Computing Trends in Nuclear Physics Software & Technologies: A perspective and look forward

Jérôme Lauret
RHIC/STAR Software and Computing Leader

BROOKHAVEN NATIONAL LABORATORY

Disclaimer

- Talk not aimed to go in depth into technology or computer science but will tell a “story”: from past success to now – my view on *a few* Computing trends – and the challenges ahead.
- My talks will be US centric but examples from High Energy Physics (HEP) / Nuclear Physics (NP) will be provided
- I may have some biases in the examples ☺️

Outline

- The NP landscape – paradigm of the past
  - Success model … and first signs of changes
- The NP Ecosystem – beyond particle physics experimental
  - External drivers and influences – emerging trends
- From Triggered to "Streaming DAQ" datasets
  - From events to frame and consequences …
- Tackling the complexities
  - Solutions? What is being worked on …
- Beyond science – the human resource factor
  - Tomorrow's collaboration landscape – an hopefully diverse environment composed of many valuable (and valued) skills
- Concluding remarks

The NP landscape – paradigm of the past

- Lesson from the past – approaches and frameworks
- Outcomes
- Cloud computing
- More changes on the horizon

BROOKHAVEN NATIONAL LABORATORY

Lesson from the past
Approaches and frameworks

Our communities have developed its computing strategy around a few principles

- (a) Reliance on Moore's law and (b) principle that memory is "free", Flops (machines) are expensive
  - Moore's law: After Gordon Moore, co-founder of Intel "the number of transistors per square inch on integrated circuits doubles every year"
  - Memory: 16-core Xeon, 20MB L2 Cache, 64GB RAM **2+ GB/process not uncommon**
  - Drove "data heavy" designs, global data-models (data structure accessible at any point in time)

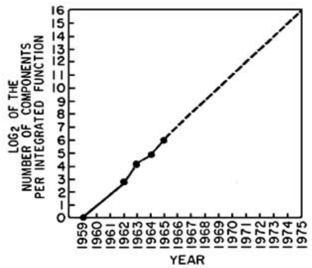

- A heavy dependence on C++, seen as the language of the future
- The development of Monolithic framework (root4star, ALIROOT, ..)
  - Aligned with C++ philosophy – Allows for code re-use, well designed code, rapid addition of new modules, ...
  - Easy integration of simulation, reconstructions, ...

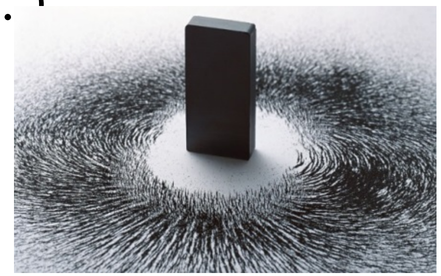

Lesson from the past
Approaches and frameworks (cont.)

- Dependence on a few packages that have become community standards (ROOT, Geant. …)
  - ROOT provided all we ever wanted (histograms, Ntuples, IO, …)
- Grid middleware and distributed computing strategies developed early & data access and distribution layers
- Often separation of online (DAQ) and offline software – DAQ had "reasonable" rates
  - GLUEX until last year, 35 kHz and 700 MB/s
  - RHIC STAR - started at a few 100 Hz, proud to exceed LHC experiments in 2010 at 600 MB/sec (CHEP 2009)

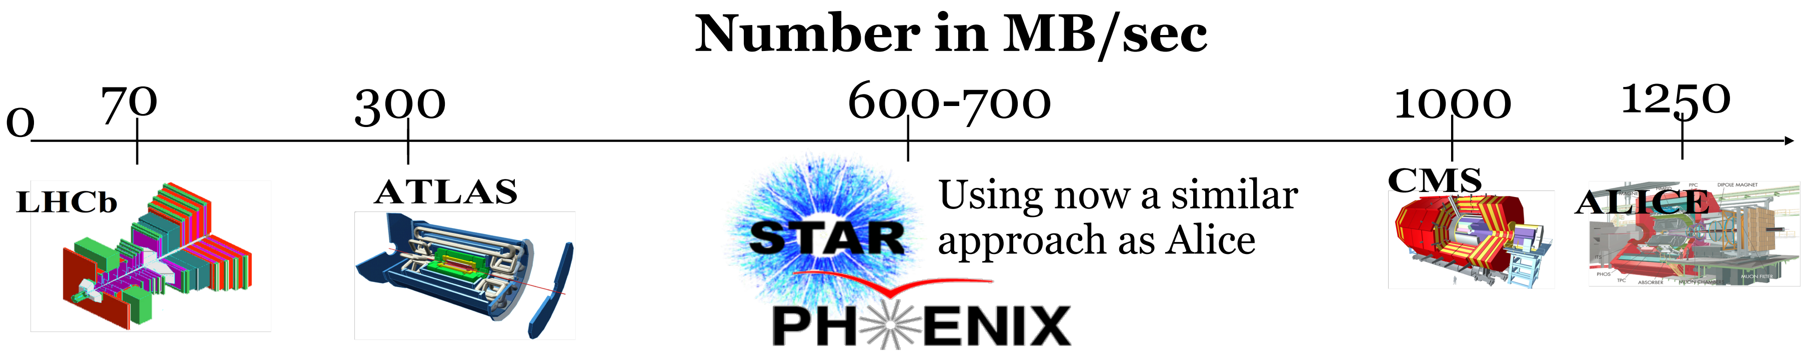

Modest DAQ evolution

- As data demands increased, did not shy off continuing on the same path
  - DAQ were modified to include L3 tracking and decision making (STAR) – data reduction by x100 (NIM 453, 397 (**2000**)
  - L3 evolved into L4 (High Level Trigger / HLT) at ALICE back to STAR
  - Algorithm were tested to do Online Cluster Finding (OCF objective = trash raw data), some tracking online to reduce data or tag events …
- Diverse level of success with HLT/OCF approach
  - "*Trigger Mode*": reduce data volume by selecting events of interest or rejecting events of no interest
  - "*Partial Reconstruction Mode*": reduce data volume by replacing raw data with processed data (TPC clusters)
  - "*Express Streams*": tag events for prioritized offline reconstruction / analysis into separate file streams
- Even more daring moves at the LHC – real-time calibrations, injection of simulation events into the data stream for on-the-fly efficiency corrections, …

Outcomes

- Rapid delivery of corner stone physics results at RHIC
  - Most processing done locally, data and IO, analysis all scaled well
  - HLT @ STAR leaded the way for a while (Anti-HyperNuclei in 2011)
- Huge success at the LHC with near-immediate discoveries
  - LHC-Grid showed to be a huge help
    - The world-wide participation in resource sharing brought together scientists from all over
    - Single site / single point of failure thought impossible
- Distributed / Federated resources and Grid seemed to be a real game changer (LCG, OSG, NorduGrid, ...)
  High Throughput Computing (HTC) ruled!
- Despite resource demand increase, full speed ahead with a well-greased model!

Small "cloud" on the horizons (pun intended)

- Cloud computing appeared from industry
  - Evaluated as soon as **2007/2008** @ STAR, Amazon/EC2 resources usable from day 1
  - Scalability tests followed …
- Resources were "standard"
  - Followed our usual model "Commodity Hardware / Off-the-Shelf" (COTS)
  - Business as usual – pleasantly parallel workflows (job parallelization)
  - Evolving project to seamlessly integrate new wave of resources (OSG, …)

- Hidden game changer
  - Cloud brought virtualization and Containers to the front scene (many positives)
  - PaaS, SaaS, …

Workflow recap

- Detector signals were digitized by electronics in front-end crates, and data readout initiated by the trigger electronics. Data was transported to an event builder, and the subsequent built events distributed for filtering, monitoring, display, etc… The event stream were also stored to MSS for offline analysis.
- *Data throughput was determined by the physics event rate*

- Frameworks were monolithic, all-purpose
- Detector simulations were performed by using standard codes such as Geant, Fluka, and virtual Monte Carlo (VMC).
- Geant, ROOT used as common blocks

- The compute infrastructure for data analysis and data storage/archiving usually resides within the host facility (Tier model)
- Distributed storage and/or computing for reconstruction (and/or analysis) – High throughput computing is king1 (HTC)

First sign of change?

- As the complexity and size of the experiments grew so did the complexity of analysis environment
  - Time dealing with the analysis infrastructure grew
  - Not due to framework design alone but framework + data handling+ larger datasets (abstracted as well / no "ls" and I see it all) + distributed computing
  - Too many abstract layers slowed framework down

- Cloud should have been a good sign that industry influences our field
  - Both hardware and new tools appeared, influencing ideas
  - Intel & Nvidia releasing more and more hardware – OpenLab created at CERN to deal with flow of novelty and create a two way communication

- Low hanging fruit Physics rapidly over
  - From RHIC – within a few years, data increased x10, x100 now in the Billions.
  - Rare probes are demanding

- And more important … busy with success, "we "lost sight of "*other branches in the Nuclear Physics fields*" and other external driving force s

The NP Ecosystem – beyond particle physics experimental

Moore's law and so what?
The other fields and their needs
HPC era and new challenges

BROOKHAVEN
NATIONAL LABORATORY

Moore's law and its impact

- **The collapse of Moore's law** is not a late discovery to anyone …
  - Clock speed plateaued as soon as 2004
  - Path was however to (still) move ahead by squeezing more cores per chip and improved architectures
  - *Experiments did not change their framework*
    - Whether more core or pseudo-cores (aka hyper-threading), pleasantly parallel still worked
- Manufacturers are struggling to shrink transistors.
  - How much further can Moore's law continue?

