

PROOF Basics

ROOT Tutorials

2013



- To start a PROOF-Lite session

```
$ root -l
root [0] TProof::Open("")
+++ Starting PROOF-Lite with 2 workers +++
Opening connections to workers: OK (2 workers)
Setting up worker servers: OK (2 workers)
PROOF set to parallel mode (2 workers)
(class TProof*)0x8330140
root [1] gProof
(class TProof*)0x8330140
```

- Now we are ready to go
- **gProof** is a global instance of TProof
- ... **but what's TProof ?**



- TProof is the interface class to interact with the PROOF session
- Everything in the session is done via TProof
 - `Print()`, gives information about the session
 - `Process()`, allows to run a TSelector
 - `GetOutputList()`, returns the list of output objects
 - `DrawSelect()`, draws distributions
 - ...
- ... and many others. Check:

<http://root.cern.ch/root/html/TProof.html>



- Gives information about the session

```
$ root [1] gProof->Print()
*** PROOF-Lite cluster (parallel mode, 2 workers):
Host name:                macphsft12.local
User:                     ganis
ROOT version|rev|tag:     5.32/02|r43514
Architecture-Compiler:   macosx64-gcc421
Protocol version:        33
Working directory:       /Users/ganis/local/root/opt/root
Communication path:
/var/folders/uC/uC0RGjQUF1mzR689bg+JJU++0gQ/-Tmp-/plite-38583
Log level:                0
Number of workers:       2
Number of active workers: 2
Number of unique workers: 1
Number of inactive workers: 0
Number of bad workers:   0
Total MB's processed:    0.00
Total real time used (s): 0.000
Total CPU time used (s): 0.000
```



Ordinal numbers

- Unique identifier
 - 0.n for workers (master is always 0)

- The sandbox
 - Location controlled by Proof.Sandbox
 - Proof.Sandbox */my/special/sandbox*
 - Defaults
 - \$HOME/.proof (standard PROOF)
 - \$HOME/.proof/path-where-we-started (PROOF-Lite)


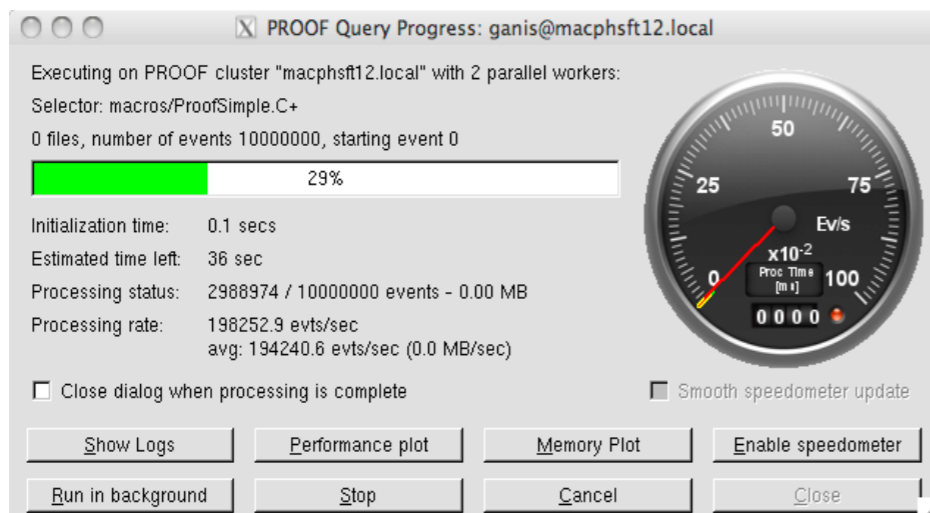


- **cache**
 - for package tarballs, code, binaries
- **packages**
 - where packages are actually build
- **queries** (master or PROOF-Lite only)
 - where the results from queries are stored
- **datasets** (master or PROOF-Lite only)
 - where information about datasets is stored
- **session-*SessionUniqueID***
 - Working area (logs, ...) for session *SessionUniqueID*



- The file `ProofSimple.C,.h` define a TSelector which fills 100 histograms with gaussian random numbers

```
root[] gProof->SetParameter("ProofSimple_Nhist", (Long_t)100)
root[] gProof->Process("ProofSimple.C+", 10000)
```

