

The MINERvA Operations Report All Experimenters Meeting

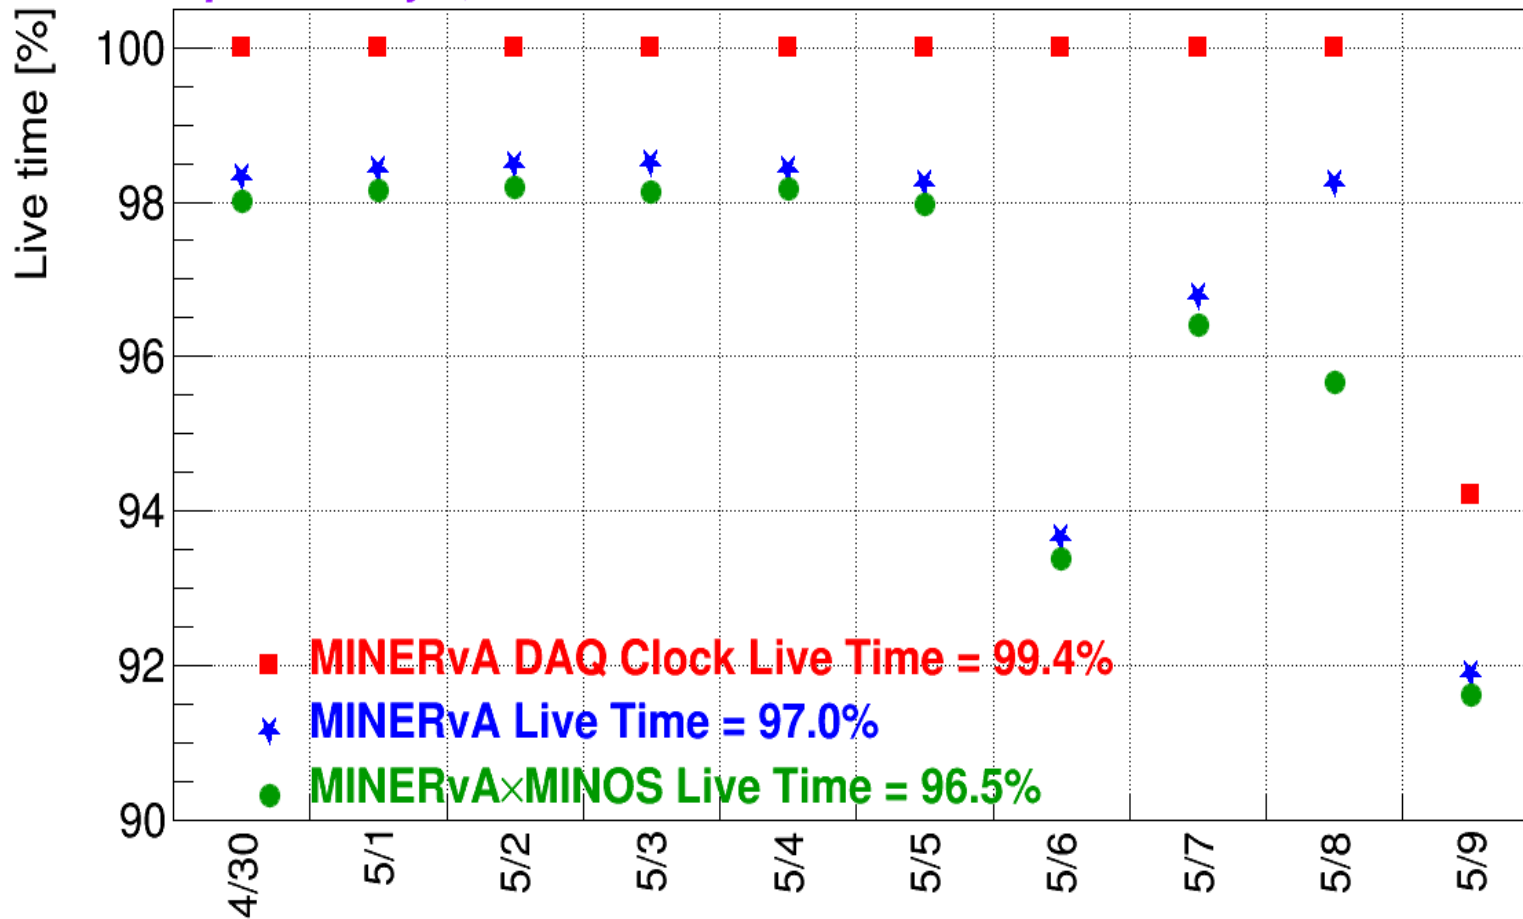
Howard Budd, University of Rochester
May 14, 2018



Anti- ν Data



Apr 30 - May 9, 2018: POT Delivered = 2.54×10^{19}





Anti- ν Data



- May 6 & 7, 95.0% MINERvA live
 - Not all the subruns were processed
- May 9, 91.9% MINERvA live
 - A CROCE hardware problem caused a problem in the data frame which has the HV. The DAQ had a problem starting subruns, as it was trying to determine if the HV was OK. The DAQ was finding it could not unpack the frame.
 - In addition, this affected the data. Donatella's unpacking changes fixed this data frame problem, so data analysis is unaffected.
 - The CROCE was replaced.
- During the May 11 shutdown we replaced 2 MINDER boards.

Average Jobs Running Concurrently

1003



Total Jobs Run

189608**May 7-13**

Average Time Spent Waiting in Queue (Production)

16.2 min

Running Jobs




Completing and Efficiency Stats

Job Success Rate




Job Success & Failures per Day



Overall CPU Efficiency



Total Time Wasted by Running Jobs



New row

New Data Cataloged

1.6 TB

Total Data Cataloged

2.3 PB

- Average concurrent jobs is lower than the average quota (average quota is ~1600)
- Job success rate is 100%
- Overall CPU efficiency is 62%
 - Small fraction of jobs (~7%) spent more time to access input files



Computing



- GENIE
 - GENIE job is very short (20~30 min level), so the efficiency is underestimated in FIFEMon.
 - GENIE job accesses flux files through StashCache and the performance of StashCache is not good due to dCache issue.
 - FIFE group is working on investigating why dCache doesn't work well for GENIE.