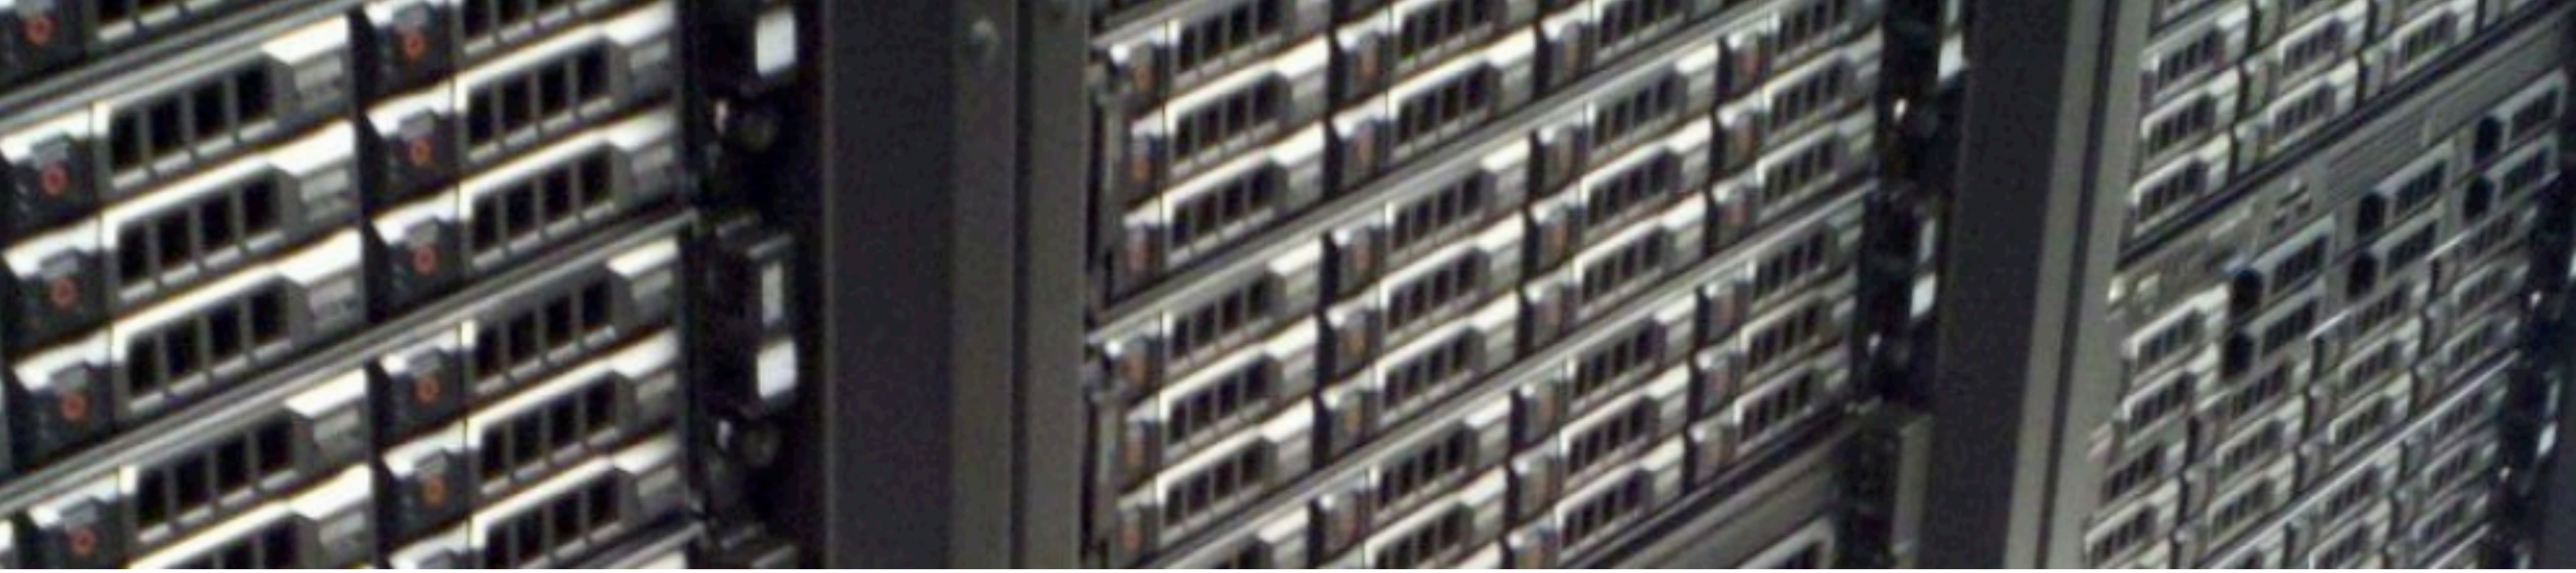

Grid Computing @ Duke
- Fall 2013 Campus Infrastructures Workshop -

sponsored by the Information Initiative @ Duke (iiD)

Organizers:
Steffen A. Bass (Duke University, Department of Physics)
Jeff Chase (Duke University, Department of Computer Science)
Rob Gardner (University of Chicago)
John McGee (The RENCI Institute at the University of North Carolina)
John Pormann (Duke University, Office of Information Technology)

Recreating the Big Bang in the Laboratory: Computational Challenges in High Energy Nuclear Physics

Steffen A. Bass

Department of Physics
Duke University

- Scientific Motivation
- Computational Challenge
- Data and Storage Management

Open Science Grid

Renci Renaissance Computing Institute at Duke University

Scientific Challenge:

