Computing Strategy of the AMS-02 Experiment

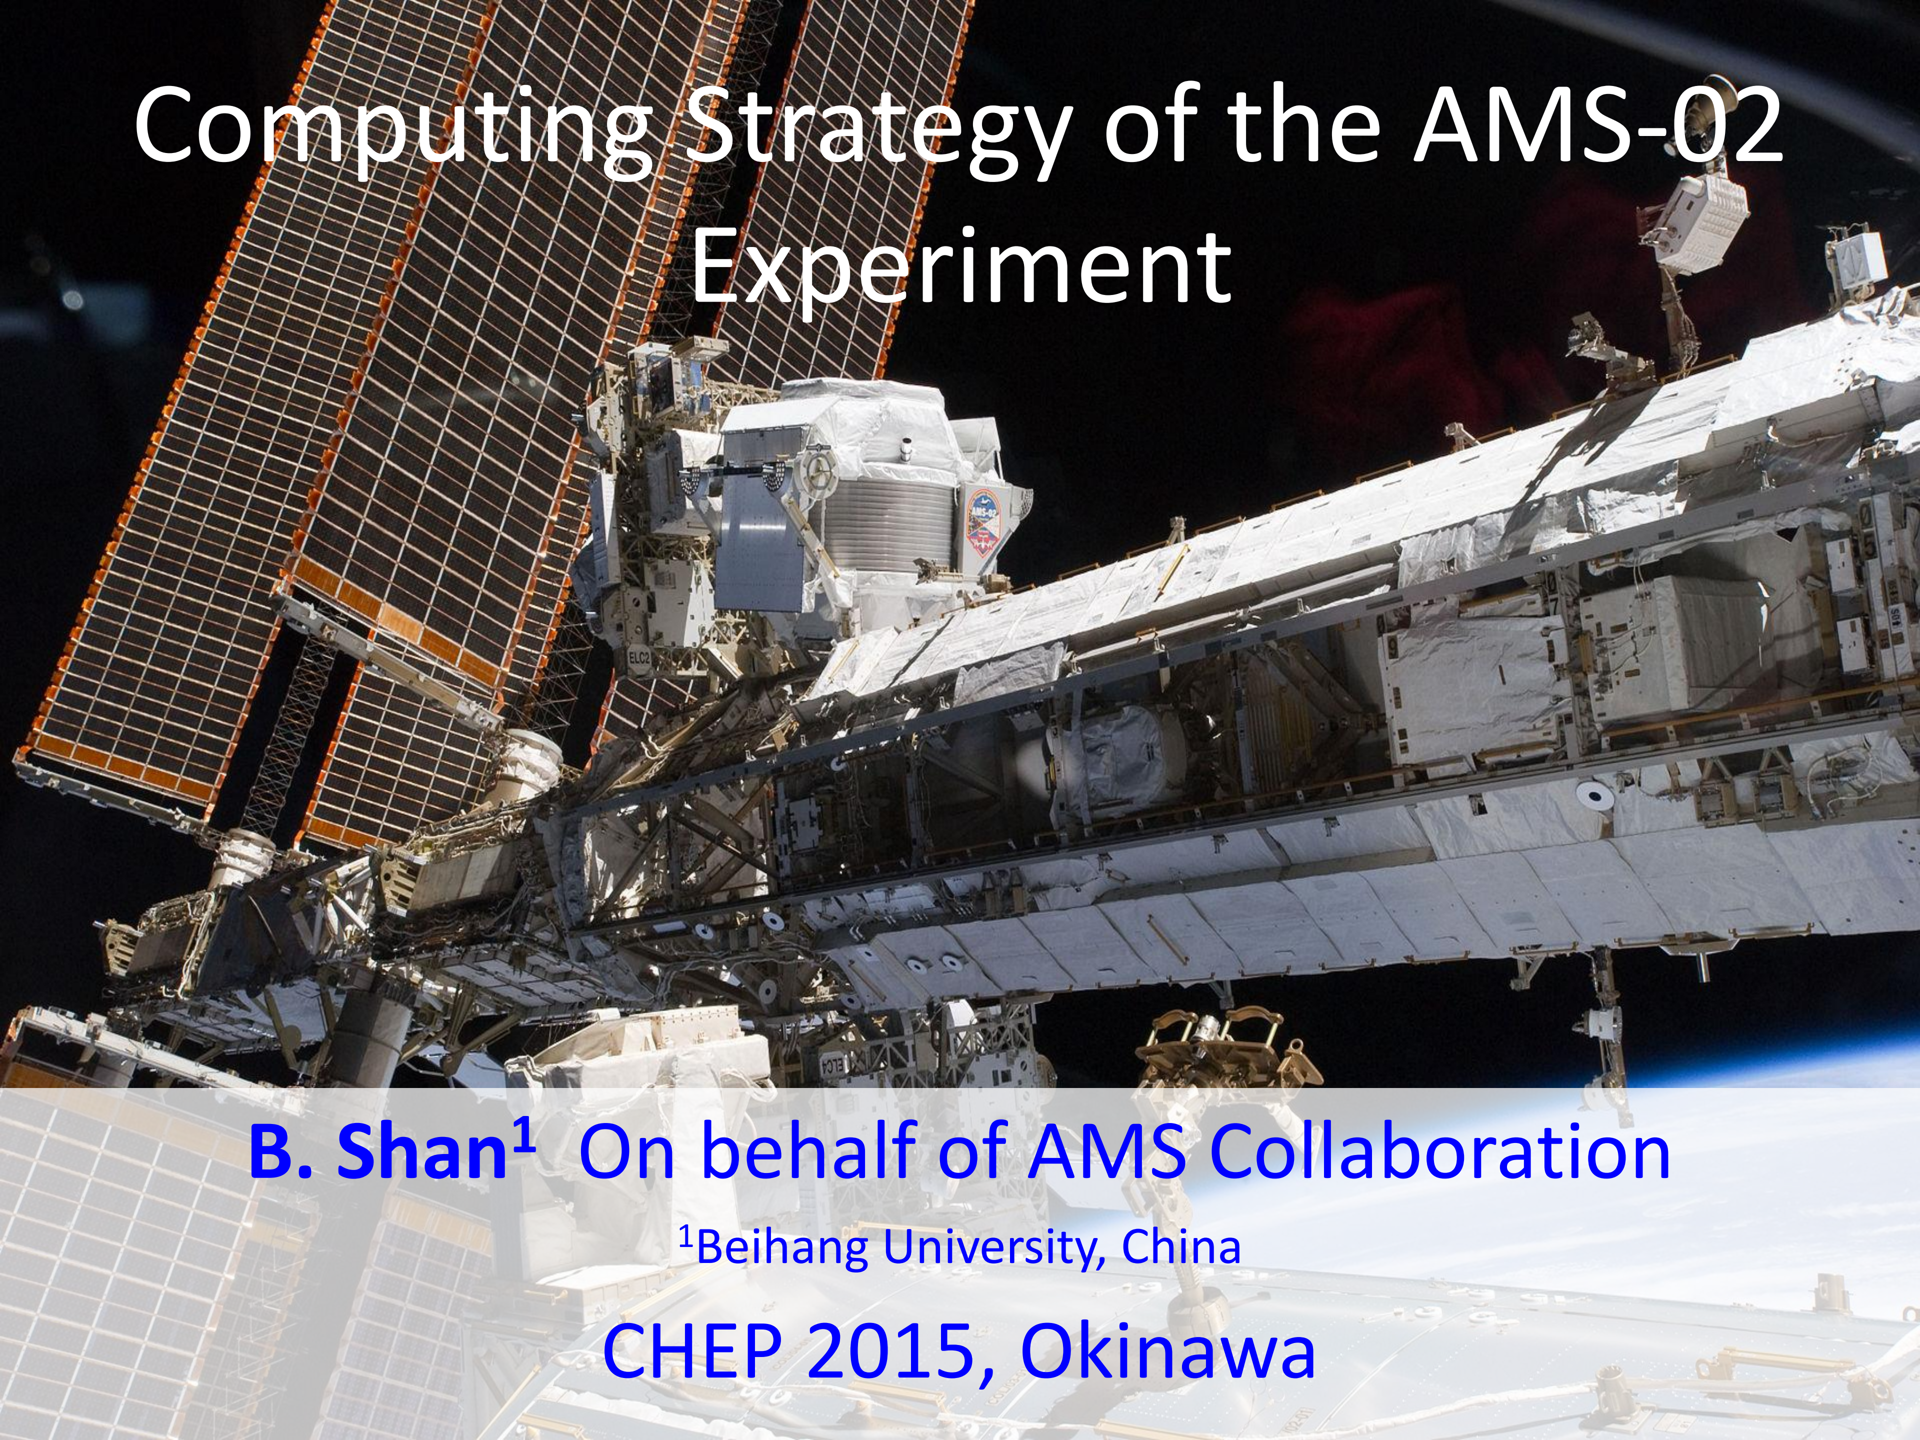

B. Shan[1] On behalf of AMS Collaboration

[1]Beihang University, China

CHEP 2015, Okinawa

AMS Experiment

- The Alpha Magnetic Spectrometer (AMS) is a high energy physics experiment on board the International Space Station (ISS), featured:
  - Geometrical acceptance: **0.5 m^2•sr**
  - Number of ReadOut Channels: ≈**200K**
  - Main payload of Space Shuttle **Endeavour's last flight** (May 16, 2011)
  - **Installed on ISS on May 19, 2011**
  - 7x24 running
  - Up to now, over **60 billion events** collected
  - Max. event rate: **2KHz**

