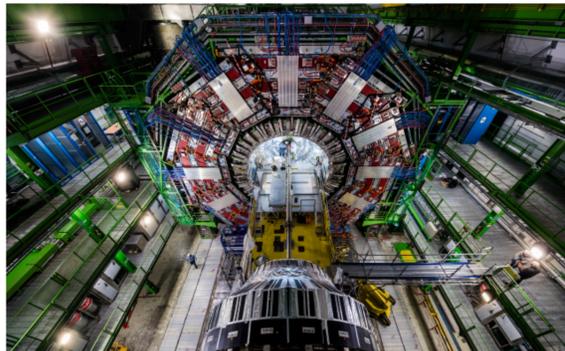


# Concurrent Conditions Access across Validity Intervals in CMSSW

Dr W David Dagenhart, Dr Christopher D Jones - Fermi National Accelerator Lab  
FERMILAB-POSTER-19-129-CMS-SCD

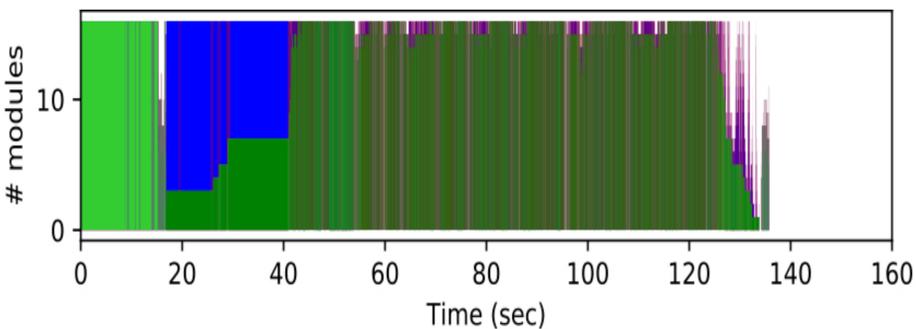
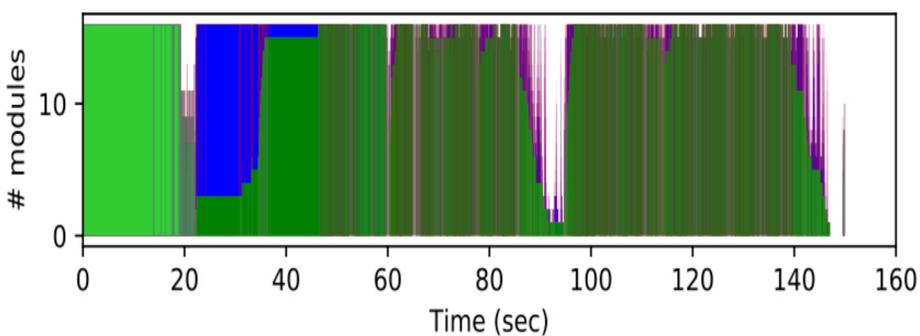
## Multithreaded Framework of the CMS Experiment at LHC



- CMS has been running a multithreaded Framework in its production executables since 2016. This Framework is built using Intel's *Threading Building Blocks (TBB)*.
- In the past, the Framework has supported multithreaded processing of events, luminosity blocks containing events, and modules within an event.
- Conditions data is typically read from a database or derived from data read from a database and includes calibration and alignment data and geometry information.
- Conditions data typically has a range over which it is valid (an *Interval of Validity* or **IOV**). With recent improvements, the Framework can process IOVs concurrently.

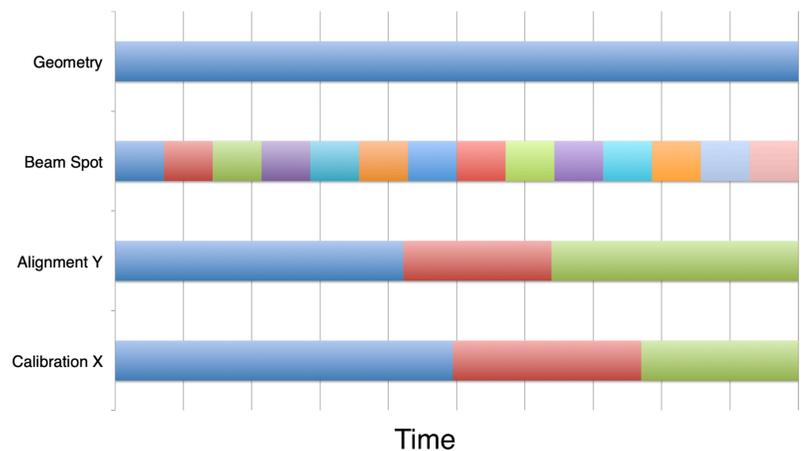
## Eliminate Inefficiency at Interval Boundaries