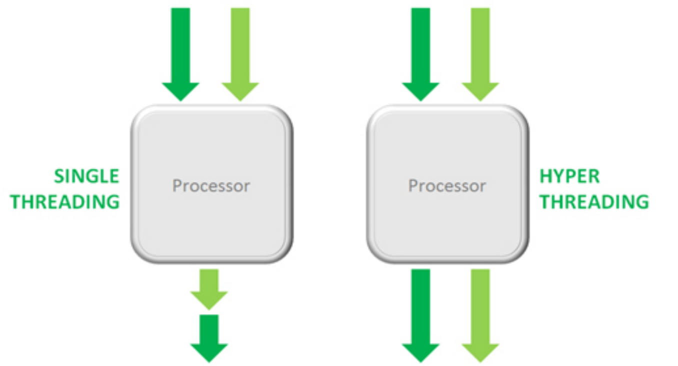

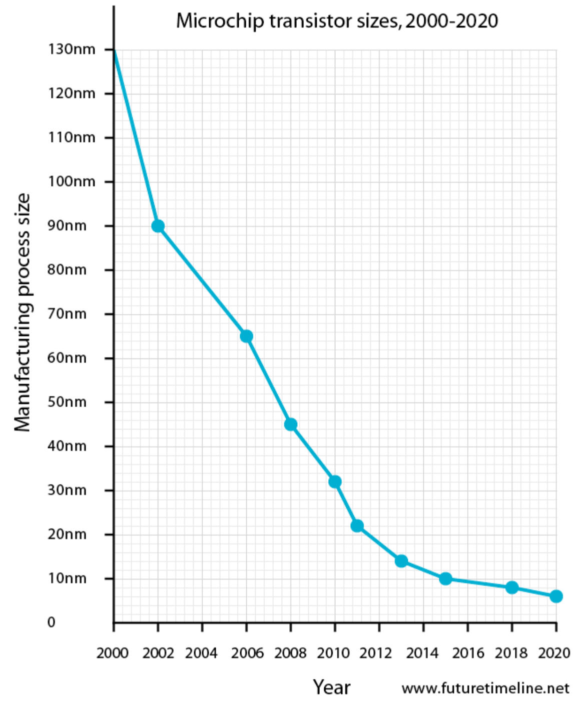

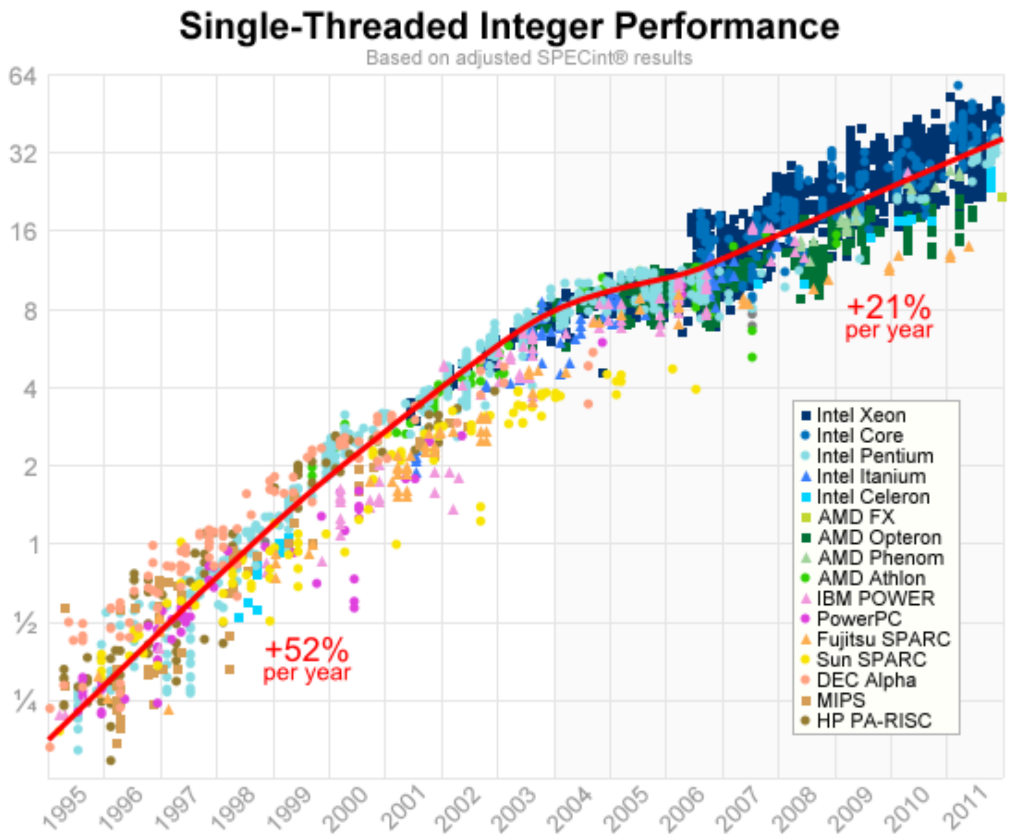

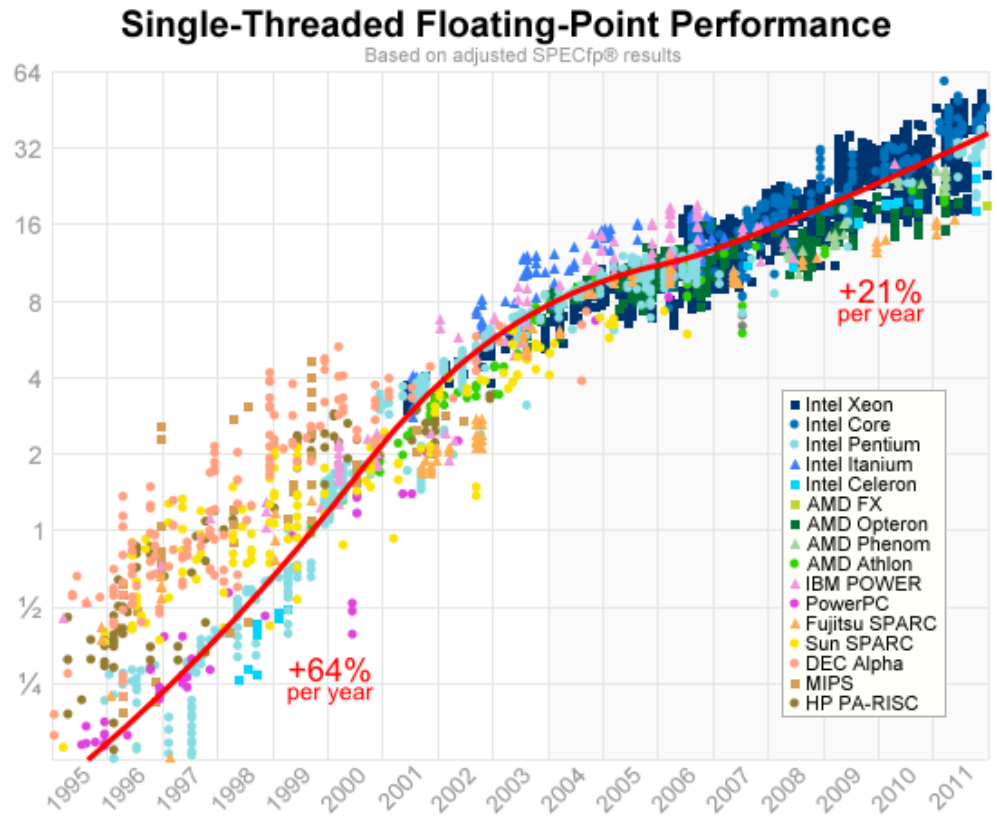

Figures from "A Look back at Single Thread CPU performance"

- The Y axes represent a compendium of multiple standard benchmarking test such as SPECInt (for Integer operations) and SPECfp (for floating point).
  - Before t2004, FP speed climbed at +64% per year: a doubling period of 73 weeks. Now, levelled at 21% per year.

Industry response – (a) try new architectures (b) push toward High Performance Computing (HPC)

Hardware diversity ⇔ porting challenge

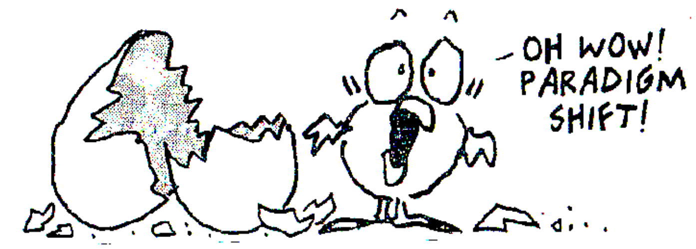

- Different "cache memory" architecture
  - code optimization has to be rethought
- HPC (usually associated to super-computers) ⇔ Full use of parallelization for TeraFlops capability (10^{12} floating-point operations per second)
  - **From "memory is cheap, CPU is expensive" move to the *reverse* paradigm: Flops are free, Memory is expensive & must be managed**
    Example: 68-core Xeon KNL, 34MB L2 Cache, 16GB HBM (High Bandwidth Memory), 96GB RAM – **now memory is < 2 GB/core!!!!**
- IO on HPC systems is known to be poor
  - Beyond networking or disk IO, abysmal performance + no scaling beyond O(1k)

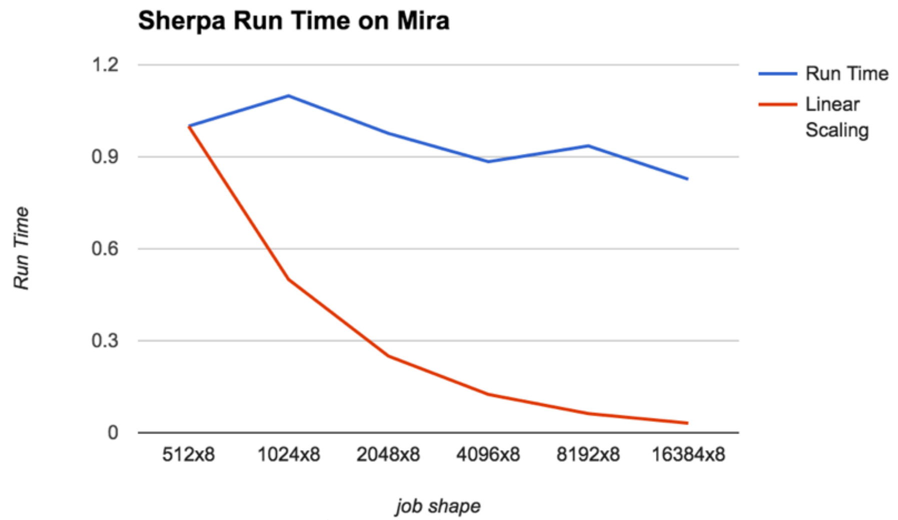

MIRA (ANL) = BlueGeneQ PowerPC

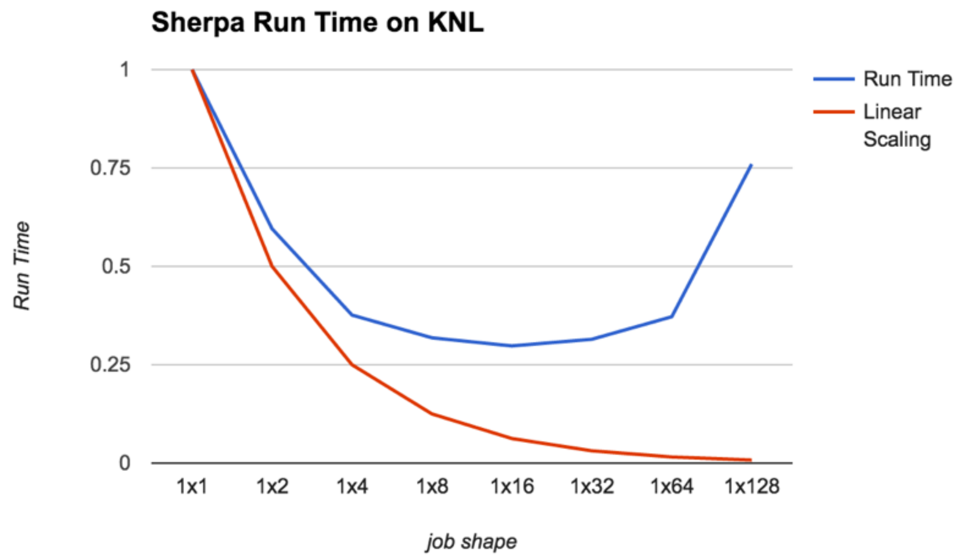

THETA / KNL (ANL) == Xeon Phi

Example: Sherpa on KNL (Taylor Childer, CHEP2016)

- Refactoring code - performed thousands of file operations at startup (hadron decay information, process information & libraries, ...)
- Removed unnecessary stat() calls + clean-up call => still cannot harvest full performance

And this is for a simulation package.

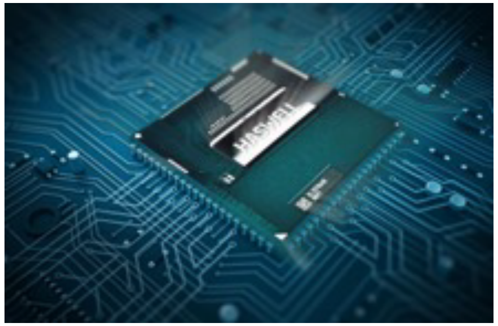

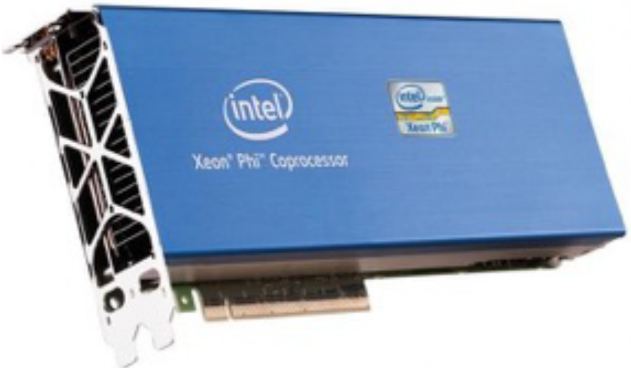

OLCF

ALCF

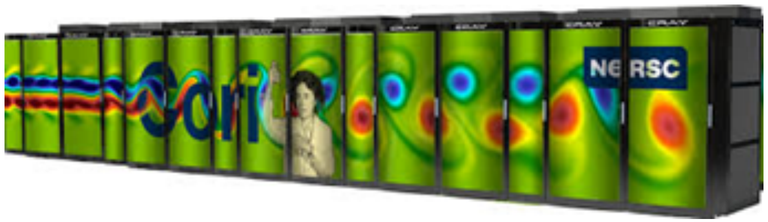

NERSC/Cori

Imagine the Code Portability issue for experimental frameworks that were NOT designed to be parallel!

Portability is a currently a big issue. The problem will get harder …

Code/frameworks are emerging as a major showstopper

Closer inspection …

- *"Detector simulations are performed by using standard codes such as Geant, Fluka, and virtual Monte Carlo (VMC)"*
  "Geant, ROOT as common blocks"
  - Geant4 is parallelized not vectorized
    Geant V (next generation design) > a decade away
    ROOT, used in production for more than 20 years as analysis frameworks – has not showed to scale to Exascale
- *"Frameworks are monolithic, all-purpose"*
  - Meaning we **solve all problems or brace for impact (or re-design?)**

Already does not look good BUT do we (NPP experiments) need to go "there" (HPC)?

