



HTCondor update and FOOT simulations

R. Zarrella

FOOT XVI General Meeting 24-26/06/2024

HTCondor update



Many things changed! (*but mainly for me...*)

- HTCondor updated to version 23 from May
- New file for HTC environment setup
- Very easy to set-up and use
- Simplifies the interaction with the scheduler
- Recommended to use SHOE scripts for job submission
- Changed handling of condor auxiliary files (see next slides)
- **New script for FOOT simulations!**

```
source /opt/exp_software/foot/set_HTC_env.sh
```

HTCondor jobs monitoring and handling



From last GM!!

Jobs can be monitored through the “condor_q” command:

```
condor_q -name sn-02
```

```
-- Schedd: sn-02.cr.cnaf.infn.it : <131.154.192.42:9618?... @ 11/27/23 17:23:48
OWNER      BATCH_NAME      SUBMITTED   DONE    RUN    IDLE  TOTAL JOB_IDS
zarrellafoott3 ID: 9934116 11/27 17:22    -     90     -    122 9934116.0-121
zarrellafoott3 ID: 9934118 11/27 17:22    -     -     1     1 9934118.0
zarrellafoott3 ID: 9934119 11/27 17:22    -    93     -    93 9934119.0-92
zarrellafoott3 ID: 9934120 11/27 17:22    -     -     1     1 9934120.0
zarrellafoott3 ID: 9934121 11/27 17:22    -   125     8   133 9934121.0-132
zarrellafoott3 ID: 9934123 11/27 17:22    -     -     1     1 9934123.0

Total for query: 351 jobs; 32 completed, 0 removed, 11 idle, 308 running, 0 held, 0 suspended
Total for zarrellafoott3: 351 jobs; 32 completed, 0 removed, 11 idle, 308 running, 0 held, 0 suspended
Total for all users: 29350 jobs; 27389 completed, 0 removed, 305 idle, 998 running, 658 held, 0 suspended
```

```
condor_q -name sn-02 -nobatch
```



Do not group jobs on same cluster

```
condor_q -name sn-02 -run
```



Display only running jobs with their runtime

You can check the output of a single running job in real time via “condor_tail”

```
condor_tail -name sn-02 -f jobId
```

```
jobId = 9934116.0
        9934116.3
        9934121.130
```

The “-f” option lets you follow the output, you can omit it if you want a single printout

HTCondor jobs monitoring and handling



From last GM!!

Jobs can be monitored through the “condor_q” command:

```
condor_q -name sn-02
```

```
-- Schedd: sn-02.cr.cnaf.infn.it : <131.154.192.42:9618?... @ 11/27/23 17:23:48
OWNER      BATCH_NAME      SUBMITTED   DONE    RUN    IDLE  TOTAL JOB_IDS
zarrellafoott3 ID: 9934116  11/27 17:22    -     90     -    122 9934116.0-121
zarrellafoott3 ID: 9934118  11/27 17:22    -     -     1     1 9934118.0
zarrellafoott3 ID: 9934119  11/27 17:22    -    93     -    93 9934119.0-92
zarrellafoott3 ID: 9934120  11/27 17:22    -     -     1     1 9934120.0
zarrellafoott3 ID: 9934121  11/27 17:22    -   125     8   133 9934121.0-132
zarrellafoott3 ID: 9934123  11/27 17:22    -     -     1     1 9934123.0

Total for query: 351 jobs; 32 completed, 0 removed, 11 idle, 308 running, 0 held, 0 suspended
Total for zarrellafoott3: 351 jobs; 32 completed, 0 removed, 11 idle, 308 running, 0 held, 0 suspended
Total for all users: 29350 jobs; 27389 completed, 0 removed, 305 idle, 998 running, 658 held, 0 suspended
```

```
condor_q -name sn-02 -nobatch
```

Do not group jobs on same cluster

```
condor_q -name sn-02 -run
```

Display only running jobs with their runtime

You can check the output of a single running job in real time via “condor_tail”

```
condor_tail -name sn-02 -f jobId
```

```
jobId = 9934116.0
        9934116.3
        9934121.130
```

The “-f” option lets you follow the output, you can omit it if you want a single printout

HTCondor jobs monitoring and handling



From last GM!!