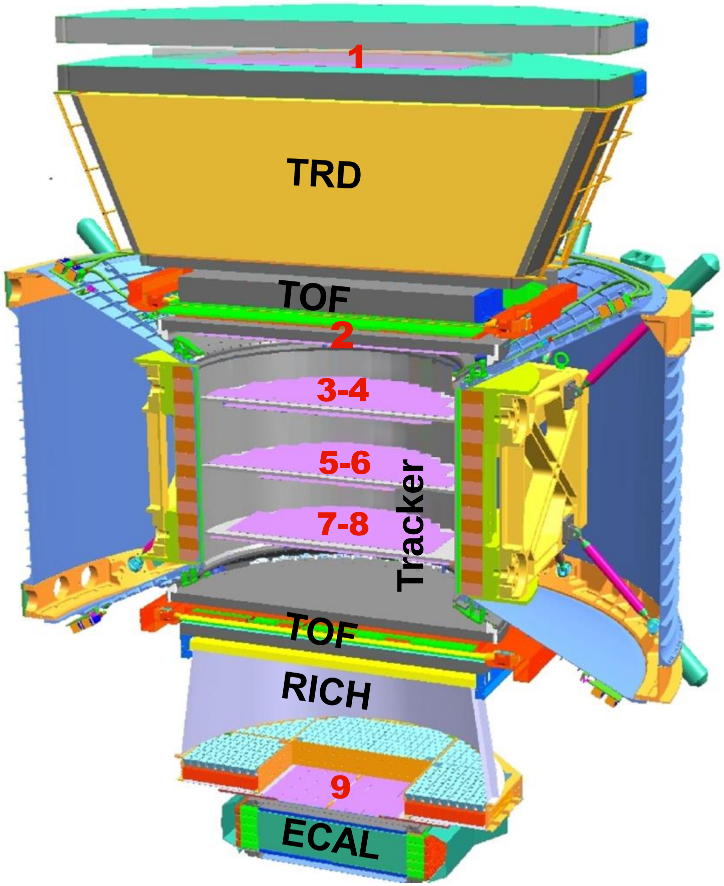

AMS Data Flow

- Data transferred via relay satellites to Marshall Space Flight Center, then to CERN, nearly real-time, in form of one-minute-frame
- Frames are then built to runs (RAW): 1 run = ¼ orbit (~23 minutes)
- RAW files are reconstructed to ROOT files

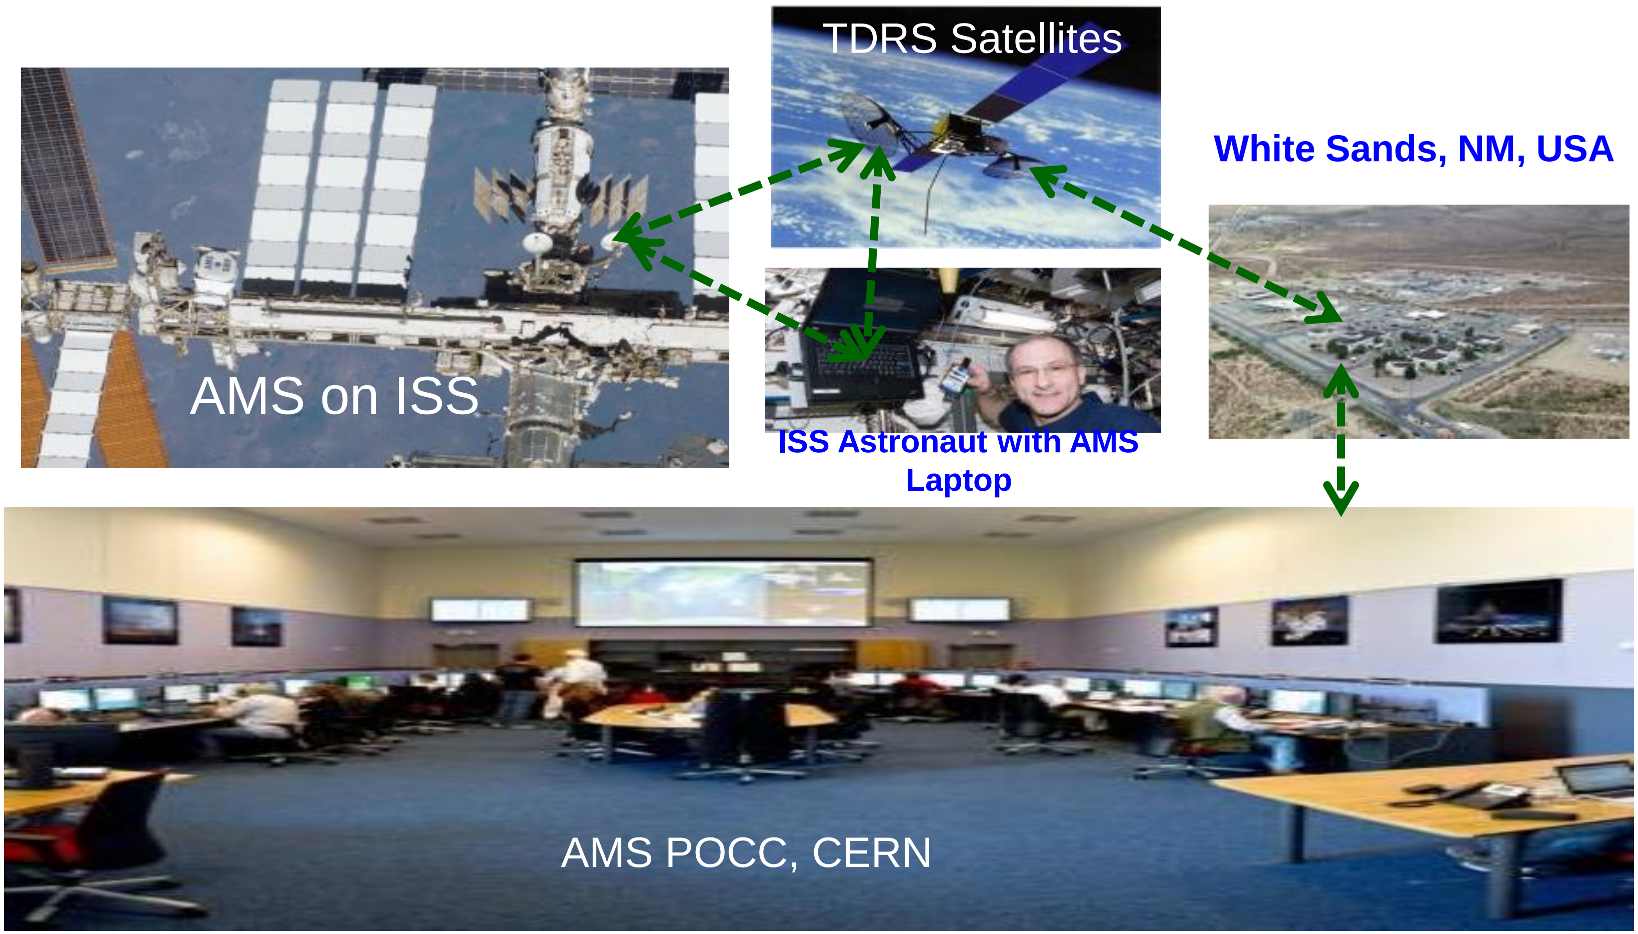

AMS Flight Data Reconstruction

- First Production
  - Runs 7x24 on freshly arrived data
  - Initial data validation and indexing
  - Produces Data Summary Files and Event Tags (ROOT) for fast events selection
  - Usually be available within 2 hours after flight data arriving
  - Used to produce various calibrations for the second production as well as quick performance evaluation
- Second Production
  - Every 3-6 months incremental
  - Full reconstruction in case of software major update
  - To use all the available calibrations, alignments, ancillary data from ISS, and slow control data to produce physics analysis ready set of reconstructed data