Race toward HPC

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	National Supercomputing Center in Wuxi China	Sunway TaihuLight Sunway MPP, Sunway SW26010 260C 1.45GHz Sunway NRCPC	10,649,600	93014.6	125435.9	15371
2	National Super Computer Center in Guangzhou China	Tianhe-2 (MilkyWay-2) TH-IVB-FEP Cluster Intel Xeon E5-2692 12C 2.200GHz TH Express-2, Intel Xeon Phi 31S1P NUDT	3,120,000	33862.7	54902.4	17808
3	DOE/SC/Oak Ridge National Laboratory United States	Titan- Cray XK7 Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc.	560,640	17590	27112.5	8209
4	DOE/NNSA/LLNL United States	Sequoia- BlueGene/Q, Power BQC 16C 1.60 Ghz Custom IBM	1,572,864	17173.2	20132.7	7890
5	DOE/SC/LBNL/NERSC United States	Cori- Cray XC40 Intel Xeon Phi 7250 68C 1.4GHz Aries interconnect Cray Inc.	622,336	14014.7	27880.7	3939

~125 Petaflops, ~15 Mwatts ~11 million compute cores, ~1.3 PB ram
Expecting 1 ExaFlop machines by 2022

Preparedness

- US Agencies have carried many workshops to identify the challenges and help prepare for the Exascale computing era
  - Community slow reacting
  - Last report looked at up to the 2025 timescale

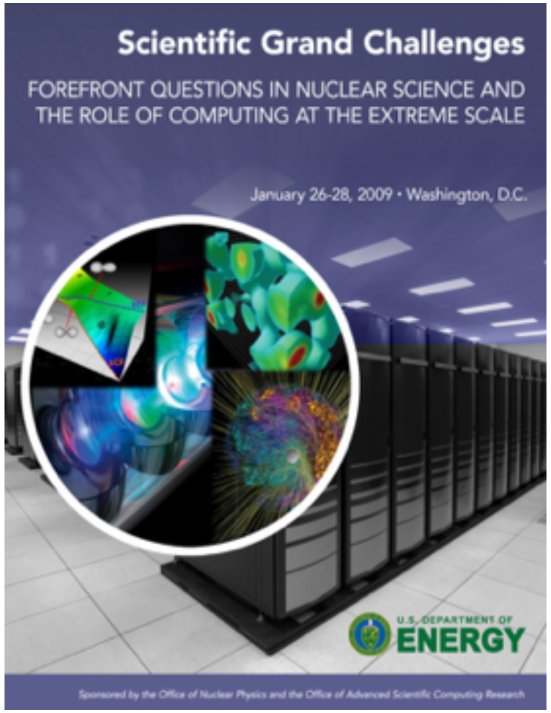

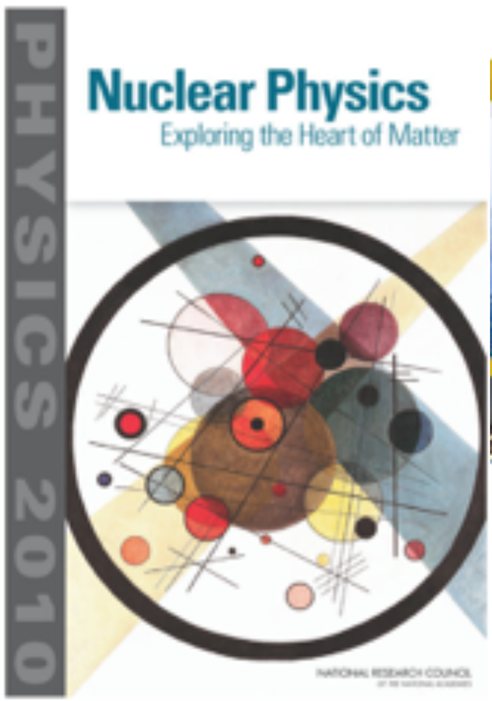

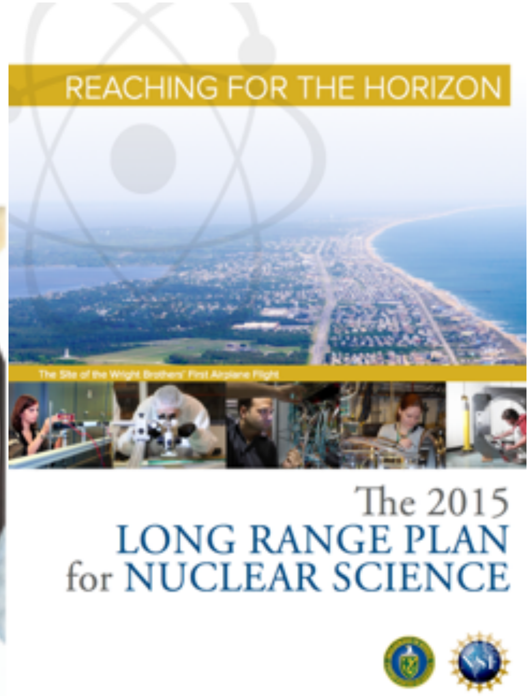

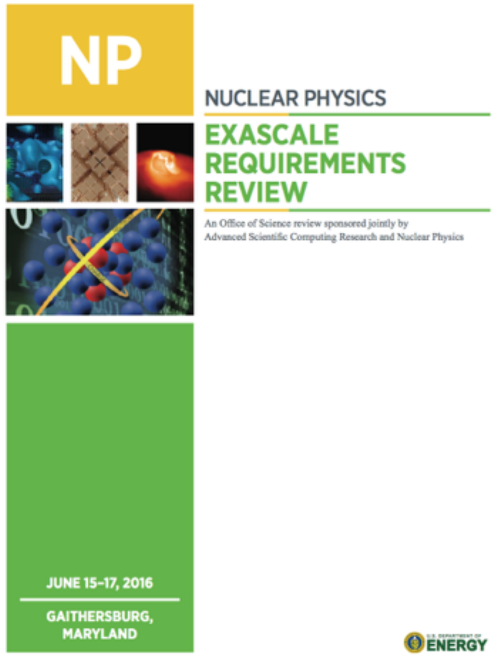

Exascale requirements in NP

The communities

- Nuclear Astrophysics
  - formation of heavy elements in the universe, the fundamental properties and role of neutrinos in supernovae
  - Gravitational waves
- Experiment and Data
  - Data streaming from detectors,
  - Complex work flows that can accommodate experimental and simulations data, and
  - Advanced networking and data sharing
- Nuclear Structure and Reactions
  - FRIB
- Cold Quantum Chromodynamics
  - Precise lattice QCD calculations
  - Gluonic structure with unprecedented precision
  - EIC simulations
- Hot Quantum Chromodynamics
  - Establishing the phase diagram of QCD
  - BES at RHIC

Grand Challenge in an Exascale ecosystem

- Essential detector and accelerator simulation
- Data acquisition, storage, retrieval, and analysis
- Precision calculations of nuclei and their reactions, with quantified uncertainties
- High-fidelity studies of the nuclear astrophysical objects in our universe
- Theoretical predictions of critical quantities directly from quantum chromodynamics (QCD) or through nuclear many-body techniques
- Exploration of new and exotic states of hadronic matter.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025 EXPERIMENT AND DATA

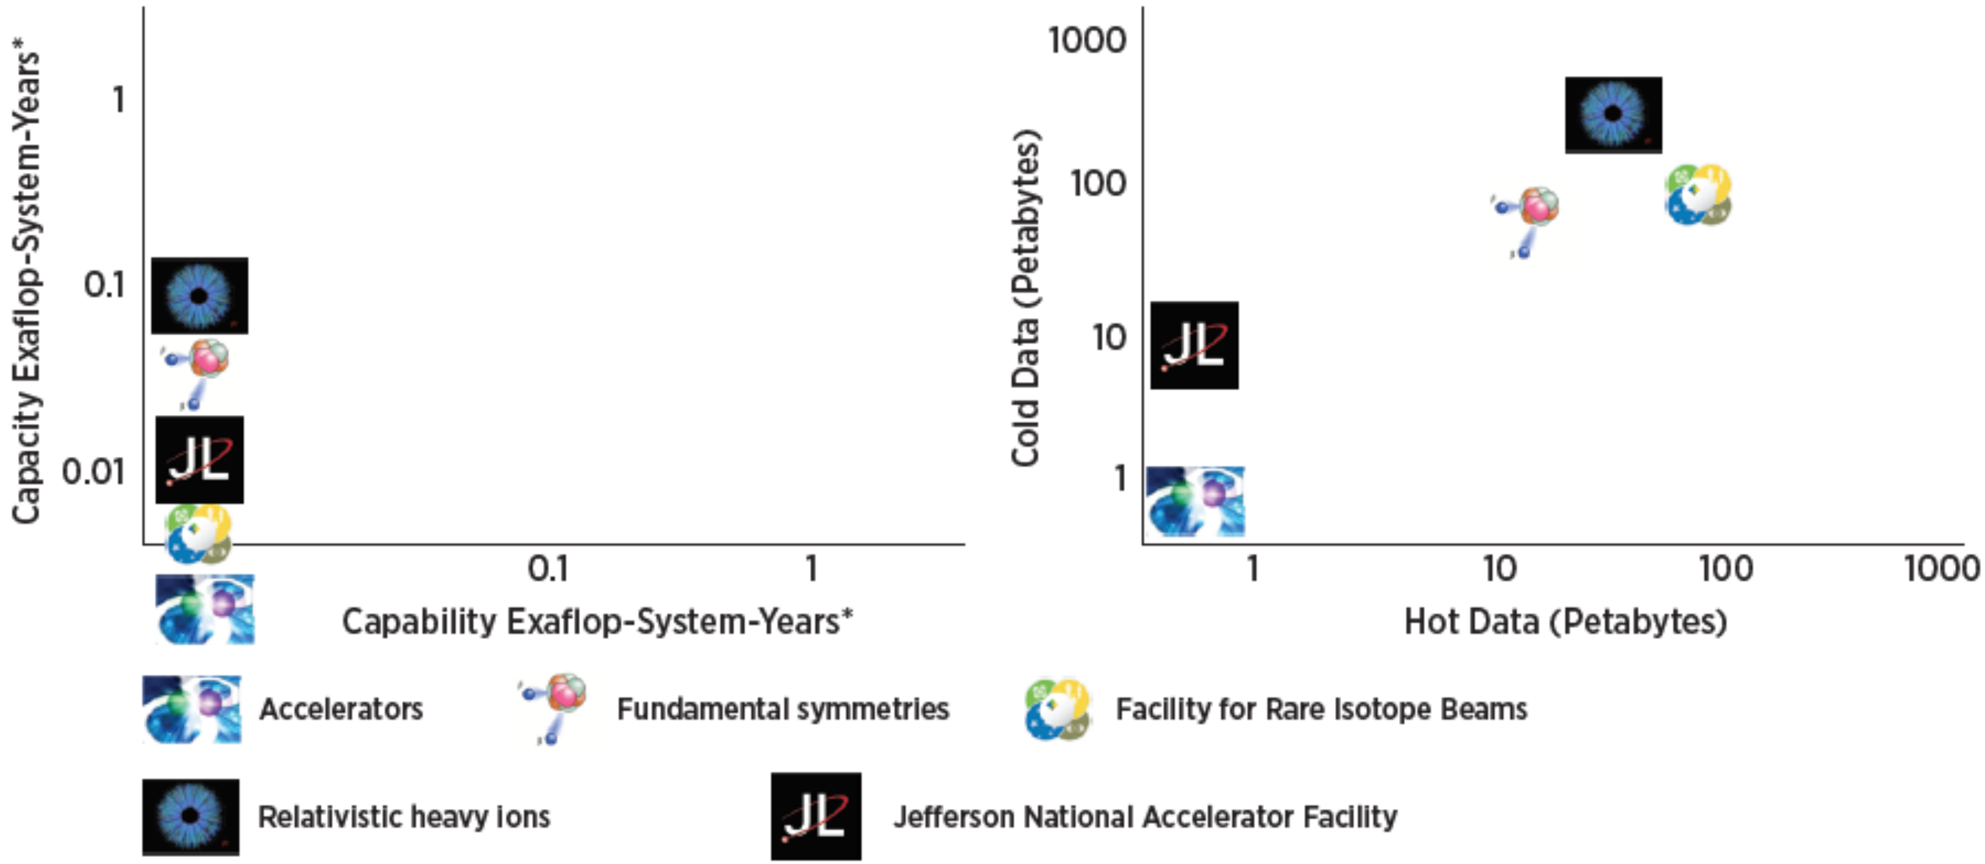

“hot data” = data available for immediate access for a short to an extended period of time
“cold data” = data that can be retrieved from storage in a reasonable time frame to data archived on tape with the expectation of being touched at intervals of a year or possibly more

Capability computing = large machines, super-computers, HPC needs
Capability computing = Hardware and storage (COTS), clusters