Once all your jobs are completed, you **need** to remove them from the queue

```
condor_rm -name sn-02 $USER
```

```
condor_rm -name sn-02 -all
```

The “-name sn-02” can be omitted by adding the following line to your ~/.bash_profile

```
export _condor_SCHEDD_HOST=sn-02.cr.cnaf.infn.it
```

If for some reason your jobs run for too long (>~1 hour), it is likely that something broke

- ☀ Usually SHOE’s “fault” → check if output files are created + condor_tail
- ☀ Some jobs **will** go on “hold” after 2-3h → intended to free resources
- ☀ Other possible issues can be pointed out in the auxiliary files of each job (.out/.err/.log)

```
condor_transfer_data -name sn-02 -all
```

Further info and commands in the [CNAF-Tier1](#) and the [HTCondor](#) documentations

HTCondor jobs monitoring and handling



From last GM!!

Once all your jobs are completed, you **need** to remove them from the queue

```
condor_rm -name sn-02 $USER
```

```
condor_rm -name sn-02 -all
```

The “-name sn-02” can be omitted by adding the following line to your ~/.bash_profile

```
export condor_SCHEDD_HOST=sn-02 cr.cnaf.infn.it
```

If for some reason your jobs run for too long (>~1 hour), it is likely that something broke

- ☀ Usually SHOE’s “fault” → check if output files are created + condor_tail
- ☀ Some jobs **will** go on “hold” after 2-3h → intended to free resources
- ☀ Other possible issues can be pointed out in the auxiliary files of each job (.out/.err/.log)

```
condor_transfer_data -name sn-02 -all
```

Further info and commands in the [CNAF-Tier1](#) and the [HTCondor](#) documentations

HTCondor interaction: summary



Set HTCondor environment

```
source /opt/exp_software/foot/set_HTC_env.sh
```

Operation	Command
Check job queue	condor_q condor_q -nobatch condor_q -run
Check job output	condor_tail jobId condor_tail -f jobId
Transfer auxiliary files	condor_transfer_data jobId condor_transfer_data -all
Remove jobs from queue	condor_rm \$USER condor_rm -all condor_rm -constraint "jobStatus=={id}"

1 pending	2 running
3 removed	4 completed
5 held	6 submission_err

{id}

HTCondor interaction: summary



Set HTCondor environment

```
source /opt/exp_software/foot/set_HTC_env.sh
```

Operation	Command
Check job queue	condor_q condor_q -nobatch condor_q -run
Check job output	condor_tail jobId condor_tail -f jobId
Transfer auxiliary files	condor_transfer_data jobId condor_transfer_data -all
Remove jobs from queue	condor_rm \$USER condor_rm -all condor_rm -constraint "jobStatus=={id}"

*Usually not needed
when using SHOE scripts!*

*(please check your queue sometimes or
probably I will tell you to do it...)*



{id}

1 pending	2 running
3 removed	4 completed
5 held	6 submission_err

Tier1 storage reminder



Each folder on the Tier1 has a specific task

→ Optimized for large files

`/storage/gpfs_data/foot` : parent directory (~90 TB)
`/storage/gpfs_data/foot/shared` : Experimental data
`/storage/gpfs_data/foot/shared/SimulatedData` : MC simulation data
`/storage/gpfs_data/foot/${USER}` : User data (create if not present!)

DATA

→ Optimized for small files

`/opt/exp_software/foot` : parent directory
`/opt/exp_software/foot/${USER}` : User software (create if not present!)

SW

→ Optimized for small files

→ Daily backup but limited and shared by all the collaboration

`/home/FOOT-T3` : parent directory (~110 GB)
`/home/FOOT-T3/${USER}` : User \$HOME

HOME

Tier1 storage reminder

From last GM!!



Each folder on the Tier1 has a specific task

→ Optimized for large files

Save your ROOTfiles here

DATA

- `/storage/gpfs_data/foot` : parent directory (~90 TB)
- `/storage/gpfs_data/foot/shared` : Experimental data
- `/storage/gpfs_data/foot/shared/SimulatedData` : MC simulation data
- `/storage/gpfs_data/foot/${USER}` : User data (create if not present!)

→ Optimized for small files

Install your software here! (also SHOE)

SW

- `/opt/exp_software/foot` : parent directory
- `/opt/exp_software/foot/${USER}` : User software (create if not present!)

