# CS848 Paper Submission Site

Go To CS848 Site

Submit or Update A Review For Paper #12

Go To CS848 Submissions

It is currently Thursday 28th of October 2004 10:35:39 PM EDT

| **Paper # 12 (Download paper of type application/pdf, 2170500 bytes)** | |
|---|---|
| Title: | Extensible Query Processing in Starburst |
| 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 describes the Starburst system and how it optimizes queries. The authors begin with a description of the two major components, Corona and Core, as well as related work to Starburst. Starburst is extensible, i.e. it supports adding new functions, data storage techniques, access methods, data types, etc.<br><br>The paper next presents Hydrogen, the query language used in Starburst. Hydrogen includes and extends the functions provided by SQL. The notion of a table expression is discussed. Table expressions allow a query to call a function that returns a table. This can be used to recursively generate input data as well as provide dynamic views. The authors continue to discuss their language and ways it may be extended by database administrators.<br><br>The method used to transform a query into a query execution plan is next presented. This process begins with parsing the query, and translating it into the query graph model. From there, the query is rewritten, to increase performance, according to a set of rules. The plan is then optimized according to cost estimations, and refined. The plan is then stored or executed. The QGM is then presented in detail as a basis for discussions regarding Starburst optimization.<br><br>The authors next turn to the query reqrite phase. A discussion of how rules are applied is given. Rules are composed of a condition that must be met, and an action that is performed on a graph meeting that condition. Rules are written in C. The rule system must be highly extensible in order to allow database administrators the ability to add more rules to optimize their added operations and functions. Some built-in rules are discussed.<br><br>The authors next discuss query optimization in Starburst. Queries are optimized on a cost basis. So, the optimizer begins by enumerating plans that access base relations. The plan space is pruned at each step. Enumeration and pruning of join plans is then performed to |

| | |
|---|---|
| | find an optimal plan to run.<br><br>Before concluding, the authors provide an overview of the query evaluation system and how a query would be run through it. Some summary points are presented. |
| What is the strength of the paper? (1-3 sentences) | This paper presents a rule-based optimization algorithm that is useful in reducing a query. |
| What is the weakness of the paper? (1-3 sentences) | This system requires the learning of a new query language, Hyrdrogen. It also appears to run slowly, even though a budget can be specified to restrict the amount of time spent optimizing. |
| 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** | **Unacceptable** |
| Novelty of paper | This is a new contribution to an established area |
| **Overall paper merit** | **The paper is a novel or new contribution with average/weak methodology, or an incremental contribution that has good methodology. Someone in the area should read it** |
| In your opinion, will this paper be important over time? | Average |
| Provide additional detailed comments to the author | Your paper presents some interesting material. However, I have several issues with it.<br><br>-Your system requires users to learn a new language. This can be a deterrant to using Starburst in an organization<br>-How long does it take to generate the query graph and perform this optimization? It seems like this could take large amounts of time if the query performs many joins, etc., or requires the query graph to be altered by many rules. Although you allow a budget, what happens if the budget is exhausted? Is there potential for a really bad plan to be chosen anyway?<br>-This paper does not present any performance evaluation of the system. How long does processing and optimization take? What overhead does this impose on queries? How accurate are the optimizations?<br><br>Overall, your discussions are easy to read and your paper is well organized. However, much work remains to be completed before this paper is publication quality. |
| **Additional comments to PC (not seen by author)** | **Although this paper presents a good concept, it does not appear as though the research is complete enough to present results. The authors should finish their implementation, run some performance evaluation, and resubmit their findings for publication.** |

Goto Main Index  Close Window