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025
NUCLEAR ASTROPHYSICS

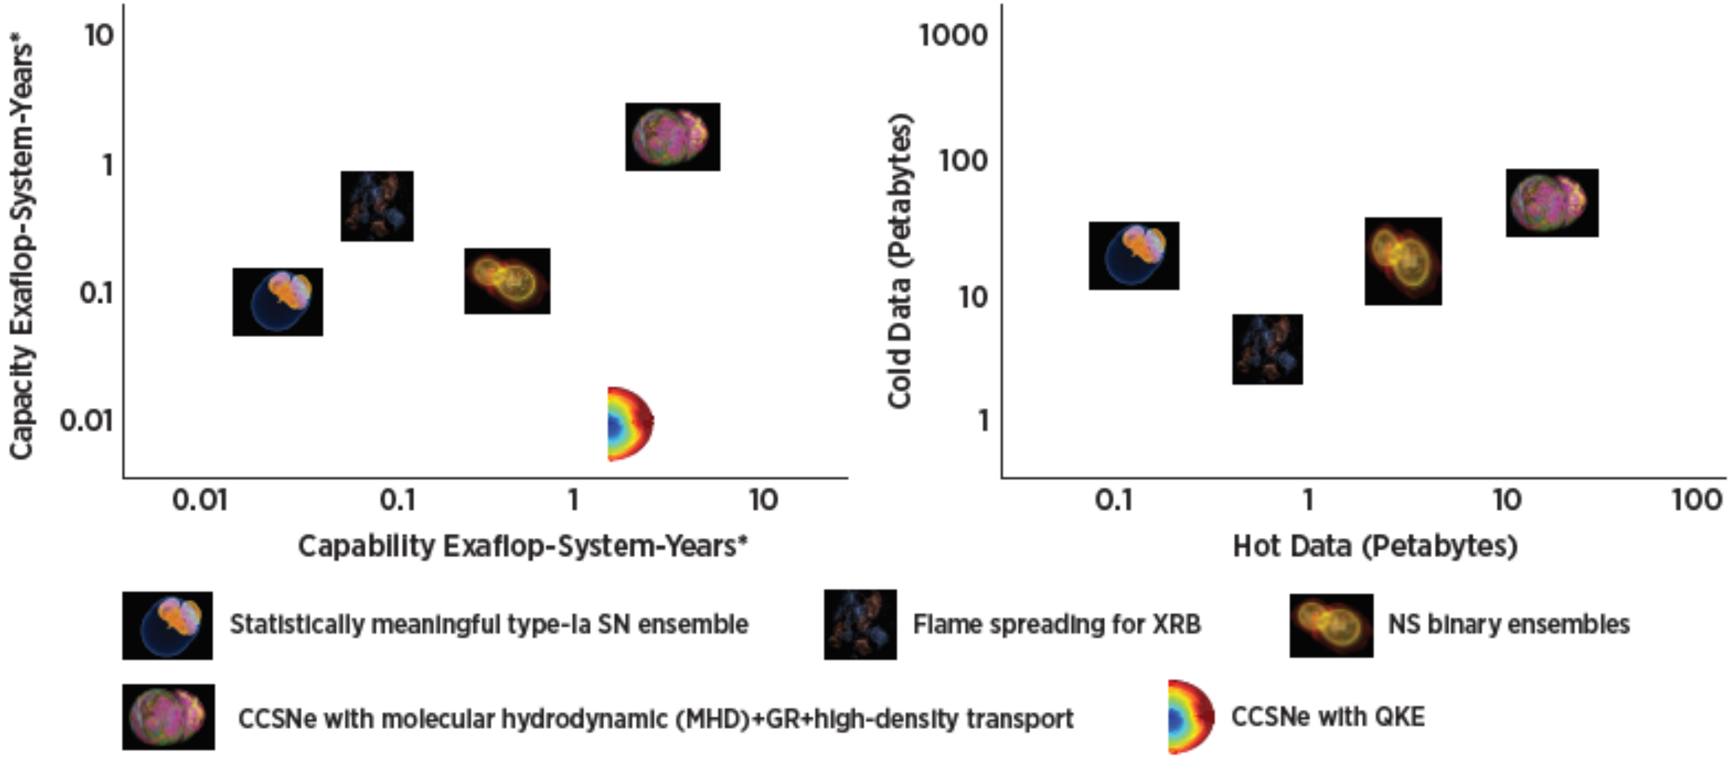

* Exaflop-system-year refers to the total amount of computation produced by an exascale computer in 1 year.

CAPABILITY/CAPACITY RESOURCES VS. HOT/COLD DATA RESOURCES IN 2025
COLD QCD

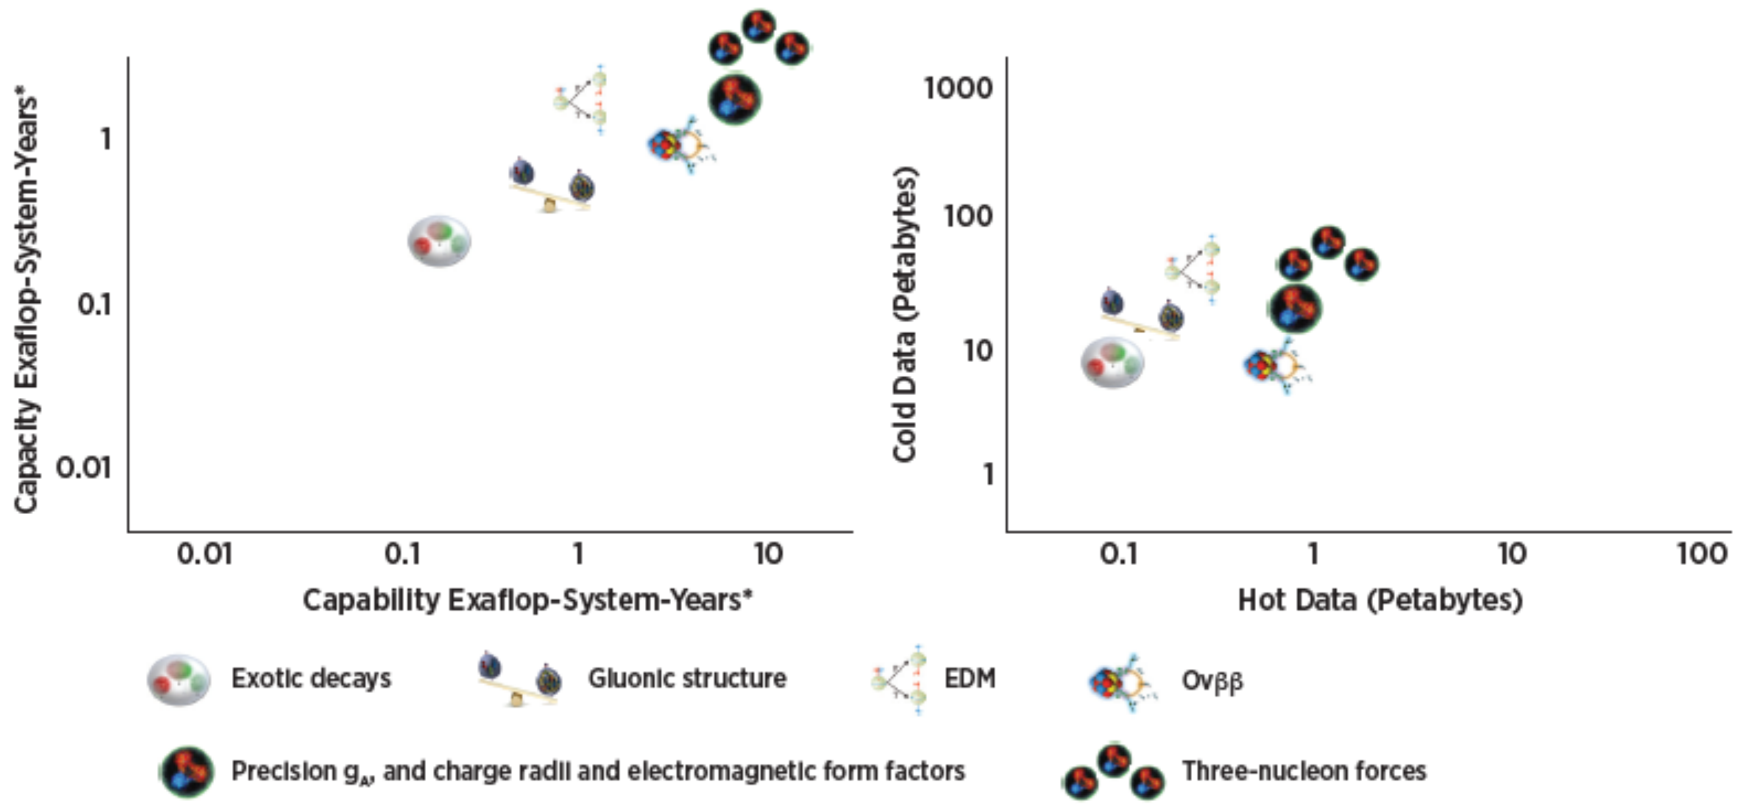

In short – all NP fields (but experiments) have large HPC requirements
Large investment in HPC is near certain …

Emerging trends

- **HPC is coming!** (actually, it is here)
  Future complexity and its daunting challenges
  How do we design "something" within the current parameters?
  - Most powerful future computers will likely be very different from the kind of computers currently used in Nuclear Physics (Exascale Computing a beginning)

- **How to design today a framework that would survive >> 1 decade?** How to make it robust against likely changes in computing environment?
  - Possibly: Apply modular design: changes in underlying code can be handled without an entire overhaul of the structure

From Triggered to “Streaming DAQ” datasets

ALICE
LHCb

BROOKHAVEN NATIONAL LABORATORY

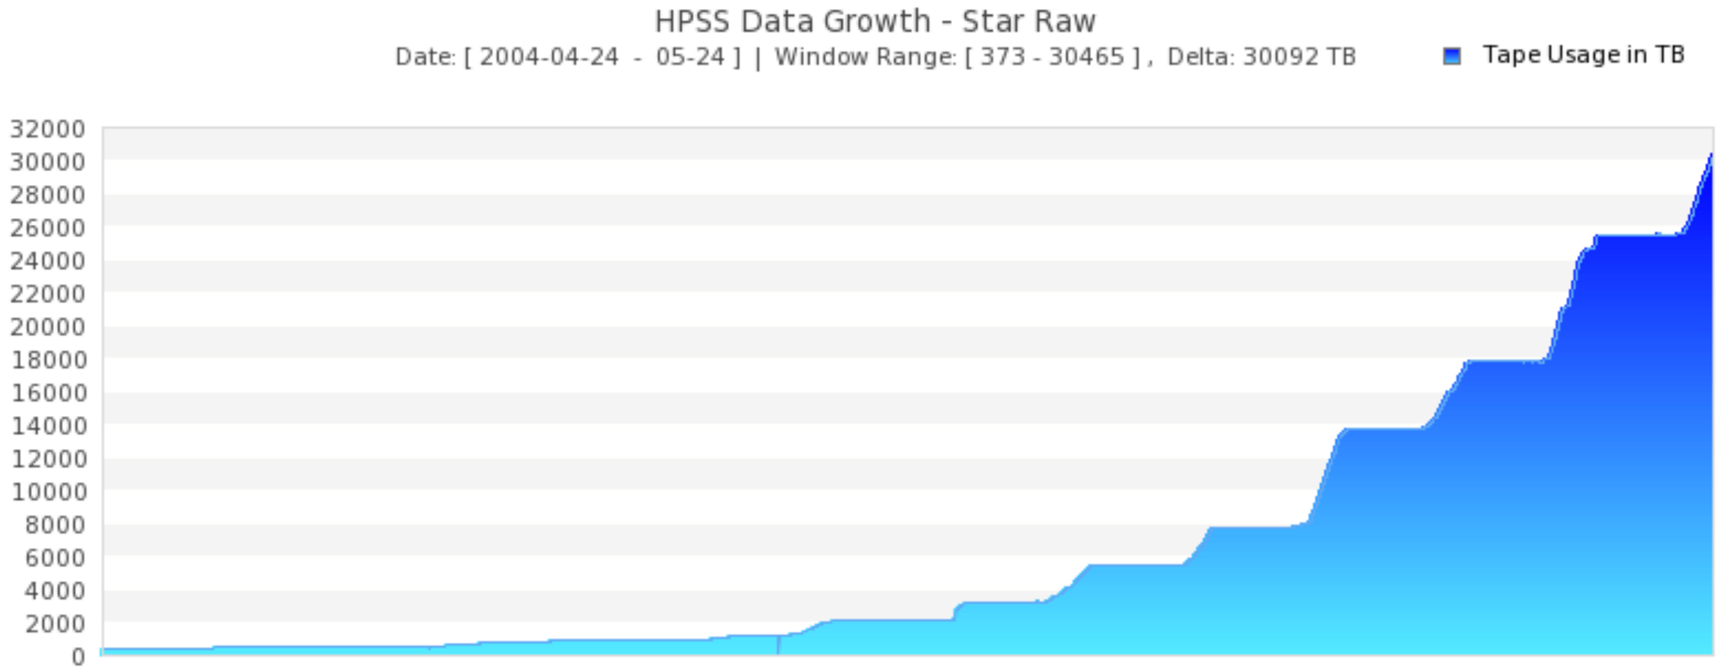

"Low hanging fruit Physics rapidly over "

Large datasets needed for statistically challenging Physics. Triger / DAQ improvements?