→ Optimized for small files

→ Daily backup but limited and shared by all the collaboration

HOME

- `/home/FOOT-T3` : parent directory (~110 GB)
- `/home/FOOT-T3/${USER}` : User \$HOME



Auxiliary files handling

All jobs produce some auxiliary files containing output/error/logs

- Up to now, saved under output directory → /storage/gpfs_data/foot/*
- **STORAGE not suited for small files!!**
- After long discussion with CNAF, moved everything to user \$HOME
- Dedicated folders for different scripts (“HTC_SHOEreco”, “HTC_SHOEanalysis”, etc.)
 - Campaign and run number in name for simplicity

➔ **Optimized for small files**

➔ **Daily backup but limited and shared by all the collaboration**

/home/FOOT-T3 : parent directory (~110 GB)
/home/FOOT-T3/\${USER} : User \$HOME

HOME



Auxiliary files handling

All jobs produce some auxiliary files containing output/error/logs

- Up to now, saved under output directory → /storage/gpfs_data/foot/*
- **STORAGE not suited for small files!!**
- After long discussion with CNAF, moved everything to user \$HOME
- Dedicated folders for different scripts (“HTC_SHOEreco”, “HTC_SHOEdanalysis”, etc.)
 - Campaign and run number in name for simplicity
- **\$HOME disk space is shared!!!! Make sure to remove your aux files if your jobs are successful!!**
- **Please remove asap all your old “HTCfiles” folders under /storage/**

➔ Optimized for small files

➔ Daily backup but limited and shared by all the collaboration

/home/FOOT-T3 : parent directory (~110 GB)
/home/FOOT-T3/\${USER} : User \$HOME

HOME

New script for FOOT simulations!



shoe/Simulation/runFOOTSimulation.sh

Arguments

- Run simulations directly on Tier1!
- **Can be run from any directory**
(just don't "mv" the script from where it is)
- Possibility to save all FLUKA aux files
 - ✓ Check quality of simulation output
 - ✓ Increase statistics if needed
 - ✓ Removed once everything is fine

- o **Output directory** (in /storage/gpfs_data/foot/\${USER})
- c **Campaign name**
- r **Run number**
- n **Total number of events to simulate**
- e **Events per job**
- k **Keep TXT files in zip archive** (optional, default "0")

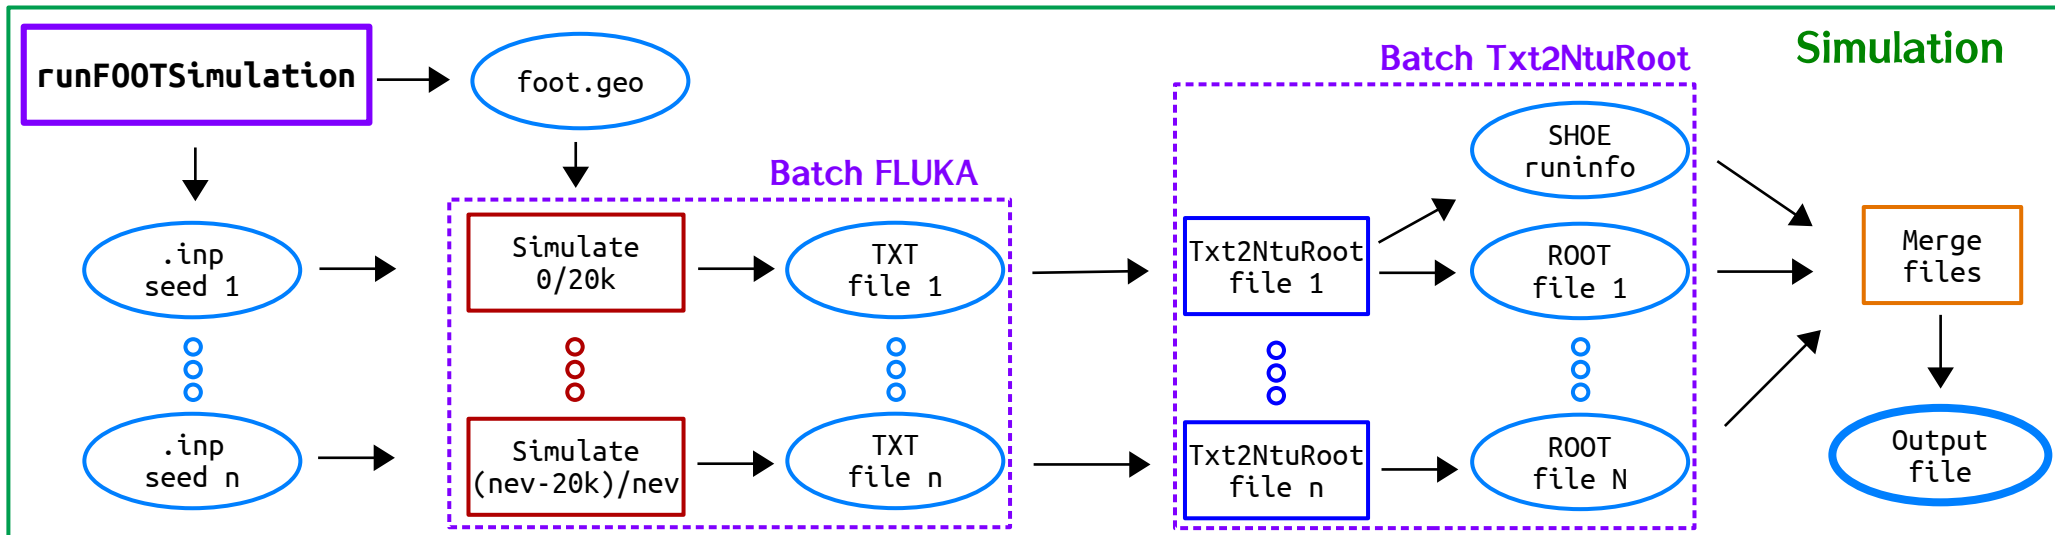
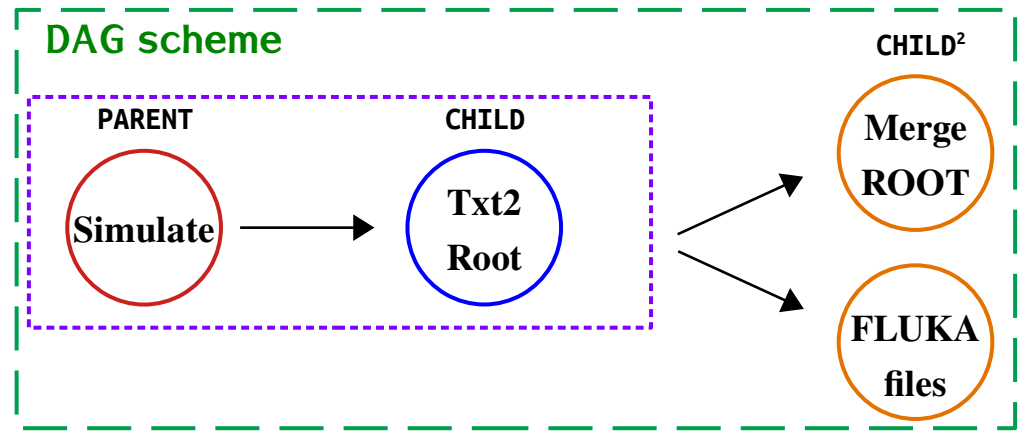
Example

```
./runFOOTSimulation.sh -o /storage/gpfs_data/foot/${USER}/CNA02023_MC_200 -r CNA02023_MC -r 200 -n 1000000 -e 1000
```



