

Basic query failure

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

```

The following setup demonstrates the problem:
<br/>

<br/>
<code><a href="local:///topazproject#test">drop </a></code>
<br/>
<code><a href="local:///topazproject#test">create </a></code>
<br/>
<br/>
<code><a href="local:///topazproject#test">insert </a></code>
<br/>
<code><a href="local:///topazproject#test">foo:bar</a> <a href="local:///topazproject#test">foo:set</a> <a href="local:///topazproject#test">user:joe</a></code>
<br/>
<code><a href="local:///topazproject#test">foo:set</a> <a href="local:///topazproject#test">topaz:implies</a> <a href="local:///topazproject#test">bar:set</a></code>
<br/>
<code><a href="local:///topazproject#test">foo:set</a> <a href="local:///topazproject#test">topaz:implies</a> <a href="local:///topazproject#test">bar:get</a></code>
<br/>
<code><a href="local:///topazproject#test">into </a></code>
<br/>
<br/>
<code><a href="local:///topazproject#test">select $s $p $o from </a><a href="local:///topazproject#test">where ( </a></code>
<br/>
<code><a href="local:///topazproject#test">$s $p $o </a></code>
<br/>
<code><a href="local:///topazproject#test">and ($s <a href="local:///topazproject#test">mulgara:is</a> <a href="local:///topazproject#test">foo:bar</a>) </a></code>
<br/>
<code><a href="local:///topazproject#test">) or ( </a></code>
<br/>
<code><a href="local:///topazproject#test">($s <a href="local:///topazproject#test">impliedBy $o and $impliedBy <a href="local:///topazproject#test">topaz:implies</a> $p) </a></code>
<br/>
<code><a href="local:///topazproject#test">and ($s <a href="local:///topazproject#test">mulgara:is</a> <a href="local:///topazproject#test">foo:bar</a>) </a></code>
<br/>
<code><a href="local:///topazproject#test">); </a></code>
<br/>
<br/>
<code><a href="local:///topazproject#test">select $s $p $o from </a><a href="local:///topazproject#test">where ( </a></code>
<br/>
<code><a href="local:///topazproject#test">$s $p $o </a></code>
<br/>
<code><a href="local:///topazproject#test">or ($s <a href="local:///topazproject#test">impliedBy $o and $impliedBy <a href="local:///topazproject#test">topaz:implies</a> $p) </a></code>
<br/>
<code><a href="local:///topazproject#test">) and ($s <a href="local:///topazproject#test">mulgara:is</a> <a href="local:///topazproject#test">foo:bar</a>); </a></code>
<br/>

```

```

<br/>
The two queries are logically equivalent (the <mulgara:is>
term is duplicated in the first query, and factored out in the second query). However, while the
first query returns three statements, the second query only returns one. Specifically, the first q
uery returns
<br/>

<br/>
<foo:bar> <foo:set> <user:joe>
<br/>
<foo:bar> <bar:set> <user:joe>
<br/>
<foo:bar> <bar:get> <user:joe>
<br/>

<br/>
but the second returns only
<br/>

<br/>
<foo:bar> <foo:set> <user:joe>
<br/>

<br/>

```

History

#1 - 11/21/2006 05:00 AM - Paula Gearon

```

I need confirm this, but I have a theory...
<br/>
In the second query the initial part is:
<br/>

<br/>
$s $p $o or
<br/>
($s $impliedBy $o and $impliedBy <topaz:implies> $p)
<br/>

<br/>
Since the ($s $p $o) triple is all that is being selected then the part in parentheses would appear redundant.
Perhaps there is code that recognizes this and does the optimization before the mulgara:is predicate is appl
ied?

```

#2 - 11/21/2006 05:14 AM - Paula Gearon

```

The theory isn't looking too good. However, the problem appears to be resolution order, if the constraint ord
er is reversed, then the bug does not appear:
<br/>

<br/>
select $s $p $o
<br/>
from <a href="rmi://localhost/server1#test">rmi://localhost/server1#test</a>>
<br/>
where (
<br/>
    <mulgara:is> $s $p $o and <mulgara:is> $p $o
<br/>
    or $s $p $o
<br/>
) and $s <mulgara:is> <foo:bar>;
<br/>

<br/>
Result:
<br/>
[ foo:bar, foo:set, user:joe ]
<br/>
[ foo:bar, bar:get, user:joe ]

```


[foo:bar, bar:set, user:joe]

#3 - 11/21/2006 05:35 AM - Andrae Muys -

Looks like non-union compatible disjunctions strike again.

Note that the second query is attempting a join between the assignment and a NUC-disjunction.

Paul, note the use of the \$p in the object position - this does mean the parenthesised expression is not redundant.

#4 - 11/21/2006 07:07 AM - ronald -

Thanks to both of you for checking this out so quickly. So should we

just be avoiding using 'or'? That would be a bit difficult, though.

#5 - 11/21/2006 02:54 PM - Paula Gearon

I had a typo in the first comment....

When I said "would appear redundant", I should have said "could appear redundant".
. The current behaviour is clearly a bug. I wondered if the resolver made this mistake. I find that scenario less likely now that I know that it works if the terms are reversed.

I was speculating out loud in case it prompted anyone to have an "A-Ha" moment. :-)

#6 - 12/11/2006 06:05 AM - Andrae Muys -

No need to avoid 'or', but it is worth avoiding non-union-compatible 'or'.

What this means is that there are variables in one term of the disjunction that do not appear in all of the others.

So

\$a <::is> <foo> or \$a <::is> <foo> is fine. but

\$a <::is> <foo> or \$a <pred> \$b is not.

The work around is to use the distributive-law to convert any POSOP... form query into SOP as per your first example above.

The only problem with this work around is that the query fragment will not join against other constraints containing \$b properly as unbound should join successfully with anything, while with an explicit placeholder it will eliminate. However for SOP form queries handle UNBOUND correctly.

#7 - 02/24/2007 02:13 AM - Andrae Muys -

Downgraded from critical to major as workaround does exist (just provide suitable <mulgara:is> bindings to avoid non-union-compatible disjunction).

#8 - 04/17/2007 05:42 AM - Andrae Muys -

Got it - this is not a non-union-compatible disjunction bug after all.

This is a bug in [[UnorderedProjection]] which can only be triggered in the presence of a conjunction/disjunction combination.

Specifically projection is being used to reorder the variables of its operand, without taking into consideration the affect this has on prefixing. To see this, look at the implementation of beforeFirst() in [[UnorderedProjection]], while considering that the order of variables in the projection may not be the same as the order in its operand.

An initial fix is to materialise the operand when a reordering of variables is required. The ultimate fix is to do this only if the reordering will invalidate a defined-prefix, as we don't care if the defined-prefix is reordered as long as any prefix variable is not reordered with respect to non-prefix variables.

It would be good to take the time to fully exploit the Annotations to provide transparent access to defineIndex and reresolve.

#9 - 04/17/2007 10:13 AM - Andrae Muys -

Resolved in branches/nuc-disj revision 231.

Please test - if there are no reports of regressions in the next few days I'll merge over to trunk.

#10 - 04/20/2007 05:27 AM - ronald -

This works for us. We tested the original query that lead to this bug report. Many thanks.

#11 - 04/20/2007 05:29 AM - ronald -

Ooops, scratch that last comment - wrong bug report. Sorry. We're still in the process of verifying this one.

#12 - 04/20/2007 10:38 PM - ronald -

Ok, for real now: we tested our full original query and the seems to be fixed. Thanks!