# Nikhef Multicore Observations

Jeff Templon

Multicore TF #2 2014.01.28

J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Data set

- Jobs running on 26 and 27 january

J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Waiting time

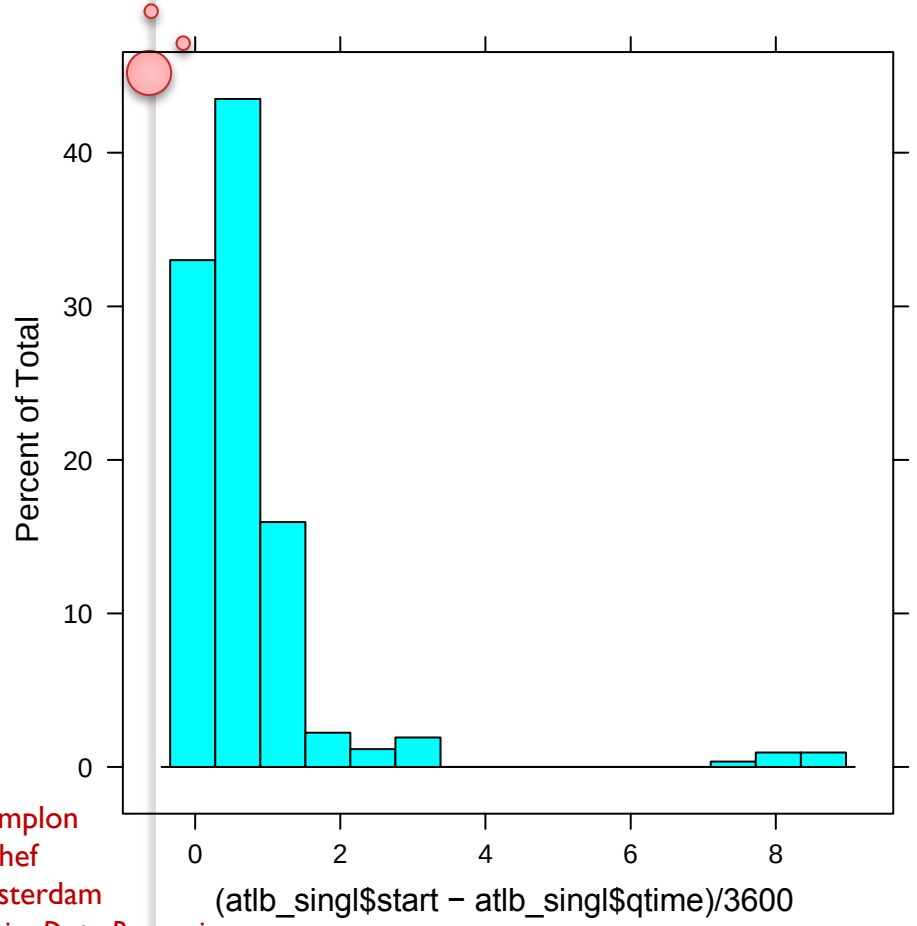


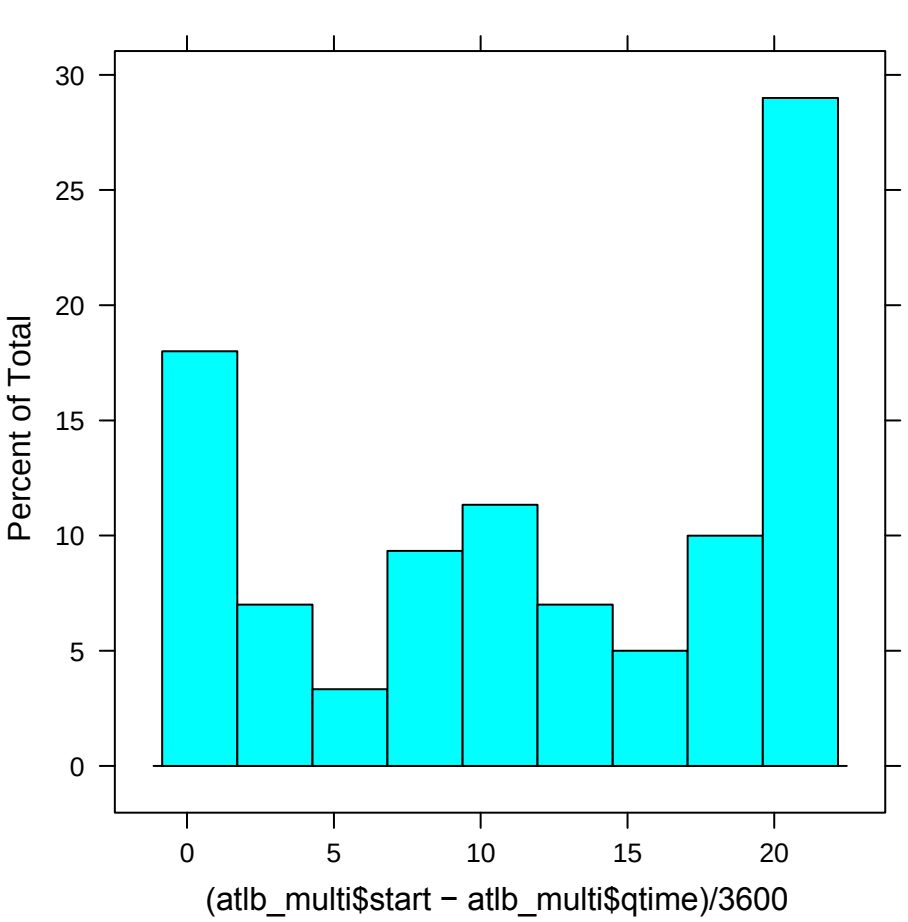


