KIT
Karlsruhe Institute of Technology

HELMHOLTZ
RESEARCH FOR GRAND CHALLENGES

R&D for High Energy and Astroparticle Physics

Data Science and Quantum Computing Workshop, Vancouver

Manuel Giffels, 27.06.2018

Institute for Particle Physics (ITP) & Steinbuch Centre for Computing (SCC)

Computing Challenges in High Energy Physics

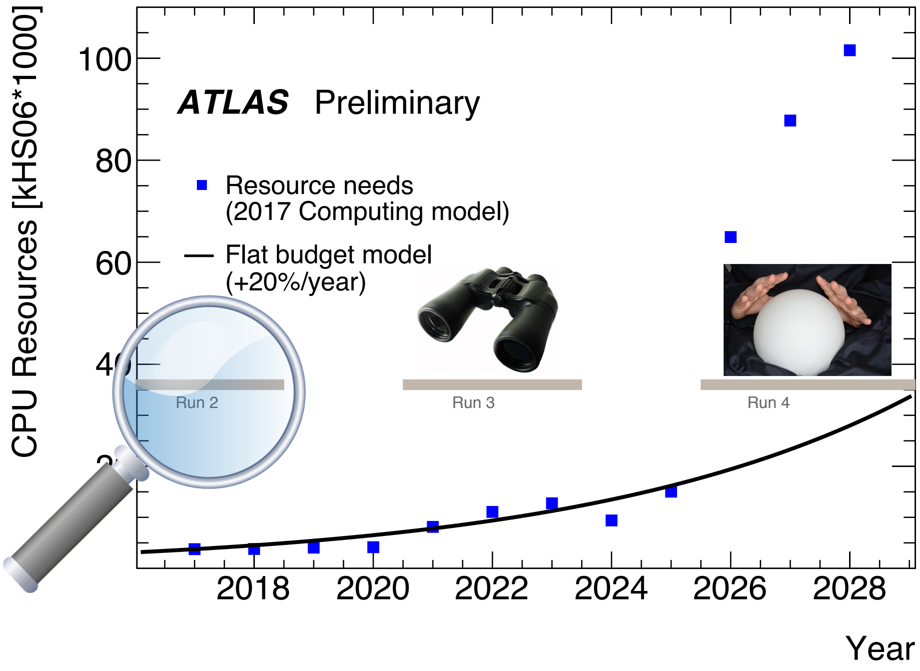

Alves et al., A Roadmap for HEP Software and Computing R&D for the 2020s, *HSF-CWP-2017-001* (2017)

- ATLAS CPU resource estimates (similar for CMS)
- Assuming flat budget and 20% technology advance per year
- CPU shortfall between needs and technology gains is about factor 4 in 2027
- Situation for disk storage is even worse (Shortfall is about factor 7 in 2027)

⇒ Critical conditions for HL-LHC

Physics will be limited by computing!