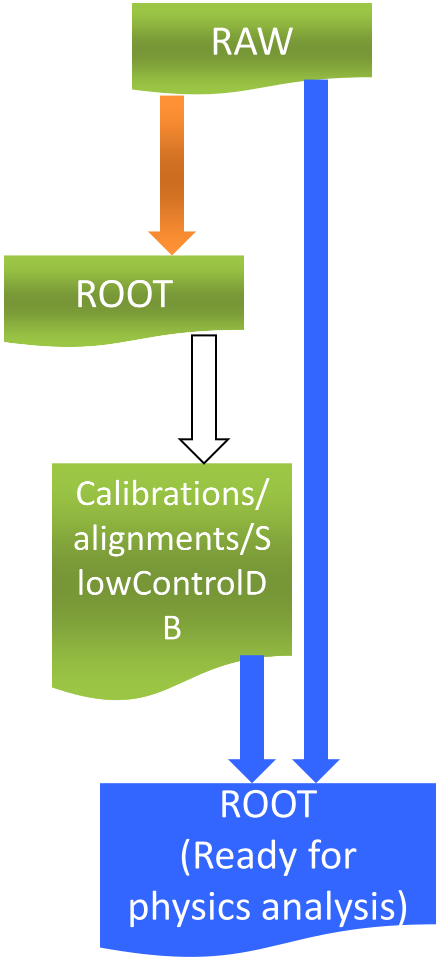

AMS-02 Major Computing Resources

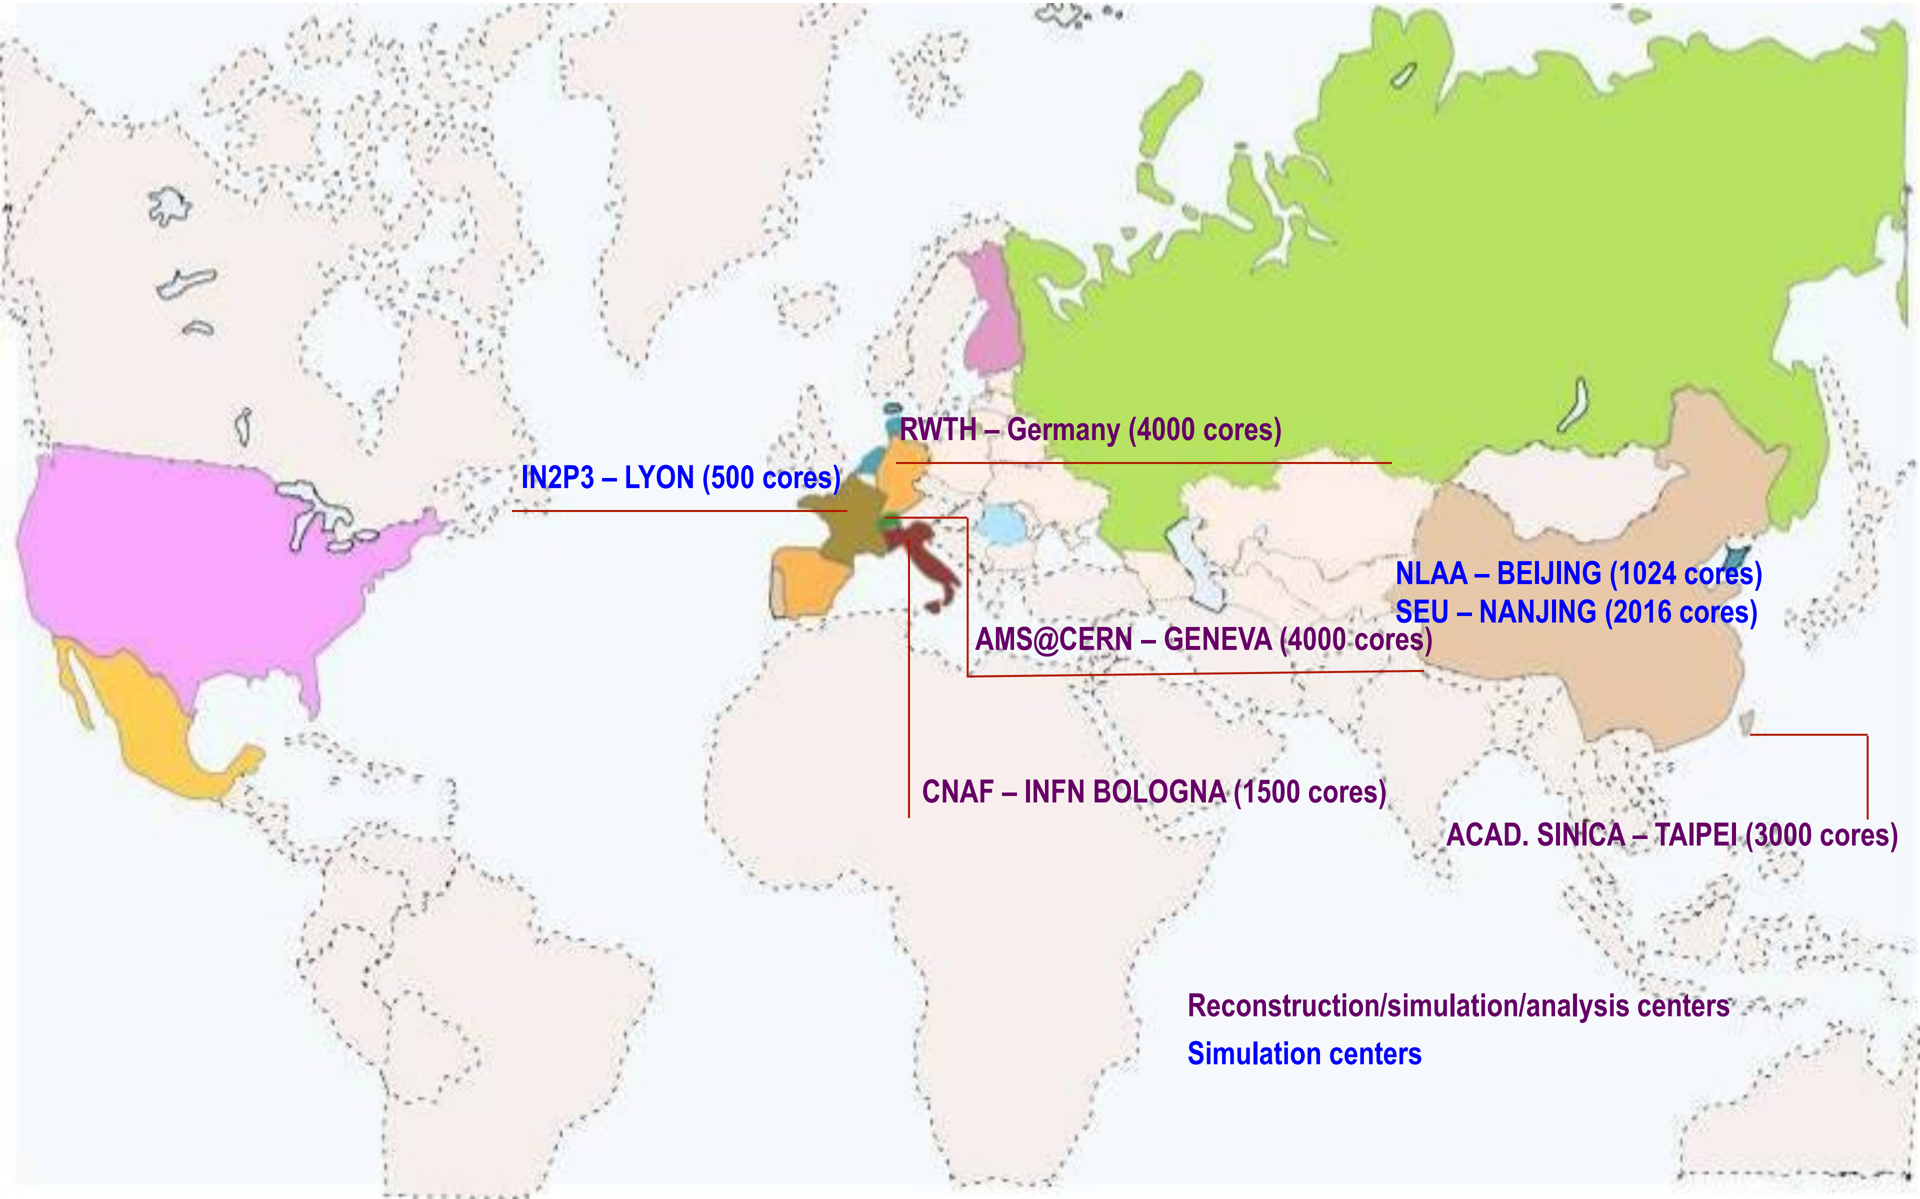

Light-weight production platform

- Fully-automated production cycle
  - Job acquiring, submission, monitoring, validation, transferring, and (optional) scratching
- Easy to deploy
  - Based on Perl/Python/sqlite3
- Customizable
  - Batch system, storage, transferring, etc.
- Running at:
  - JUROPA, CNAF, IN2P3, NLAA