J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Short pilots

- Submission to short queue (4 hrs) helped multicore throughput a lot

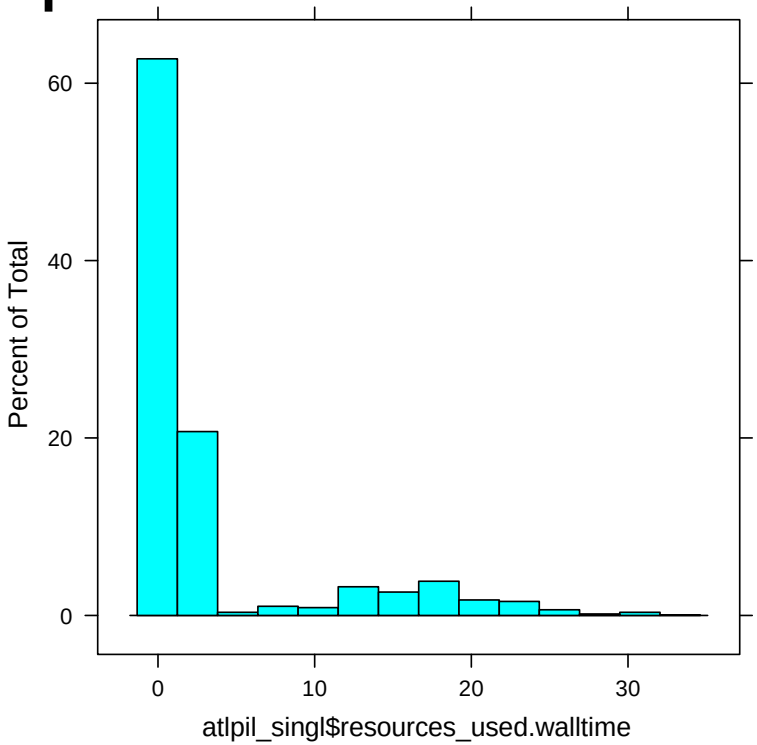


J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Production pilots

- Still could win some ground here with walltime limits

ATLAS prod submits normally to ‘atlas’ or ‘medium’ queues: 52 and 38 hours resp.

