, 2002	
International Joint Conferences on Artificial Intelligence (IJCAI)	no		1.58
National Conference on Artificial Intelligence (AAAI).	no		
Pacific Symposium on Bioinformatics	no		
Pattern Recognition in Information System (PRIS)	no		
Proceedings of Genetic and Evolutionary Computation Conference (GECCO)	no		0.88
Structural, Syntactic, and Statistical Pattern Recognition (SSPR)	partially	2002	
World Wide Web Conference	partially	as from 2001	
- Other			
Advances in Intelligent Data Analysis	partially	2001	
Advances in Neural Information Processing Systems	no		
Advances in Soft Computing	partially	2002	
Springer Lecture Notes on Economical and Mathematical Systems (LNEMS)	partially		

Logic

- Journals	*Coverage*	
Nordic Journal of Computing	no	
ACM Transactions on Computational Logic	no	
Argumentation	no	
Artificial Intelligence & Law	no	
Cognitive Science Quarterly	no	
Formal Aspects of Computing	partially	as from 2004
International Journal of Electronic Commerce	partially	as from 2000
Journal of Autonomous Agents and Multi Agent Systems	partially	as from 2000
Journal of Logic, Language and Information	no	
Logical Methods in Computer Science	no	
Mathematical Structures in Computer Science	partially	as from 2005
Studia Logica	no	
- Proceedings		
International Colloquium on Automata, Languages and Programming (ICALP)	partially	1999
International Conference on Concurrency Theory (CONCUR)	partially	2000-2002

9. Differences among subfields of Computer Science

Many respondents underlined differences in publication practices among subfields of Computer Science. Especially respondents from the subfields *logic* and *artificial intelligence* stated that the publication practices in their subfields are different from those in for instance *computing and imaging* or *programming research and algorithmics*.

An appropriate evaluation system of scientific-scholarly research quality has to take into account differences among (sub)fields of science and scholarship. In fact, the principal aim of the current study was to develop a methodology that takes into account specific properties of scientific communication in Computer Science – particularly the important role of conference proceedings. The new methodology differs from that normally applied in physics, chemistry or biometrical sciences. On the other hand, differences among subfields within Computer Science have not received full attention in the current study. The subfield logic is a highly multi-disciplinary field, that can be located on the interface between exact or technical sciences on the one hand, and humanities on the other. This subfield certainly has characteristics of a humanities field, and a follow-up study could focus on special indicators in this subfield.

10. Use of the bibliometric outcomes in the policy domain

Several respondents underlined that the outcomes of the bibliometric study at the level of individual groups or even universities should not be used in the policy domain.

The study presented in this report was methodological and exploratory. The report does not present outcomes of individual groups, and not even of individual universities, other than in a purely anonymous or aggregate way.

9 Discussion and Conclusions

9.1 Data collection (Chapter 2)

The conclusion from Chapter 2 is that expanding the *WoS* database with conference proceedings sources is technically feasible, provided that meta-data (including cited reference lists) are available in electronic form, but it involves a lot of elementary data processing. The amount of work depends upon the nature and quality of the relevant meta-data (including cited references) of articles from these sources. If the meta-data have to be extracted from PDF documents, the process of data collection can be qualified as cumbersome.

In this process, a part of the relevant data is lost. In our study, it was estimated that from about 17 per cent of source articles in PDF format, the cited reference lists were not extracted at all, due to the fact that the text was not or only partially extracted, no reference list section was identified, or no reference separator was recognised. At this moment we do not see how these problems can be easily solved. But it could be useful to collect in a follow up study even more expert knowledge on these processes of data extraction, and examine whether their recall can be enhanced.

As outlined in Section 2.1, this problem of missed cited reference lists is much smaller for papers from 2004 than it is for those published in the beginning years of the time period considered. Moreover, in the citation analysis carried out in this study recent papers make a larger contribution to the citation counts than earlier papers. Whether or not these missed citations substantially affect the accuracy of the analyses based on the extracted papers, depends upon the type of analysis carried out and the research question addressed.

For instance, if one aims to measure the strength of citation links among proceedings volumes, the outcomes for those from which a large part of cited references is missing may be inaccurate. But it is unlikely that the missed cited references substantially affect the citation analysis of all target articles published by a research group or institution, carried out in a citation universe containing hundreds of journal and conference proceedings volumes.

Chapter 2.2 discussed the problem of accuracy of citation matching and citation counts. It is concluded that more work needs to be done in order to tackle this problem and thus further increase the accuracy of citation counts. In principle there are two lines along which one could proceed.

1. *A 'classical' approach*. The methodology applied in this report is an extension of the one developed at *CWTS* in the analysis of the *Web of Science*, which covers mainly – though not exclusively – scientific/scholarly *journals*. The basic unit in the citation

analysis carried out in this database is the individual (journal) paper. Cited references are linked to target articles on a paper-by-paper basis. Citation rates of the papers of a particular group are compared to the average citation rates of all articles published worldwide in the subfields covered by that group. Based on this idea, the collection of sources covered by the database could be expanded even further, and the citation matching algorithms could be further developed. Perhaps this approach can be denoted as the 'classical' approach.

2. *A new approach*. But there is also a more fundamental, conceptional problem at stake. A basic issue is the extent to which the various versions of publications have indeed the same contents and thus represent the same paper. It can perhaps be assumed that a series of papers from the same authors with identical or similar titles all relate to a particular *concept* – for instance, an idea, or a methodology. During the process of its development, published papers may all bear this concept in their titles, but their contents may have changed, reflecting the various maturing stages of the concept. In an alternative approach, the basic unit of analysis would be a concept, – embodied in a *series* of publications – rather than an *individual* publication. This concept should be linked in some way with authors and research groups. A statistic as a citation per publication ratio that plays a key role in the 'classical' approach would have less significance than in this new, concept-based approach.