Computing Challenges in High Energy Physics

Karlsruhe Institute of Technology

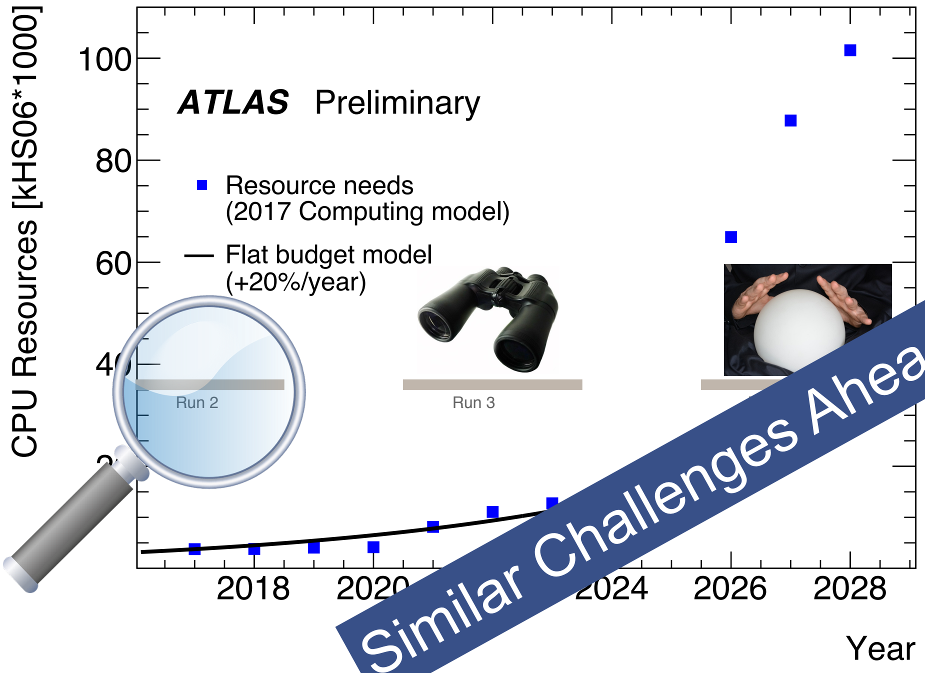

Alves et al., A Roadmap for HEP Software and Computing R&D for the 2020s, *HSF-CWP-2017-001* (2017)

Similar Challenges Ahead in Astroparticle Physics

- ATLAS CPU resou… stimates (similar for CM…
- Assumi… budget and 20% tec… advance per year
- …U shortfall between needs and technology gains is about factor 4 in 2027
- Situation for disk storage is even worse (Shortfall is about factor 7 in 2027)

⇒ Critical conditions for HL-LHC

Physics will be limited by computing!

