

Mulgara - Feature #21

HybridTuples should defer sort and support prefix-define/reresolve

09/26/2006 03:51 AM - Andrae Muys -

Status:	New	Start date:	
Priority:	High	Due date:	
Assignee:	Andrae Muys -	% Done:	0%
Category:	Mulgara	Estimated time:	0.00 hour
Target version:			
Resolution:			

Description

Currently `[[HybridTuples]]` performs it's sort when it is constructed. This forces the evaluation of its argument before join-optimisation has an opportunity to specify a desired prefix or reresolve. `[[HybridTuples]]` should defer it's sort until at least the first call to `beforeFirst`.

Considerations:

Reresolve - If `[[HybridTuples]]` argument is `Reresolvable` `[[HybridTuples]]` should support passing this argument through.

`[[DefinablePrefix]]` - If our argument is unordered then a projection is free, so we should expose this to the join optimisation. However note that if we are only sorting once, then we cannot pass this call through to the argument. Rather we should project the argument to maximise the prefix length available.

Single vs. Multiple sorts - A sort is $O(n \log n)$, and a n from m prefix is $O((0 \text{ from } m)^{(n - m) / m})$. If we can increase the prefix length sufficiently to pay for the cost of the extra sorts we should resort on each call to `beforeFirst()`. Note that if the argument is itself prefix-definable then the cost of a sort is in terms of the estimated reduced length.