

Morpho - Bug #1777

Selecting columns is very slow

11/18/2004 05:01 PM - Andrea Andrea

Status:	Resolved	Start date:	11/18/2004
Priority:	Normal	Due date:	
Assignee:	ben leinfelder	% Done:	0%
Category:	morpho - general	Estimated time:	0.00 hour
Target version:	1.8.0	Spent time:	0.00 hour
Bugzilla-Id:	1777		
Description It takes an excruciatingly long time to select a column in a data table. On Andrea's computer and on Veronique's computer, it takes up to 20 seconds for a column to be selected after you click on the column header. On Dan's computer, it takes about 4 seconds. All are PCs. The search and open functions of morpho also take a long time to complete, although not as long as the column selection takes.			
Related issues: Blocked by Morpho - Bug #2366: Selecting attributes in a table is slow			
		Resolved	02/22/2006

History

#1 - 12/09/2004 02:46 PM - Veronique Connolly

Andrea and I both noticed that selecting columns is no longer as slow. This bug seems to come and go - we should at least try and figure out what causes it to happen.

#2 - 12/17/2009 05:31 PM - Jim Regetz

Selecting columns definitely gets slower after repeatedly opening and closing a data package. To reproduce, start morpho, then open the jscientist sample data package. Note that selecting columns is reasonably fast (not instantaneous, but definitely <0.5 s). Now simply close and reopen this data package 10 times without doing anything else. Then open once again and try clicking columns again. Response time is now ~2 seconds. It gets slower still after opening and closing another handful times.

Similar slowdowns result when opening and closing different data packages, not just the same one multiple times. Moreover, performance seems to degrade more rapidly when the packages in question have more tables and/or attributes, but I haven't thoroughly tested this.

I've reproduced this on two platforms, both with 1.7.0 and with a recent dev build (svn r4704):
Ubuntu 9.10, java version "1.6.0_16"
OS X 10.6.2, java version "1.6.0_17"

#3 - 12/17/2009 05:49 PM - Jim Regetz

As another test, I created two fake DPs with attribute descriptions (but no attached data). One DP described 1 table with 20 columns, and the other described 10 tables each with only 2 columns. In other words, both DPs have descriptions for 20 total attributes.

Degradation in column click responsiveness after repeatedly opening and closing these DPs 10 times was about the same in both cases. So at least based on this very simple test, there was no clear evidence of a difference between having many small tables versus one larger table.

On a tangential performance note, opening the DP with more tables is noticeable slower than opening the one with a single table.

#4 - 12/18/2009 10:47 AM - ben leinfelder

well it's not the actual column selection (see 4617 about background thread) - it's displaying the attribute metadata.
Now when I test with this fix in (opening jscientist DP 10x) the column is selected quickly, but the lag is substantial for displaying the metadata...
It appears to be parsing the full XML and applying a XSLT to it. There is a history mechanism that I see in the code, but it's unclear if that is caching the display information and/or working as expected.

#5 - 12/18/2009 03:20 PM - ben leinfelder

I've made the datapackage frame tighter in a few ways:
1. Only listen/handle data selection events that originate in this same frame
2. Remove yourself as a listener when you are closed
3. Display attribute metadata in a separate background thread (bug 2366)

1 and 2 are crucial, whereas 3 is possibly confusing when lag is very long (though I now believe it won't be very long)

#6 - 12/18/2009 05:20 PM - ben leinfelder

Now that the listeners are behaving themselves...closing

#7 - 03/27/2013 02:18 PM - Redmine Admin

Original Bugzilla ID was 1777