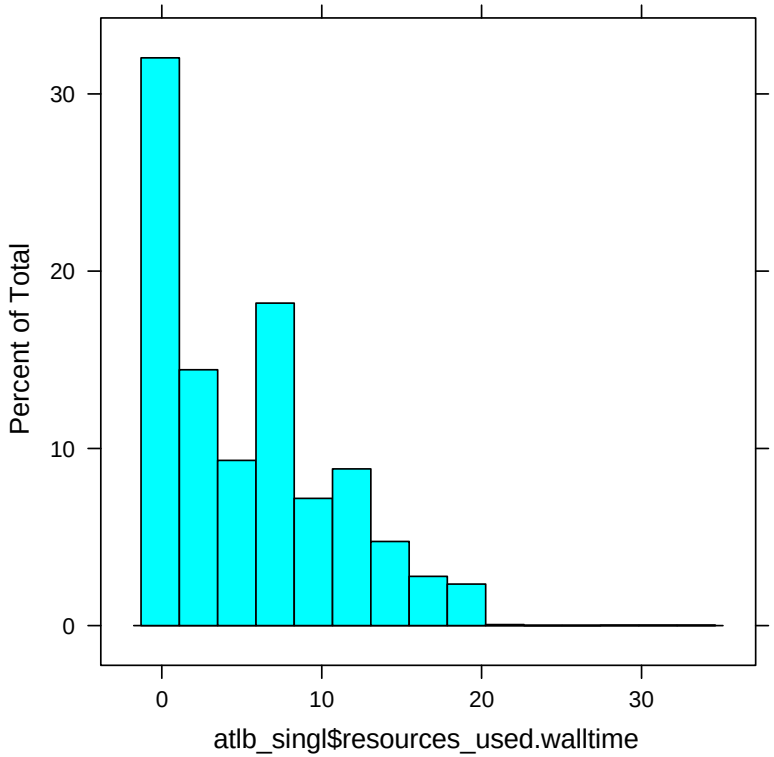


J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Multicore memory

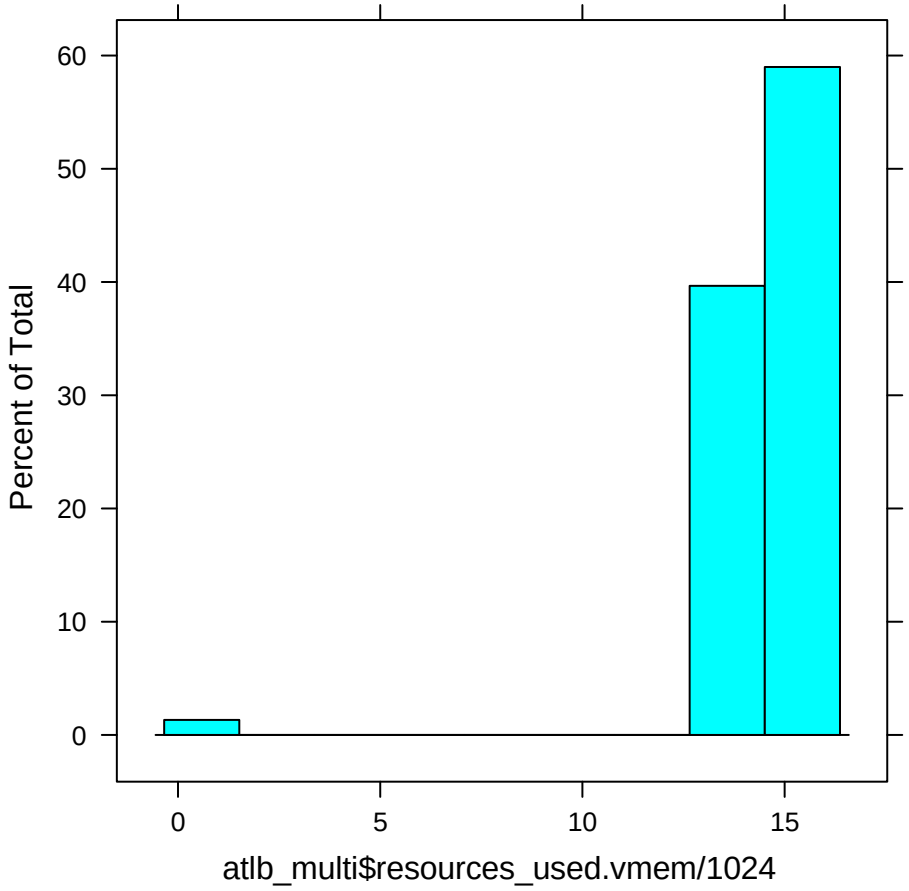


J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Memory and maui

- Maui checks mem; recall we use pvmem
- Maui assumes each proc needs 'pvmem'
- Submit filter assumes "multicore" : 8 cores = 8 times mem …
- Hence ATLAS was "asking for" 256 GB per job!!!
- Changed by hand to 32 GB per job; seems ok (jobs consume 16)

J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Maui & Memory II

- My preferred solution:
  - Ask for nr cores: assumes independent cores
  - Hence pvmem = 4096mb
  - Need true multicore? Ask for mem too
  - But then Maui will over-ask because it does not understand multicore
- "Not a problem" at moment? All ATLAS jobs use the indep. Core model