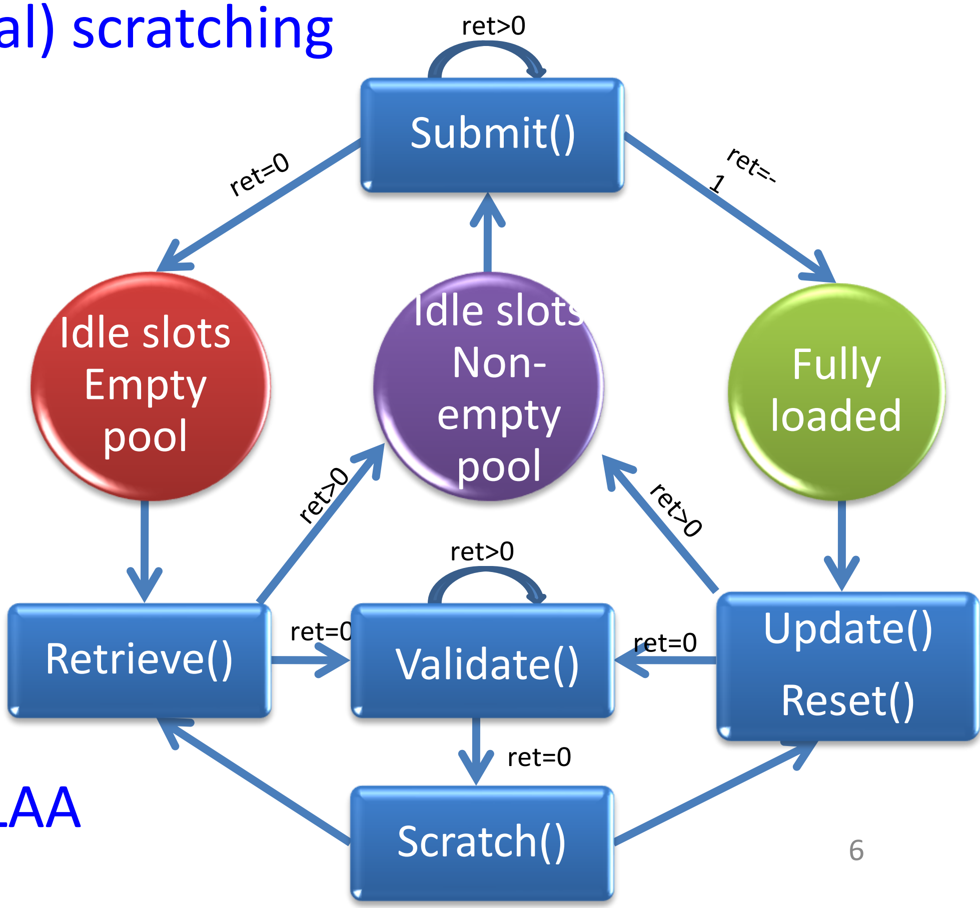

Latest Full Data Reconstruction

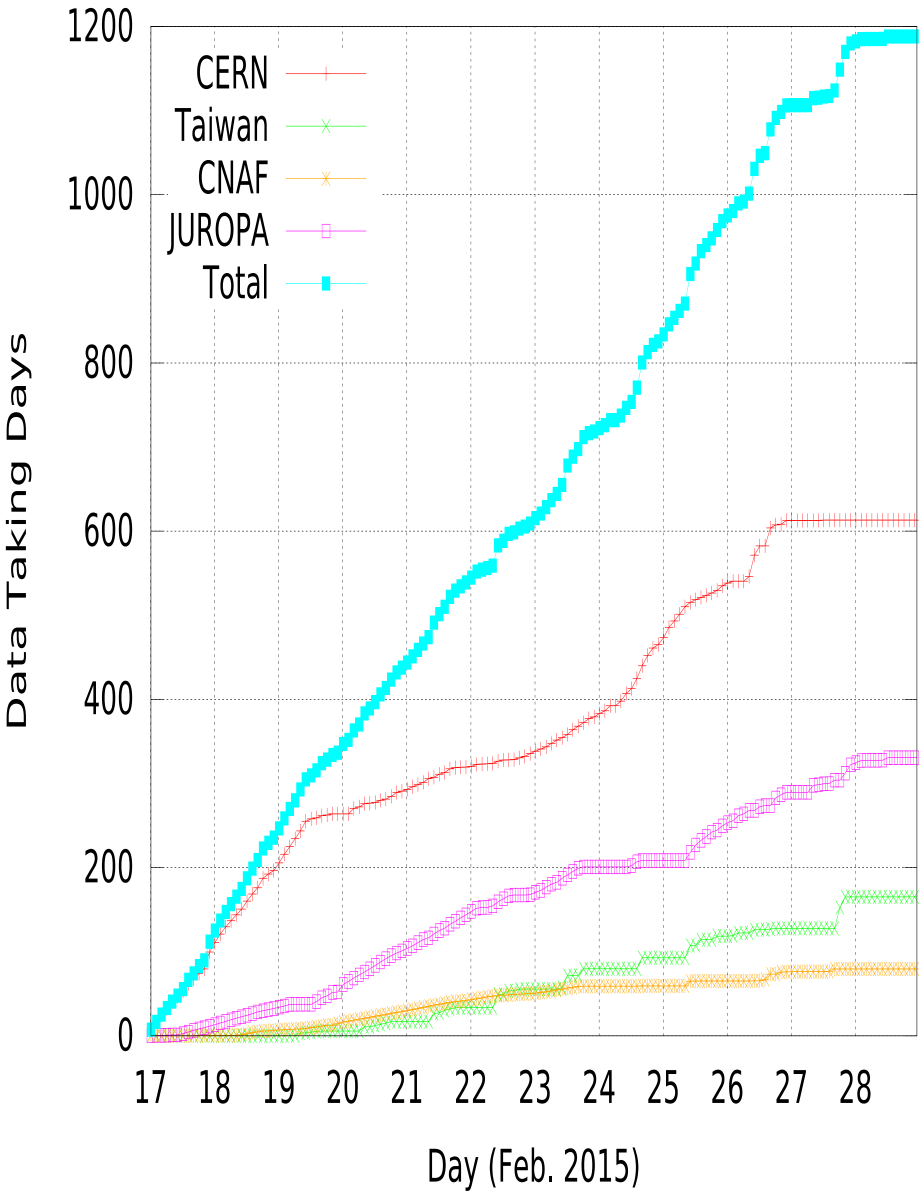

- Second half of Feb. 2015
- Performed at CERN, JUROPA, CNAF, ACAD. SINICA, and all provided us exclusive resources
- 40 months' flight data
- Reconstruction completed essentially in two weeks
- 100 times data collection speed
- 10^{-5} failing rate

Parallelization of Reconstruction

- Motivations
  - To shorten the (elapsed) time of data reconstruction
  - To increase productivity by using multi-core/hyper-thread enabled processors
  - To reduce memory usage per core
  - To ease the production jobs management
- Tool: OPENMP
  - No explicit thread programming
  - Nearly no code change except few pragmas