- Knowledge Extraction from Heavy-Ion Collisions

Phases of Matter

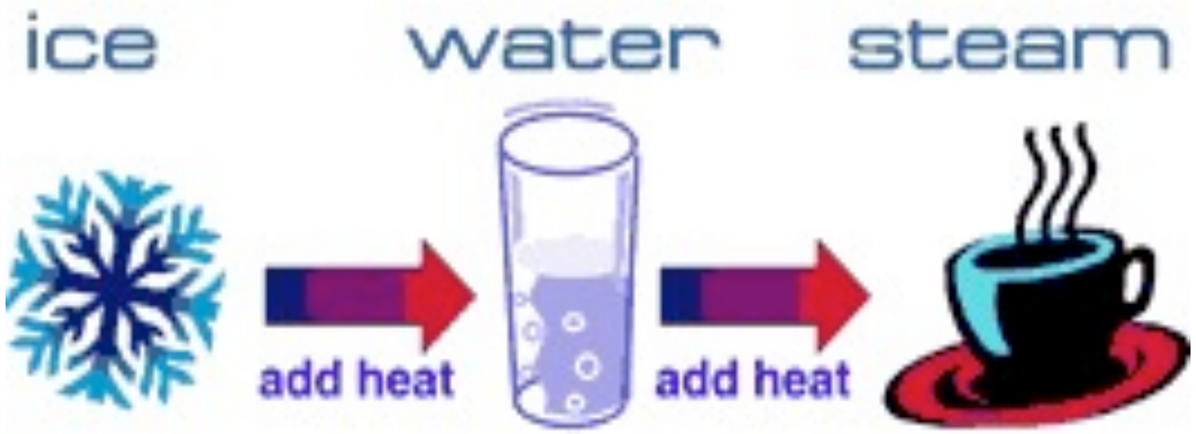

solid liquid gas

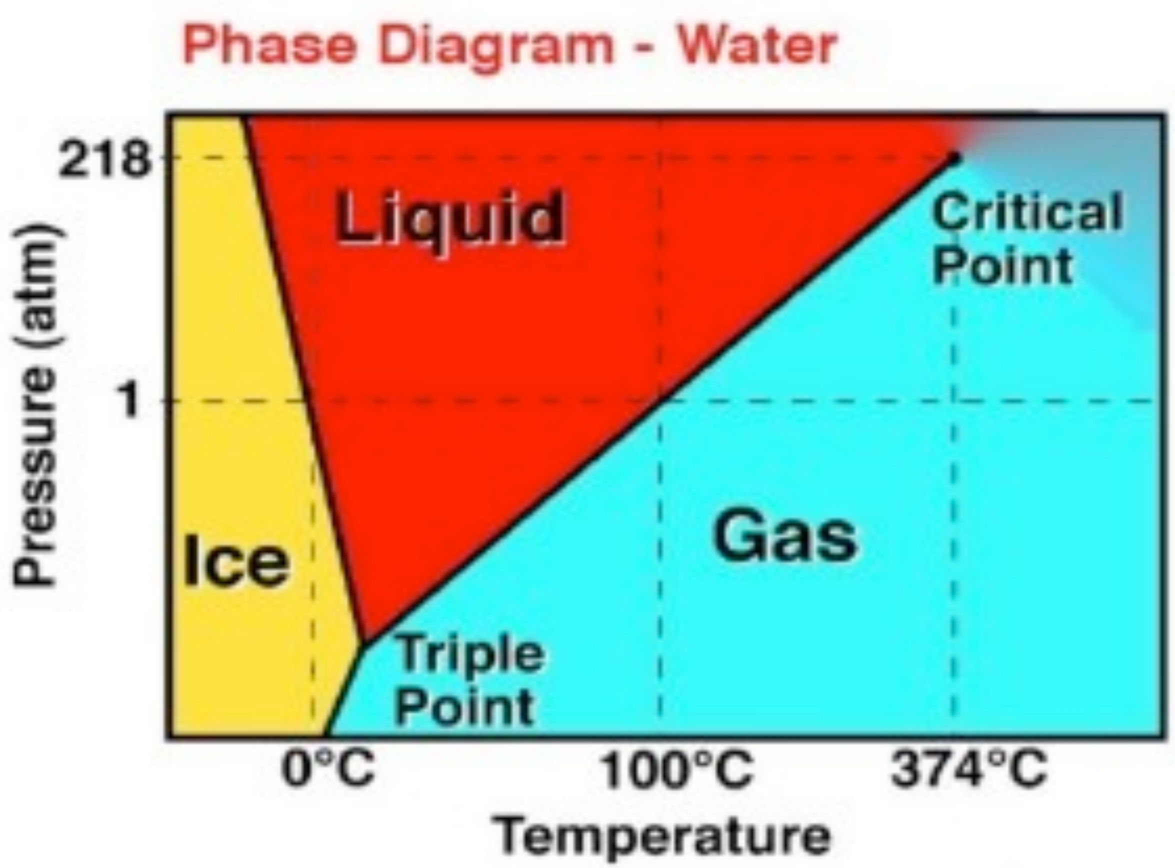

QCD matter analogues of familiar phases:

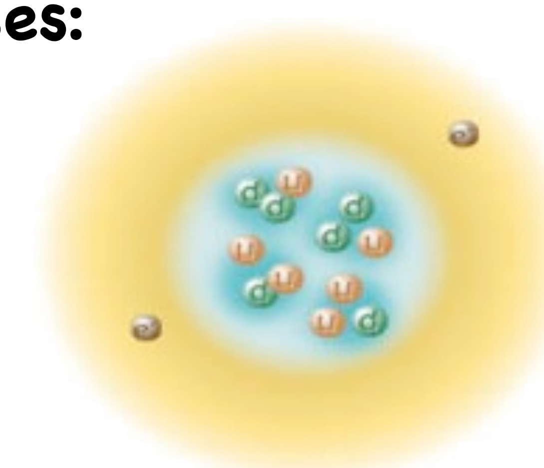

- Nuclei behave like a liquid
  - Nucleons are like molecules
- **Quark Gluon Plasma:**
  - "ionize" nucleons with heat
  - "compress" them with pressure

➢ new state of matter!