New script for FOOT simulations!

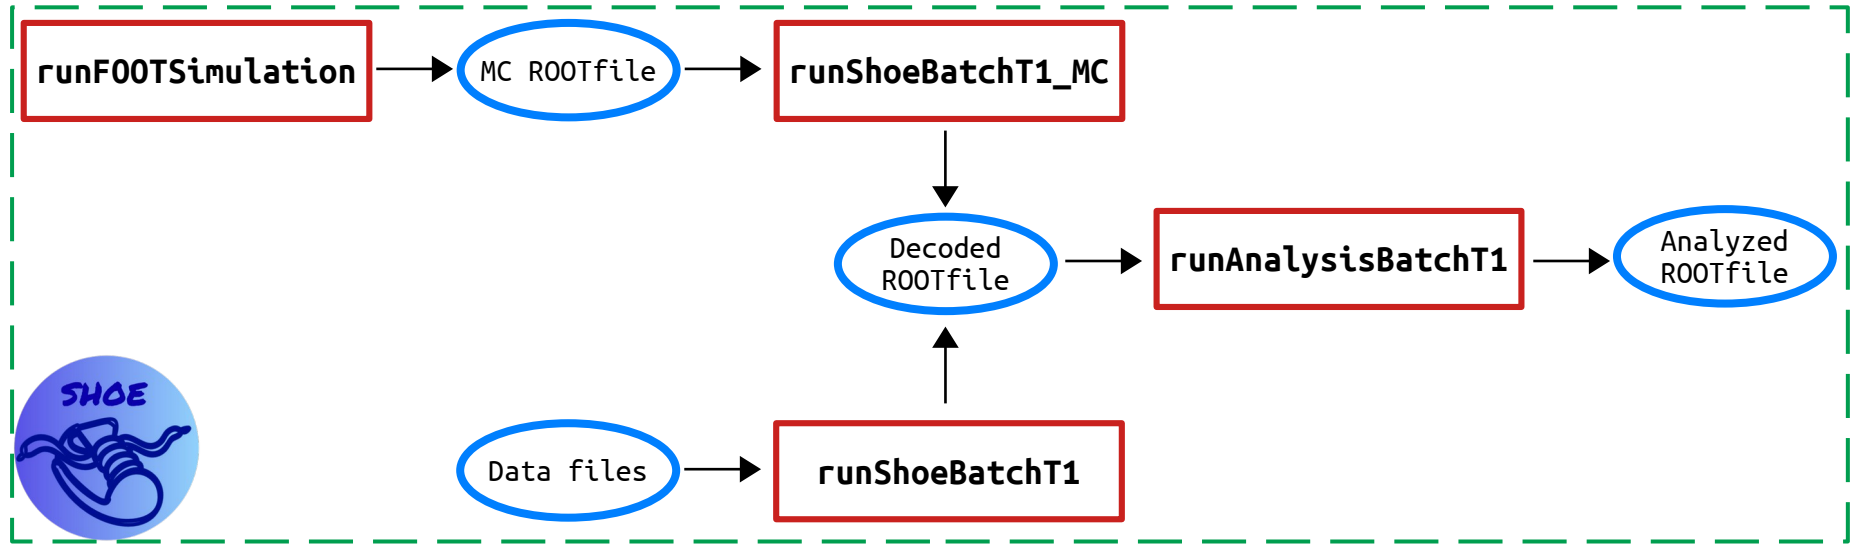
- 1) Create geometry with SHOE (“makeGeo”)
- 2) Create FLUKA input files w/ different seed
- 3) Run simulations in parallel
- 4) Convert to ROOT-files
- 5) Merge ROOT-files and save FLUKA aux files