Parallel Processing

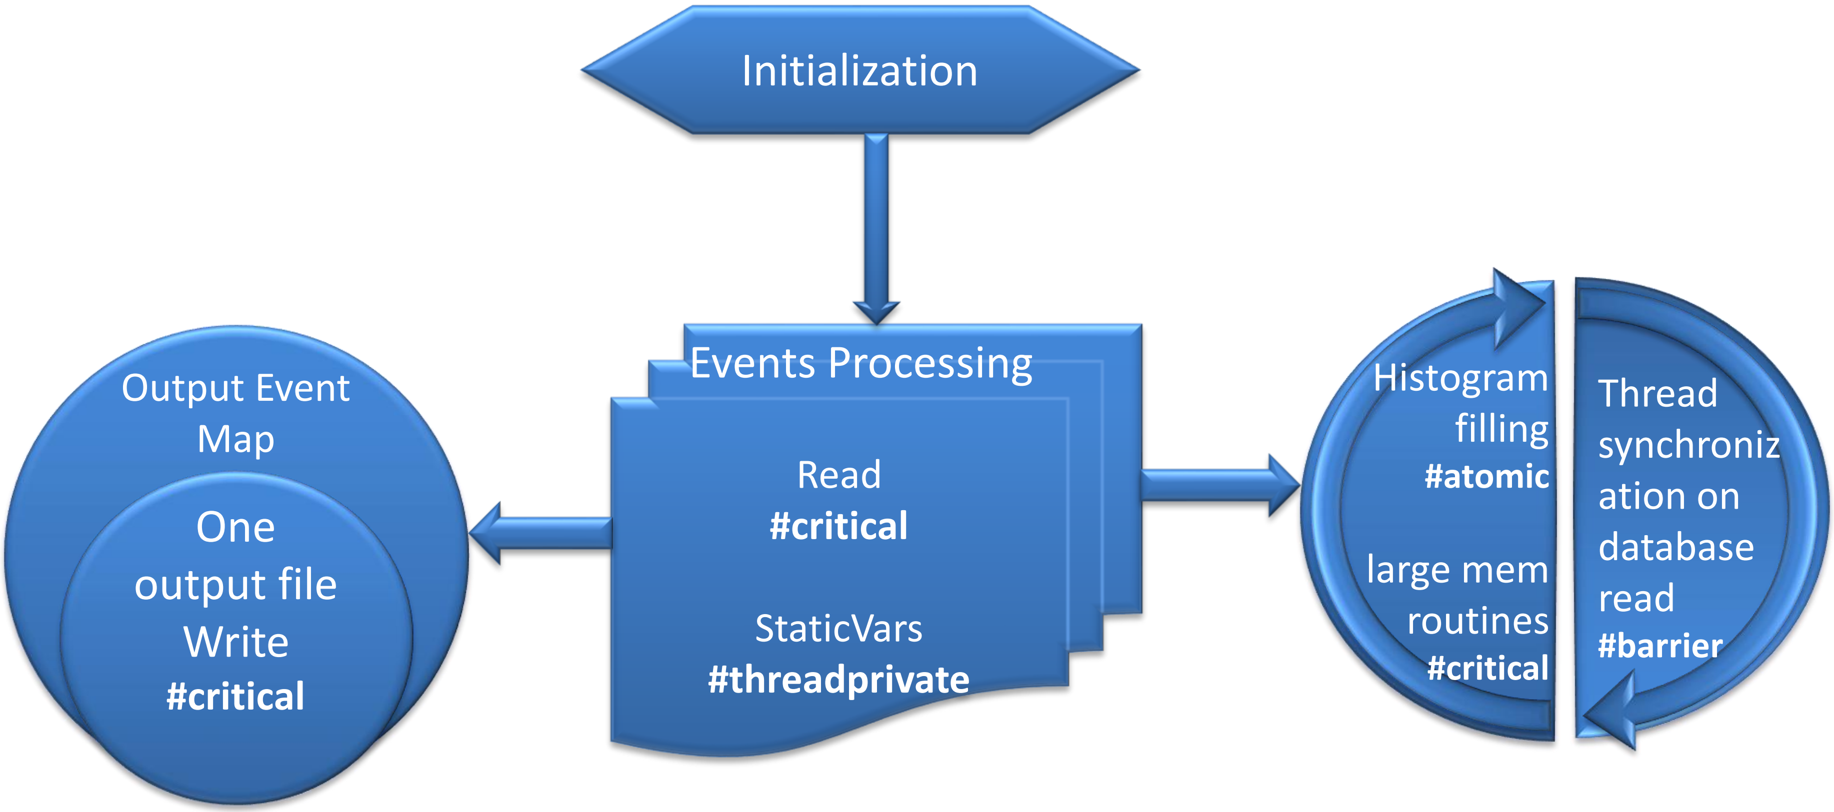

Parallel Reconstruction Performance
(Xeon E5-2699 v3 2.30GHz, max. 18 threads per job)

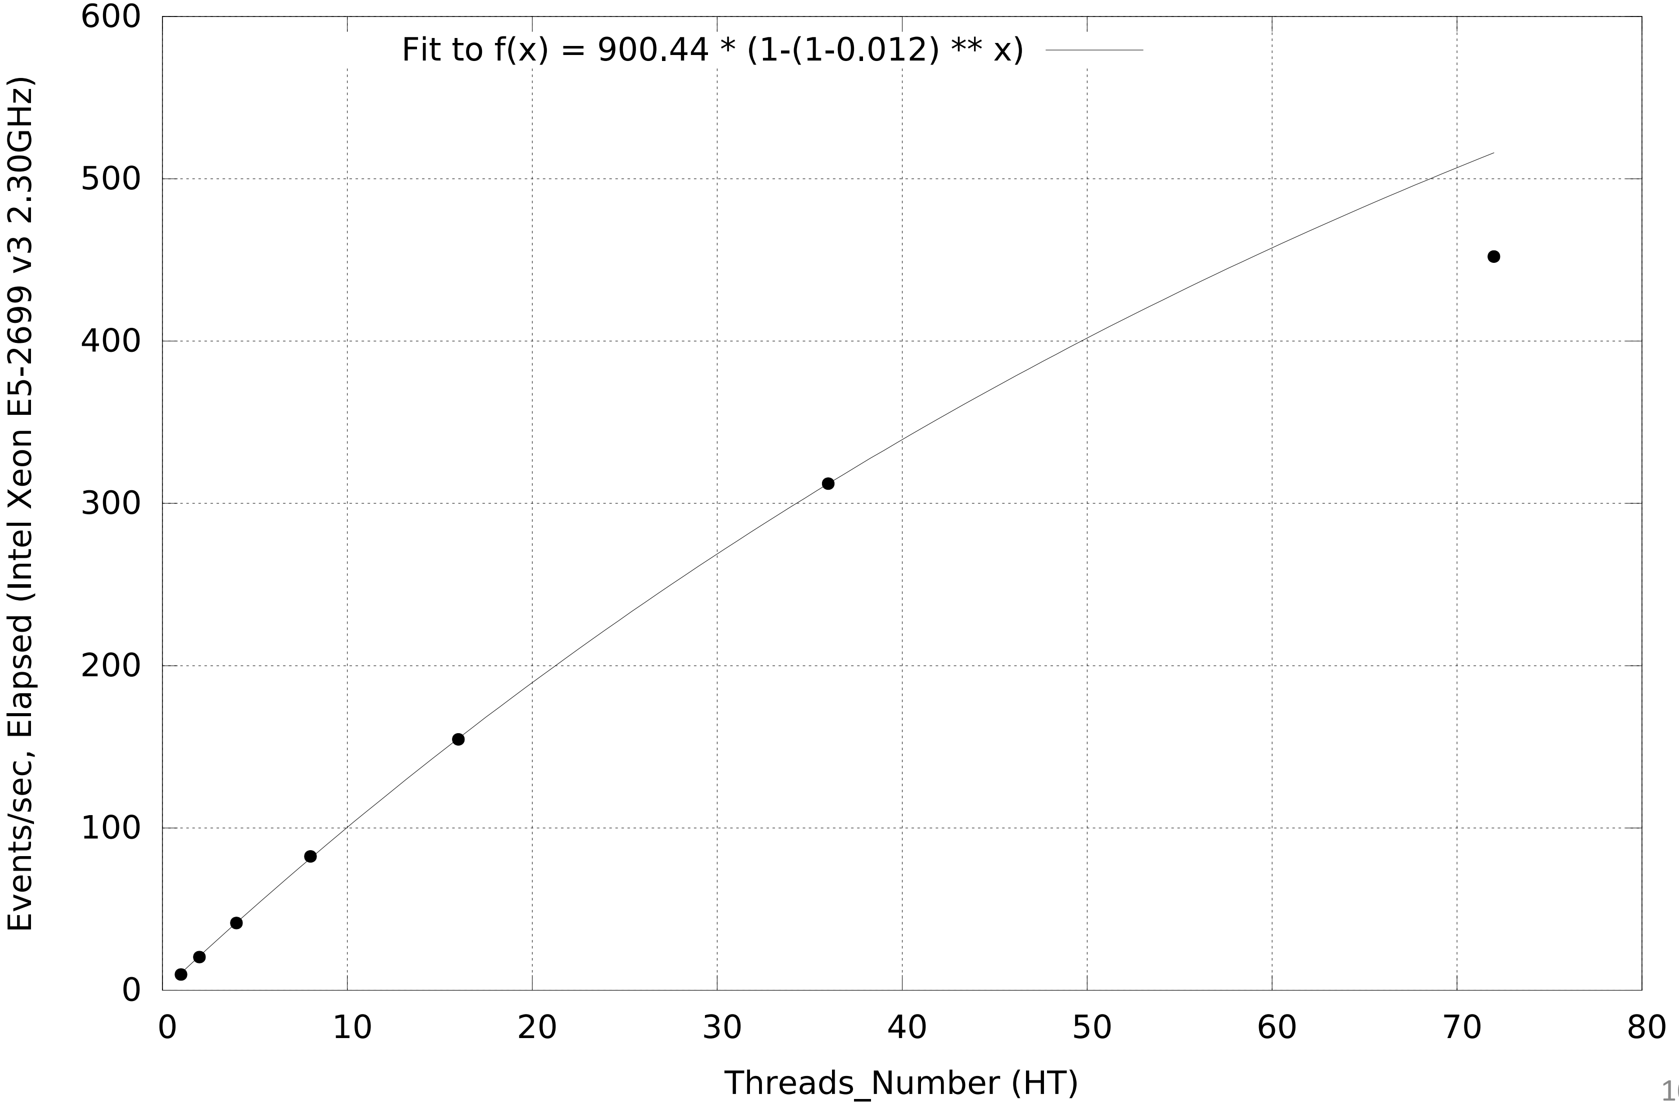

Parallel Reconstruction Performance (Intel MIC, no software change)