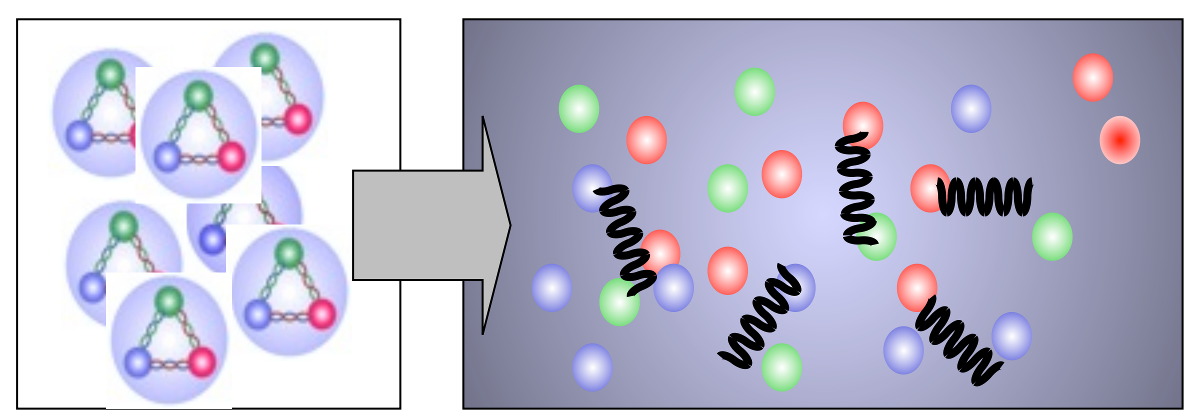

QGP and the Early Universe

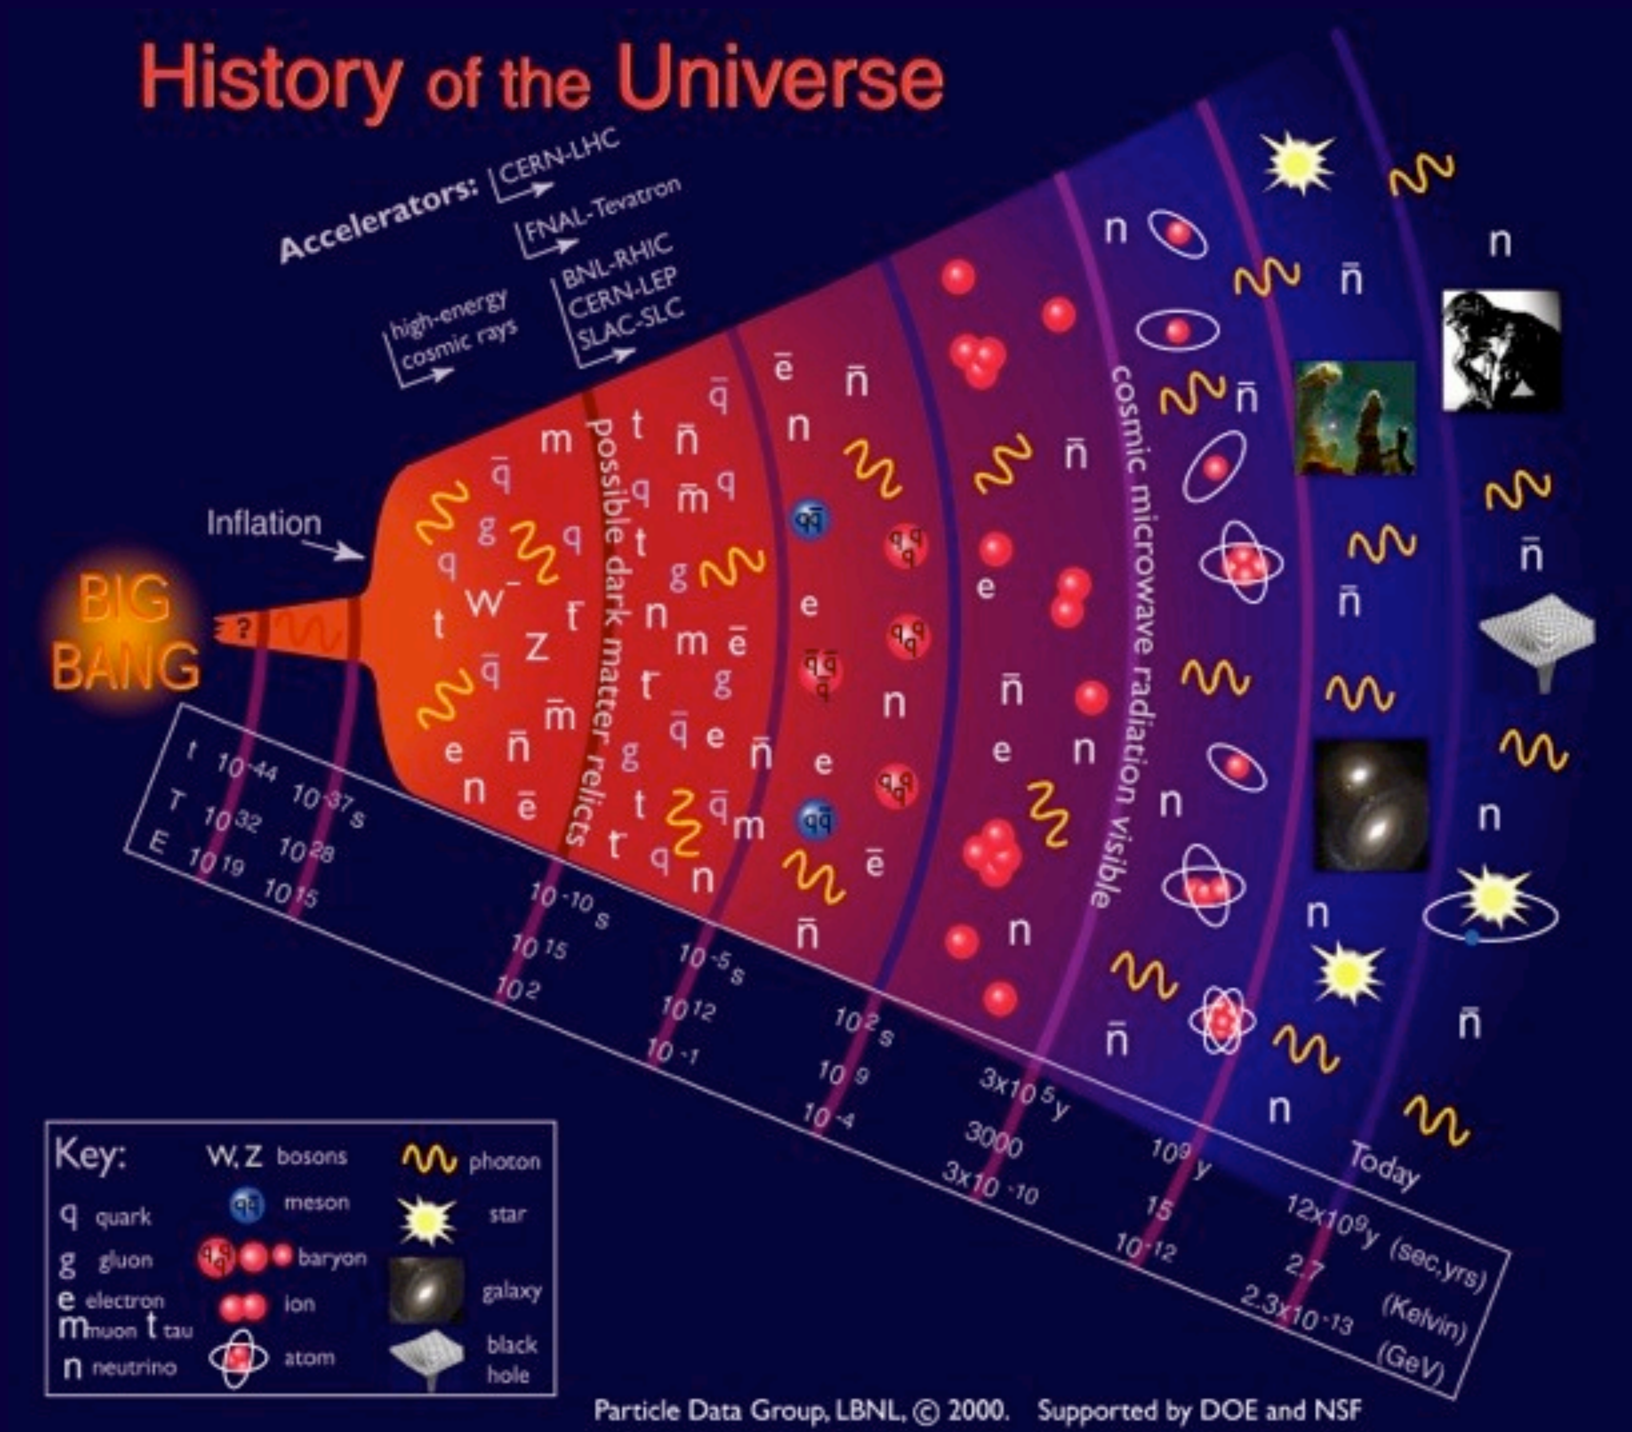

- a few microseconds after the Big Bang the entire Universe was in a QGP state
- compressing & heating nuclear matter allows to investigate the history of the Universe
- the only means of recreating temperatures and densities of the early Universe is by colliding beams of ultra-relativistic heavy-ions