Experiments' new needs

- FPGA -> CPU (included Phi) -> CPU/GPU mix … and back to FPGA
  - "Song and dance" driven by cost/performance

- The highest trigger rate experiments increase rate by factor of ~10 every 10 years.
  - RHIC/STAR Maximum rate = 2.1 GB/s, average = **1.6 GB/s** (600 MB/sec in 2009, 100 MB in 2005)
  - FRIB - GRETA 4000 channel gamma detector, 120 MB/s/channel (2025 timescale) ⇔ **~ 480 GB/s**
  - LHCb – RAW rates @ **TB/sec**
    - Storage wise would be a huge cost impact - online processing a MUST
    - Final rates foreseen at 50 GB/sec
  - ALICE upgrade (2020-2025 timeframe) – **also in multi TB/sec** detector readout
    - Can only be done by Data reduction by (partial) online reconstruction and compression

Rates are so high that "streaming DAQ" are on the horizon (triggerless / no longer a concept of "event")

ALICE / O² project

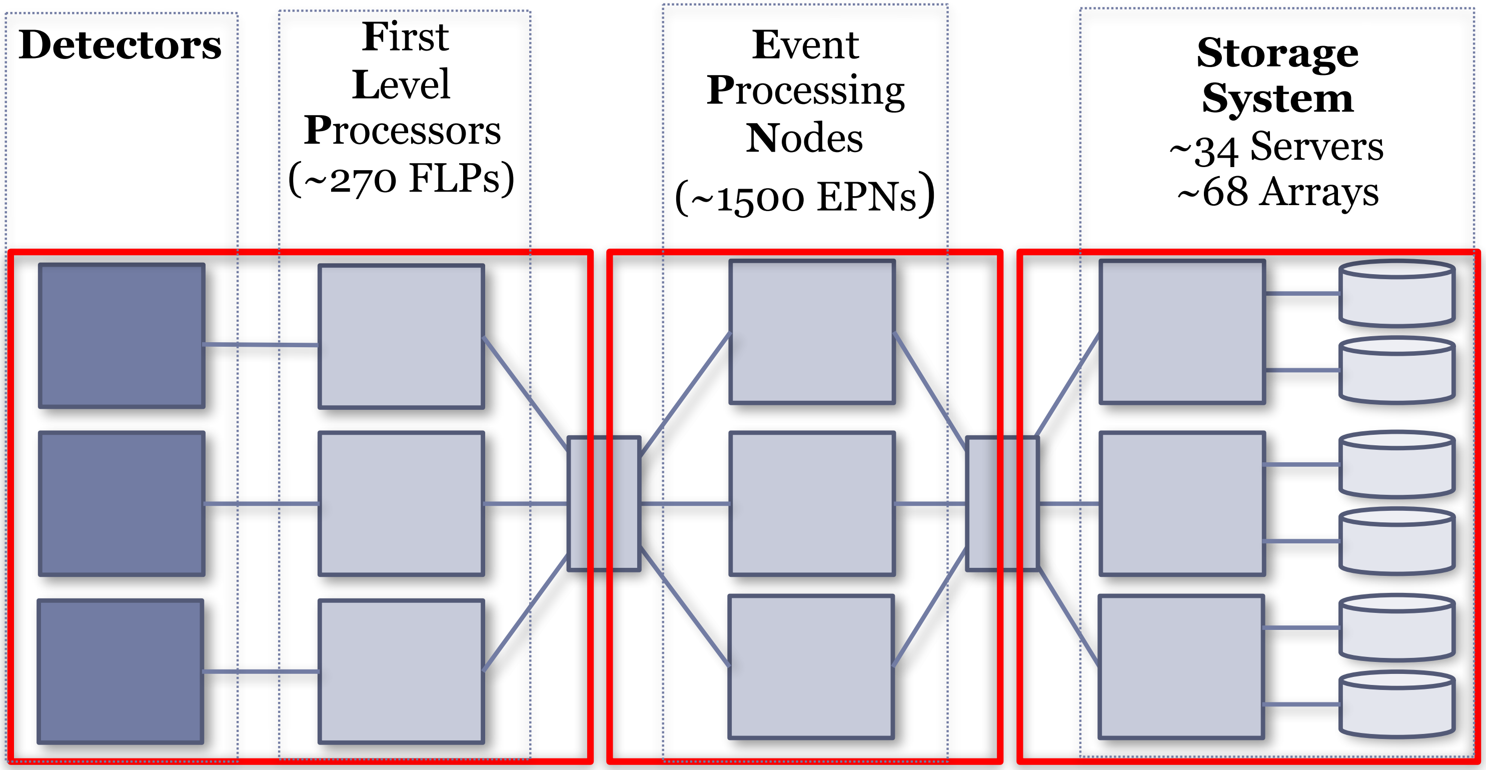

Versatile Links/GBT ~9000 → ~3.4 TB/s

COTS Network → ~500 GB/s

COTS Network ← ~270 GB/s → ~170 GB/s

→ ~20 GB/s

- Triggerless (streaming) readout-electronics
- HLT performs calibration online
- Data compressed and reduced
- ALICE O² ⇔ conceptually, no difference between online and offline
- Large farm moved closer to the experiment ...

Thorsten Kolleger, From todays HLTs to Streaming DAQs and Online Reconstruction: challenges and limits, Software & Technology Forum

LHCb

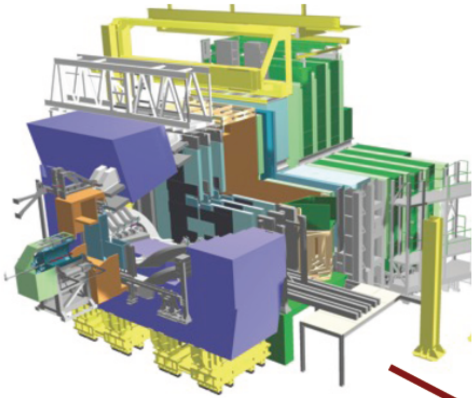

Simple feature-building in custom electronics (e.g. FPGAs) required to reduce the data volume to a transferable rate.

TB/s
(post zero suppression)

Online computing farm processes 250 PB / year -- can only persist 1% of this.

50 GB/s

LHCb will move to a triggerless-readout system for LHC Run 3 (2021-2023), and process 5 TB/s in real time on the CPU farm.

- Full data calibration in the online system in near real time
- Run reconstruction in real time in the online system at the full 40 MHz rate (now 1 MHz)
- How is this possible?
  - Use of Machine Learning pattern recognition all the way!
  - Trained on MC data

V.Gligorov, MW, JINST 8 (2012) P02013 (Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree)

Michael Williams, Machine Learning and Real-Time Analysis at LHCb)

Envisioned to have large processing farm "closer" to the experiment to provide the power to process in Real-Time!

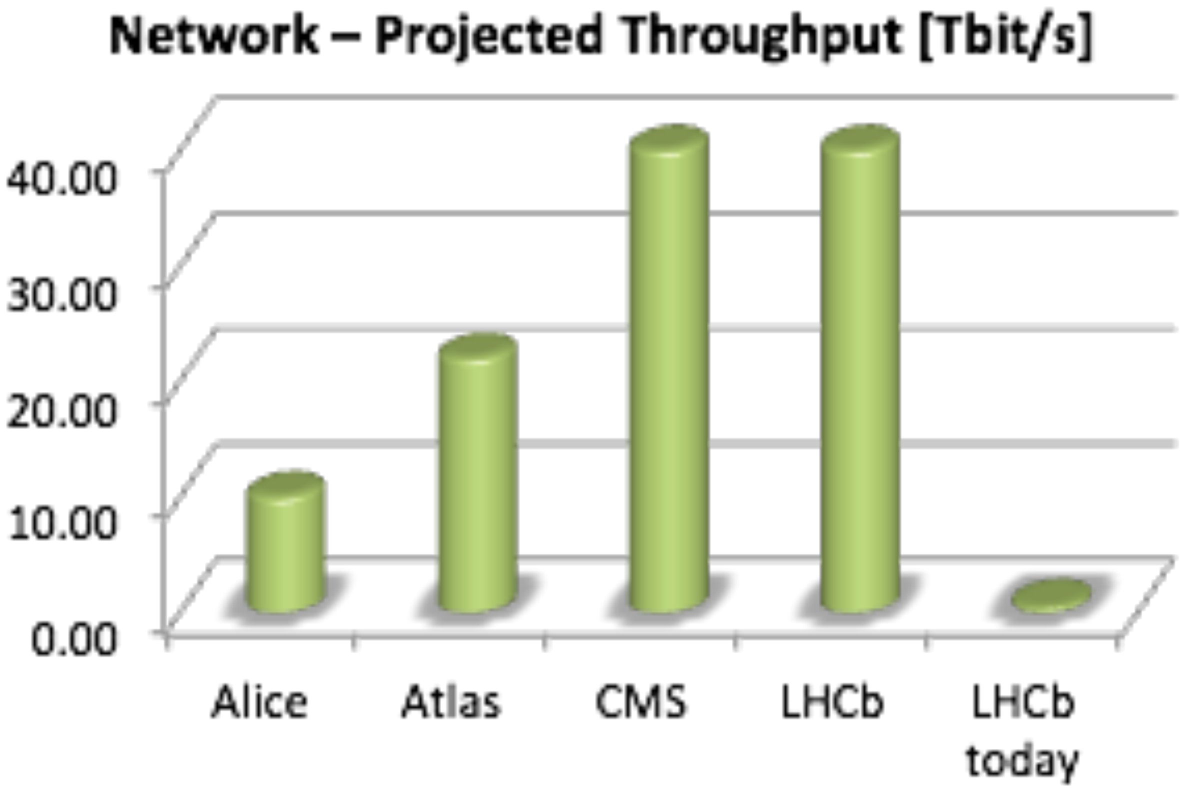
Network – Projected Throughput [Tbit/s]
40.00
30.00
20.00
10.00
0.00
Alice
Atlas
CMS
LHCb
LHCb today

RECAP

Trends in DAQ

- Rates are so high that the concept of events is broken - "streaming DAQ" on the horizon (triggerless)
- The real-time processing demands are very high – data compression and online analysis. Several strategies used
  - Write to buffer, re-analyze and classify / data reduction
  - On-the-fly data compression and reduction
  - On the fly calibrations, classification, …
- LHCb/ALICE envisioned to have large processing farm "closer" to the experiment to provide the power to process in Real-Time!
- LHCb already making heavy used of Machine Learning to classify events, anomaly detection, PID (DNN), Jet Tagging (2D BDT), …)
- Blurring of online/offline (LHCb speaks of online analysis, ALICE O2, FAIR, …)

Tackling the complexities

Job scheduling and heterogeneous resources

Framework design?

Facility efforts

Community efforts

IoT and Apps

BROOKHAVEN NATIONAL LABORATORY

Job scheduling and heterogeneous resources

- How do I handle the complexity / heterogeneity of resources?
  - We now have a mix of local COTS (Clusters), Grid, Clouds and possibly HPC

- The Job Broker way (ATLAS/PANDA, CMS/CRAB, …)
  - Submit "pilot job" that checks the slot capability
  - Contact a central server and pull a job in depending on what is available and matches
  - May combined with "*event service*" for simulations
    - Pros: relatively easy to design, pilot would not run a task if resource is faulty
    - Cons: can a global optimization really be done? Input pulled in only after the facts …

- The **A**rtificial **I**ntelligence way – A reasoner and planner
  - All resources are looked at simultaneously. A "plan" for a production campaign is done based on available resources (CPUs, storage, network bandwdith, connections between resources, QoS)
  - A partial plan is initiated – its reaction is feedback to the planner "learning"
  - Next iteration is reactive
    - Pros: Global optimization IS possible, industry standard (Amazon deliver, Ambulance scheduling, Home Depot deliveries, …), can add constraints such as "cost"
    - Cons: job starts the usual way and may not run if the input to the plan are not accurate

Example

- D. Makatun et al., Provenance-aware optimization of workload for distributed data production, CHEP 2016
- Network ow maximization approach with polynomial complexity (good)
- Looked at multiple experiments job pattern and network complexity pattern to get proper input to the model
  - Showed up to 30% improvements on simple networks and resources
  - Overall makespan improvement is 7% on a most complex Tier structure
  - AI always wins!

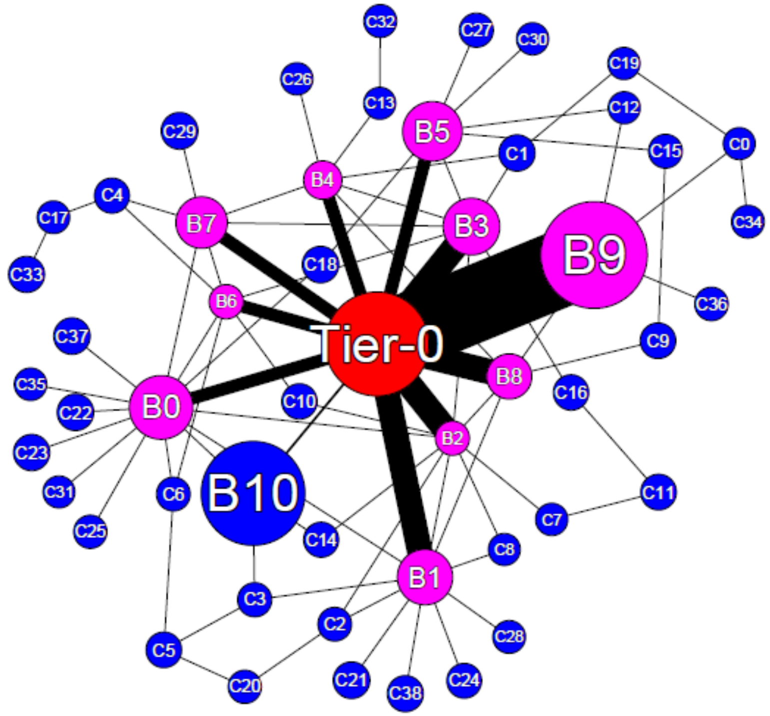

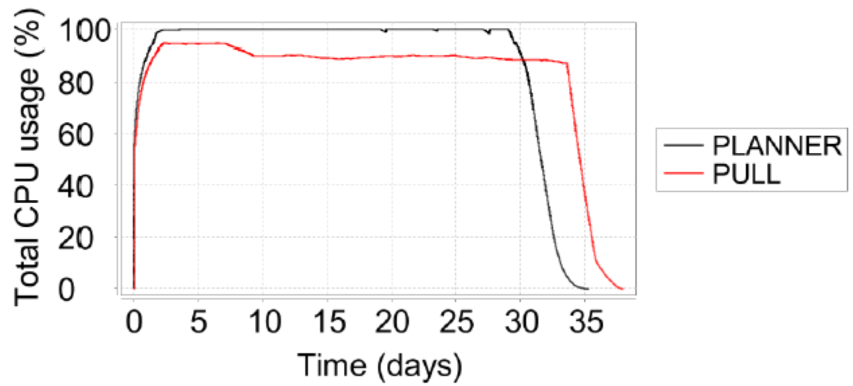

Framework (re)design?