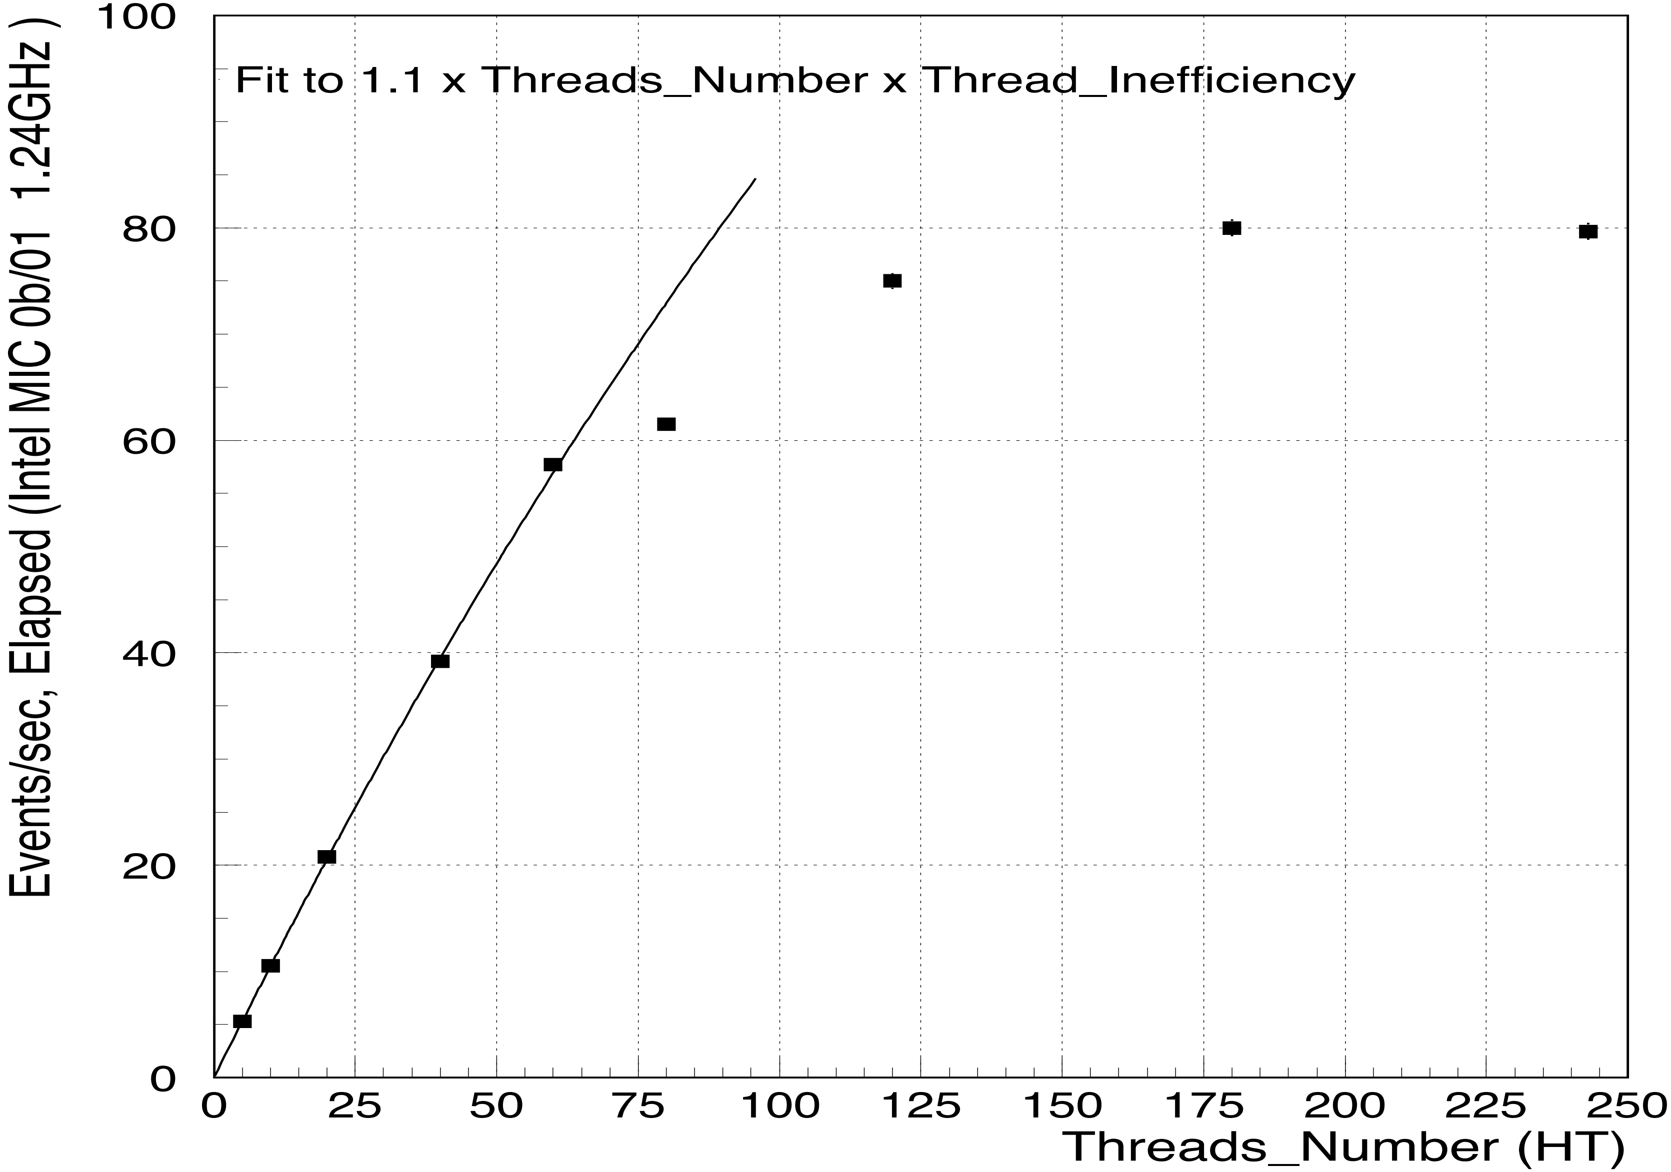

Parallelization of Simulation Software

- Motivations
  - To decrease required memory amount per core
  - To shorten elapsed running time so as to allow faster Monte-Carlo parameter/software tuning
- Minimal software change:
  Original Geant4.10.1 MT model + OPENMP (ICC 15)
  - Thread ID, thread number, and core number are taken from Geant4.10.1, instead of from OPENMP
  - Barrier from OPENMP does not work, so barriers are implemented
  - Other OPENMP pragmas work (minimal changes to AMS software)
  - No changes to original G4.10.1

Parallel Simulation Performance (Xeon E5-2699 v3 2.30GHz)