Inefficiency can occur at IOV boundaries when CPUs are waiting for the next IOV to start and have no work to do.



Number of modules running during reconstruction of 151 events with 16 threads. In top plot, events from only 1 interval of validity (IOV) are processed at a time. In the bottom plot, events from 2 IOVs are processed at the same time. There is an IOV boundary between 80 and 100 seconds. The disappearance of the dip is significant and the purpose of the plots is to show this. (The overall change in efficiency cannot be inferred from these plots both because it depends highly on the use case, and because these tests were not designed to measure it. It will become more significant as the number of threads increases and the length between IOV boundaries decreases.)

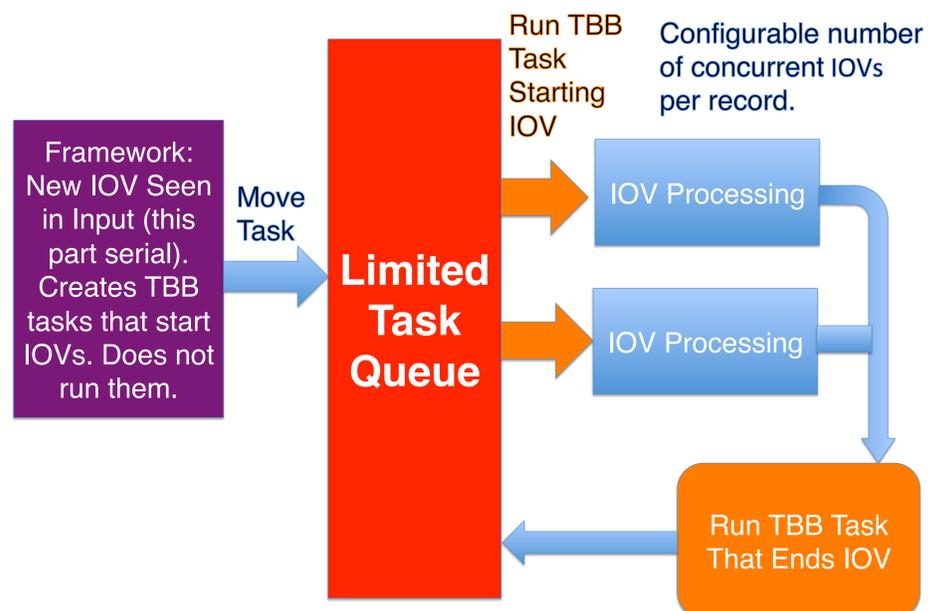
## Interval of Validity (IOV)



- **IOV**: Time period over which data is valid
- Data sharing same IOVs are grouped into a **record**
- 100's of records defined in CMS
- IOVs for different records vary widely in duration
- Memory required for a record can vary widely
- **Number of concurrent IOVs configurable by record!**

## Limited Task Queue

- The first arrow depicts TBB Tasks being moved into a LimitedTaskQueue. Each of these tasks starts one IOV.
- Second set of 2 arrows depicts running the TBB tasks that start processing for an IOV. Each of these tasks creates a different TBB task that runs after waiting for all data for that IOV to finish processing.
- The last set of arrows depicts running the TBB task that finishes an IOV. Among other things, this task tells the limited queue that an IOV slot is available.



CHEP 2019 – Also see these CMS Framework Presentations in Track 5 on Monday, Nov. 4:  
*Using OpenMP for HEP Framework Algorithm Scheduling*, Chris Jones, 11:45  
*Bringing heterogeneity to the CMS software framework*, Oliver Gutsche, 14:30