STAR Software

G. Van Buren

September 4, 2019

(with contributions from J. Lauret)

STAR

- Large acceptance with PID tracking & calorimetry, moderate rate capabilities
- Spanning interests from hot QCD (quark gluon plasma) to cold QCD (spin structure of the nucleon)
- Flexible enough to enable several unique studies in the field, like CME (isobars) and critical point search (beam energy scans)
- Collaboration size: currently ~600 people from ~60 institutions
- Expecting to run the entire lifetime of RHIC, from the very first collisions in 2000, to the very last in the mid-2020s (never a year off!)
  - ~100 different collision setups so far! (species, energy, fixed target)
  - ~2000 (very roughly) collaborators so far!

The Detector

- Numerous upgrades over the years, with more still to come!

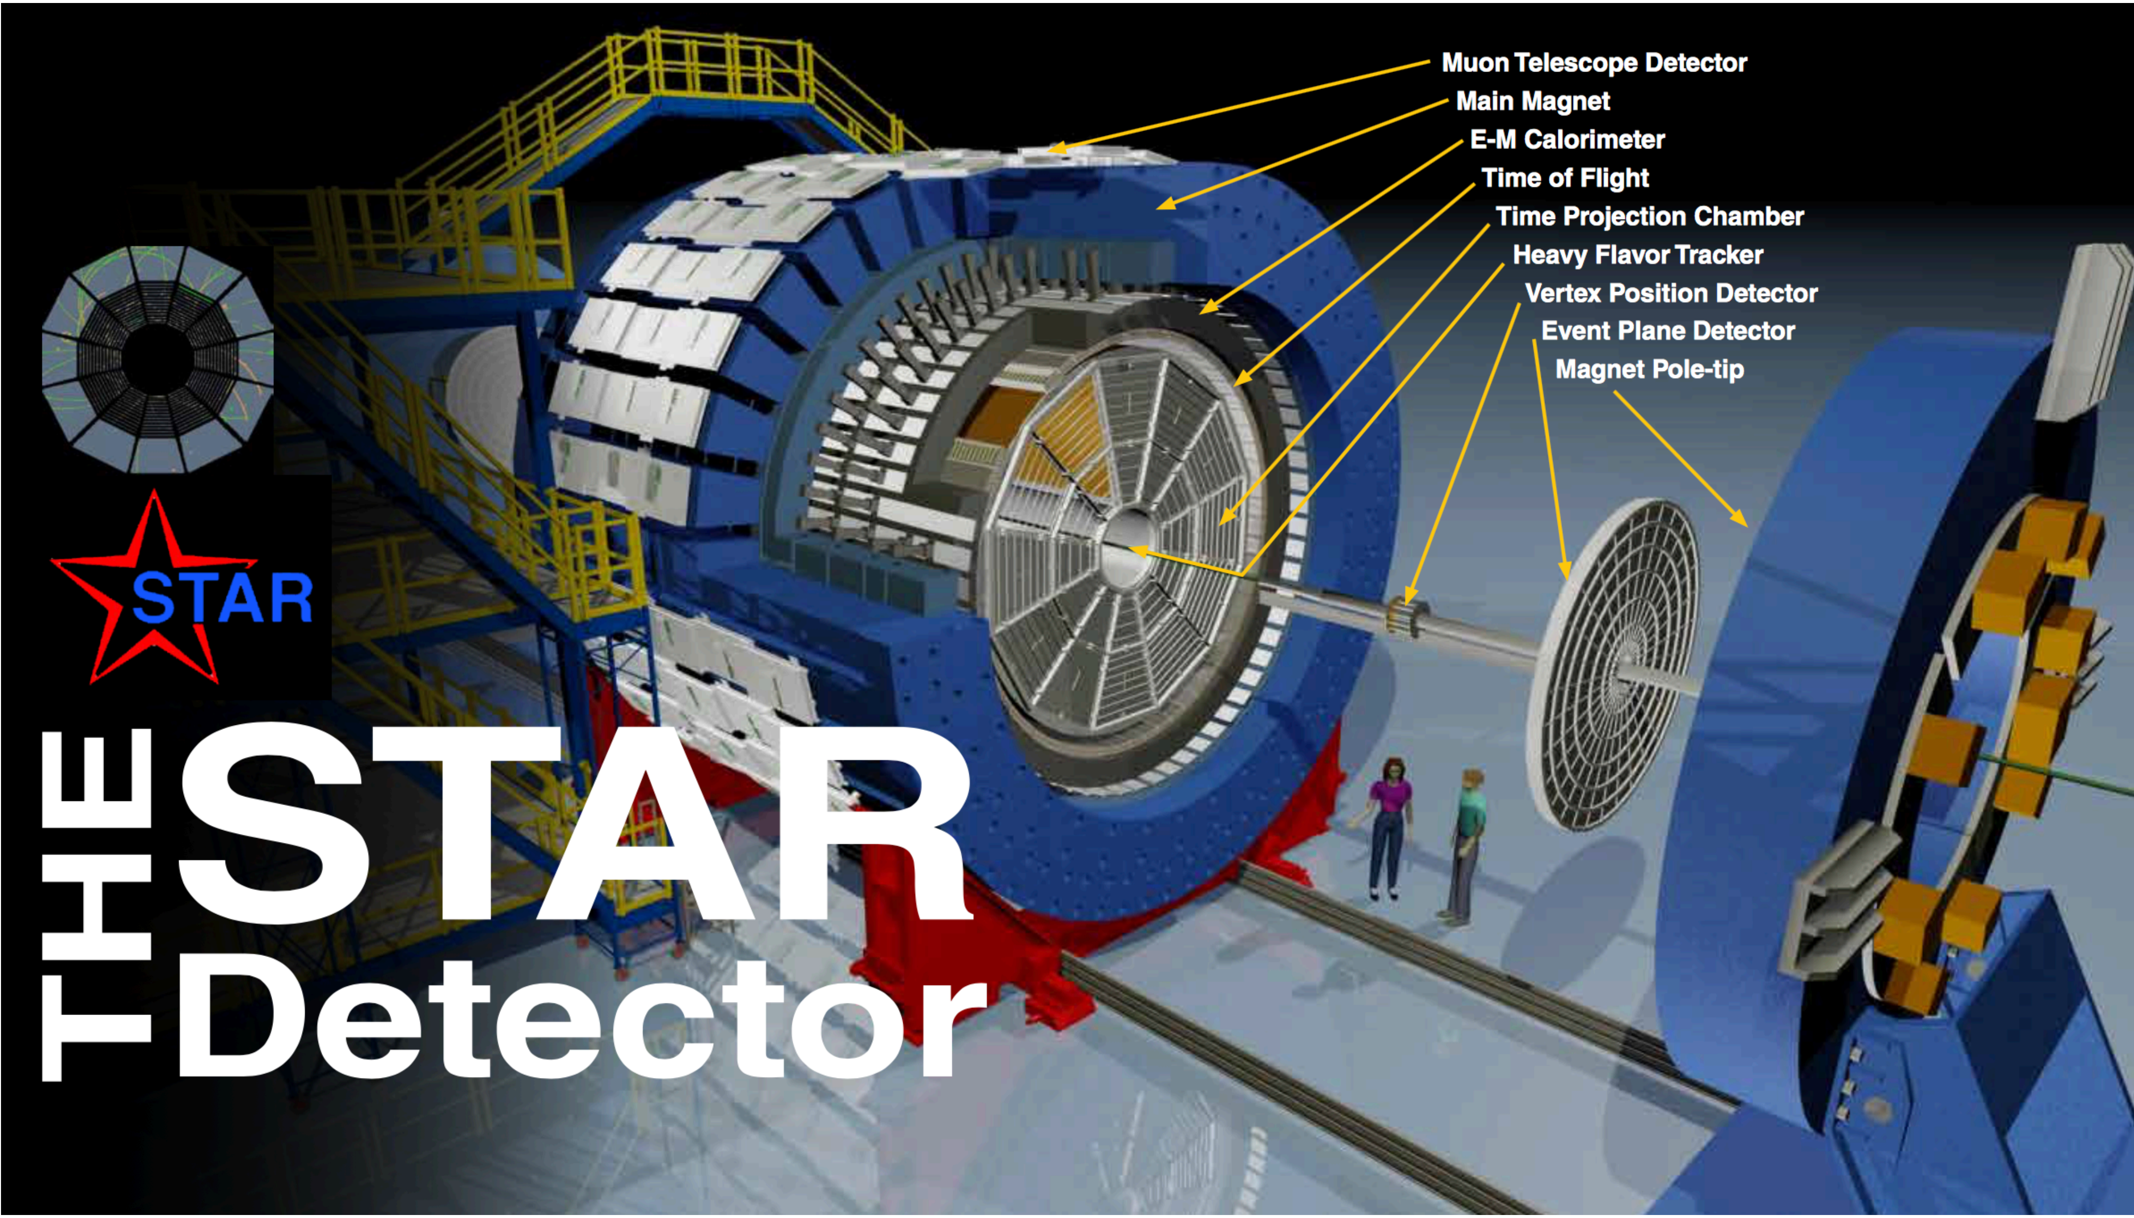

(Meta-)Data flow

- An outline for the rest of the slides...
  (NB: this will NOT be fully comprehensive, and more details may come from the presentations of others in NPSS who work on STAR)

- Acquired data files are stored in HPSS at BNL RACF / SDCC
  - ~40 PB of integrated raw data