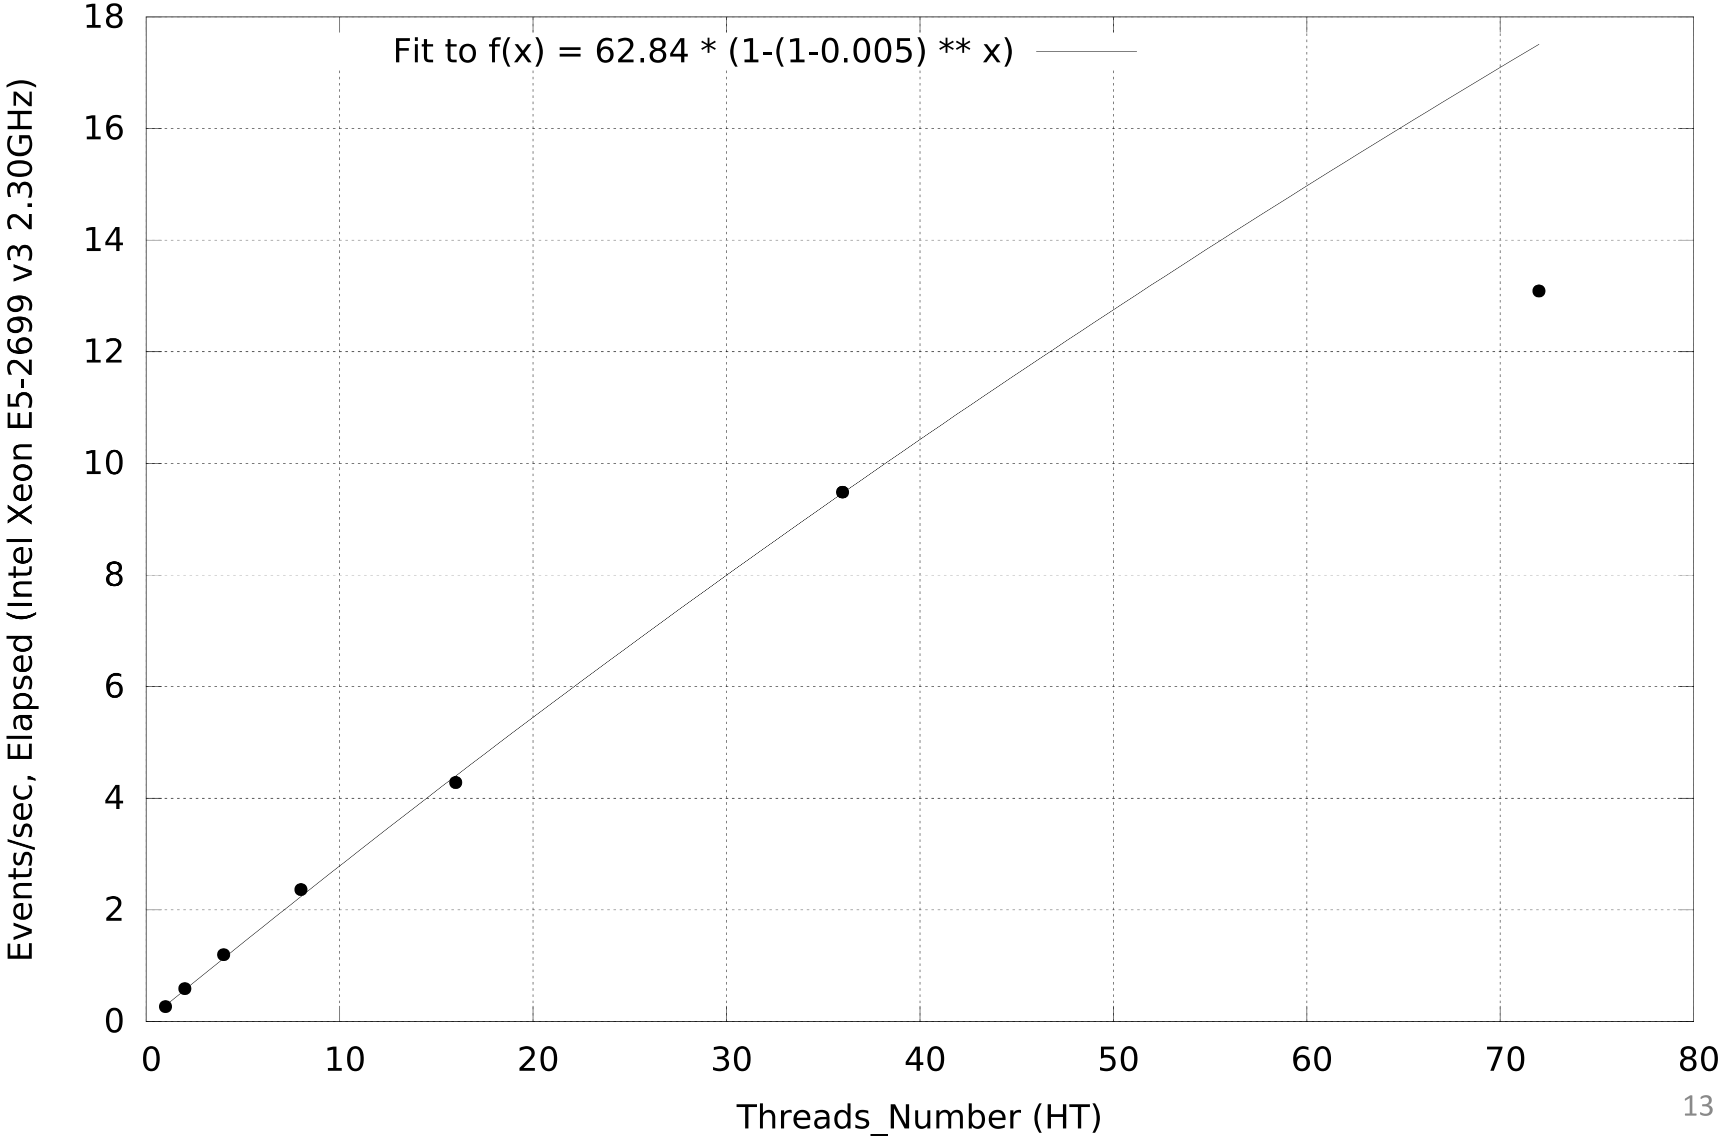

Parallel Simulation Performance (Intel MIC, no software change)

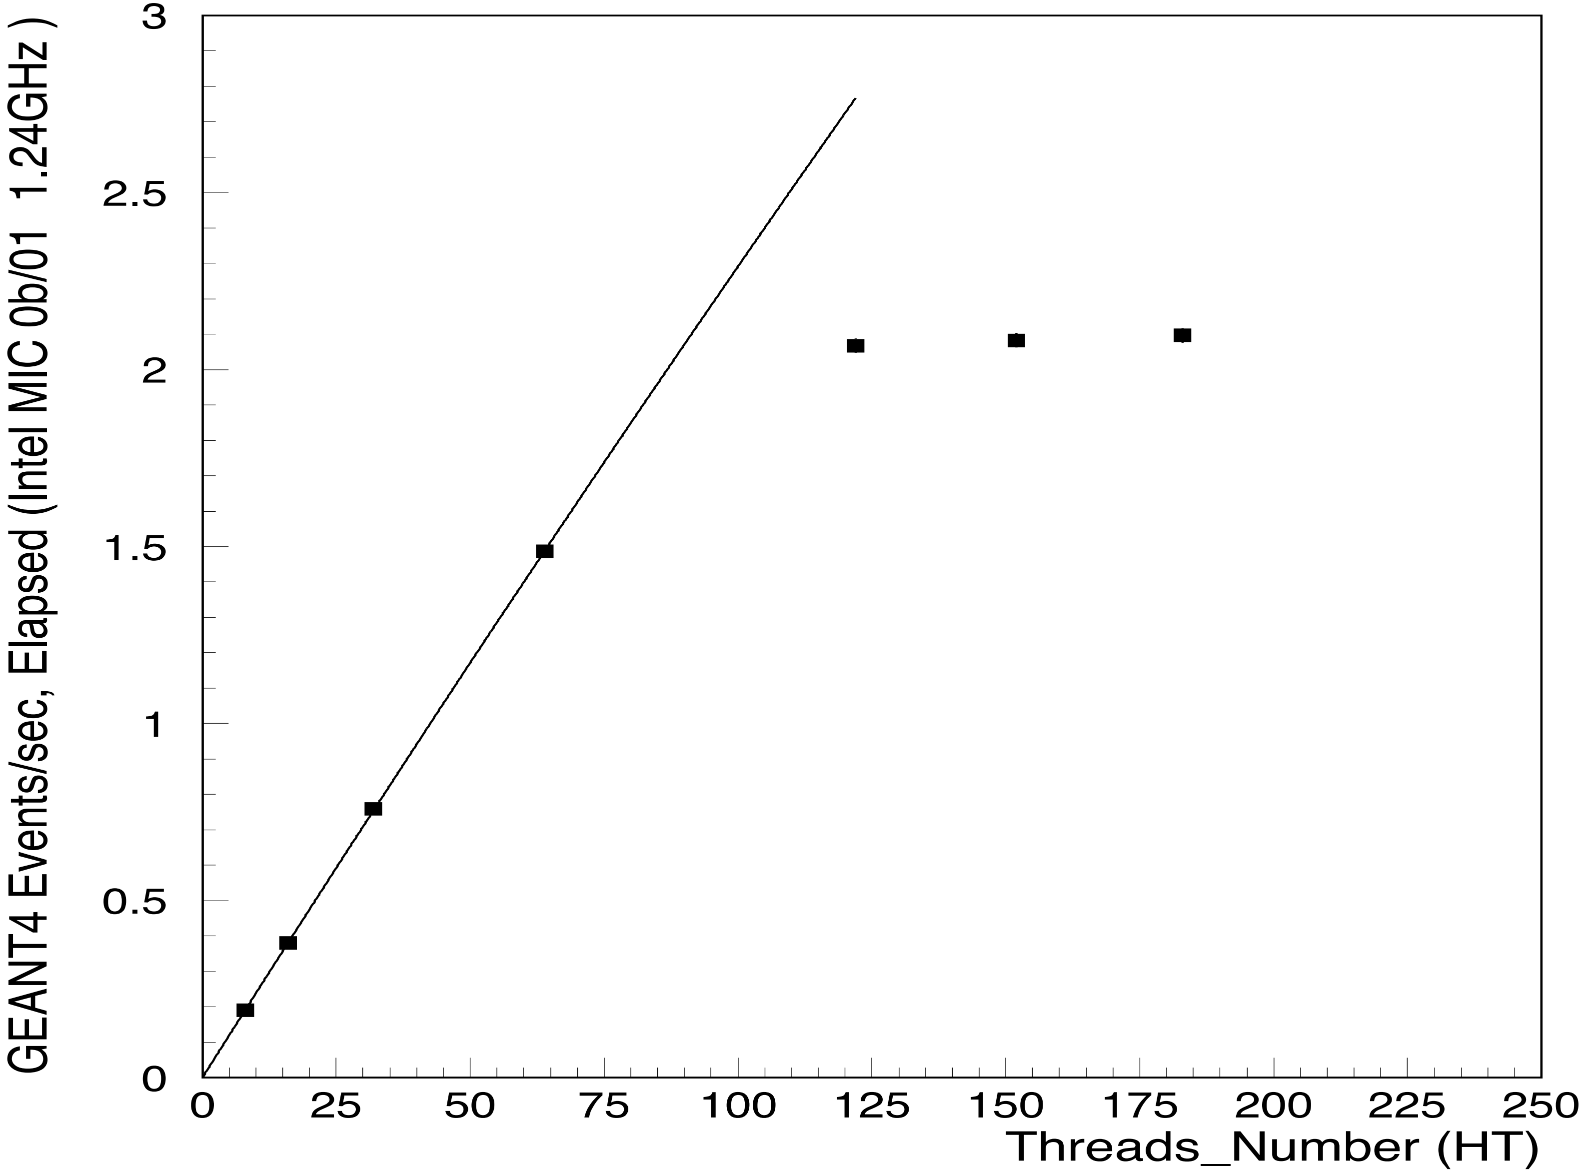

Memory Optimization for Geant4.10.1

- Garbage collection in G4AllocatorPool
  - We added garbage collection mechanism in G4Allocator
  - Decreased the memory consumption by factor 2 and more for long jobs (W.R.T. original G4AllocatorPool)