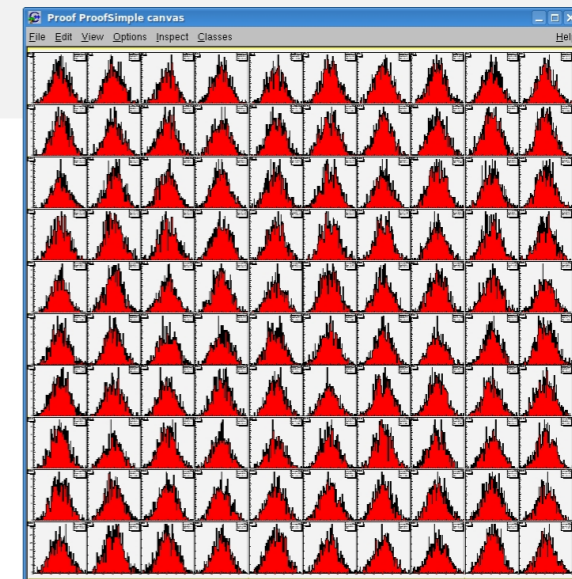




- Repeat on the local session

```

root[] TProofPlayerLocal *p = new TProofPlayerLocal()
root[] p->AddInput( \
           new TParameter<Long_t>("ProofSimple_Nhist", 100)
root[] p->Process("ProofSimple.C+", 10000)
    
```



- Try to measure the time with `gROOT->Time()`
 - What do you find?



- **TProofBench** is a tool to benchmark a PROOF cluster (lite or full installation)
- Framework for CPU-intensive and I/O-intensive scalability tests
- Default benchmarks
 - **CPU-intensive**
 - Random number generation
 - **I/O-intensive**
 - Readout TTrees based on \$ROOTSYS/test/Event.C,.h
- **Exercise 14b**



Selector being run

active workers

Progress bar

Executing on PROOF cluster "macphsft12.local" with 2 parallel workers
Selector: macros/ProofSimple.C+
0 files, number of events 10000000, starting event 0

29%

Initialization time: 0.1 secs
Estimated time left: 36 sec
Processing status: 2988974 / 10000000 events - 0.00 MB
Processing rate: 198252.9 evts/sec
avg: 194240.6 evts/sec (0.0 MB/sec)

Close dialog when processing is complete Smooth speedometer update

Show Logs Performance plot Memory Plot Enable speedometer
Run in background Stop Cancel Close

stats

Log dialog box



Select logs to display

PROOF - Processing logs for session 'macphsft12.local-1334613619-39527', started on Tue Apr 17 00:00:19 2012 at

Enter cluster URL:
proof://ganis@_lite_//

Enter session: 0 Get logs info

Choose workers: Clear All

0
0.0
0.1

```
// ----- Start of element log -----  
// Ordinal: 0.0 (role: worker)  
// Path: /Users/ganis/.proof/dropbox-Private-Tutorial-root-tutorial-tutorial/session-macphsft12.local-1334613619-39527  
// # of retrieved lines: 4  
// -----  
00:00:19 39532 Wrk-0.0 | Info in <TProofServLite::Setup>: fWorkDir: /Users/ganis/.proof  
cp: ./ProofSimple.C and /Users/ganis/.proof/cache/ProofSimple.C are identical (not copied).  
cp: ./ProofSimple.h and /Users/ganis/.proof/cache/ProofSimple.h are identical (not copied).  
00:00:50 39532 Wrk-0.0 | Info in <TUnixSystem::ACLiC>: creating shared library /Users/ganis/.proof/dropbox-Private-  
// ----- End of element log -----  
  
// ----- Start of element log -----  
// Ordinal: 0.1 (role: worker)  
// Path: /Users/ganis/.proof/dropbox-Private-Tutorial-root-tutorial-tutorial/session-macphsft12.local-1334613619-39527  
// # of retrieved lines: 3
```

Display

Lines: all svcmsg From -100 to 0

Grep for: Grep file: <session-tag>.log Save Close

Can be started also with
TProof::fLogViewer

Grep functionality

Save to a file



- In the output list

```
root [] gProof->GetOutputList()
(class Tlist*)0x89eae58
root [] gProof->GetOutputList()->ls()
OBJ: TStatus      PROOF_Status      : 0 at: 0x8a264a8
OBJ: TH1F         h0          h0 : 0 at: 0x89d5b48
OBJ: TH1F         h1          h1 : 0 at: 0x8a22de0
OBJ: TH1F         h2          h2 : 0 at: 0x8a21f88
OBJ: TH1F         h3          h3 : 0 at: 0x8a215f8
OBJ: TH1F         h4          h4 : 0 at: 0x8a24100
OBJ: TH1F         h5          h5 : 0 at: 0x8a288b8
OBJ: TH1F         h6          h6 : 0 at: 0x8a31c20
...
OBJ: TH1F         h96         h96 : 0 at: 0x89e9b38
OBJ: TH1F         h97         h97 : 0 at: 0x89ea000
OBJ: TH1F         h98         h98 : 0 at: 0x89ea4c8
OBJ: TH1F         h99         h99 : 0 at: 0x89ea990
root []
```