It must be noted that this approach would not only be relevant for studies in the field Computer Science, but in many other fields as well, in view of the increasing importance of electronic publishing in general, and particularly the creation of preprint archives and document repositories.

Evaluated researchers justly demand that citation counts of their publications are accurate, especially since citation distributions are highly skewed. It appears to be extremely difficult to determine in a fully automated way accurate citation counts to publications that are not published in journals or in well formatted and structured conference proceedings, or, if papers are published in different versions (e.g., as technical report, proceedings article, book chapter and journal paper), to allocate citations to the intentionally cited version of such papers. The authors of this report therefore recommend that a citation analysis of these types of papers published by researchers under evaluation is carried out – or at least checked – *manually*.

9.2 The importance of conference proceedings (Chapter 4)

The results presented in Chapter 4 underline the importance of conference proceedings, – particularly *ACM*, *LNCS* and *IEEE* proceedings volumes –, in terms of the citation impact these volumes generated both upon proceedings and upon journals. The *Source Citation Rate*, an 'impact factor-like' citation impact measure, is for the proceedings volumes analysed in this study on average somewhat higher than for annual journal volumes.

Journal articles give about two out of three of their references to other journal articles, and one out of three to proceedings papers. For proceedings papers, it is the other way around. They give about two out of three references to other proceedings articles, and one out of three to journals. In other words, each source type (journal or proceedings) tends to show the same preference ratio for citing papers of the own type.

It was found that in the sub-universe analysed in Chapter 4, papers in proceedings volumes contain on average more linked references to other articles than journal papers do. As a result, at the cited side, proceedings papers receive a higher proportion of their impact from other proceedings papers than journal articles gain from other journal articles: 73 versus 54 per cent.

But one should keep in mind that these outcomes relate to a sub-universe of articles in Computer Science journals and proceedings volumes published during 1996–2004. References given to articles published before 1996 are not included in this universe. It should also be noted that the values presented for citations and references are averages over a range of available years. It should also be taken into account that the age of citing conference papers is on average younger than that of citing journal articles. As a result, cited references in conference papers have a higher chance of being linked within the sub-universe than cited references in journals have. However, it can be shown that journal articles contain on average a higher number of cited references than proceedings papers do. On the other hand, cited references in the latter tend to be published more recently than those in the former (Visser and Moed, 2005).

Proceedings volumes tend to show a somewhat higher variability in their citation impact rates than annual journal volumes do. In Computer Science, both in the top and the bottom of the distribution of citation impact among sources, proceedings volumes are somewhat overrepresented. In other words, there are relatively more highly cited *and also* more poorly cited proceedings volumes than there are annual journal volumes.

The citation links among proceedings volumes of recurring (e.g., annual) conferences tend to be as strong as those among annual volumes of the same journal. These proceedings series reveal citation patterns that are statistically similar to those shown by a journal's annual volumes, if not stronger. It should be noted that the outcomes presented in this chapter represent a first analysis of the citation patterns in Computer Science. Follow-up studies could carry out more detailed analyses of these patterns, including the creation of maps of sources on the basis of their citation links.

These findings corroborate outcomes from earlier studies and claims made by computer scientists as regards the importance of conference proceedings as channels of written communication in their field. Using citation impact of a publication source as an indicator of its importance, it follows that in the *WoS* database a number of important conference proceedings volumes is missing. These volumes tend to be as important as the journals that are covered by this database. For proceedings of recurring conferences, successive volumes tend to be as important for one another as successive annual volumes of journals are for one another. Their inclusion in the *Expanded WoS* enhances the coverage of the important channels of written communication, and therefore provides a more accurate and a more valid bibliometric assessment of research performance in Computer Science.

9.3 The coverage of the Expanded WoS database (Chapters 3, 5 and 8)

Adequacy of coverage of the publication database is a crucial issue. The fraction of papers of Netherlands computer scientists published during the time period 1996-2001 in journals and proceedings included by the *Expanded CWTS-WoS* database – termed as *external* coverage – was found to be 35 per cent. In other words, slightly more than one out of three Netherlands publications was published in these sources. For the '*Pure WoS*' database, excluding the proceedings volumes that were added, this coverage percentage amounts to 25. These percentages were calculated relative to the total number of publications submitted to the QANU/VSNU review or listed in CWI annual research reports, excluding research reports. It is an overall statistic, related to the total time period 1996-2001, and to groups from all subfields of Computer Science.

The overall coverage of the *Expanded WoS* database for Netherlands papers increased from 22 per cent in 1996 to 41 per cent in 2001. The major part of the papers from ACM, LNCS and IEEE sources included in the *Expanded WoS* database are from later years. Therefore, the coverage rates of Netherlands Computer Science papers can be expected to further increase during de time period 2002-2006.

At the level of subfields, and for Netherlands papers published in 2001, the *Expanded WoS* coverage was highest in *Computing and Imaging* (53%) and *Programming Research and Algorithmics* (41%). For *Information and Knowledge Systems*, it was 32%, and for *Logic* and *Telematics* 36 and 24 per cent, respectively. *Mathematics*, covering the more mathematically oriented groups, showed a coverage rate of 53%.

An examination of the database's internal coverage, based on an analysis of cited references in papers included in the database revealed the same differences among subfields. It also showed differences among groups: 25 per cent of groups had a publication coverage percentage above 46 per cent, and 25 per cent below 28 per cent. These differences can be expected to be highly subfield-dependent.