- How to make sure frameworks are still valid a decade down the road?
  - Guess? Or address tasks one by one but remain within single-framework?

- ALICE/FAIR teaming for a new framework: ALFA
  - Has data-flow based model (Message Queues based multi-processing)
  - Each process assumes limited communication and reliance on other processes
  - ⇔ each process can be a "tasks" that run on ANY hardware or the hardware where it fits best, communication is MQ based – network is a true resource
    - Can be multithreaded, SIMDized, …etc.
    - Runs on different hardware (CPU, GPU, …, etc.)
    - Can be written in an any supported language (Bindings for 30+ languages)
  - Orchestration – Dynamic Deployment System

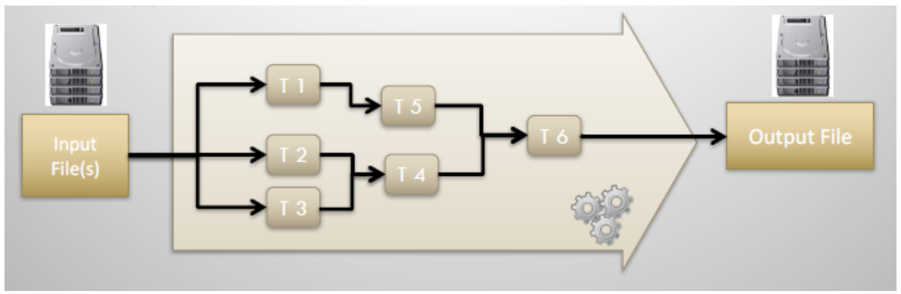

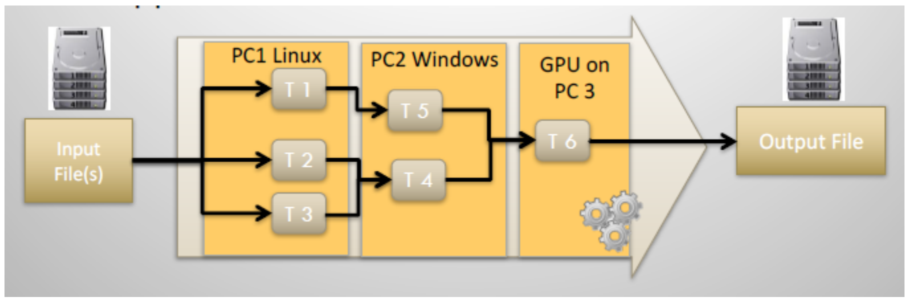

Facility efforts

National Energy Research Scientific Computing Center

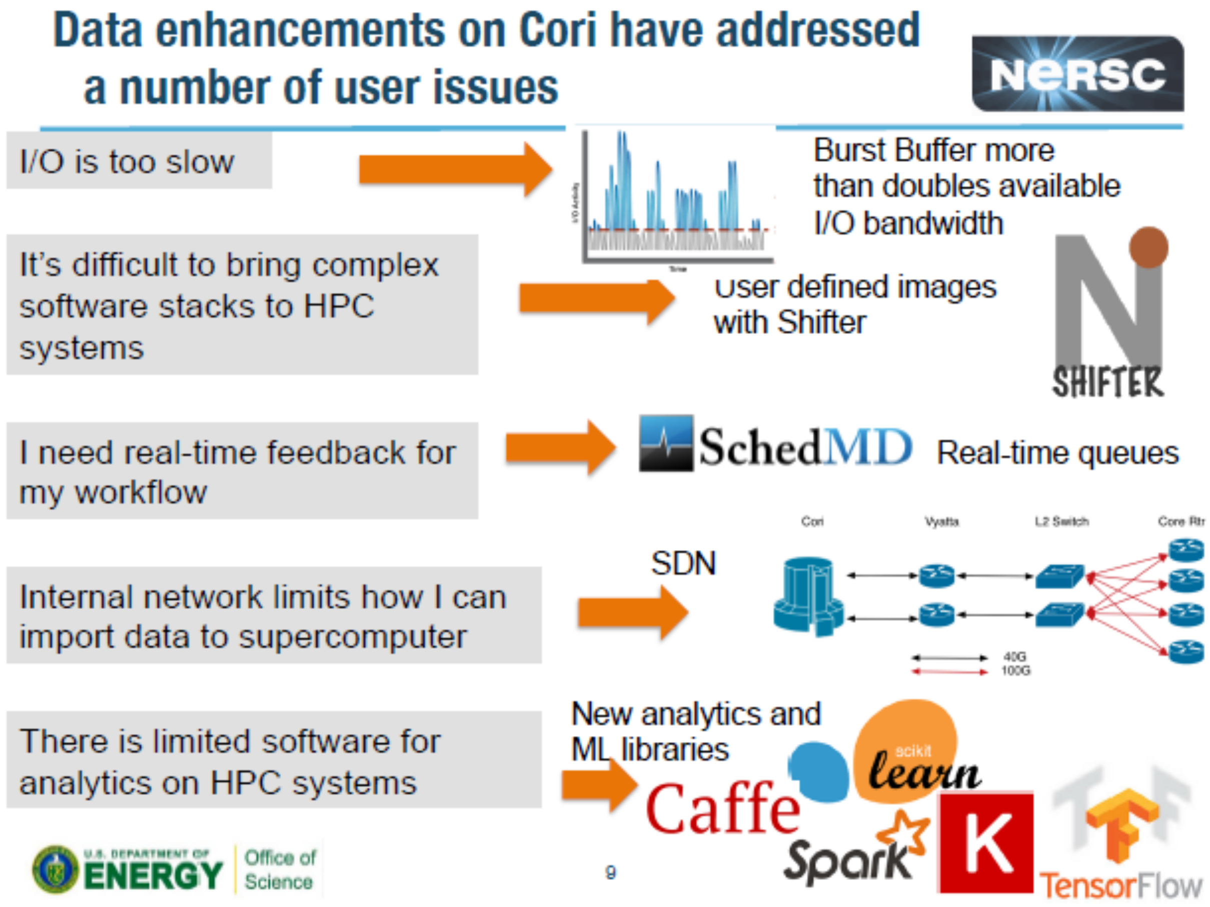

Katie Antypas (NERSC, Lawrence Berkeley National Laboratory)

Community efforts
HEP Software Foundation

HEP Software Foundation

Following the "Free Software Foundation" idea, can HEP (NP) do the same and …

- Share expertise
  - White papers, peer reviews, topical workshops
- Promote commonality
  - Raise awareness of existing software and solutions
  - Catalyze new common projects
  - Promote collaboration on new projects to make the most of limited resources
- Support common software
  - Aid developers & users in creating, discovering, using, and sustaining common software
  - Act as a framework for attracting support to S&C common projects
- Support careers
  - Support career development for software and computing specialists
  - Serve as a training resource
- Facilitate wider connections with other sciences & communities

Good set of goals with daunting sociological challenges

Component design a powerful concept

- At the core of the Internet Of Thing (IoT)
  - Web based tools?
  - Phone App based monitoring?
- All is already done in NP
  - D. Arkhipkin, Online Meta-Data Collection Framework
  - Integration with the Pre-existing Controls Infrastructure, CHEP 2016
  - MQTT at the core

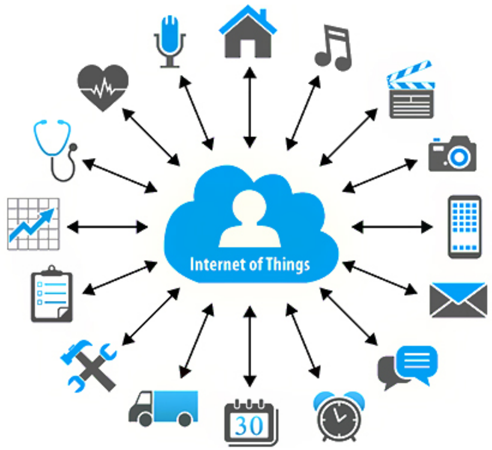

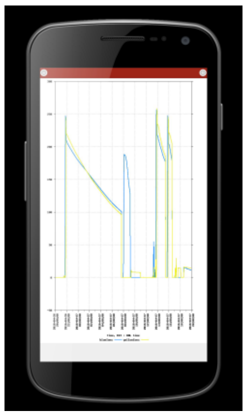

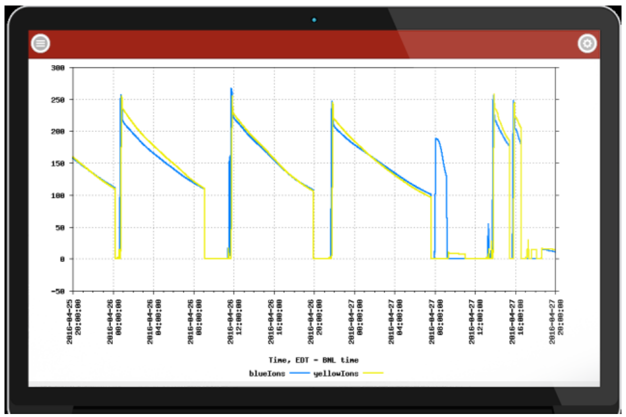

20.8 billion devices on the Internet of things by 2020 according to Gartner, 30 Billion estimate from ABI Research. Those includes embedded and wearable – all talk to each other.
Can we learn from this?