Heating & Compressing QCD Matter

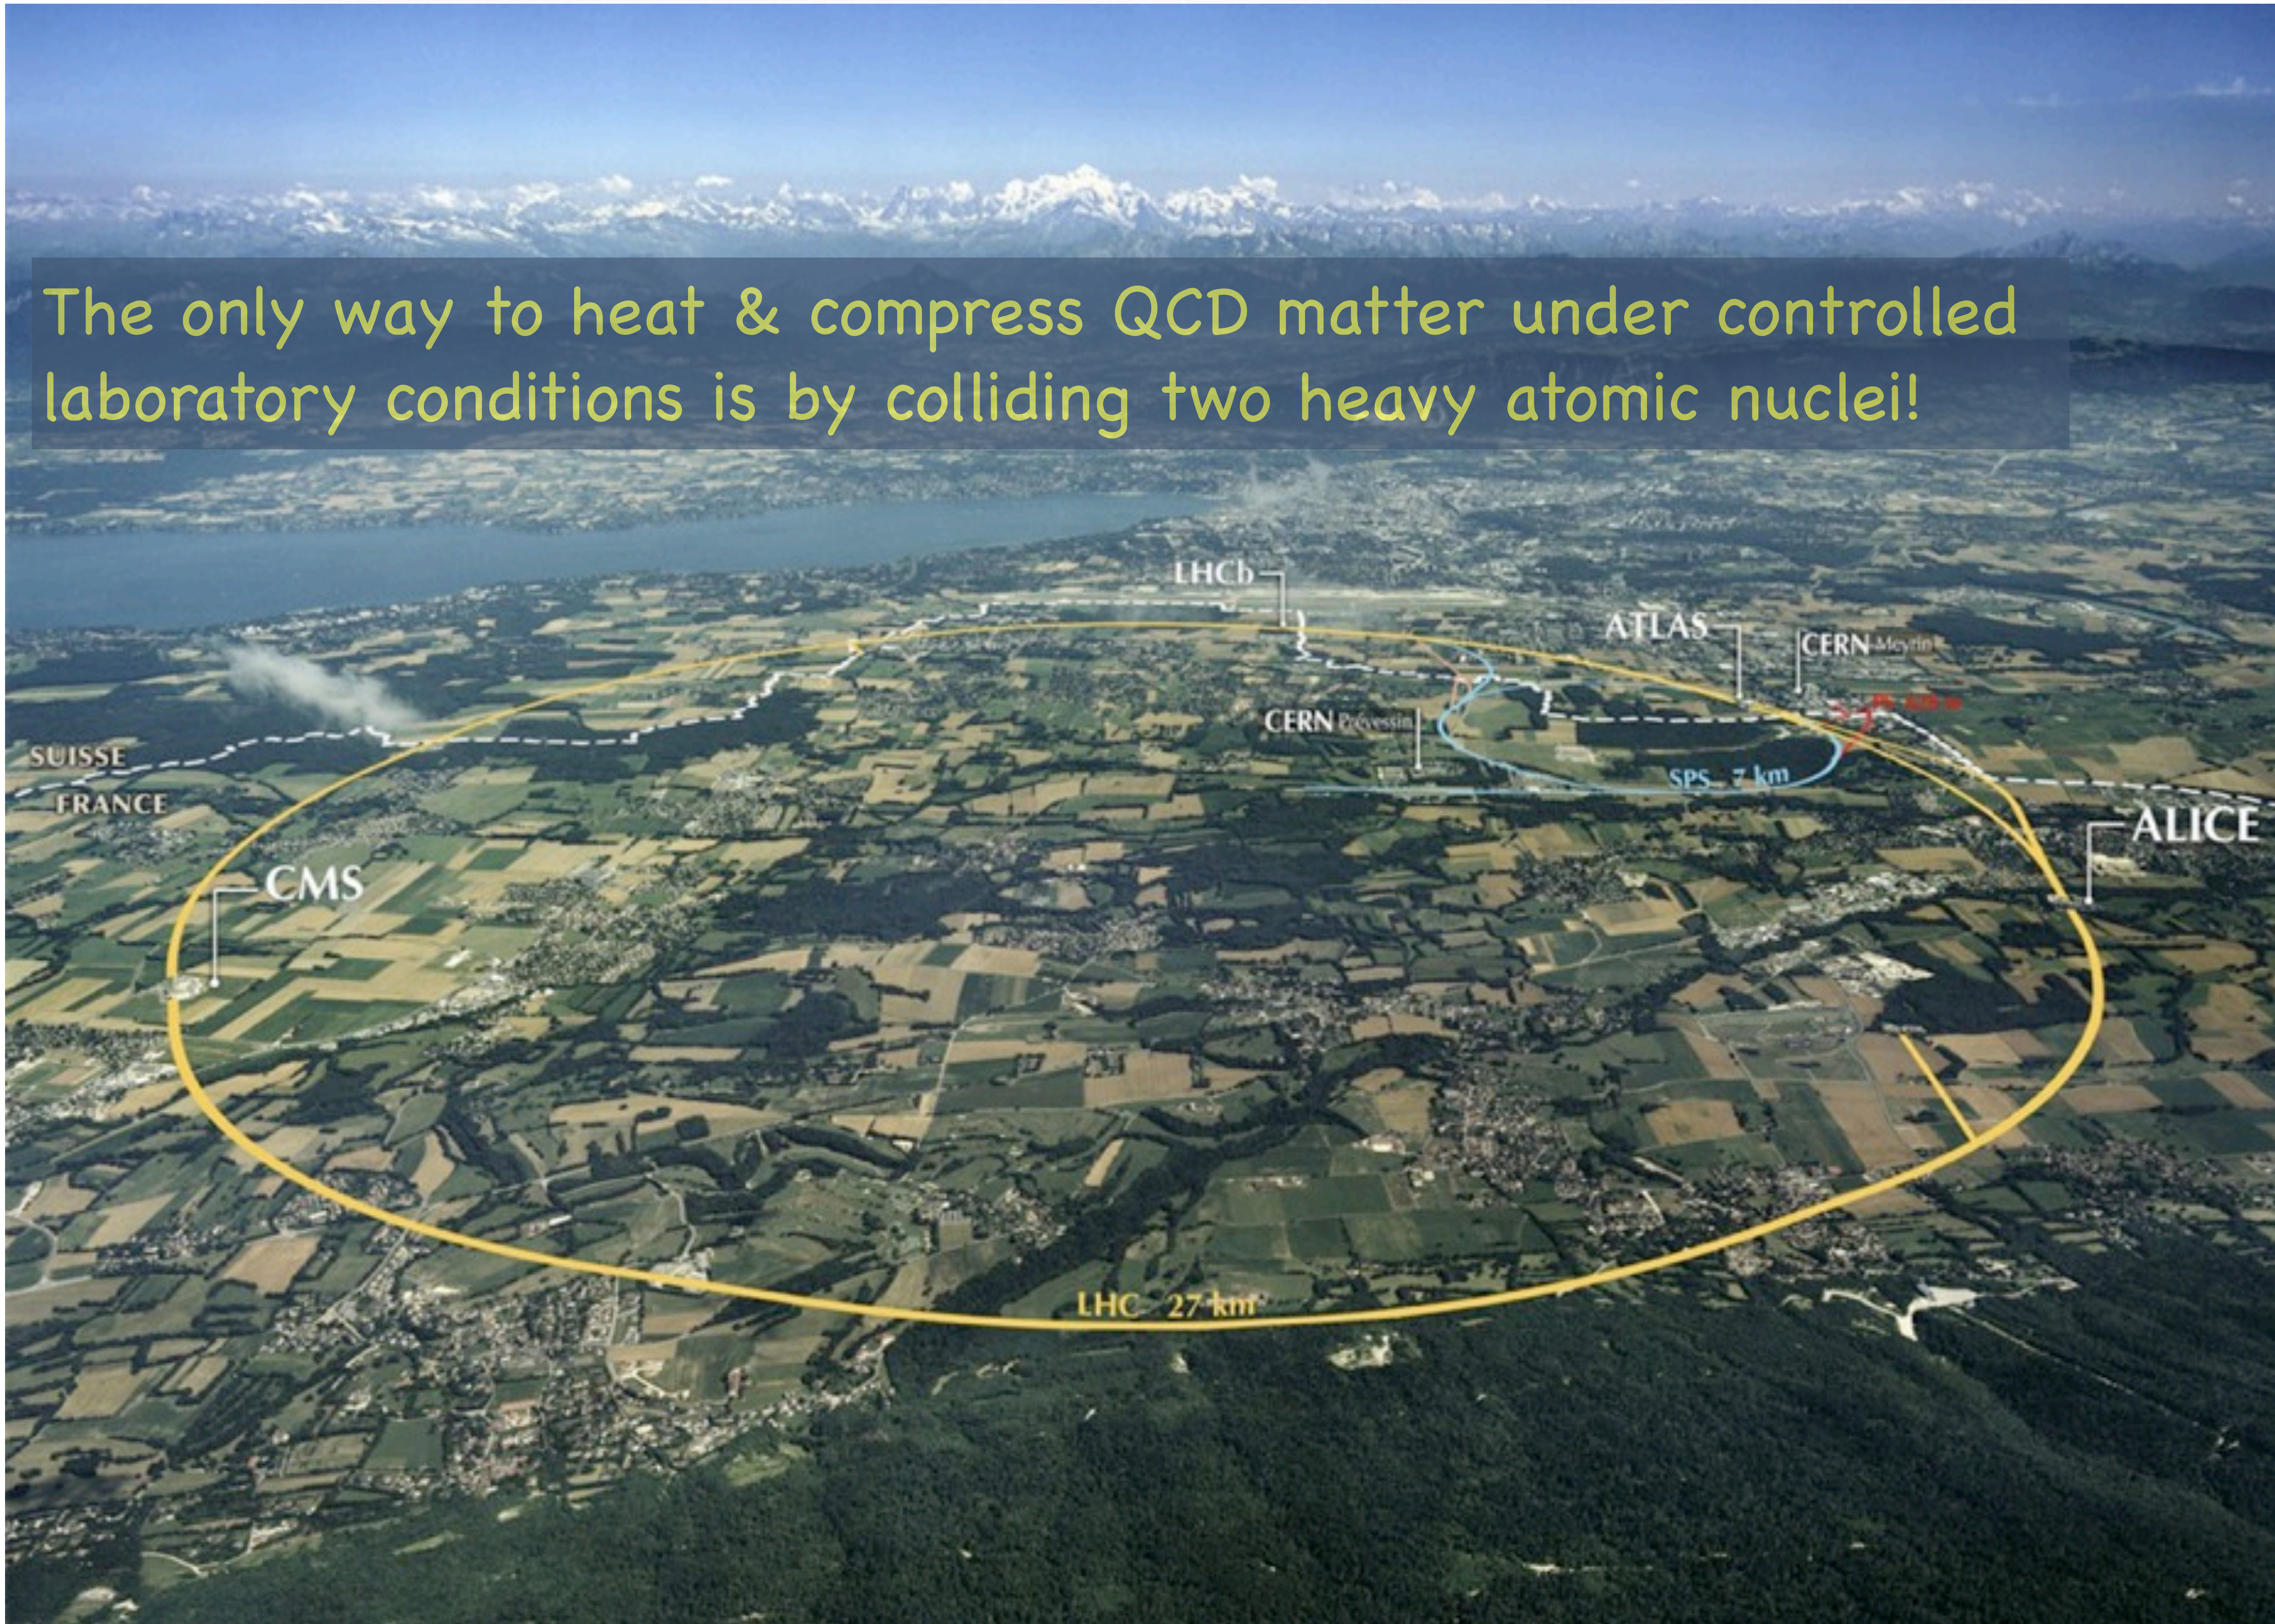

Heating & Compressing QCD Matter

ALICE experiment @ CERN:

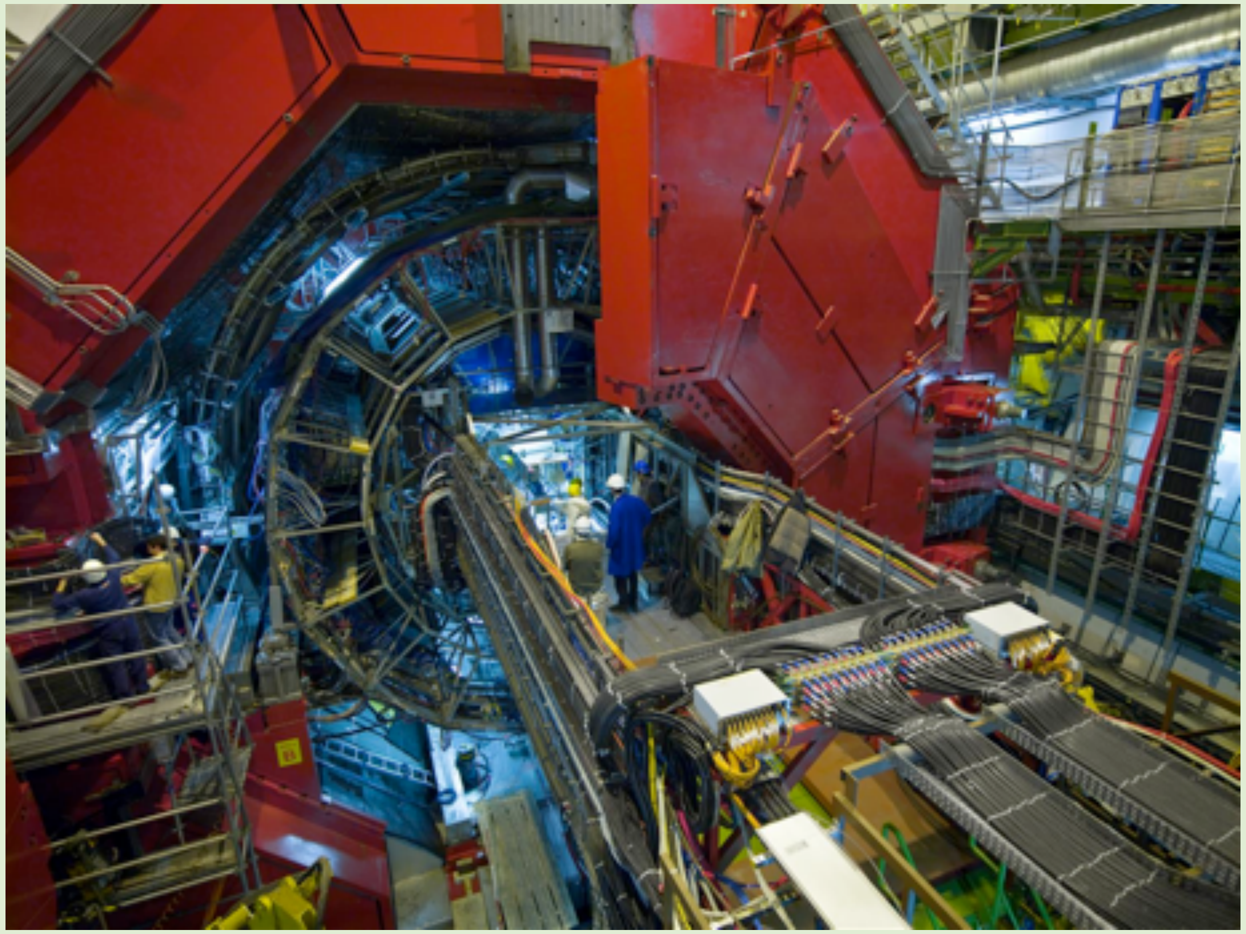

- 1000+ scientists from 105+ institutions
- dimensions: 26m long, 16m high, 16m wide
- weight: 10,000 tons

two more experiments w/ Heavy-Ions:

- CMS, ATLAS

typical Pb+Pb collision @ LHC:

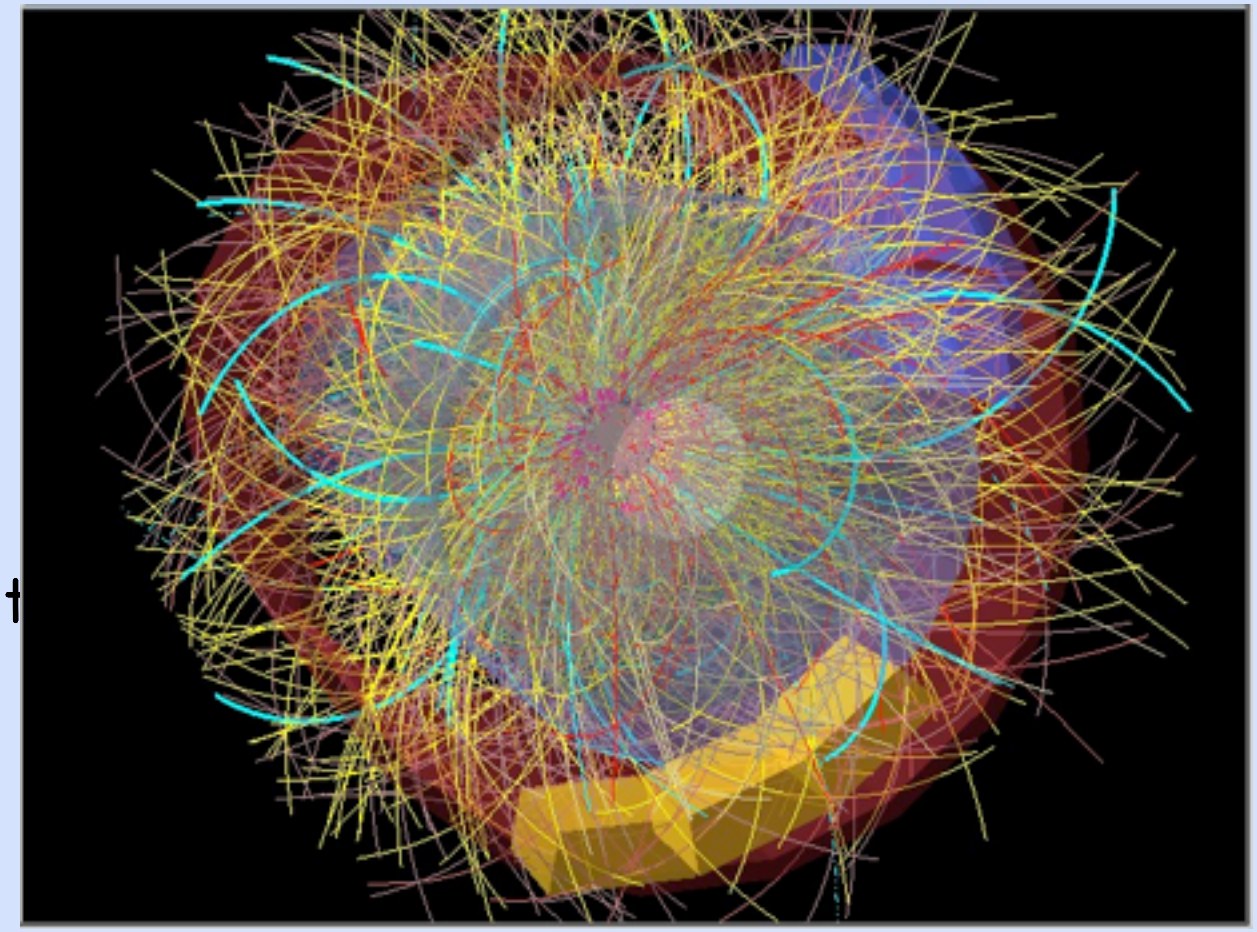

- 1000s of tracks
- task: reconstruction of final state to characterize matter created in collision

LHC Data Challenge

LHC Data:

- 15 Petabyte per year raw data
- WLCG: World-wide LHC Computing Grid: 170 computing centers across the world analyze 25 Petabyte of data & simulations annually

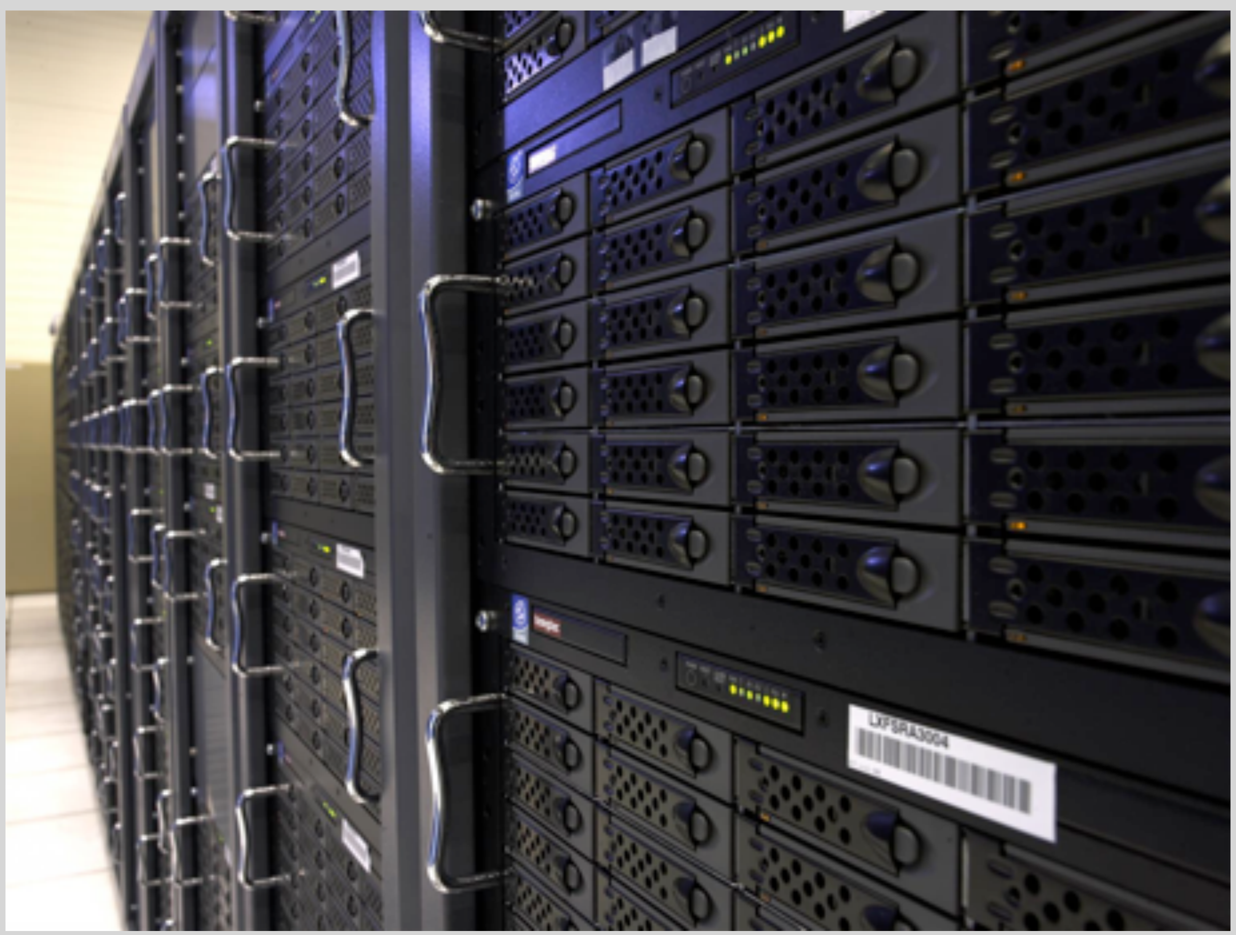

Alice Data Flow:

- factor of 100 higher data acquisition rate than RHIC experiments

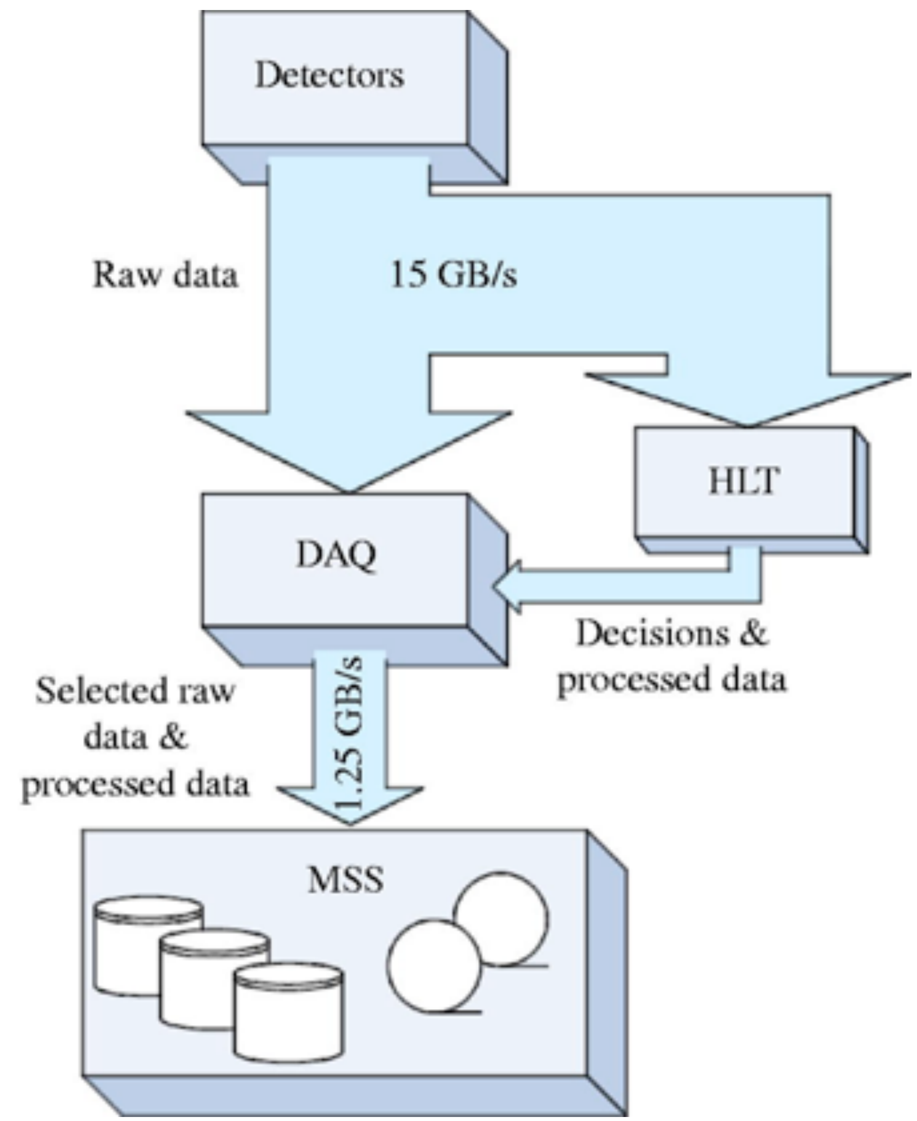

Knowledge Extraction: Need for Modeling

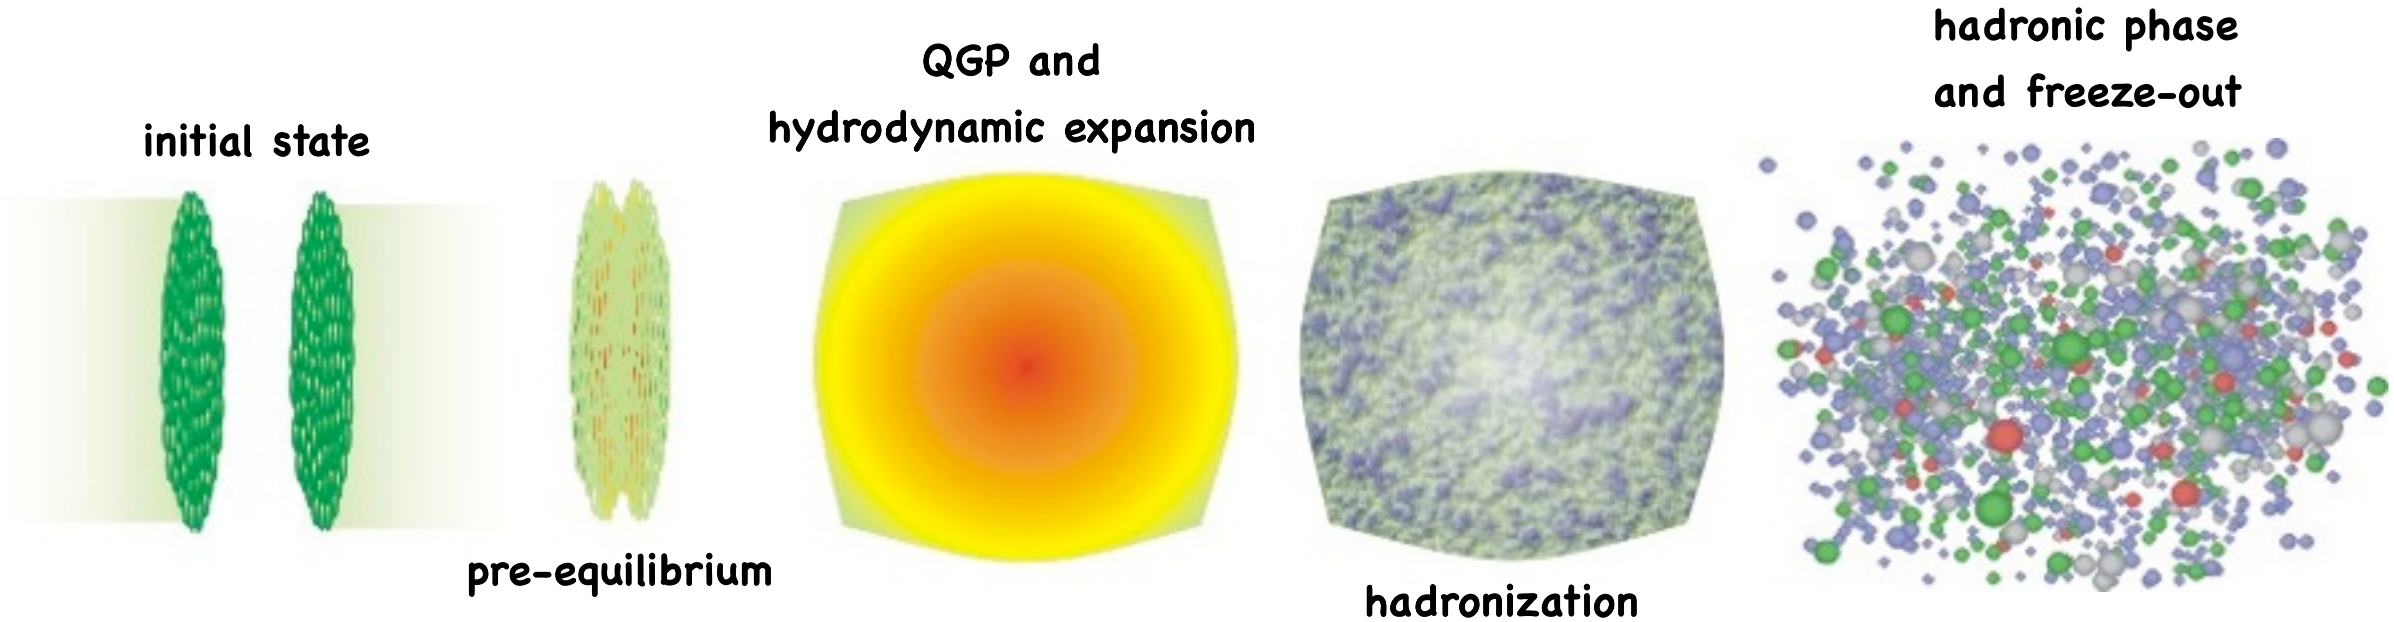

Challenges:

- time-scale of the collision process: 10^{-24} seconds! [too short to resolve]
- characteristic length scale: 10^{-15} meters! [too small to resolve]
- confinement: quarks & gluons form bound states @ hadronization, experiments don't observe them directly

Experiments:

- observe only the final state
- rely on QGP signatures predicted by Theory

Transport-Models:

- simulates the outcome of the experiment
- gives access to not directly observable quantities
- parameters encode the science to be extracted

Transport Models for RHIC

microscopic transport models based on the **Boltzmann Equation**:

- transport of a system of microscopic particles
- all interactions are based on binary scattering

$$\left[\frac{\partial}{\partial t} + \frac{\vec{p}}{E} \times \frac{\partial}{\partial \vec{r}}\right] f_1(\vec{p}, \vec{r}, t) = \sum_{processes} C(\vec{p}, \vec{r}, t)$$

diffusive transport models based on the **Langevin Equation**:

- transport of a system of microscopic particles in a thermal medium
- interactions contain a drag term related to the properties of the medium and a noise term representing random collisions

$$\vec{p}(t + \Delta t) = \vec{p}(t) - \frac{\kappa}{2T} \vec{v} \cdot \Delta t + \vec{\xi}(t) \Delta t$$

(viscous) relativistic fluid dynamics:

- transport of macroscopic degrees of freedom
- based on conservation laws:

$$\partial_\mu T^{\mu\nu} = 0$$

$$\begin{aligned} T_{ik} &= \varepsilon u_i u_k + P\left(\delta_{ik} + u_i u_k\right) \\ &- \eta\left(\nabla_i u_k + \nabla_k u_i - \frac{2}{3}\delta_{ik}\nabla \cdot u\right) \\ &+ \varsigma\, \delta_{ik} \nabla \cdot u \end{aligned}$$

(plus an additional 9 eqns. for dissipative flows)

hybrid transport models:

- combine microscopic & macroscopic degrees of freedom
- current state of the art for RHIC modeling

Each transport model relies on roughly a dozen physics parameters to describe the time-evolution of the collision and its final state. These physics parameters act as a representation of the information we wish to extract from RHIC.