It must be concluded that, despite the enormous effort that was made to expand the Web of Science with conference proceedings from three important sources (ACM, LNCS and IEEE), substantial differences exist in the coverage of the *Expanded WoS* database among subfields of Computer Science. The coverage of the Netherlands publication output in *Computing and Imaging*, *Programming Research and Algorithmics* and *Mathematics* is higher than that for *Information and Knowledge Systems*, *Logic* and *Telematics*.

Even though external coverage percentages are most relevant and informative, they depend upon what researchers include in their publication lists. For instance, evidence was obtained that, although some groups listed all their articles in the ASCI proceedings, other groups decided not to do so, since they considered their articles in these proceedings of limited significance. A similar statement can be made for the BNAIC proceedings. In other words, coverage percentages depend upon the degree of self selection researchers have imposed themselves when compiling their publication lists. This observation also explains why it is difficult to compare coverage percentages of research output in different research fields. For instance, researchers in chemistry do publish articles in conference proceedings or in books, but they often do not even include these articles in their publication lists submitted in an evaluation process.

Therefore, it is also relevant and informative to examine one by one the sources in which Netherlands computer scientists published their articles and that were *not* covered by the *Expanded WoS* database. Table 5.4 in Chapter 5 presented a list of the most frequently used (recurring) conference proceedings, as well as the average citation impact of the papers published therein. An evaluation of journals or conference proceedings requires a detailed knowledge of the fields they cover, and their communication networks. The authors of this report do not have such knowledge. On the other hand, citation impact measures do provide useful tools to carry out such an evaluation.

It was found that the Netherlands papers in the *Proceedings of the Conference of the Advanced School for Computing and Imaging (ASCI)* and especially in the *Proceedings of the Belgium-Netherlands Conference on Artificial Intelligence (BNAIC)* have rather low citation impact. As outlined in Chapter 5, citations given in papers in these proceedings to other papers in the same proceedings were not included in the counts. If one assumes that such citations account for around 20 per cent of citations to these proceedings, the citation impact of the Netherlands papers in these proceedings are still rather low.

But several other proceedings listed in Table 5.4 have a substantial citation impact, and some of these were also qualified as important in the verification round. Conference proceedings in which the Netherlands papers have a relatively high impact but that were not included in the *Expanded WoS* database are (for a complete list of important sources mentioned in the verification round the reader is referred to Table 8.1 in Chapter 8):

Conference of the IEEE Communications Society (INFOCOM)
Conference on Parallel Problem Solving from Nature (PPSN)
Eurographics incl. Various Workshops
European Conference on Artificial Intelligence (ECAI)
European Conference on Software Maintenance and Reengineering (CSMR)
Genetic and Evolutionary Computation Conference (GECCO)
IEEE International Conference on Robotics and Automation (ICRA)
International Joint Conference on Artificial Intelligence (IJCAI)
International Conference on Very Large Databases (VLDB)
Text Retrieval Conference (TREC)
Proceedings of Recherche d’Information Assistée par Ordinateur (RIAO)

It is interesting to note that respondents in the *verification round* did not only indicate important conference proceedings to be added to the database, but also a number of scientific-scholarly journals that are not (yet) covered by the Web of Science. But it needs emphasizing that many journals reported as missing in the *Expanded WoS* database appear to be fully covered by the *WoS in recent years*. Apparently, the *WoS* coverage of the journal literature in Computer Science has increased in recent years.

A second observation is that many sources respondents qualified as important and missing, and listed in Table 8.1, are *not* in the list presented in Table 5.4 of sources not covered by the *Expanded WoS* database and most frequently used by Netherlands computer scientists. The standardization of source titles underlying this table has not been perfect, and possibly some conference proceedings would appear on the list after an additional de-duplication process. In addition, the publication lists related to the time period 1996-2001, while respondents may have indicated sources that were founded in recent years, or sources in which they started publishing after 2001. But even if this were so, there seem to be major discrepancies between the two lists of sources.

Section 3.2 presented empirical data on internal coverage percentages of the *WoS* in all domains of scholarship, based on the extent to which (source) articles included in a publication database cite other (source) articles in that database. In addition, it distinguished four types of bibliometric studies. The principal criterion in deciding which type of study is the most appropriate in a particular field is based on the internal coverage of the publication database for that field. As explained in Section 3.2, in a citation impact analysis one should distinguish a citing or source side and a cited or target side. Target articles are those that are subjected to a citation analysis. Source articles are documents from which cited references are extracted. The following tentative classification is proposed:

- In fields with an excellent *WoS* coverage, for which the internal coverage percentage is typically between 80 and 100, it is generally sufficient in a citation impact analysis to take into account as target articles only those that are published in *WoS* source journals, and to use the total collection of cited references in *WoS* source journals as citation universe. This type of analysis is labelled above as the 'Pure' *WoS* analysis or Type I study.

- If *WoS* coverage in a field is not excellent, but can nevertheless be qualified as good, with internal coverage percentages typically between 60 and 80, the scheme suggests to expand the collection of target articles analyzed in the 'pure' *WoS* analysis by including target articles that are not published in *WoS* source journals (a ***target expanded*** citation analysis or Type 2 study).

- If *WoS* coverage of a field is moderate, with internal coverage percentages typically between 40 and 60, it is proposed to expand the universe of citing sources with articles in proceedings volumes from a range of subsequent years. Such an approach is labelled as a ***source expanded*** citation analysis (Type 3 study).

- Finally, if *WoS* coverage in a field is poor, showing internal coverage percentages below 40, it is questionable whether it is useful conducting a citation analysis based on *WoS* data, even if target or source universes are expanded (Type 4 study).

In an analysis presented in Section 5.3 of the internal coverage of the 'Pure WoS' database of the field computer science, an internal coverage percentage of almost 38 per cent was found. This outcome constituted the principle reason to carry out a source expanded (or type III) study, and to expand the *WoS* with a number of proceedings volumes of important international conferences. Expansion of the database with ACM, IEEE and additional LNCS volumes not included in the 'Pure' *WoS* raised the internal coverage to 51 per cent in 2004. Although the increase in the internal coverage is substantial, it should be noted that, according to the tentative scheme presented above, it is still moderate.