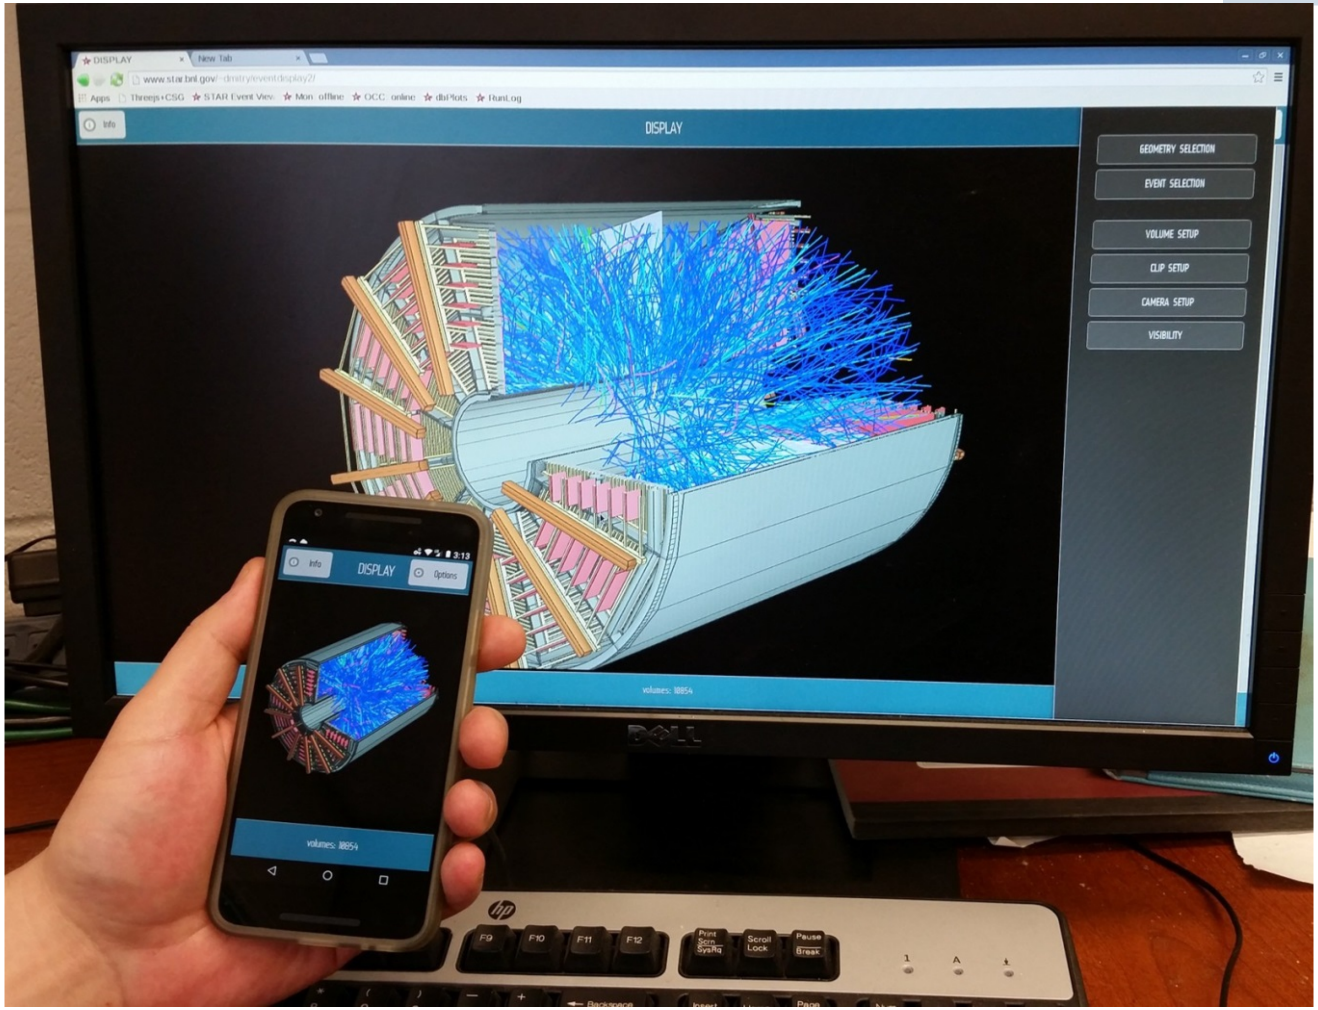

Real-Time event display Phone/Web App

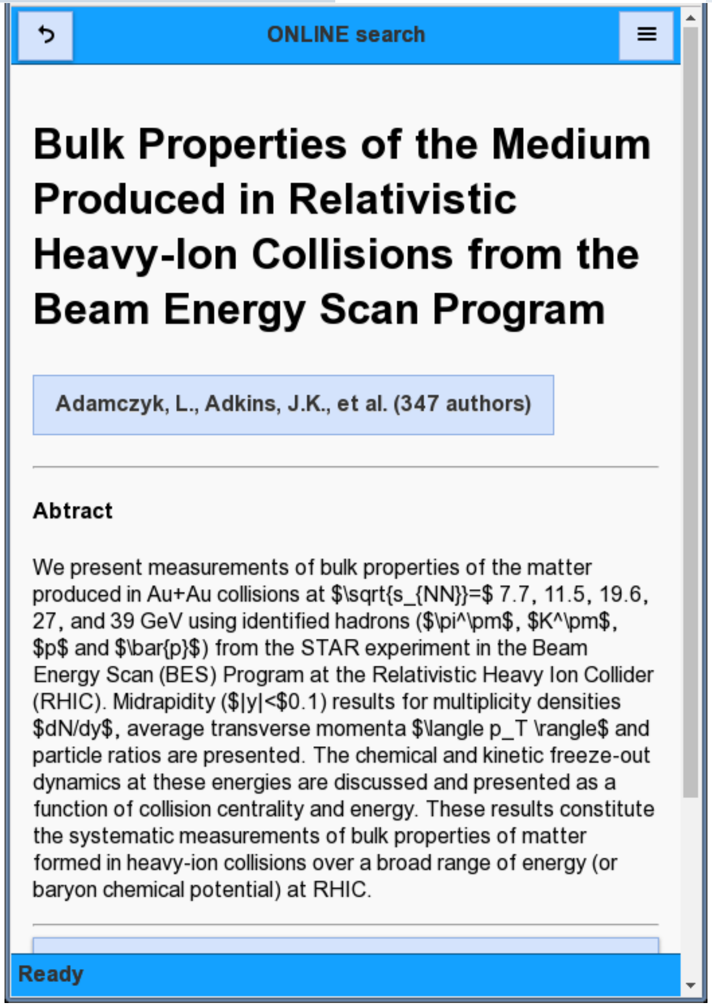

INSPIRE Paper Browser App (beta)

[a] Modernize environment [b] brings us closer to component based framework [c] slowly equipped with tools allowing morphing into IoT.

Use of Message Queuing and using networks as a resource is a key ingredient to distributed computing / IoT paradigm. It is happening NOW!

Beyond science – the human resource factor

Tomorrow's collaboration
Workforce attraction and retention
Ripples ahead (the “unknown”)

BROOKHAVEN NATIONAL LABORATORY

A need for stronger collaborations

- A very complex environment requires fostering all knowledge and talents with a positive attitude toward collaboration
  - *We should drop the thinking that Physicists can resolve it all*

- To tackle the most challenging "Big Data" problems, new teams will need
  - A cross-section of many fields is needed: Experimentalists & Theorists, Applied Mathematics, Computer Science
  - Out-reach to industry : Vendors (Intel, …), Google (analytics) and beyond –they are at forefront of Machine Learning
  - To learn from other sciences – Plasma, Fusion, … and *integrate facilities as part of "the team" (Ex: KISTI and STAR)*

- A daunting sociological and funding challenge
  - A "buy in" from above (agencies) with priority given to collaborative work
  - *Acknowledgement and recognitions*
  - Support travel, communication and exchanges: cross disciplines workshops, "rencontres", bootcamps, big data shops, … we will see more of those
  - *Support (both ways) common projects – Geant4.*
    What about FTE paid by experiments? Consulting $?

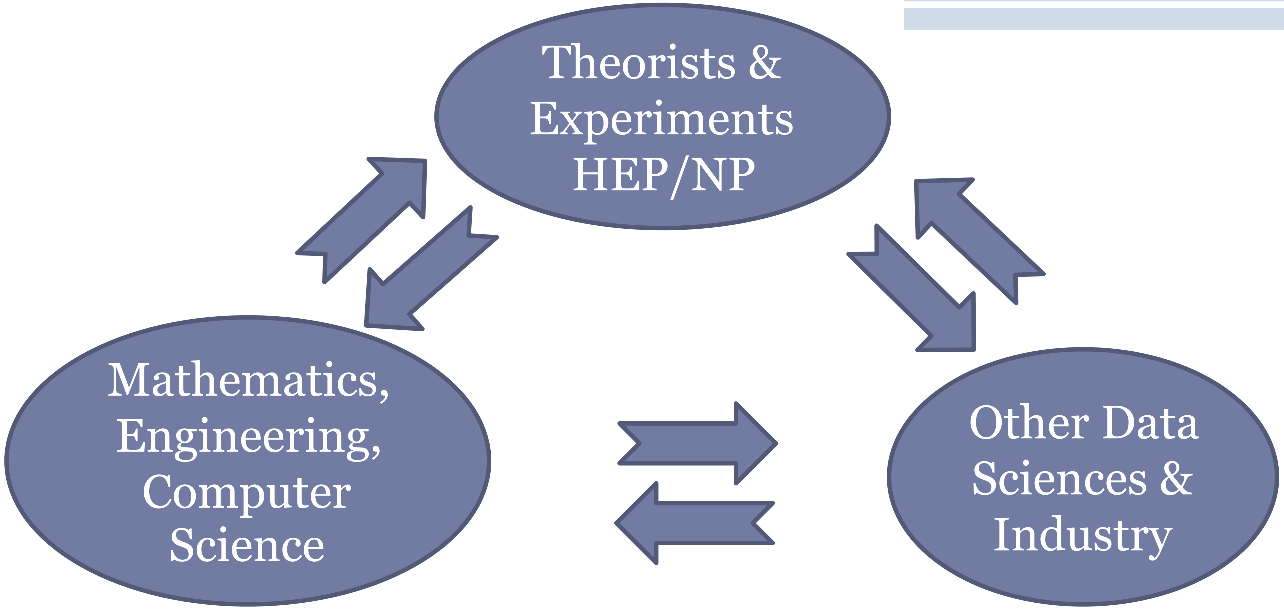

SciDAC
Scientific Discovery through Advanced Computing

Software institutes were popular in the mid-2000 . Institutes encouraged collaborations across science domains but agency did not support activities beyond R&D (hardening & integration a challenge)

Open Science Grid

Jointly funded by the Department of Energy and the National Science Foundation, includes many science domains, computer scientists and resource providers.

CMSE

The Department of Computational Mathematics, Science and Engineering (CMSE) at Michigan State University as created such “power team” across section of scientists.

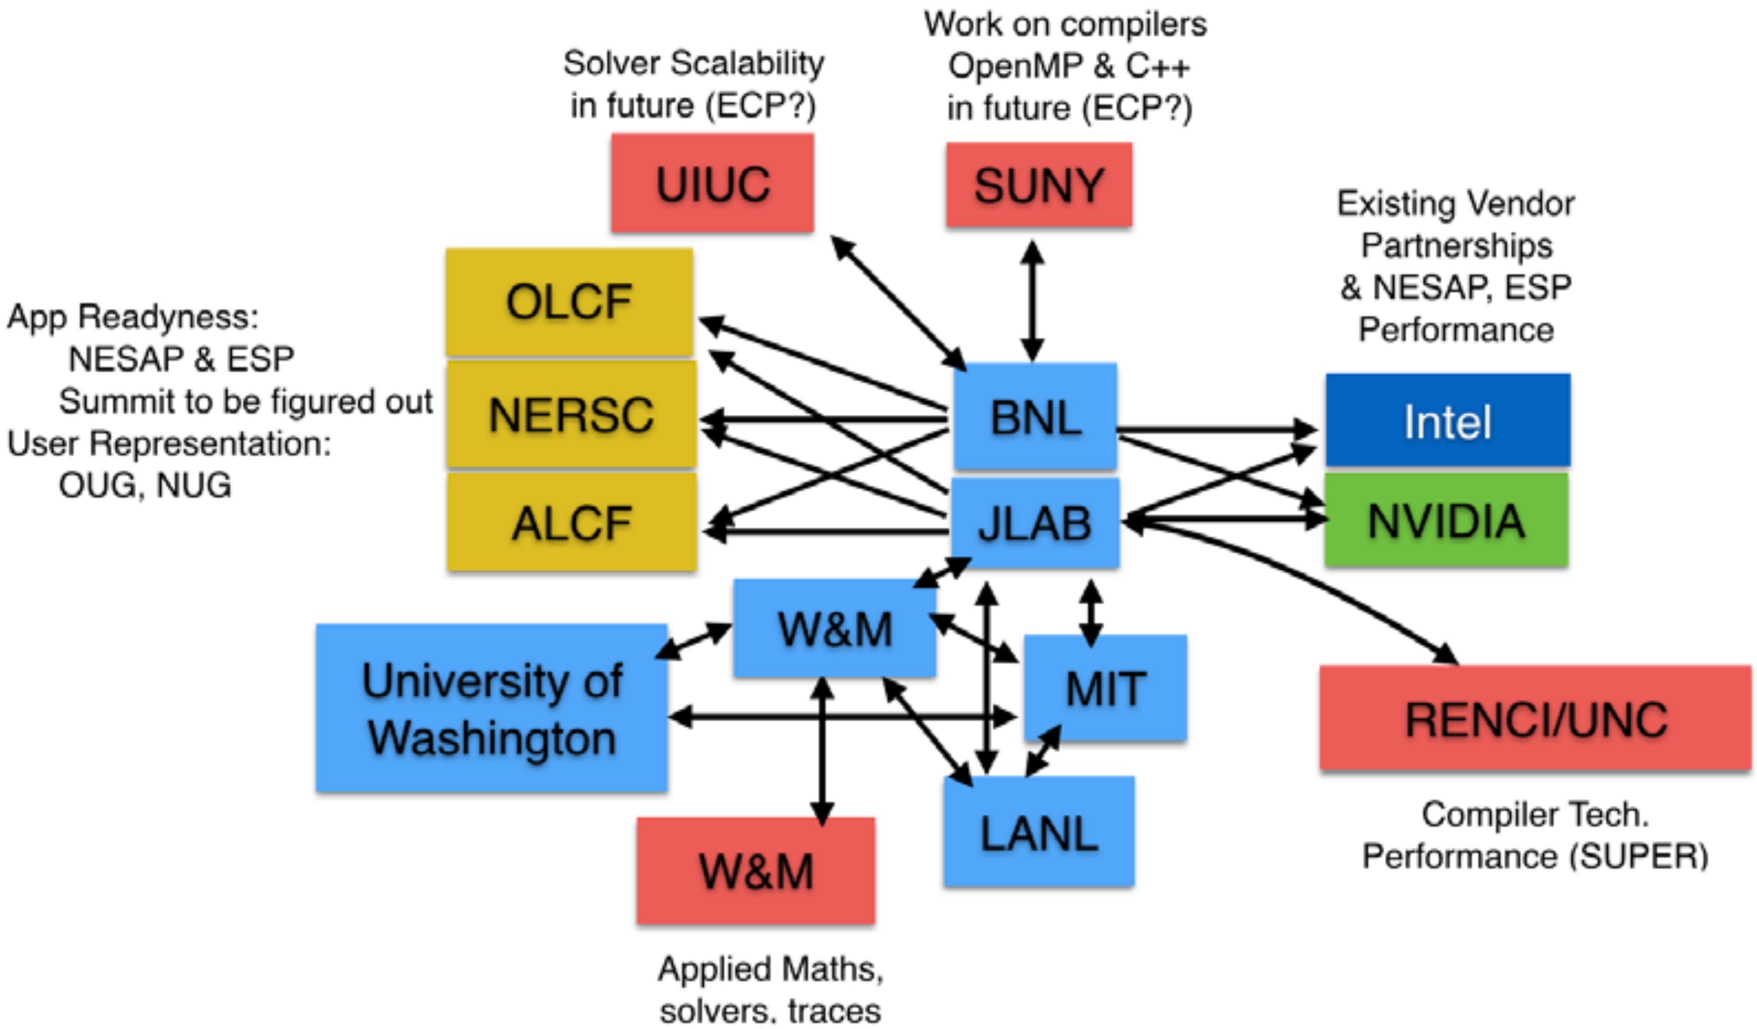

Collaborations are hard (especially within a competitive field)

Wisdom will be required at agency and individual's levels

If we do all of this, we would have a bright future ahead.
But if we do not succeed in keeping the best talent from multiple-disciplines …

Workforce attraction and retention

- The harsh reality for Graduate students and PhDs in the US
  - Only 1 out of 6 physicists continues to a PhD degree (American Institutes of Physics), only 1/3rd of PhD will have a job in academia
  - Majority of the permanent jobs is outside of academic research

- We have a responsibility to train, prepare and secure our youth for a career within our outside academia
  - IF this is done, the field *will continue to attract and retain* the best individuals with many needed skills
    - NB: In my opinion, we are currently losing our best students, realizing they have "options" outside academia – we must create an attractive and rewarding environment and stop the "brain leak"

- QOTD: How can we attract & prepare young scientists for their most likely career and align skills to all eventualities?