The Challenge of a Model to Data Comparison

Each computational model relies on multiple physics parameters to describe the outcome of the experiment it simulates. These parameters act as a representation of the information we wish to extract from the model to data comparison.

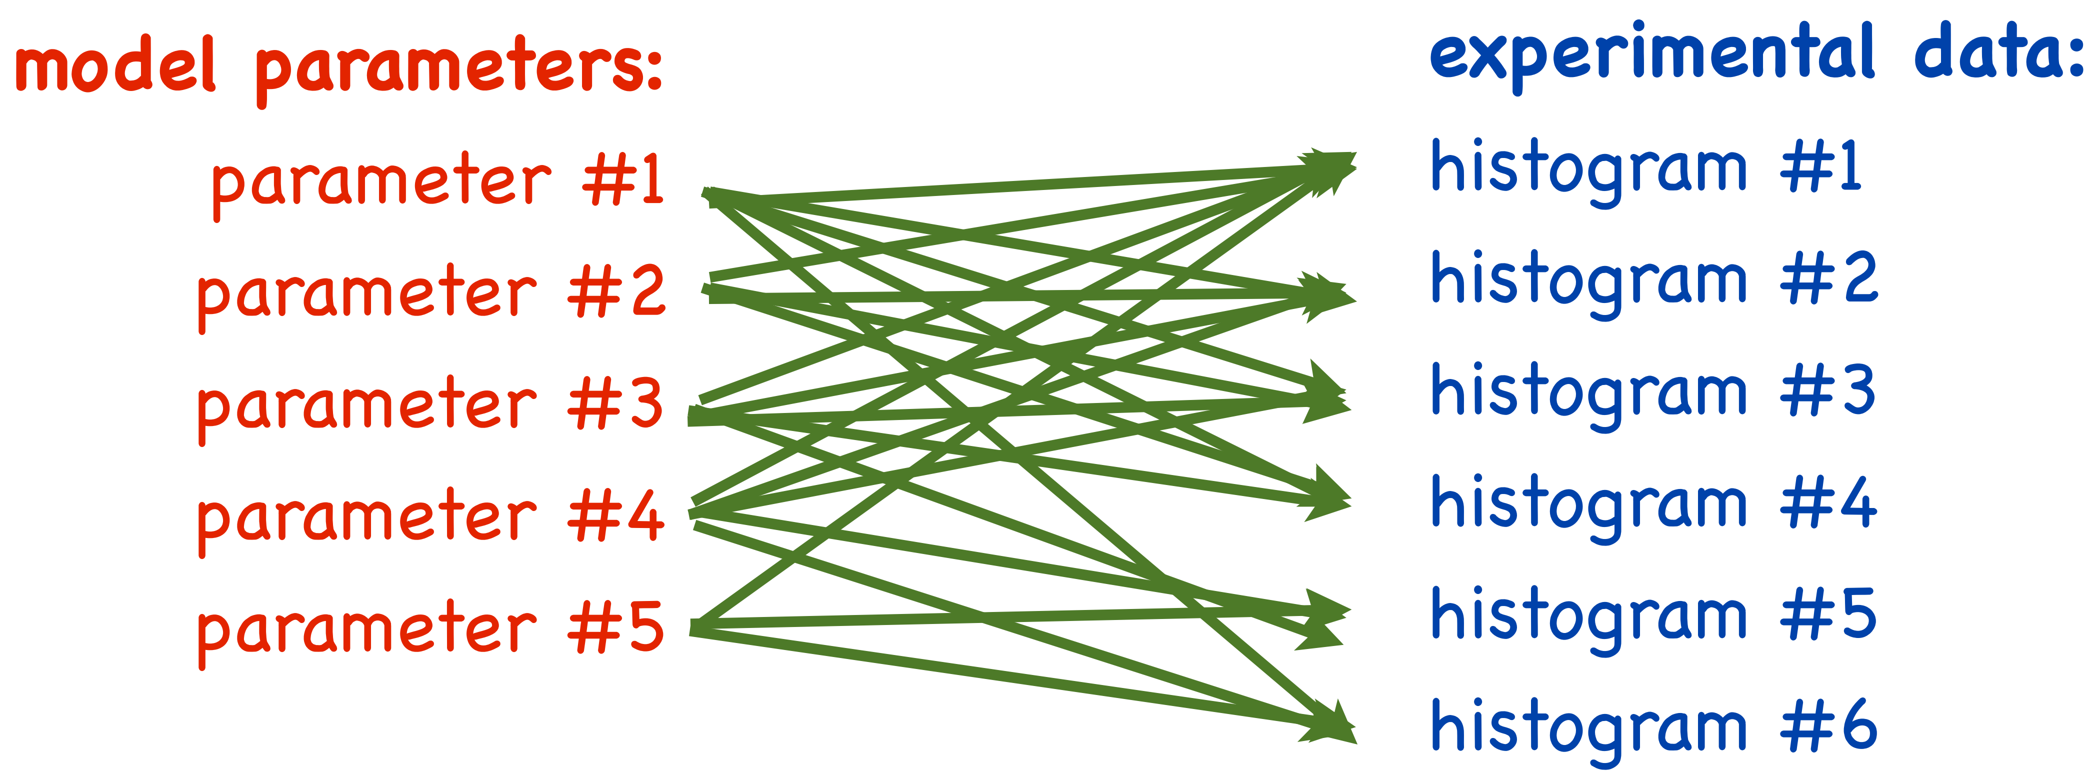

- large number of interconnected parameters w/ non-factorizable data dependencies
- data have correlated uncertainties
- develop novel optimization techniques: Bayesian Statistics and MCMC methods
- models are computationally expensive: need Gaussian Process Emulators
- general problem, not restricted to HEP/Nuclear Physics

→ **interesting challenge for Statistical Sciences**

Computational Challenges

- Open Science Grid
- Storage & Archival of Data

The Computational Challenge

- use selection of models which best cover the relevant physics
- accumulate sufficient statistics for each parameter set

Model	CPU/event	# of events	# of parameters	total CPU
Hydro	4 h	1	12	480 h
Hydro +afterburner	4 h (hydro) + 1 h per afterburner event	10,000	18	9,000,480 h
EbE-Hydro +pre/ afterburner	5 h	10,000	12	6,000,000 h
microscopic transport	1 h	10,000	6-18	1,200,000 h

assumption: use 10 bins for each parameter set

creating an event database reasonably covering the parameter-space for a selection of most promising transport approaches will require significant computing resources in excess of 15 Million CPU-hours (almost two years on a 1000 node CPU-farm!)

OSG Workflow: Duke QCD Group

RENCI Engage VO submit node:
- CONDOR script transmits job to OSG nodes
- job may run on 1-100,000 nodes independently

local desktop at Duke:
- prepare executable & input files
- configure job for 10-20 cpu-hours

Open Science Grid

Open Science Grid is best suited for:
- trivially parallelizable MC-based simulations
- tasks which can be completed within 10-30 cpu hours

GridFTP protocol

Nukeserv: 100 TB array at Duke

compute cluster @ Duke:
- combine individual job outputs
- run analysis on output files
- perform visualization tasks