Therefore, it is recommended to *further expand* the publication database with a number of important conference proceedings volumes. Suggestions for sources to be included were given above. As outlined above, one needs to have detailed knowledge of a (sub-) field in order to identify its principal sources. The authors of this report recommend that experts in the various subfields compile a list of the most important journals of conference proceedings that should be added to the *Expanded WoS* database. The citation impact rates of sources and their importance as perceived by respondents in the verification round would constitute valuable sources of information in this process. But it needs emphasizing that, from a practical point of view, the availability of sources in electronic format should also be an important criterion.

It is unlikely that the expansion of the database with these additional proceedings volumes would raise the internal coverage to a level above 80 per cent. Assuming realistically that it will be below 80 per cent, our tentative classification scheme of types of bibliometric studies proposes to carry out a Type II or target expanded citation analysis. Therefore we recommend to carry out *within a further expanded WoS database* such a target expanded citation analysis, in which in principle *all* research articles published by a group are taken into account, not only those that are published in journals or proceedings included in the database. In this way, citations to articles published in journals or proceedings that are not covered by the database are also counted, and contribute to the citation impact rates of a group's publication output. The same is true for book chapters. Whether or not citations to monographs or technical reports should also be taken into account is an open question. From a pragmatic point of view, citations to this type of publications could be tabulated separately.

Evaluated researchers justly demand that citation counts of their publications are accurate, especially since citation distributions are highly skewed. Missing all citations to one highly cited article may seriously affect the values of citation impact indicators. The experiences with citation matching collected in this study have shown that it is extremely difficult to determine in a fully automated way accurate citation counts to individual publications that are not published in journals or in well formatted and structured conference proceedings. In addition, if papers are published in different versions (e.g., as technical report, proceedings article, book chapter and journal paper), it is extremely difficult to allocate merely by means of computer programs citations to the intentionally cited version of such papers. It is therefore recommended that a citation analysis of these types of papers published by researchers under evaluation is carried out – or at least checked – *manually*. The insights obtained in these manual searches and checks can be used to further improve the computerized citation matching algorithms.

As outlined in Section 3.3, the indicators calculated in a Type 1 study – regardless of whether the database is expanded or not – allow one to compare from an international viewpoint the citation impact of a group's publication output with that of other groups active in the same subfields, particularly with the world citation average in the subfields in which a group is active. This indicator is important, since in a bibliometric assessment of Netherlands research groups one is not merely interested in comparisons among Netherlands groups, but also – if not primarily – in comparing the performance of Netherlands researchers with that of their international peers. It needs emphasising that in the target expanded (Type 2) citation analysis recommended above, there is as of yet *no* 'standard' methodology available that allows for such international comparisons, although several suggestions were made in earlier studies. In a follow up study, further research into the development of such a methodology is necessary.

The web application created for the verification round (see Chapter 8) only included citations to articles published in journals or proceedings included in the Expanded WoS database. It thus presented the raw counts of a Type I citation analysis, but within a citation universe of an *expanded* WoS database. The analyses presented in Chapters 6 and 7 of this report are also based on this type of citation analysis. The conclusions reached in these chapters are 'tentative'.

Finally, we make a comment on the process of identifying important journals or conference proceedings outlined above. Participants in this process could address the following issues. The lists of conference proceedings in which Netherlands computer scientists published reveal that there is a huge scattering of published articles among sources. Is this a basic characteristic of the field (or a specific subfield), or does it perhaps at least partly indicate that many groups in Netherlands academic computer scientists have not (yet) developed a well-considered strategy as regards the participation at international conferences and the selection of publication sources?

The authors of this report wish to emphasize that they do not claim that BNAIC conferences and similar national conferences have no value or are in all respects unimportant. On the contrary, such national conferences and their proceedings may play a most important role in creating and maintaining national networks of researchers. However, from the perspective of assessing the contribution these researchers make at the international research front, – the central perspective in this study and this report –, they seem to be less important.

In each field, researchers are more and more stimulated – if not forced – to create and apply quality standards. One way to come to such standards could be to identify particular conferences and publication sources that one can normally only enter if the work presented is of high quality. If one has finished an excellent paper, presenting it at an AAAI conference rather than at the BNAIC, and publishing it in the AAAI proceedings rather than in the BNAIC proceedings, does not make the paper 'better', in the sense that its contents would be different. But one would in this way contribute to the maintenance of a quality control system, aimed at discriminating between research of the highest quality and less significant research, as papers of lower quality tend to be rejected for publication in AAAI proceedings. In addition, more selectivity in the selection of conferences and publication sources could increase the visibility and international standing of Netherlands academic Computer Science.

9.4 Comparison with QANU peer ratings (Chapter 6)

Chapter 6 compared the outcomes of peer reviews of the past research performance of research departments with bibliometric indicators of their publication output. The main objective of this analysis is to subject the bibliometric indicators to a first test, by comparing their values to peer ratings. The analysis does *not* aim at re-doing the QANU evaluation or suggesting that correlations should be perfect.

As outlined by Moed (2005) such a comparison can be made from two distinct points of view. The first is that of validation of bibliometric indicators. These indicators are tested using peer judgements as a benchmark. Thus, it is assumed that peer review provides a more direct measure of research quality, and to the extent the bibliometric indicators correlate with peer judgements, they are validated. A second point of view critically examines peer review processes, especially peer ratings of evaluation panels. In this analysis bibliometric indicators are applied as a benchmark. The basic assumption holds that bibliometric indicators are valid indicators of research performance, and thus can be used to assess peer judgements, and raise questions about how peers evaluated, which criteria they applied, and whether their judgements were biased.