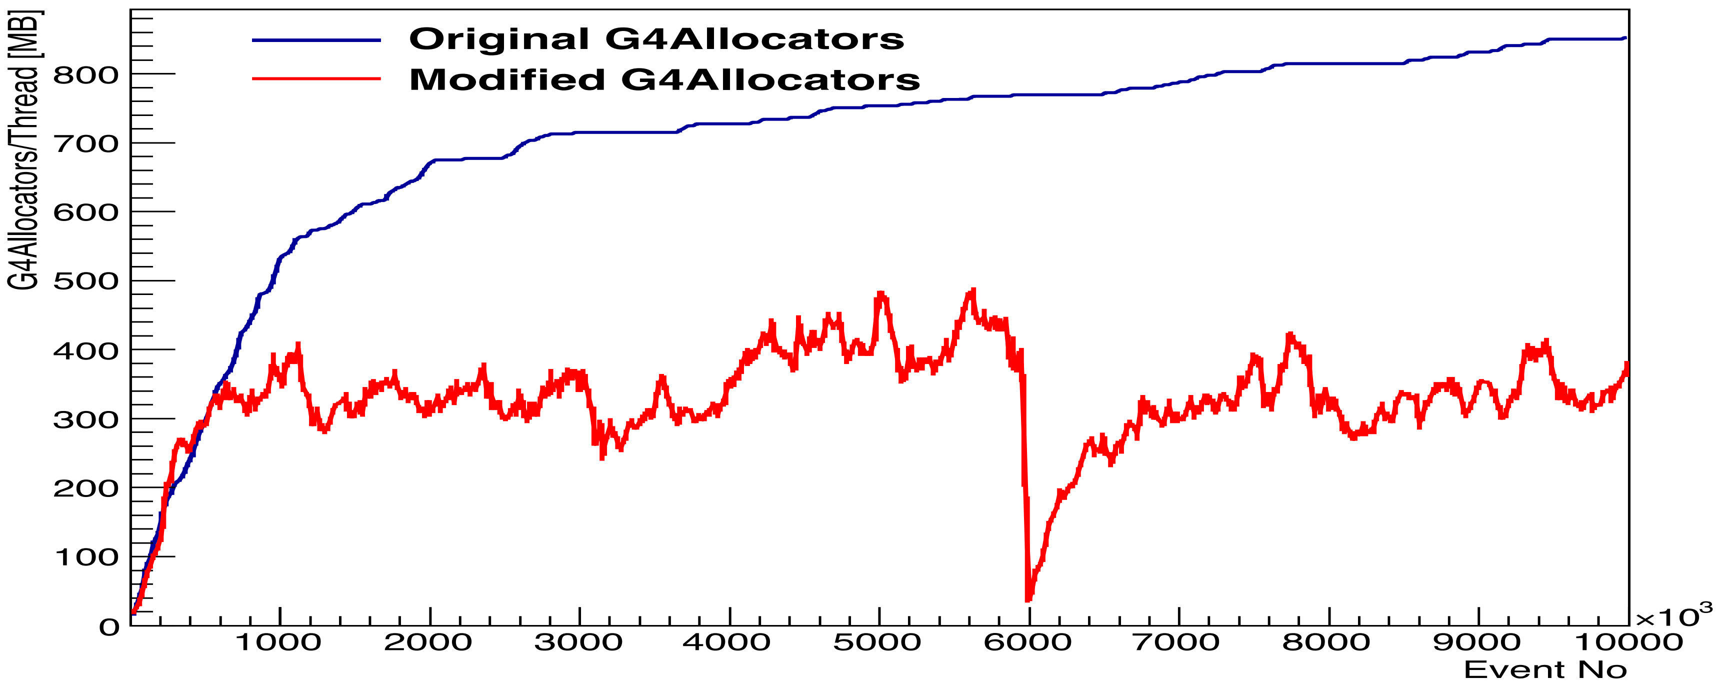

- Decreased the memory consumption by factor 5 W.R.T. sequential simulation job

Summary

- With production supporting software, data reconstruction rate reached 100 times of data collection rate.
- Parallelization of the AMS reconstruction software allows to increase the productivity of the reconstruction program to 30-50% using the HT enabled processors as well as to shorten per run reconstruction time and calibration job execution time by factor of ~50 (72 threads on HT).
- Parallelization of the AMS simulation software and optimization on G4AllocatorPool allow to decrease the total memory consumption per core by factor 5, which makes it much easier to run long simulation jobs for light nuclei.

AMS Days at CERN
The Future of Cosmic Ray Physics and Latest Results

CERN, Main Auditorium,
April 15-17, 2015

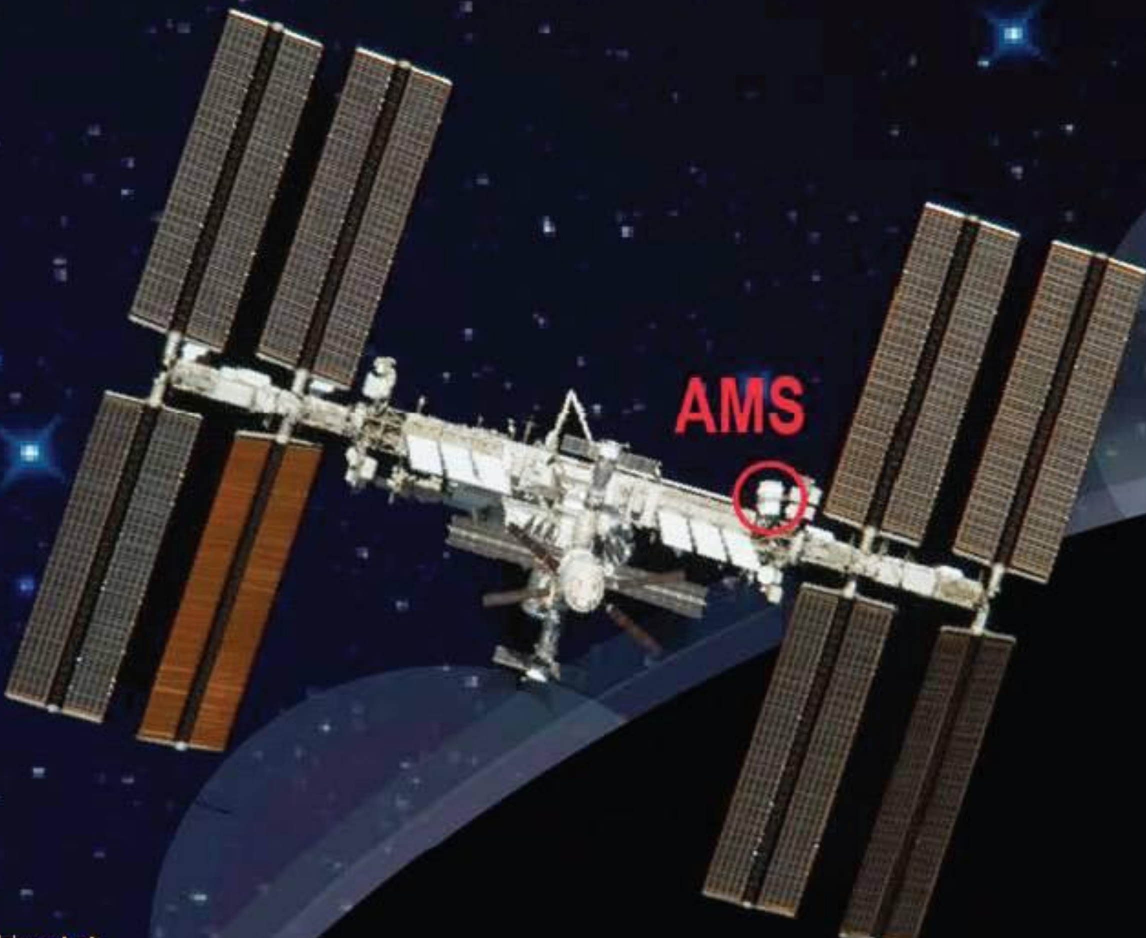

Speakers:
Roberto BATTISTON, ASI, Trento
Kfir BLUM, IAS, Princeton
John ELLIS, King's College, London, CERN
Jonathan FENG, UC Irvine
William GERSTENMAIER, NASA
John M. GRUNSFELD, NASA
Francis HALZEN, Wisconsin
Werner HOFMANN, MPI Heidelberg
Gordon KANE, Michigan
Peter F. MICHELSON, Stanford
Igor V. MOSKALENKO, Stanford
Angela OLINTO, Chicago
Piergiorgio PICOZZA, INFN, Tor Vergata
Vladimir S. PTUSKIN, IZMIRAN, Moscow
Lisa RANDALL, Harvard
Michael SALAMON, DOE
Subir SARKAR, Oxford, Niels Bohr Inst.
Eun-Suk SEO, Maryland
Tracy SLATYER, MIT
Edward C. STONE, Caltech
Michael TURNER, Chicago
Alan A. WATSON, Leeds
Yue-Liang WU, UCAS/ITP, CAS
Fabio ZWIRNER, Padua, CERN
and
presentations on the AMS latest results

Event Webcast at
http://cern.ch

Contact: Ms. Laurence Barrin <laurence.barrin@cern.ch>