OSG Throughput: Duke QCD Group

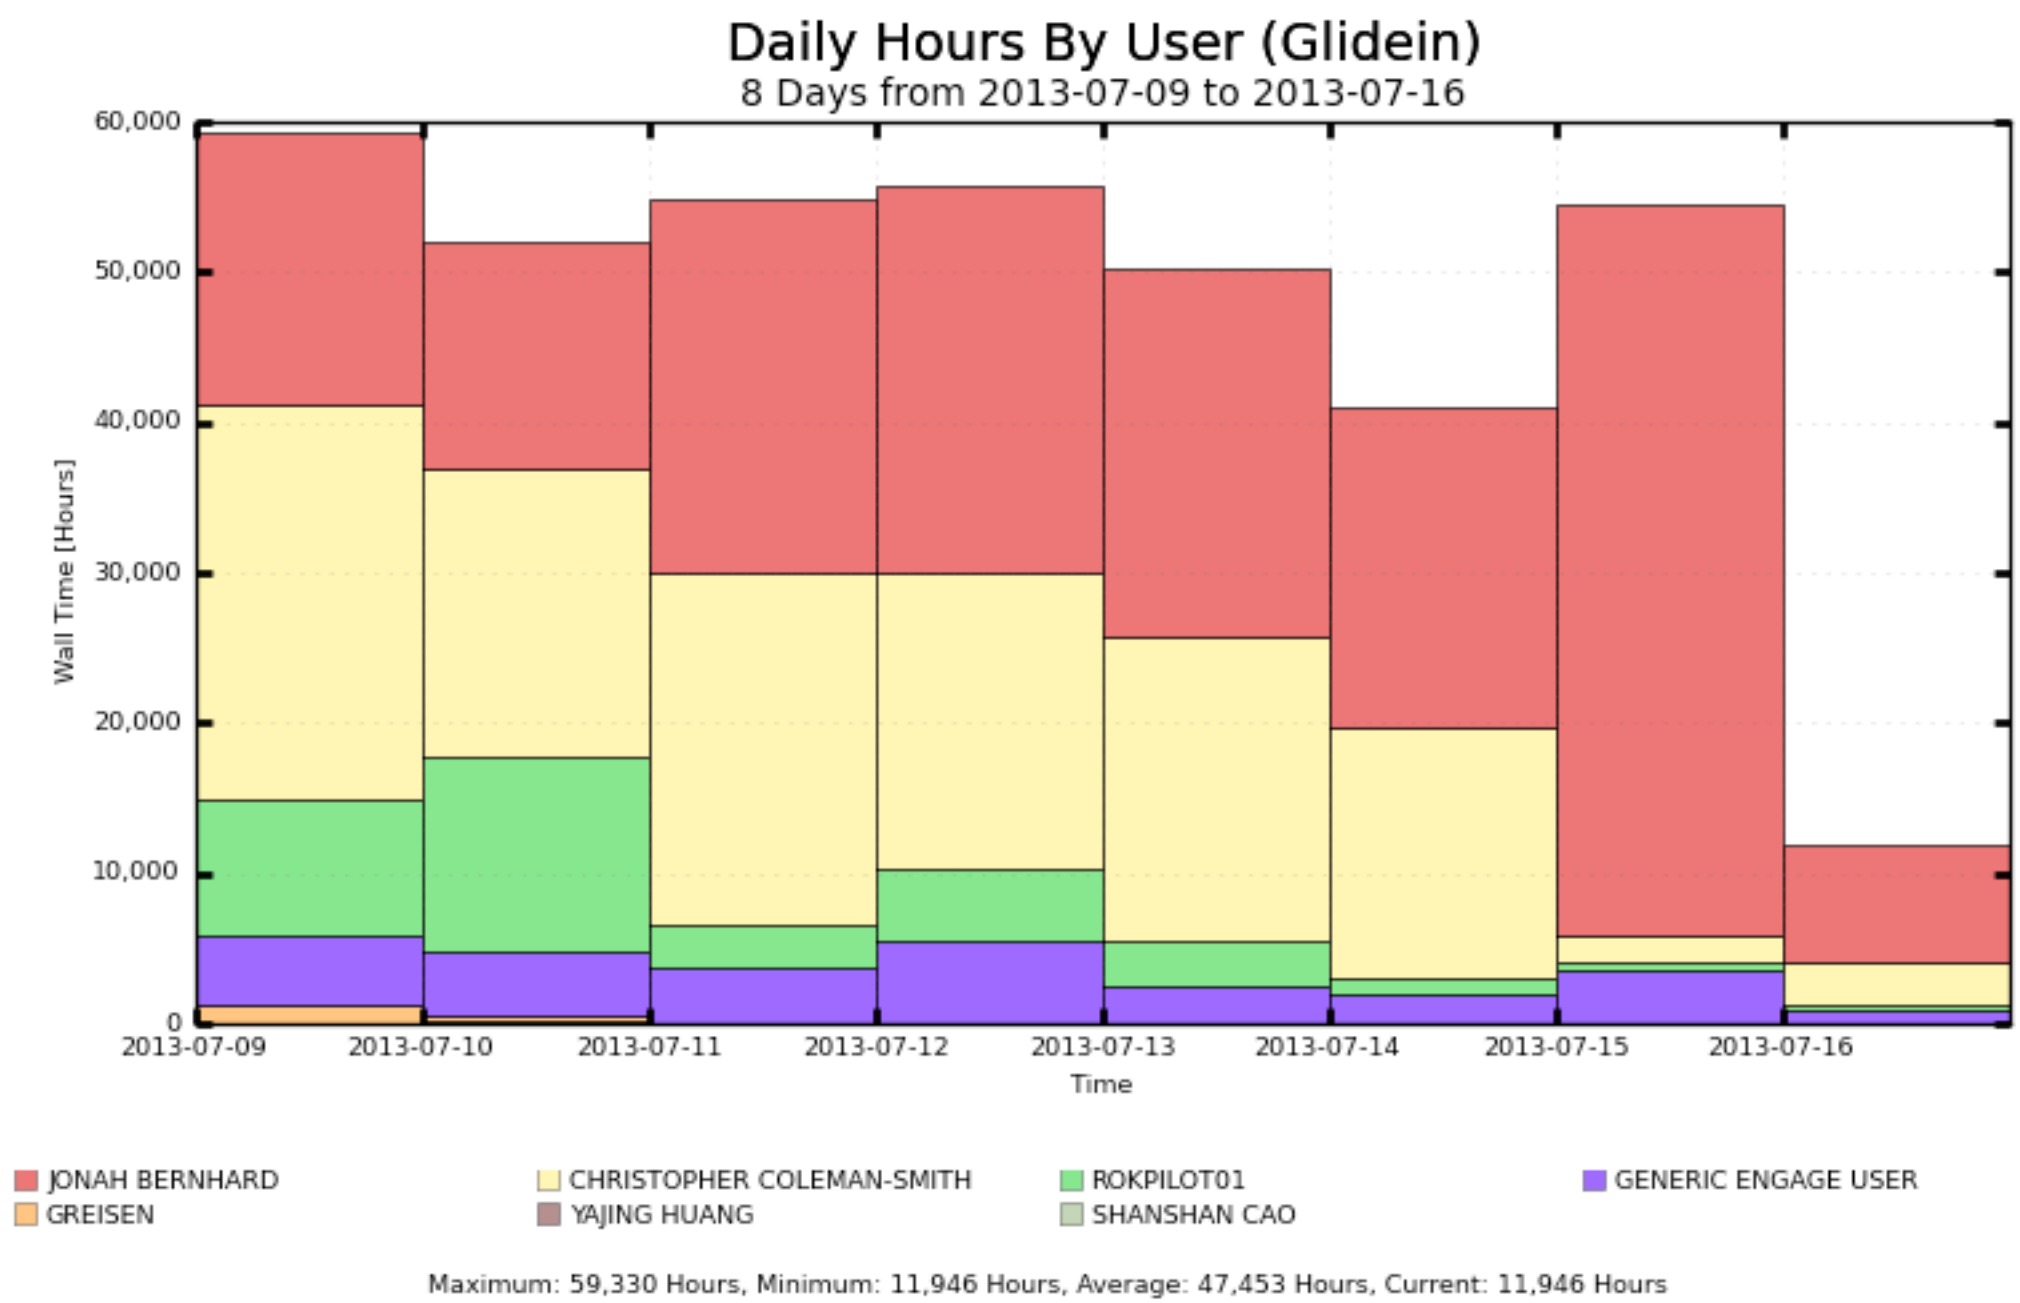

- up to 50,000 cpu hours per day from OSG resources!
  - computational projects previously thought unfeasible are becoming doable
  - still need to utilize statistical tools such as Gaussian Process Emulators to reduce the computational footprint

Data Storage & Management

- simulation data contains information of the full time-evolution of the collision: necessary for correlating hard probes with medium evolution

Model	storage/event	# of events	# of parameters	total storage
Hydro	2 GB	1	12	240 GB
Hydro +afterburner	2 GB (hydro) + 100 MB per event (micro)	10,000	18	180 TB
EbE-Hydro +pre/ afterburner	1.1 GB	10,000	12	1,212 TB
microscopic transport	100 MB	10,000	6-18	120 TB

assumption: use 10 bins for each parameter set

creating an event database for the simulation data requires about the same storage capacity as 1-2 years of running of RHIC experiments (1.5 PB). Requires significant resources and management tools previously not known or available to the Theory Community

Data Storage: HPSS at NERSC

HPSS Capabilities:

- theoretical capacity: 200 Petabytes
- buffer (disk) cache: 288 Terabytes
- theoretical maximum throughput: 6.4 GB/sec

Problem:

- network bandwidth/capacity to transfer data to/from NERSC

Outlook & Challenges

Heavy-Ion collisions at RHIC and LHC have produced a novel state of matter which is called the **Quark-Gluon-Plasma:**

- frontier-science w/ a vibrant community: most discoveries have only been made during the last decade and many more are expected for the next decade
- experiments and model calculations are generating an unprecedented amount of data (Petabytes) and require extreme-scale computing resources (Petaflops)
- grid computing has been an extremely successful tool for providing the required compute cycles
  - main challenge: data storage and archival needs of the modeling community outpace current resources and/or available bandwidth for transfer to archival locations
  - need the OSG equivalent for data storage!!

The End