In this discussion section *both* viewpoints will be further developed. The discussion is essentially open. It will not assume the primacy of one of the two methodologies to assess research performance – peer review or bibliometric analysis – above the other. Instead, it aims to show how the combination of the two methodologies may provide insight into the validity of both peer review and bibliometric indicators, by presenting a number of significant observations, and raising critical questions regarding the construction of bibliometric indicators, the publication and citation practices from which these are derived, and the ways peers evaluated the past performance of research groups.

First, it should be noted that the correlation analysis between computer science peer ratings and bibliometric indicators has a limited significance. There is little variation in the peer quality ratings (see below), and the correlation analysis in a sense magnifies the small differences that there are. The finding that the number of articles in the *WoS*-Expanded database is the indicator showing the highest correlation with the quality peer ratings of the Netherlands academic Computer Science groups can be interpreted as evidence that the extent to which groups published in refereed international journals and in important conference proceedings (*ACM, LNCS, IEEE*) has been an important criterion of research quality for the Review Committee.

This interpretation would also explain the apparent lack of correlation between the quality and productivity ratings made by the peers. The productivity rating takes into account the total publication output, including all types publications, whereas the quality rating is based on a screening of a complete lists of publications and assessment of selected top publications. Following this interpretation, the methodology developed in this study is on the right track. Any attempt to identify research quality or excellence should discriminate between truly important publications, and less significant ones.

In the field of Netherlands Chemistry, publishing in journals processed for the *WoS* is hardly discriminatory, since Netherlands chemists tend to publish the overwhelming part of their publication output in these journals. Therefore, it is not surprising that for Netherlands Chemistry groups a very weak correlation was found between the quality peer rating and total number of published articles.

In the Computer Science Review, the peer quality rating of a group shows a positive correlation with the total number of citations the group's publications received. It should be noted that the total number of citations received by a group depends strongly upon the number of papers it published – in fact, these two variables show a very strong correlation. Therefore, this outcome does not tell us much about the relationship between a group's peer rating and its citation impact. It is more informative to consider relative or normalised citation impact indicators. It was found that the peer quality rating revealed a weak, positive, correlation with the citation impact compared to the world average, and with the number of citations per FTE research time. In the Chemistry review, the first is somewhat higher.

From the perspective of validation of the bibliometric indicators calculated in this study, the observed positive correlation between peer quality rating and relative or normalised citation impact can be interpreted as empirical evidence of the validity of these types of indicators as measures of research quality in the field Computer Science. On the other hand, it can be argued that the correlations between quality ratings and relative/normalised citation impact indicators, though positive, are rather weak.

In order to discuss this outcome further, it seems fruitful to compare the ratings for Computer Science to those assigned in the Chemistry Review. It should be noted that the two peer reviews were carried out by different committees, and related to different fields in different phases of development. Computer Science is a relatively young and expanding field, whereas Chemistry is a classical field, subjected to severe retrenchments (particularly, though not exclusively, in the Netherlands). Moreover, in the review on Chemistry bibliometric, citation based indicators constituted one of the sources of information to the Committee, whereas in the review for Computer Science they did not.

Compared to the Chemistry groups, the quality ratings of the academic Computer Science groups show a higher level and less variability, whereas their citation impact tends to be lower, and, – taking the ratio of standard deviation and mean as a standard –, shows more variability. One can argue that it is remarkable that the quality ratings of the various academic Computer Science groups are so similar one to another, whereas their citation impact reveals so much variation, and focus on the peer review process in which the ratings were generated. Following this line of thought, this finding suggests that the Review Committee for Computer Science gave a lower weight to citation impact than the Review Committee for Chemistry did.

The authors of this report do *not* claim that the Review Committee for Computer Science (or any other review committee) *should* have based their judgments merely upon citation analysis, or that there should be a perfect or even a strong correlation between its ratings and citation impact. From the apparent weak correlation between peer ratings and relative or normalised citation impact indicators, one can *not* conclude that the peer ratings are basically invalid. But one could at least raise the question how the Review Committee for Computer Science evaluated aspects as ‘international recognition and innovative potential’, the extent to which the work of group is ‘at the forefront internationally’ and generates an ‘important and substantial impact in the field’. In the view of the authors of this report, citation analysis, when used properly, is a valuable, *additional tool* in the assessment of these aspects of research performance. It provides in a quantitative framework a condensed representation of citation patterns in an entire field’s literature from a range of years, and can be used as such to sharpen or even correct a peer’s own impression of an entity’s research quality.

The Review Committee for Computer Science did *not* use the outcomes of the citation analysis presented in this report. A further development and validation of the indicators computed in this study could take place in a future peer review of the field. Application – be it experimental – in a peer review context ensures that background knowledge about the groups to be evaluated and the subfields in which they are active is taken into account in the interpretation of bibliometric indicators, and in this way establishes necessary conditions for their proper use.

9.5 Discussion of preliminary results (Chapter 7)

In this study we applied a methodology that is most similar to the one we applied in many other studies of research performance. The principal difference is that in the current study the *WoS* database was expanded with *ACM*, *LNCS* and *IEEE* proceedings. The study is *exploratory*. The application of bibliometric indicators is *experimental*. Assuming that the publication database and methodology provide a valid reflection of the research performance of *Netherlands* academic Computer Science, which *tentative* conclusions could be drawn from the analyses presented in Chapter 7?

Focusing on the citation impact of the papers published during 1996–2001, the outcomes suggest that the impact of the Netherlands academic Computer Science is significantly above world average. Interestingly, during the time period considered, the impact increased substantially: the impact of articles published in the last years of the time period is higher than that of papers published in the beginning.

