

Update on MJF vs simulation in LHCb



Ph. Charpentier
LHCb - CERN

- At September GDB:
 - First comparison between MJF and simulation event rates
 - Only 2 sites at the time:
 - ☆ CERN and GRIDKA
 - Several differences:
 - ☆ Not same absolute normalization
 - * Shows that Dirac benchmark is off in absolute scale by ~60%
 - ☆ Consistent results at GRIDKA
 - ☆ Less consistent at CERN
 - * Many dependencies on WN types
 - * Memory performance not taken into account
 - * In some cases, wrong information (nb of slots or cores)
 - Some effects from boosting of WNs
 - ☆ Event throughput > expected
 - * Not a problem for job matching
 - * ... nor for accounting (bonus)
- Today: update with more jobs and more sites
 - LPNHE, Imperial and LAL (since December 2nd, less stats)



- **New sites**
 - GRIF sites: LPNHE and LAL (since December 2nd)
 - Imperial College
 - Any other site?

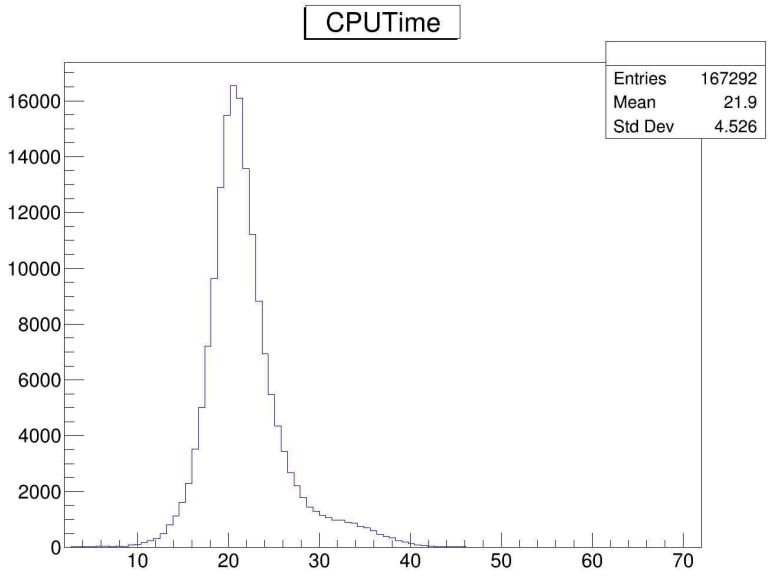
- **Known issues have been / are being fixed**
 - **Number of slots at GRIDKA**
 - **Still some minor issues at CERN (inconsistencies between LSF and MJF)**
 - ☆ **Number of slots**
 - ☆ **HS06 rating (should relate to LSF parameter cpuf)**
 - * $Cpuf = HS06perSlot / 3.9$
 - * CERN decided 1 kSI2k = 3.9 HS06
 - ☆ **Issue being fixed on HTCondor nodes**
 - * `$MACHINEFEATURES` not pointing to the standard location (`/etc/machinefeatures`), directory existing but with wrong information...



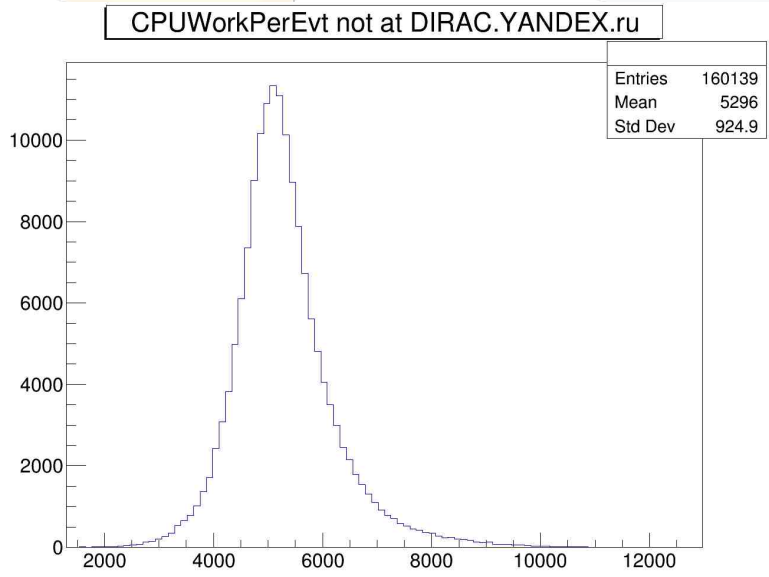
Reminder on MC production used

- Large MC production (> 200,000 jobs) targeting the whole Grid, single event type, therefore stable CPU work
- Each job reports:
 - CPU time, Dirac rating, MJF if available, number of events and many other parameters