- Meta-data (conditions) recorded in online databases
  - Migrated to offline databases in near-real time
- Portions of the data used for Quality Assurance & Calibrations
- Physics productions deliver multiple reduced data types
  - ~30 PB of integrated DSTs archived in HPSS
- Analyzers have access to prioritized DSTs on live storage
- STAR Data Management System automatically catalogues produced files and maintains consistency of live datasets

See intro of...
NPPS/STAR: Dmitry A.

Quality Assurance

- Online QA [by shift + detector subsystems]:
event pool (a few % of data) generally for detector functionality and feedback to collider
- Offline QA [by shift + detector subsystems + analyzers]:
automatic full reconstruction (~50% of files: "FastOffline") and histogram generation for all-purpose QA
    - Automated QA analysis (vs. a reference) serves as a non-definitive guide to the shift
        - References need to be human-vetted for each collision/trigger setup
    - Observations are recorded/tracked & reported
- Productions [analyzers]:
    - Real data QA: no technical issues, but hasn't been well-organized nor enforced (analyzers prone to wait instead of following as it goes) and we're moving to address this
    - Simulations/embedding QA: working well with commitments from the analysis side to contribute to QA (a good model)

Code

- Nearly 30 years worth! And some of the initial codes were borrowed (e.g. NA49 & ALICE TPC tracking)
- Currently in the official stack, we have ~5.3M lines of code:
    - C++ source code: ~3.2M lines (61.2%), ~600 MB
    - Fortran source code: ~1.8M lines (35.0%), ~100 MB
    - The rest: ansi C, python, perl, php, sh, csh, tcl, yacc, java, f90
    - (...excludes retired software packages)
- Code Verification
    - Peer review (2 people) for any new software
    - Comprehensive nightly test suite (all years' data + simulations)
    - Code sanity dashboards (cppcheck, coverity)
    - Test suite for software library releases
- Libraries released 4-12 times per year historically
    - Numerous developers from across the collaboration, but releases are managed by a small, efficient effort within the core S&C team

New dashboard from...

NPPS/STAR: Amol J.