Section 2.2 reported that the number of *ACM*, *LNCS* and *IEEE* proceedings articles included in the *Expanded CWTS-WoS* database has also increased substantially. Is the increase in impact of the Netherlands papers perhaps a database artefact? We believe that this is not the case. On the one hand, the number of conference proceedings included in the database increased substantially during 1996-2004. Especially LNCS published in recent years many more volumes than they did during the 1990s. The citation per article ratio for Netherlands papers did increase under the influence of the database expansion, but the same is true for the world citation average. The relative citation impact indicator (citation impact compared to the world average) corrects for the observed increase in database coverage. It is also noteworthy that the increase in normalised citation impact is also visible in an analysis of papers in *journals* covered by the Web of Science, and for articles from each major *subfield*. We are therefore inclined to conclude that the outcomes indicate a genuine increase in the average citation impact of Netherlands Computer Science papers.

We found an *overall citation impact* ratio of 1.30 (presented in Table 7.1) and an increase of this ratio to a level of 1.41 and 1.56 for papers published in the last two years (Figure 7.2). The bibliometric study on academic Chemistry (*VSNU*, 2002) reported an overall impact of Netherlands groups in this field of 1.45. Although this outcome related to articles published during a longer, somewhat earlier time period (1991-2000), the findings presented in this report suggest that the citation impact of Netherlands academic Computer Science has the same level of that of Netherlands academic research activities in other disciplines from the exact sciences. But it was also found that among the top 10 per cent most frequently cited articles published world-wide in Computer Science, the number of papers by Netherlands academic computer scientists is 50 per cent higher than expected on the basis of the total volume of Netherlands publication output in the field.

The outcomes of the bibliometric analyses presented in this report, be it of an exploratory and preliminary nature, indicate a rather strong variability in the citation impact of Netherlands computer science, both at the level of individual papers and that of research groups, a variability that is not reflected in the quality ratings of the QANU Peer Review Committee evaluating a large segment of Netherlands academic Computer Science groups in 2003/2004. Even though variability in citation impact and lack of correlation with peer ratings were also found *within* subfields, the extent to which these findings are due to differences in coverage of the Expanded WoS database created and explored in this study, awaits further research.

9.6 Concluding remark

Each methodology has its strengths and limitations, and is associated with a certain risk of arriving at invalid outcomes in individual cases. A methodology, even if it provides invalid outcomes in individual cases, may be beneficial to the scholarly system as a whole. This is true both for bibliometric analysis and for peer review. It is primarily the task of members from the scientific/scholarly community and the domain of research policy, and not of the authors of this report to decide whether or not these risks of using a particular method of citation analysis are acceptable and whether its benefits prevail. This task may also comprise an assessment of whether the extra costs of an advanced, sophisticated bibliometric analysis, compared to those of a less sophisticated one, match its surplus value in a research evaluation process. This report aims at providing information about the potentialities and limits of the various types of citation analysis that help scholars and policy makers to carry out such a delicate task.

References

Drott, M.C. (1995) Reexamining the Role of Conference Papers in Scholarly Communication. *Journal of the American Society of Information Science & Technology* 46, 299–305.

Garfield, E. (1974) Some comments on pure and applied research – stimulated by a list of works cited by applied chemists. *Current Contents*, December 4, 1974. In: *Essays of an information Scientist*, 2, 184–197. Philadelphia: *ISI* Press.

Garfield, E. (1979). *Citation Indexing – its Theory and Applications in Science, Technology and Humanities*. Wiley, New York.

Goodrum, A.A., McCain, K.W., Lawrence, S. and Giles, L. (2002). Scholarly publishing in the Internet age: a citation analysis of computer science literature *Information Processing and Management,* 37, 661–675.

Martin, B.R. and Irvine, J. (1983). Assessing Basic Research. Some Partial Indicators of Scientific Progress in Radio Astronomy. *Research Policy,* 12, 61–90.

Moed, H.F, De Bruin, R.E. and Van Leeuwen, Th. N. (1995). New bibliometric Tools for the Assessment of National Research Performance. *Scientometrics*, 33, 381–422.

Moed, H.F. (2005), *Citation Analysis in Research Evaluation*. Dordrecht (Netherlands): Springer.

Narin, F. (1976). *Evaluative Bibliometrics. The Use of Publication and Citation Data in the Evaluation of Scientific Activity*. Computer Horizons, Inc., Cherry Hill, New Jersey, USA.

QANU – Quality Assistance Netherlands Universities (2004). Computer Science. Report of the Review Committee for Computer Science.

Schubert, A., and Glänzel, W. (1983). Statistical reliability of comparisons based on the citation impact of scientific publication. *Scientometrics* 5, 59-73.

Schubert, A., Glänzel, W. and Braun, T. (1989). Scientometric Datafiles. A Comprehensive Set of Indicators on 2649 Journals and 96 Countries in All Major Science Fields and Subfields 1981–1985. *Scientometrics* 16, 3–478.

Van Raan, A.F.J. (1996). Advanced bibliometric methods as quantitative core of peer review based evaluation and foresight exercises. Scientometrics, 36, 397–420.

Visser, M.S. and Moed, H.F. (2004), Measuring the impact of non–ISI source items, Presentation at the 8th Science and Technology Indicators Conf 2004. Slides available at: http://conference.CWTS.nl/Downloads/ppt/128_Visser.pdf.

Visser, M.S. and Moed, H.F. (2005), Developing Bibliometric Indicators of Research Performance in Computer Science. In: P. Ingwersen and B. Larsen (eds.). Proceedings of ISSI 2005. Vol 1, 275–279.

VSNU – Association of universities in the Netherlands (2002). Chemistry and Chemical Engineering. Report of the Review Committee on Chemistry and Chemical Engineering.

