

## ctools - Change request #1339

### Reduce memory usage to ctools

10/23/2014 05:33 PM - Knödlseider Jürgen

<b>Status:</b>	Closed	<b>Start date:</b>	10/23/2014
<b>Priority:</b>	Normal	<b>Due date:</b>	
<b>Assigned To:</b>	Knödlseider Jürgen	<b>% Done:</b>	100%
<b>Category:</b>		<b>Estimated time:</b>	0.00 hour
<b>Target version:</b>	00-08-00		
<b>Description</b>			
<p>ctobssim can take a large amount of memory if multiple runs are simulated at once. This is related to the fact that first all events are simulated and only after all simulations are done the results are saved. One can circumvent this by adding a save step directly after the simulation step in case that the execute() method is called. The only problem is then that ctobssim may process the runs in parallel (due to the OpenMP) support, leading still to important memory requirements. But this could be controlled by setting the</p> <pre>export OMP_NUM_THREADS=2</pre> <p>environment variable to limit for example the maximum number of threads to 2.</p>			

#### History

##### #1 - 10/24/2014 08:26 PM - Knödlseider Jürgen

- Subject changed from Reduce ctobssim memory usage to Reduce memory usage to ctools
- Status changed from New to In Progress
- Assigned To set to Knödlseider Jürgen
- Target version set to 00-08-00

##### #2 - 10/24/2014 08:47 PM - Knödlseider Jürgen

Here some information about memory usage.

The size of a GCTAEventAtom object is 264 Bytes, which is about 2.3 times larger than the event written to disk. The difference mainly comes from the GCTAInstDir member which takes 112 Bytes due to pre-computed information.

The IFAE20120510\_50h response predicts about 60000 events for 30 min, which corresponds to 16 MB in RAM and 7 MB on disk. The IFAE20120510\_50h response file takes 210 kB on disk. The binning of the response file is still relatively coarse, and it is certainly conceivable that the final response file takes a factor of 10 more on disk, i.e. 2 MB. As a rule of thumb, one can thus assume that a 30 min run will take 10 MB on disk and 20 MB in RAM. The following table summarizes the expected memory needs:

Duration	Disk	RAM
30 min	10 MB	20 MB
50 hr	1 GB	2 GB
200 hr	4 GB	8 GB
1 year	28 GB	56 GB
30 years	840 GB	1.7 TB

(1400 hr of observing time have been assumed per year).

As today laptops can accommodate a few GB of RAM, analyses of 50 hr can be done with the existing code in memory. Memory usage can be reduced by reducing the size of an GCTAEventAtom object to the information that is needed for analysis. The minimal information is Right Ascension, Declination, energy and time, which would occupy 32 Bytes for double precision values. This corresponds to a factor of 8 reduction.

To go beyond the RAM memory limitations, data may be stored on disk when not used. This goes of course at the expense of computation speed.

**#3 - 10/24/2014 10:50 PM - Knödseder Jürgen**

- % Done changed from 0 to 20

ctobssim now uses less memory when the execute() method is called. In that case, events are written immediate to disk and are not kept in memory. They are however read back when needed.

**#4 - 10/25/2014 12:34 AM - Knödseder Jürgen**

- % Done changed from 20 to 50

I implemented a new logic in GCTAObservation where the events are actually not loaded when reading information from an XML file. The read() method now just stores the event file name, and the events() method makes sure that the events are loaded if they are not yet online. To get rid of events once they have been loaded, the dispose\_events() method has been added.

With this logic, only minimal changes need to be done on the ctools side.

**#5 - 10/25/2014 12:37 AM - Knödseder Jürgen**

- Description updated

**#6 - 10/25/2014 01:40 AM - Knödseder Jürgen**

- File test.py added

- Status changed from In Progress to Feedback

- % Done changed from 50 to 90

Attached a test script that I used to test a pipeline with reduced memory usage. Looks good from my side.

attachment:test.py

**#7 - 10/30/2014 12:12 PM - Knödseder Jürgen**

- Status changed from Feedback to Closed

- % Done changed from 90 to 100

**Files**

---

test.py	6.94 KB	10/24/2014	Knödseder Jürgen
---------	---------	------------	------------------