Calibrations

- Some tasks performed automatically in near-real time using either small online cluster, or the FastOffline framework
- Initial calibrations, useful for online HLT reconstruction or for FastOffline, performed within first ~week of new collision conditions
- Full calibrations generally delivered in a few months after data-taking completes for the year
  - STAR Time Projection Chamber (TPC) calibrations encounter something different that requires further study almost every year
  - A rather limited set of calibrations not required for tracking (see next slide) can be applied post-production

Reconstruction

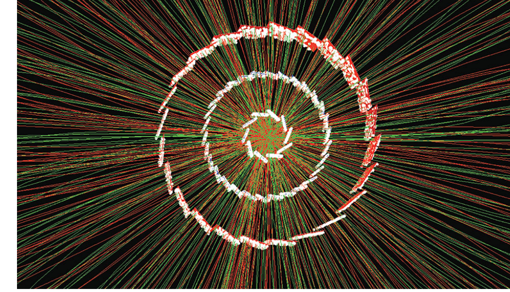

- Tracking: Cellular Automata Finding + Kalman Finding & Fitting
  - CA collaboratively contributed from GSI/FIAS (ALICE & CBM)
  - Uses simplified version of the geometry (including tools to approximate such geometry on the fly)
  - History of dealing with considerable pile-up in slow detectors
  - First heavy ion collider experiment to successfully use precision silicon tracking (STAR Heavy Flavor Tracker)
  - Handling of non-uniform magnetic fields and forward tracking for coming updates is in development
- Vertex-finding & fitting: Kalman + annealing, constrained fits
- Subsystems not used in the track finding & fitting can generally re-reconstruct as post-production afterburners if necessary (avoided)
  - Tracking dominates reco time; inefficient to do post-production
- Tracks are matched (or not) to time-of-flight, calorimetry, muon detectors (outside the primary solenoid)

Broad contributions from...

NPPS/STAR:
Victor P.
Dmitri S.

Productions

At RACF...

NPPS/STAR: Amol J.

- Highly efficient: *99+% of jobs succeed!*
- Sites:
  - RACF (BNL)
    - Includes scavenging other experiments' unused nodes
  - OSG tools:
    - KISTI (Korea), JINR/Dubna (Russia), PDSF (LBNL), STAR Online Cluster
  - Amazon Cloud
  - HPC: Cori (LBNL)
    - All data processing workflows except user analysis (for now)!
    - STAR is a part of the OSG, which has an NP activity focused on HPC
- Over a decade of utilizing VMs

Embedding

- Inject simulated data into real data at the detector level (...later called "event overlay" by some...):
  simulation + *detector response ("digitization")* + *reconstruction*
  - Full simulated events into real backgrounds ("zerobias")
  - Simulated tracks into real events
- Simulations can use misaligned geometries
- Currently, 1% of recorded data retains un-reduced, fully raw data for use in embedding (can re-use real events if larger simulation statistics are needed)
- Running with HPC providing the bulk of its resources
- Utilizing CVMFS for software distribution
- CPU needs have historically been less than 10-15% of the real-data reconstruction demands

Analyses

- Live storage for DSTs:
  - Central (NFS-mounted GPFS)
  - Distributed (RACF computing farm)
    - Xrootd access
  - Not sufficient to hold the entirety of STAR reduced data (some analyses cannot use the most compact "picoDst" reduction), but the Collaboration decides what's "hot"
- Job handling: STAR Scheduler
  - Unified interface to multiple batch systems at multiple sites
  - Match job resource demands to batch queue specifications
  - Access to the catalog to find produced data files
  - Can bring job-to-data, or data-to-job, to reduce potentially crippling network & file system traffic

Summary

- The STAR Experiment has a long history...and an active future
  - A quarter century of collecting data, after a decade of preparation (ask Torre about the [very] old days ;-)
  - Software has had a lot of evolving to do:
    - Computing platforms & operating systems have come and gone
    - Computing facilities have come and gone
    - Detector subsystems have come and gone
    - Code authors/maintainers have come and gone
  - Propagation of technical knowledge across time and across the collaboration (e.g. analyzers) are factors
- R&D has prioritized re-usability and contributions beyond STAR (e.g. ROOT, CA tracking, embedding, DataCarousel, Xrootd, TPC calibrations, MetaData & Database approach, PhoneBook, unified event display, and more)
- Throughout it all, succeeding to deliver on scientific goals within the finite available resources (computing *and* human)