Annex A1: Specification of bibliometric indicators calculated in the study

Number of articles (P)

This indicator gives the total number of papers published during the time period considered ***(P)***. These are scientific articles published in sources included in the Publication and Citation Universe. Multiple occurrences of papers – particularly when a paper is listed in the publication lists of several Netherlands groups – are excluded on higher aggregation levels. The time period for *WoS* articles relates to 'database' years: i.e., the year in which they were entered into the *WoS* database. Due to a time lag in processing articles, late papers published for instance in 2004 may enter the database in early 2005, and contribute to the counts for the year 2005. For *ACM*, additional *LNCS* and *IEEE* proceedings papers the year indicator gives the publication year. As *WoS* articles we considered only papers classified as *normal articles*, *letters*, *notes*, and *reviews*. Meeting abstracts, corrections, and editorials are *not* included. From the three proceedings groups, we included all papers with at least one author, and categorized them as (proceedings) article.

Number of citations, excl. self citations (*C)*

This indicator gives the total number of citations received during the time period considered, without self citations (***C***). A self-citation (sc) to a paper is a citation given in a publication of which at least one author (either first author or co-author) is also an author of the cited paper (either first author or co-author).

Total citations incl. self citations (*C+sc*)

This indicator gives the total number of citations received during the time period considered, with self citations included (***C+sc***). A self-citation (sc) to a paper is a citation given in a publication of which at least one author (either first author or co-author) is also an author of the cited paper (either first author or co-author).

Citations per article (*CPP*)

The *fourth* indicator is the average number of citations per publication calculated while self-citations are not included (***CPP***).

Percentage of uncited articles (*Pnc*)

A *fifth* indicator is the percentage of articles not cited during the time period considered (***Pnc***), excluding self-citations.

Mean Journal Citation Score (JCSm, intermediary variable, not included in Tables)

Two international reference values are computed. A first value represents the mean citation rate of the sources in which a research unit has published (***JCSm,*** the mean **J**ournal **C**itation **S**core), taking into account both the type of paper (e.g., normal article, review, and so on), as well as the specific years in which the papers were published. To give an example, the number of citations received during the period 2000 – 2004 by a *letter* published in 2000 in journal X is compared to the average number of citations received during the same period (2000 – 2004) by all *letters* published in the same journal (X) in the same year (2000). Generally, a research unit publishes its papers in several journals/sources rather than one. Therefore, we calculated a weighted average JCS indicated as ***JCSm***, with the weights determined by the number of papers published in each journal. With respect to the *ACM*, *LNCS* and *IEEE* conference proceedings, we calculated ***JCSm*** values for each individual conference proceeding. In the present study, self-citations are excluded from the computation of JCSm.

Mean Field Citation Score (FCSm, intermediary variable, not included in Tables)

The second reference value presents the mean citation rate of the subfields (journal categories) in which the research unit is active (***FCSm,*** the mean **F**ield **C**itation **S**core). Our definition of subfields is based on a classification of scientific journals into *categories* developed by *Thomson Scientific/ISI*. Although this classification is certainly not perfect, it is at present the only classification available to us. Within this classification system there are several journal categories covering subfields of Computer Science, including 'Computer Science, Artificial Intelligence', 'Computer Science, Software Engineering', and several others. In the present study this journal category system was adjusted in the following way: all journals assigned by *Thomson/SI* to one of the categories related to Computer Science, as well as all sources from *ACM*, *LNCS* and *IEEE* conference proceedings, were assigned to one compound 'journal' category, Computer Science. In calculating ***FCSm***, we used the same procedure as the one we applied in the calculation of ***JCSm,*** with journals/sources replaced by subfields. In most cases, a research unit is active in more than one subfield (i.e., journal category). In those cases, we calculate a weighted mean value, the weights being determined by the total number of papers published in each subfield. JCSm and FCSm are intermediate variables, and will not be presented in the data table.

Impact compared to journal average (*CPP/JCSm*)

This indicator compares the average number of citations to the oeuvre (***CPP***) to an international reference value, namely the corresponding Average Journal Citation Score (***JCSm***), by calculating the ratio of the two. In the case of conference proceedings, an 'Average Proceedings Citation Score' is calculated, also denoted as ***JCSm***. Self-citations are excluded in the calculation of this ratio, to prevent that ratios are affected by divergent self-citation behaviour.

If the ratio ***CPP/JCSm*** is above 1.0, the mean impact of a research unit's papers exceeds the mean impact of all articles published in the sources (journals or proceedings volumes) in which the particular research unit has published its papers (the research unit's source set).

Impact compared to field average (*CPP/FCSm*)

This indicator compares the average number of citations to the oeuvre (***CPP***) to an international reference value, namely the corresponding Average Field Citation Score (***FCSm***), by calculating the ratio of the two. Self-citations are excluded in the calculation of this ratio, to prevent that ratios are affected by divergent self-citation behavior.

If the ratio ***CPP/FCSm*** is above 1.0, this means that the oeuvre is cited more frequently than an 'average' publication in the subfield(s) in which the research unit is active. ***FCSm*** constitutes a *world subfield average* in a specific (combination of) subfield(s). In this way, one may obtain an indication of the international position of a research unit, in terms of its impact compared to a 'world' average. This 'world' average is calculated for the total population of articles published in the Publication and Citation Universe, assigned to a particular subfield or journal category. As a rule, about 80 percent of these papers are authored by scientists from the United States, Canada, Western Europe, and Japan. Therefore, this 'world' average is dominated by the Western world.

The indicator ***CPP/FCSm*** is the most frequently applied indicator in *CWTS* bibliometric studies, and is sometimes denoted as the '*crown indicator*'. It is considered to be the best available reflection of the average citation impact of a research unit's papers.