Full FOOT workflow on Tier1!



Possible now to run all steps of the workflow on data center resources!



Full FOOT workflow on Tier1!

Blue args are optional!!!



Simulate

runFOOTSimulation

MC ROOTfile

- Arguments**
- o Output directory (in /storage/gpfs_data/foot/\${USER})
 - c Campaign name
 - r Run number
 - n Total number of events to simulate
 - e Events per job
 - k Keep TXT files in zip archive (optional, default "0")

Example

```
./runFOOTSimulation.sh -o /storage/gpfs_data/foot/${USER}/CNA02023_MC_200  
-c CNA02023_MC -r 200 -n 1000000 -e 1000 -k 1
```

Decode simulation

MC ROOTfile

runShoeBatchT1_MC

Decoded
ROOTfile

- Arguments**
- i Input file (in /storage/gpfs_data/foot/shared/)
 - o Output directory (in /storage/gpfs_data/foot)
 - m Merge output files (optional, default "1")
 - f Use full statistics (optional, default "0")

Example

```
./runShoeBatchT1_MC.sh -i  
/storage/gpfs_data/foot/shared/SimulatedData/CNA02023_MC/12C_C_200_1.root  
-o /storage/gpfs_data/foot/${USER}/results -m 1 -f 1
```

Decode data

Data files

runShoeBatchT1

Decoded
ROOTfile

- Arguments**
- i Input directory (in /storage/gpfs_data/foot/shared)
 - o Output directory (in /storage/gpfs_data/foot/)
 - c Campaign name
 - r First run number
 - l Last run number (optional)
 - m Merge output files (optional, default "1")

Example

```
./runShoeBatchT1.sh -i /storage/gpfs_data/foot/shared/DataGSI2021sync/ -o  
/storage/gpfs_data/foot/${USER}/results -c GSI2021 -r 4306 -l 4310 -m 1
```

Analyze

Decoded
ROOTfile

runAnalysisBatchT1

Analyzed
ROOTfile

- Arguments**
- i Input file (in /storage/gpfs_data/foot/)
 - o Output file (in /storage/gpfs_data/foot/, optional)
 - m Is MC (optional, default "0")
 - n No batch (optional, default "0")

Example

```
./runAnalysisBatchT1.sh -i  
/storage/gpfs_data/foot/${USER}/Merge_CNA02023_6106.root -o  
/storage/gpfs_data/foot/${USER}/MergeAna_CNA02023_6106.root -m 0 -n 0
```


Conclusions



Upgrading HTCondor w/ SHOE!

- **HTCondor updated** → easier to setup and use!
- Moved Condor aux files to \$HOME
- **New script for running simulations on Tier1**
 - All handled inside SHOE
 - Possibility to save FLUKA aux files
- **Full FOOT workflow interfaced with Tier1!**



Conclusions



Upgrading HTCondor w/ SHOE!

- **HTCondor updated** → easier to setup and use!
- Moved Condor aux files to \$HOME
- **New script for running simulations on Tier1**
 - All handled inside SHOE
 - Possibility to save FLUKA aux files
- **Full FOOT workflow interfaced with Tier1!**



*Automatic production of Decoded
ROOTfiles will start in some time...*





Backup slides