

art - Feature #8426

Understand gradual memory growth and determine if it can be mitigated

04/24/2015 10:36 AM - Kyle Knoepfel

Status:	Accepted	Start date:	04/24/2015
Priority:	Normal	Due date:	
Assignee:		% Done:	0%
Category:	I/O	Estimated time:	16.00 hours
Target version:		Spent time:	0.00 hour
Scope:	Internal	SSI Package:	art
Experiment:	-		

Description

In issue [#4946](#), we were able to resolve the memory spike at the end of the art process that Andrei had observed. The gradual memory growth (see attached image) as documented there, however, is not yet fully understood. At this point, we know the following:

For the specific job as described in [#4946](#), a gradual physical memory growth of roughly 150 MB is observed. The breakdown is as follows:

```
FileIndex      : 60.1 MB ( 39.7%)
ParameterSetDB : 24.6 MB ( 16.2%)
Unaccounted for: 66.7 MB ( 44.1%)
-----
Total          : 151.4 MB (100.0%)
```

The FileIndex and ParameterSetDB entries correspond to metadata that is propagated to the final output stream. When no output stream is specified, the FileIndex contribution is removed, but the ParameterSetDB information is still aggregated as each input file is read. What remains to be accounted for is 66.7 MB.

As discussed at yesterday's stakeholder meeting, it is possible to mitigate the FileIndex memory growth by omitting the sorting call in RootOutputFile.cc. It was deemed that this was not yet a necessary step to take.

Analysis should be performed to determine where the remaining memory growth comes from and whether steps can/should be taken to mitigate the memory growth from each of the sources during the art process.

Related issues:

Related to art - Feature #4946: art memory use: multiple subruns

Closed

History

#1 - 04/27/2015 11:46 AM - Christopher Green

- Category set to I/O
- Status changed from New to Accepted
- Estimated time set to 16.00 h
- SSI Package art added
- SSI Package deleted ()

After some further study, about 35% of the observed memory increase remains unaccounted-for. The estimated time covers the analysis, categorization and documentation of that remaining memory growth.

#2 - 05/07/2015 09:59 AM - Kyle Knoepfel

- Related to Feature #4946: art memory use: multiple subruns added

Files

fileFormatVersion7.png	47.1 KB	04/24/2015	Kyle Knoepfel
------------------------	---------	------------	---------------