Impact journals compared to field average (*JCSm/FCSm*)

This indicator compares the impact of the sources in which the selected articles were published *(JCSm)* to the average Field Citation Score in the subfield covered by these sources (*FCSm*). If the ratio ***JCSm/FCSm,*** is above 1.0, the mean citation score of the set of sources in which the research unit has published exceeds the mean citation score of all papers published in the subfield(s) covered by these sources. In this case, one can conclude that the research unit publishes in sources with a relatively high impact.

It should be noted that the three indicators: impact compared to journal/source average (*CPP/JCSm*), Impact compared to field average (*CPP/FCSm*) and impact sources compared to field average (*JCSm/FCSm*) are not independent. The value of each one of these follows directly from the values of the other two indicators.

Statistical test, significance of differences

We applied a statistical test to establish whether the average impact of a research unit's publication oeuvre ***(CPP)*** differs significantly from the average impact of all papers in the research unit's journal/source set ***(JCSm)*** or from the world subfield average ***(FCSm)*** in the subfield(s) in which the research unit is active. If a research unit has a citation per publication ratio (***CPP***) significantly above (below) the average field (***FCSm***) or journal/source citation score (***JCSm***), this is indicated in the tables by means of a '+' ('-')

symbol directly after the numerical value of the indicators ***CPP/FCSm*** and ***CPP/JCSm.*** A '?' indicates that the test has insufficient information to interpret the result.

Because the present study does not use randomly sampled data, significance tests are not appropriate for inferential analysis. However, significance is reported here as an arbitrary criterion in deference to its widespread use in social science for exploratory analysis of non-random data.

Due to the presence of error, only the first decimal of the ratios is normally reliable, provided that they are based on a sufficiently high number of publications, as a rule of thumb, more than 50. But even for a quite high number of publications a 5 per cent difference or shift in the value of an indicator should not be regarded as a significant result.

Percentage of author self-citations (*%sc*)

We also calculated the percentage of author self-citations ***(%sc***), relative to the total number of citations received. The percentage of author self-citations to a research unit's oeuvre is influenced by a number of factors, such as the subfield in which a research unit is active; the number, type and age distribution of articles published by a research unit; the size of the research unit; and the extent to which the papers published by a research unit are cognitively related. As noted, author self-citations are excluded from most indicators in the present study.

External or publication coverage (Ext. Coverage)

This indicator gives for a particular research unit the percentage of articles published in sources included in the Publication and Citation Universe, relative to the total number of publications made by the research unit.

Internal or cited reference coverage (Int.Coverage)

For a particular research unit, this indicator gives for all articles published in sources included in the Publication and Citation Universe, the percentage of their cited references that are published in sources in the Publication and Citation Universe. This indicator measures – within the Publication and Citation Universe – how well this universe covers the documents upon which the research unit's work is based, as reflected in the cited reference lists in its papers.

Citation coverage (C/Ctot)

This indicator gives the percentage of citations to articles published in (journal or proceedings) sources included in the *Expanded WoS* database (***C***), relative to the total number of citations to all types of publications (***Ctot***), including articles in journals or proceedings not included in the *Expanded WoS* database, books, chapters, but excluding publications in the category 'other'.

Full time Equivalents research time (*FTE*)

Chapter 6 provides data on the 'input' of Netherlands groups in the field academic Computer Science. These data were taken from the Report of the Review Committee for Computer Science (*QANU*, 2004). Input was defined as the number of Full Time Equivalents research time in a group, not counting Ph.D. students.

Number of published articles per FTE research time (*P/FTE*)

This indicator gives the number of articles in the *Expanded WoS* database per FTE research time. This indicator can be termed as a publication productivity indicator, in the sense that it relates 'output' to 'input'. It is used only in Chapter 6 of this report.

Number of citations per FTE research time (*C/FTE*)

This indicator gives the number of citations (received during 1996–2004 by a group's articles published during 1996–2001), per FTE research time. It can be denoted as a citation productivity indicator. It is used only in Chapter 6 of this report.

Frequently cited publications

An additional set of impact indicators reflects the contribution to the most frequently cited papers world-wide. Two research units may have equal impact scores on the ***CPP/FCSm*** indicator, but one produces a steady stream of publications that are cited well but fails to produce really high impact publications, while the other contributes considerably to the high impact publications (and also has a larger number of less well cited publications). To examine the distribution of frequently cited papers, we have ranked publications by the number of citations it received up to four years after publication, and marked those belonging to the 10 % most frequently cited papers published word-wide in a given year. Moreover, letters were excluded, due to the fact that letters display a deviant citation pattern compared to articles and reviews. The following three indicators were calculated.

The actual number of highly cited or 'top' publications (*Ptop10%*)

The indicator ***P top10%*** gives the absolute number of papers that are represented among the top 10% most frequently cited of all papers published in a particular year, and subject category. The rank of papers is derived from on the actual impact distribution of <u>all</u> similar papers, and author self-citations are excluded.

The expected number of highly cited or 'top' publications (*E(Ptop10%)*)

E(Ptop10%) gives the expected number of highly cited papers based on the number of papers published by the research unit. It is not simply 10 per cent of the number of articles analysed, since it takes into account deviations from the 90^{th} percentile, if tied values occur due to the discrete nature of the citation distribution.

The ratio of actual and expected number of highly cited or ‘top’ publications (*A/E(Top10%)*)

Finally, the ***A/E(Ptop10%)*** indicator marks the relative contribution to the 10 % most frequently cited papers, and is calculated as the ratio of the ***Ptop10%*** and ***E(Ptop10%)***. Here, a value above (below) one indicates a relatively high (low) contribution to the 10% most frequently cited papers.