- Generate a few files, for example 10

```
root [] .L CreateEventTree.C
root [] Int_t i = 0;
root [] for(;i<10;++i){\
           CreateEventTree(Form("data/evtree_%d",i));}
```

- Create a chain

```
root [] TChain c("tree")
root [] for(i=0;i<10;++i){\
           c.AddFile(Form("file:///home/user/data/evtree_%d",i));}
```

- Process the chain inside PROOF

```
root [] c.SetProof()
root [] c.Process("EventDataSelector.C+")
```



- Use `gROOT->Time()` to measure **speed-up**
 - No PROOF

```
root [] c.SetProof(0)
root [] gROOT->Time();
root [] c.Process("EventDataSelector.C+")
```

– PROOF

```
root [] c.SetProof()
root [] gROOT->Time();
root [] c.Process("EventDataSelector.C+")
```



- The concept of **dataset** is very useful in HEP: it refers to a set of files containing **homogeneous data**
 - e.g. all the data taken during Summer 2009 under uniform detector conditions.
- Is useful to refer to a dataset **by name**
- **TFileCollection**: **named** list of TFileInfo
- **TFileInfo**: most generic way of describing a file
 - Multiple URLs, meta-information
- A TFileCollection is typically the **result of a query to a catalog**



- TProof has a set of methods to perform basic operations on datasets
 - `RegisterDataSet(const char *name, TFileCollection *)`
 - `VerifyDataSet(const char *name)`
 - `ShowDataSets()`
 - `TFileCollection *GetDataSet(const char *name)`
 - ...
- The name is in the form `/group/user/datasetname`
 - 'group' is an advanced PROOF concept: by default anybody is in group 'default'



- Second part of **Exercise 14c**
- Create a dataset for the files we used in the previous example
- Register and Verify it
- Process the dataset 'by name'



- When the selector needs additional code - for example a new class defined in the files `MyClass.C` and `MyClass.h` – PROOF provides two ways to make it available
 - `gProof->Load("MyClass.C")`
 - Equivalent of `.L` on the ROOT shell
 - Convenient for simple things
 - Package ARchives (PAR)
 - Structured archives with build and setup facilities
 - Convenient for more complex and stable things, e.g. the experiment analysis suite



- Zipped **tarballs** identified by a name and the `.par` extension, e.g. `pack.par`
- The tarball contains a structure like this
 - `./pack`
 - `./pack/PROOF-INF`
 - `./pack/PROOF-INF/BUILD.sh`
 - `./pack/PROOF-INF/SETUP.C`
- The code (`.C`, `.h`, makefiles, ...) should be put in the top level directory
- **BUILD.sh**: script to build the package, e.g. runs 'make'
- **SETUP.C**: is a macro running the final setup



- This is an example of running a real Monte Carlo simulation in PROOF
 - Pythia8 is the first usable C++ version of a famous HEP generator
- It needs a PAR file to setup the environment
- To run the example we need to link to pythia8
 - Available under `/cvmfs/sft.cern.ch/lcg/dev`



```
$ tar tzvf par/pythia8.par
drwxr-xr-x mslawins/sf          0 2008-07-03 19:45 pythia8/
-rw-r--r-- mslawins/sf      1521 2008-07-03 19:45 pythia8/main03.cmdnd
drwxr-xr-x mslawins/sf          0 2008-07-03 19:46 pythia8/PROOF-INF/
-rw-r--r-- mslawins/sf       233 2008-07-03 19:46 pythia8/PROOF-INF/SETUP.C
```

```
void SETUP()
{
    // Load the libraries

    gSystem->Load("$PYTHIA8/lib/libpythia8.so");
    gSystem->Load("libEG");
    gSystem->Load("libEGPythia8");

    // Set the include paths
    gROOT->ProcessLine(".include $PYTHIA8/include");
}
```



- Exercise 14d
- Download the `pythia8.par` and the related selectors
- Run the code to produce the following plot



- We have seen the basic concepts of PROOF
 - How to run a selector for a CPU intensive job
 - How to process a TChain
 - How to do the same from a named dataset
 - How to use a PAR file