Introductory remarks

- Students are the main code developers, data-analysis workforce, … In practice:
  - Long term scientists (faculty and staff scientists) are mostly disconnected from coding and its complexity … and complain that students spend too much time on coding or service
    - Anecdote: 50% of their time is spent on writing code instead of thinking about physics (e.g. graduate students at LHC).
  - Short term scientists (under-graduate, graduate, post-doc) are trained, some are versed in software development and have opportunity to self-develop valuable skills (mostly "learn by doing" with no formal training)

Dilemma: without training and hand-on work, students would not acquire skills to attain (good) opportunities in industry. They would not acquire the skills to understand tomorrow's computing techniques & landscape either.

#1 know your skills …
and develop what you do not have

- What skills do physicists have?
  - Have handled computer modeling and simulation, data analytics
  - Versed in big data workflows (practical or implicit knowledge)
  - Fluent with math, strong problem solvers
  - Experience with coding
- What skills are physicists (often) missing?
  - Ability to function on multi-disciplinary teams and harvest multiple skills
  - Ability to recognize value of diverse relationships (customer, management, …)
  - Ability to design a system, knowledge in software engineering and architecture based on "requirements" (tend to code first to "get something going"), think later), interface specifications, …
  - *Leadership skills*
  - Familiarity with basic business concepts (i.e. cost-benefit analysis, funding sources, leadership, project management)
  - *Awareness of career paths outside of academia*

Senior s should make sure opportunities exist, encourage it and student should not be afraid to "jump in"

#2 change the design & accommodate for all needs

- Get back to user-driven strong design
  Monolithic framework have many "cons"
  - Single language, steep learning curve of "the whole" rather than a framework philosophy, rarely have industry standard approach

- Move to "micro-service" instead – Redesign from scratch and bring back the "Unix" philosophy way
  - A chain of event is composed of "logical tasks", each responsible for a unit of work – exchange format in between (do one thing and do it well)
  - Focus on data type and formats (then algorithms can be swapped as well as programing languages – logical unis work together)
  - Huge plus: Potentially can scale with cores, adaptable to changes and allows mix of architecture

This is where an idea like ALFA can help

#3 Bring State Of the Art techniques in

- Spice it up with Industry standards …
  - Phone-App design (why not?), protocol design through service tasks – expose students to business principles and methods
  - Use industry format for data streaming – JSON, XML, Google protocol buffer , AMQP, CEPHfs, HDFs, Python and R
  - Introduce industry standard tools: Introduce packages for Data-Science and ML that are truly industry standards (Anaconda, scikit-learn) or business analytics software, use common tools and packages for Machine Learning
    - Don't role your own – nobody cares about TMVA in industry
  - Think container

- … but be pragmatic
  - ROOT, csv, hdf,… are not used widely n industry but
    - HDF (for example) has been around for 2 decades
    - Everyone use ROOT – at least, don't build layers on top!
  - Google protocol buffer would lead to knowledge outside the field but would it serve the community's need?
    - For long term projects such as the EIC, may be a good solution

Ripples ahead

- Quantum Computing is a big unknown
  - What would it mean for our field? Microsoft, Google, IBM have all indicated more development within a year

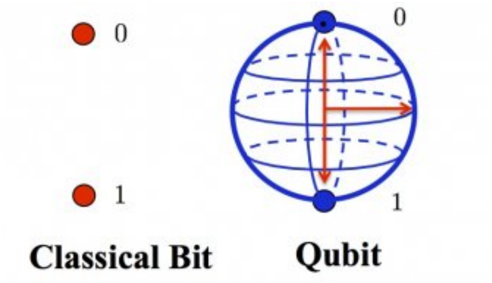

- Edge computing? Another term for IoT?
- Machine Learning
  - "Machine Learning gives computers the ability to learn without being explicitly programmed" - Arthur Samuel in 1959 (pioneer in AI and gaming)
  - ML is used in your every day life – from Google and FaceBook "tag someone" face recognition, by our phones to interpret voice commands, by internet retailers to make recommendations, and by banks to spot unusual activity on a credit or debit card.
  - Already used in our field for sophisticated pattern recognition and find correlations in data
  - Advances in industry far surpasses anything done in HEP/NP so far

We should read the signs of changes in today's technology trends ...

The current applications only scratch the surface and understanding of how powerful those techniques really are. Robots at home are scarier than AI-infused systems used for policing and healthcare. (Royal Society study)

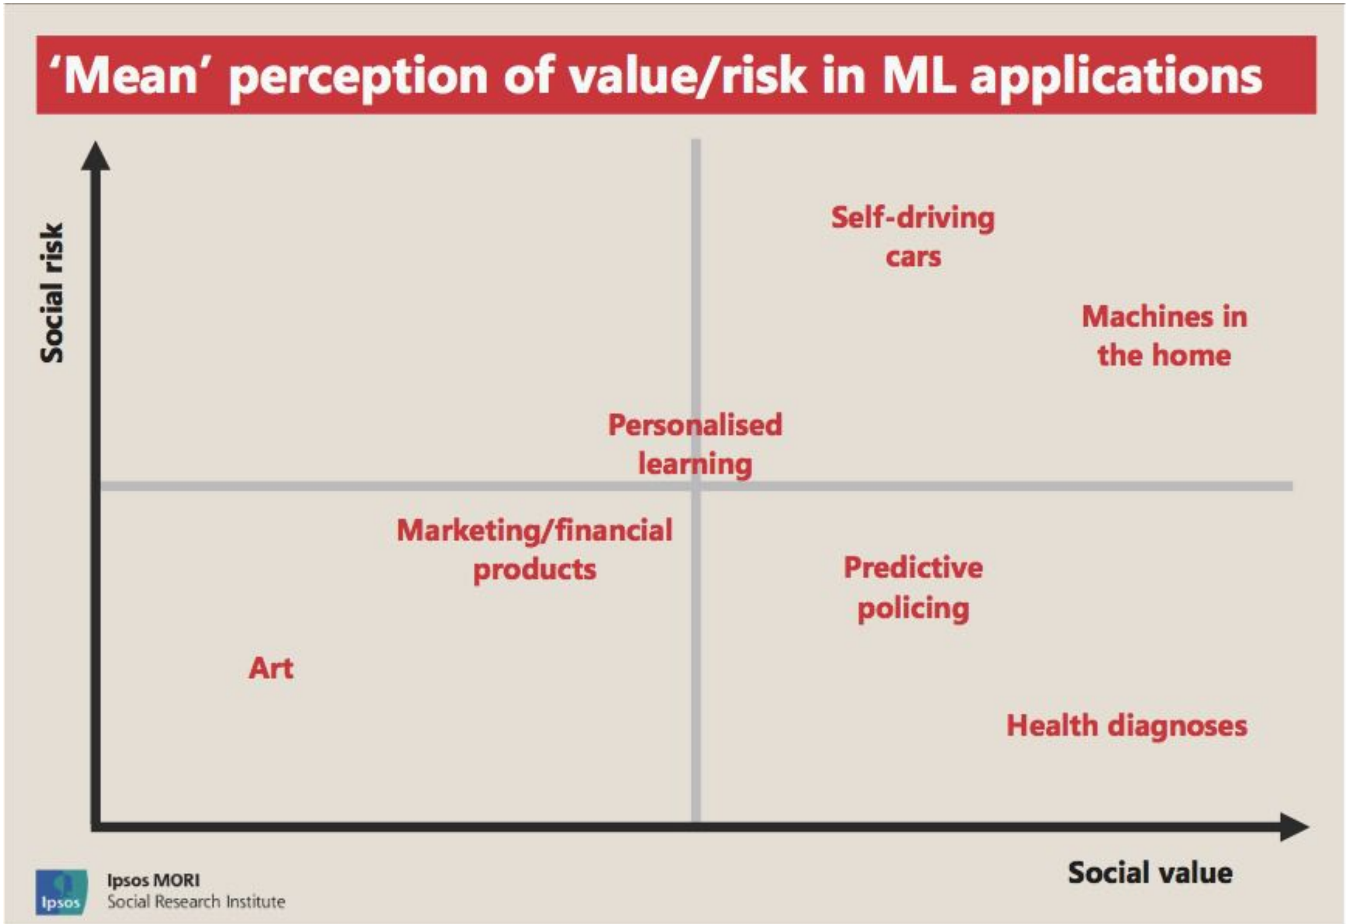

One thing for sure – **ML is appearing in our field and offers a goldmine opportunity to learn industry re-usable skills**. To attract talents, we also need to adapt our field to the world around us and the wave of time.

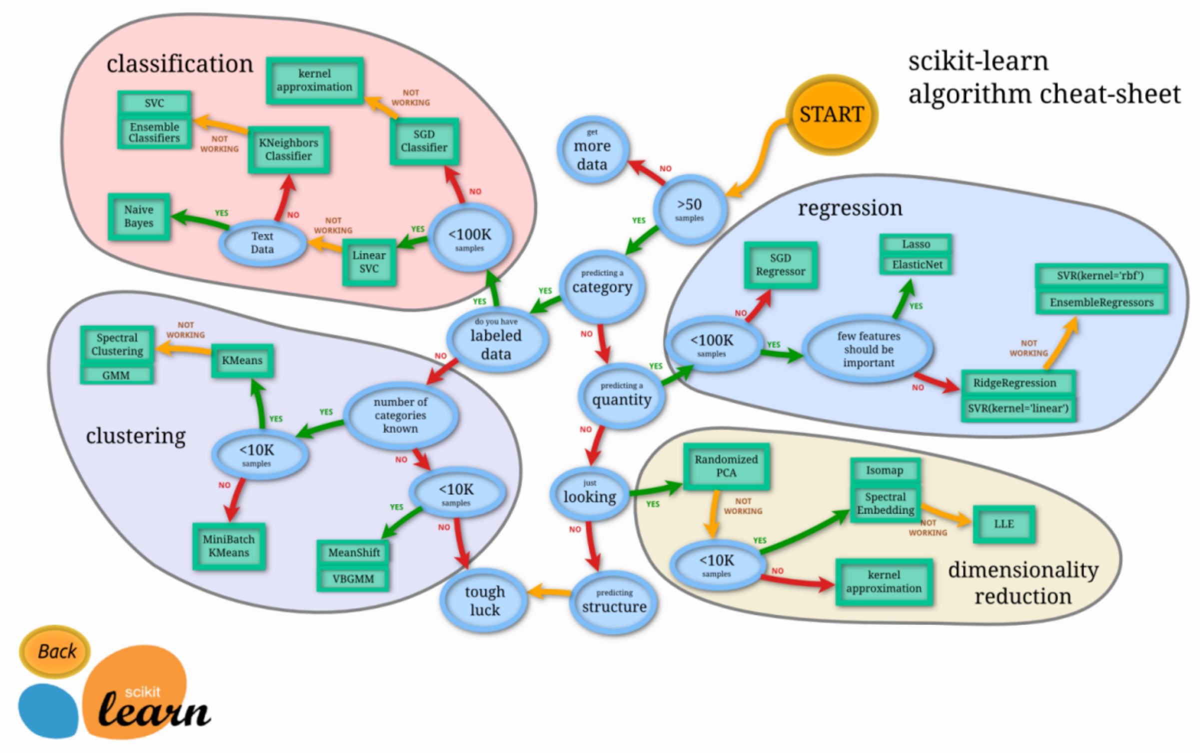

Roadmap: [How to Learn Machine Learning in 6 months]

By Zacharia Miller (ex STAR PhD ☺)

Choosing an algorithm

Concluding remarks

BROOKHAVEN NATIONAL LABORATORY

Conclusions

- The 2000 computing architecture and designed worked and served well our communities
  - Near 2 decades of amazing discoveries were possible and enabled by it
- Next level of science in Nuclear Physics requires new approaches
  - The ecosystem has changed, driven not only by increasing Physics needs but also by technology evolution, the maturing of AI approach ever-present in our lives, the race toward HPC
  - DAQ rates will reach unprecedented rate levels – IO will be challenging
- Workflow adaptations to new realities may require talents from many fields
  - Re-think framework along "component based" or "task based", the Unix way
  - Push boundaries with AI and ML, IoT
  - New collaborative teams are being shaped and seeing life
  - Cross section of Physics, Math, Computer science, Engineering, Industry, …
- As the world is changing and so is our field, *focus on people* and introduce (and endorse) potentials and opportunities for (industry) re-usable skills
  - Winning strategy to attract and retain the best from many fields of expertise …