# CS848 Paper Submission Site

Go To CS848 Site

Submit or Update A Review For Paper #1

Go To CS848 Submissions

It is currently Wednesday 13th of October 2004 07:51:45 PM EDT

| **Paper # 1 (Download paper of type application/pdf, 435692 bytes)** | |
|---|---|
| Title: | Rank-aware Query Optimization |
| Abstract: | |

**You have already finalized your review for this paper.**
**You can no longer modify it, but you may view it.**

**If you made a mistake in your review and you want it "unfinalized",**
**you may send mail to the program chair asking them to unfinalize it**

Send yourself this review by email

| Attribute | Value |
|---|---|
| Are you finished with this review? | Finalize, I am done editing |
| Provide a short summary of the paper | This paper begins with a description of the top-k ranking problem, why it is important, and some unanswered research questions surrounding it (such as how to generate rank-aware query plans, and how to provide a query cost for them during optimization). The authors then discuss the contributions that their work will have to this field.<br><br>Once the problem definition and contributions have been presented, the authors turn their attention to a discussion of some other techniques for performing result ranking. Presented techniques include positional ranking, NRA, TA, NRJN and HRJN. As well, the authors give an overview of some techniques used for generating query plans, such as the dynamic programming approach, and the interesting orders approach.<br><br>Next, attention is paid to the task of extending a dynamic programming optimizer to include new rank-join operators. The authors discuss the two major tasks that need to be accomplished in order to make the optimizer rank-join aware: including plans with these operators in the plan space, and using costing to prune the plan space after addition of these plans. Plans can be added using a set of heuristic rules, and can be pruned using enhanced cost estimations considering interesting orders and selectivities.<br><br>The next section of this paper presents an algorithm for estimating the amount of data that needs to be read in order to generate the top-k results, as well as how to minimize the depths of the two join relations. The authors present some theorems and proofs of those theorems.<br><br>Later, experimental validation of the presented estimation models is shown. The authors outline the experimental setup, and later present the results of input cardinality estimation and buffer size estimation. In both cases, the presented algorithms appear to work well. Estimations were found to contain only 30% error.<br><br>The authors conclude the paper by relating their work to others, as well as providing a |

| | |
|---|---|
| | short summary of their paper. |
| What is the strength of the paper? (1-3 sentences) | This paper presents a novel approach to estimation of depths and result cardinalities for rank-join operations. Also, algorithms are presented and many examples are used to clarify the concepts. |
| What is the weakness of the paper? (1-3 sentences) | This paper presents a good analysis of the work, but does not provide any comparison to other work or naive approaches. As well, some presented mathematical models are tough to follow and understand. |
| Your qualifications to review this paper | I know the material, but am not an expert |
| Writing Quality | Good |
| Relevance to query processing? | The paper is relevant to query processing |
| **Experimental Methodology** | **Good** |
| Novelty of paper | This is very novel |
| **Overall paper merit** | **A novel or new contribution to this area with good methodology, or an incremental contribution paper that has excellent methodology. A must read for anyone in the area.** |
| In your opinion, will this paper be important over time? | Good |
| Provide additional detailed comments to the author | This paper presents some interesting concepts as well as providing a good overview of the problem domain. However, some points for improvement are worth mentioning:<br><br>-Some words are missing in places thoughout the paper. It makes some sentences difficult to understand. However, the overall writing style is quite good.<br>-Statistical analysis is very useful, but it would have been good to see a contrast with results from a naive processor without the aid of the implemented estimation changes.<br>-Some mathematical models and proofs presented in section 4 could use clarification. Models should be presented as clearly as possible, and in as plain english as possible, to ensure the reader can understand it.<br><br>Overall, this paper is well written and presents many interesting concepts fairly clearly. |
| **Additional comments to PC (not seen by author)** | **Although there are some missing words, this paper is well written. Also, I do not believe that the missing contrasting information with naive approaches presents enough of a reason to not publish this novel paper.** |

Goto Main Index  Close Window