MJF UPDATE



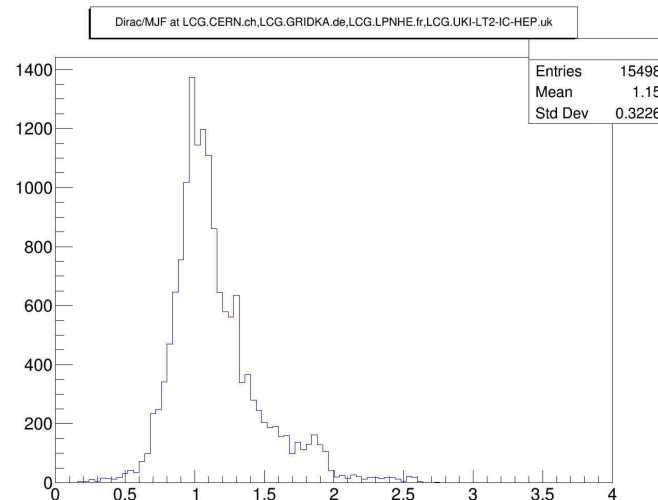
CPU time in hours



CPU-work per event in HS06.s



- CPU power:
 - DiracPower: result of Dirac's fast benchmark
 - ☆ Renormalized to match MJF "on peak" (factor 0.627)




- JobPower: renormalized event rate ($\sim 1/\langle \text{CPUTimePerEvt} \rangle$)
- MJFPower: single processor power as computed from MJF information
- CPU work:
 - CPUWork: computed from DiracPower
 - MJFWork: computed using MJF



CPU work per event per site (DiracPower)

Site vs CPUWorkPerEvt at LCG.CERN.ch, LCG.GRIDKA.de, LCG.LPNHE.fr, LCG.UKI-LT2-IC-HEP.uk




MJF UPDATE




JobPower vs MJF per CPU type

MJF UPDATE


WNModel vs Job/MJF at CERN




WNModel vs Job/MJF at GRIDKA





WNModel vs Job/MJF at LCG.UKI-LT2-IC-HEP.uk



WNModel vs Job/MJF at LCG.LAL.fr



WNModel vs Job/MJF at LCG.LPNHE.fr






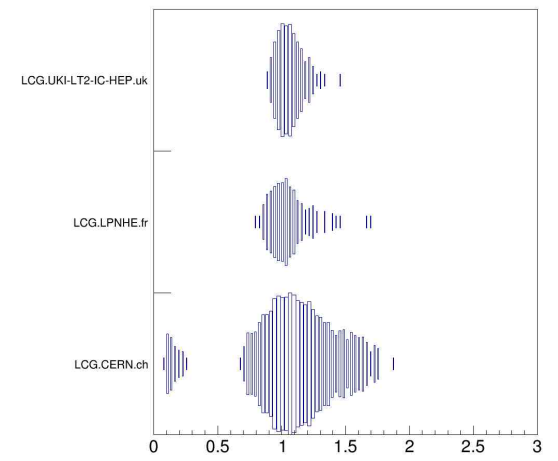
Site comparison for same model

MJF UPDATE


Job/MJF for Intel(R)Xeon(R)CPUE5-2630v3@2.40GHz




Job/MJF for Intel(R)Xeon(R)CPUE5-2650v2@2.60GHz




Job/MJF for Intel(R)Xeon(R)CPUE5420@2.50GHz




Job/MJF for Intel(R)Xeon(R)CPUE5520@2.27GHz



Job/MJF for Intel(R)Xeon(R)CPUE5-2670v3@2.30GHz



Job/MJF for Intel(R)Xeon(R)CPUE5650@2.67GHz


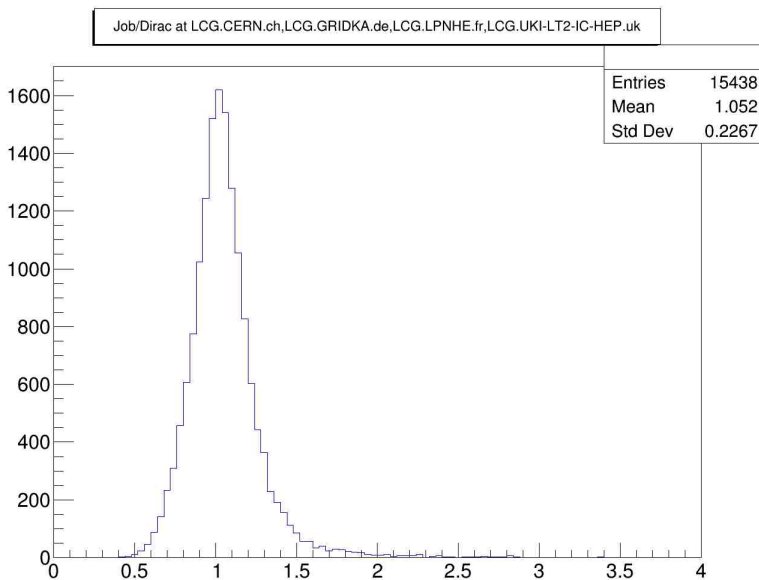




JobPower to DiracPower comparison

MJF UPDATE

- Using same sites as for MJF comparison
- Nice (almost) symmetric peak
- Small tail
- RMS 22%
 - Larger than with MJF... when it scales





- Still too few sites !
 - We must make sure HS06 is run in the same conditions
- LHCb simulation doesn't scale with HS06!
 - Big differences for different WN models
 - Mostly two/three classes
 - ☆ Up to 50% difference!

WNModel vs Job/MJF at LCG.CERN.ch, LCG.GRIDKA.de, LCG.LPNHE.fr, LCG.UKI-LT2-IC-HEP.uk