J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Holding onto cores

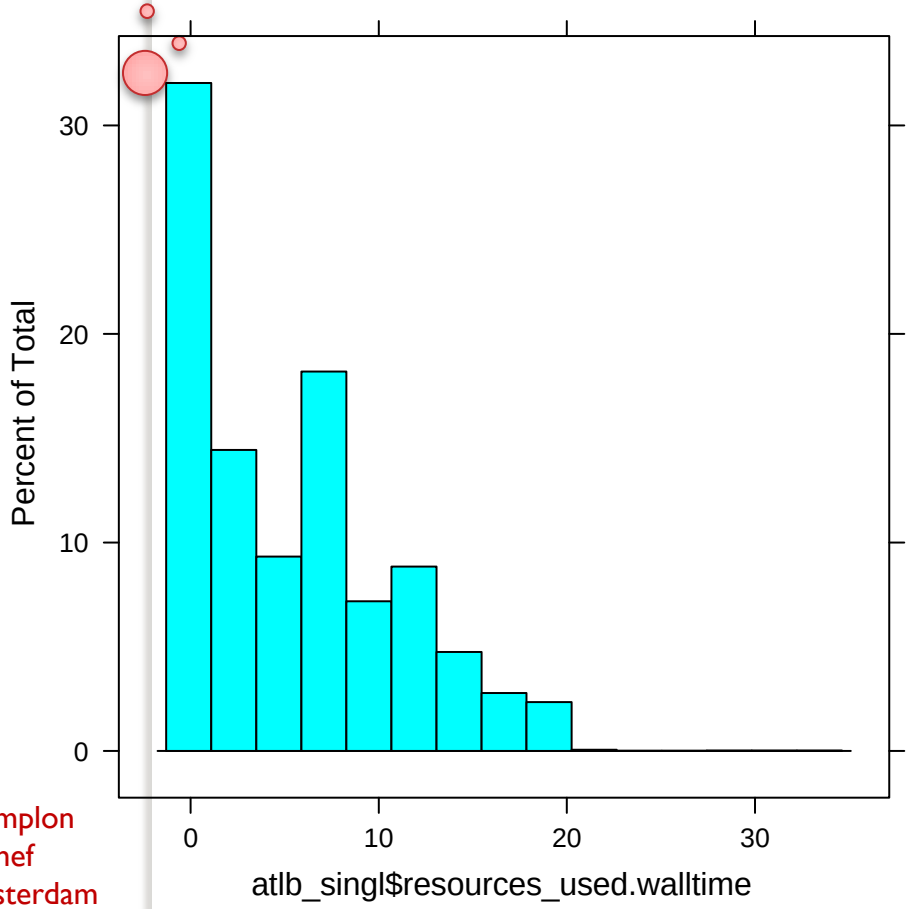


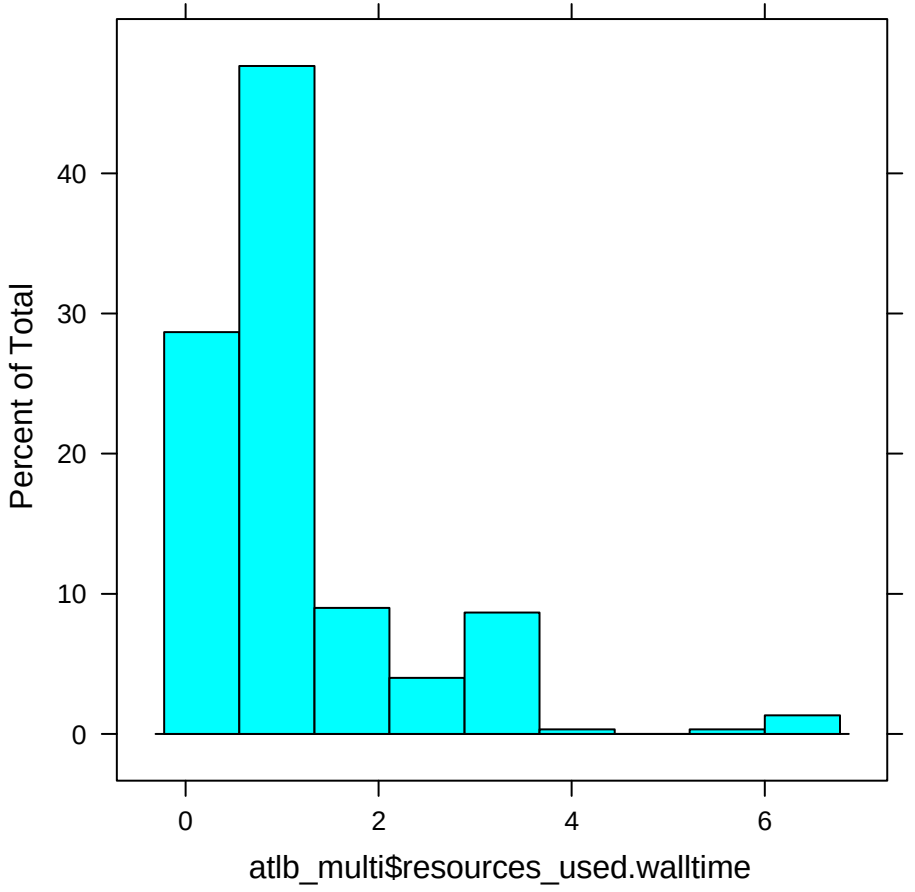


J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*

# Holding onto cores

- ATLAS doesn't
- Normal fair share fluctuations -> lose cores
- Probably need some dynamic partitioning to "multicore" segment. Dynamic means others can use it for single-core work when ATLAS workload disappears.

J. Templon
Nikhef
Amsterdam
*Physics Data Processing Group*