Rethink Software & Computing Models

Rethink Software & Computing Models

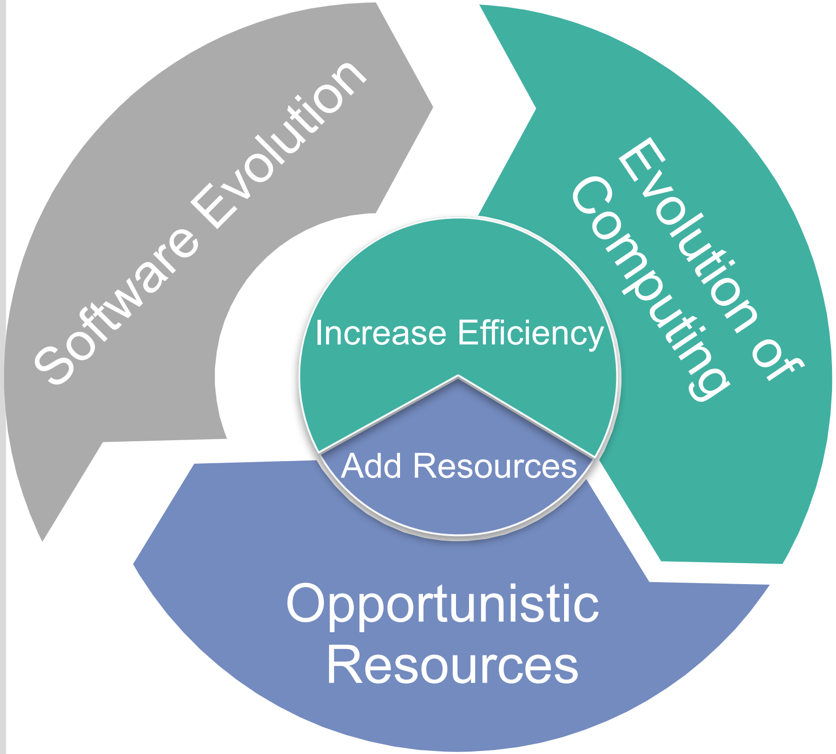

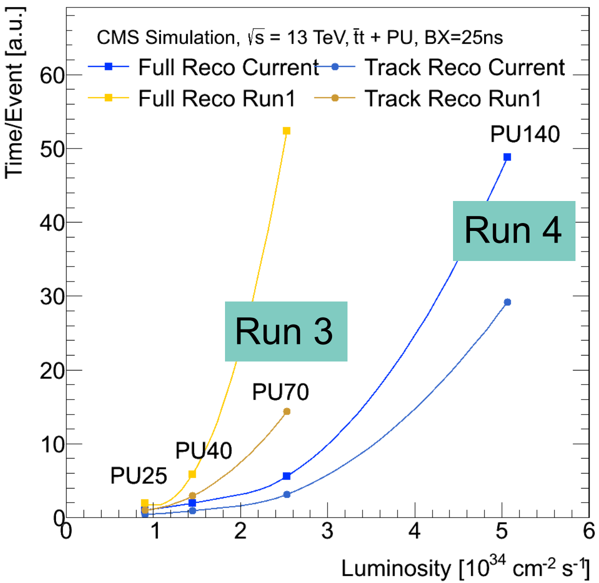

- Exploit modern technologies
- Improve algorithms / simulation
- Utilize Machine Learning

Machine Learning in HEP

- Omnipresent in offline analyses for signal/ background classification (20 years+)
- Increasingly applied in offline and online reconstruction (real-time application)
  - Flavour tagging (Classification of Jets)
  - Particle identification
  - Online event selection (LHCb trigger)
- Already utilized in core reconstruction algorithms at DUNE (Neutrino Physics)
- Full reconstruction of B mesons at Belle using NeuroBayes was developed at KIT

Nuclear Instruments and Methods in Physics Research A 654 (2011) 432–440

Contents lists available at ScienceDirect

Nuclear Instruments and Methods in Physics Research A

journal homepage: www.elsevier.com/locate/nima

A hierarchical NeuroBayes-based algorithm for full reconstruction of B mesons at B factories

M. Feindt, F. Keller, M. Kreps [1], T. Kuhr, S. Neubauer *, D. Zander, A. Zupanc

Institut für Experimentelle Kernphysik, Karlsruher Institut für Technologie, Campus Süd, Postfach 69 80, 76128 Karlsruhe, Germany

ARTICLE INFO

Article history:
Received 7 April 2011
Received in revised form
3 June 2011
Accepted 3 June 2011
Available online 17 June 2011

Keywords:
Full reconstruction
B-factory
Neural networks
Probability

ABSTRACT

We describe a new B-meson full reconstruction algorithm designed for the Belle experiment at the B-factory KEKB, an asymmetric e^+e^- collider that collected a data sample of 771.6×10^6 $B\bar{B}$ pairs during its running time. To maximize the number of reconstructed B decay channels, it utilizes a hierarchical reconstruction procedure and probabilistic calculus instead of classical selection cuts. The multivariate analysis package NeuroBayes was used extensively to hold the balance between highest possible efficiency, robustness and acceptable consumption of CPU time.

In total, 1104 exclusive decay channels were reconstructed, employing 71 neural networks altogether. Overall, we correctly reconstruct one $B^\pm$ or B^0 candidate in 0.28% or 0.18% of the $B\bar{B}$ events, respectively. Compared to the cut-based classical reconstruction algorithm used at the Belle experiment, this is an improvement in efficiency by roughly a factor of 2, depending on the analysis considered.

The new framework also features the ability to choose the desired purity or efficiency of the fully reconstructed sample freely. If the same purity as for the classical full reconstruction code is desired ($\sim$25%), the efficiency is still larger by nearly a factor of 2. If, on the other hand, the efficiency is chosen at a similar level as the classical full reconstruction, the purity rises from $\sim$25% to nearly 90%.

1. Full B meson reconstruction at B factories

1.1. The experimental setup

2. For the B^+B^- or $B^0\overline{B^0}$ pairs produced in this two-body decay, the four-momenta are related by

$$p(B_1)+p(B_2)=p(e^+)+p(e^-). \quad (1)$$

ML experience available at KIT

Software Evolution

- Detector simulation utilizes most CPU cycles
  ⇒ fast simulation (parameterized simulation)
- Increased demand for CPU cycles in reconstruction tasks at HL-LHC
  ⇒ Need an (r)evolution of algorithms
- Exploit technologies of modern CPUs
  - SIMD ⇒ Data Oriented Programming
  - Muli-core / Many-core / GPU support
- Machine Learning / Deep Learning
  - Under evaluation in fast simulation
  - Further application at reconstruction level (pattern recognition / object classification)

KIT engaged in tracking algorithms and ML

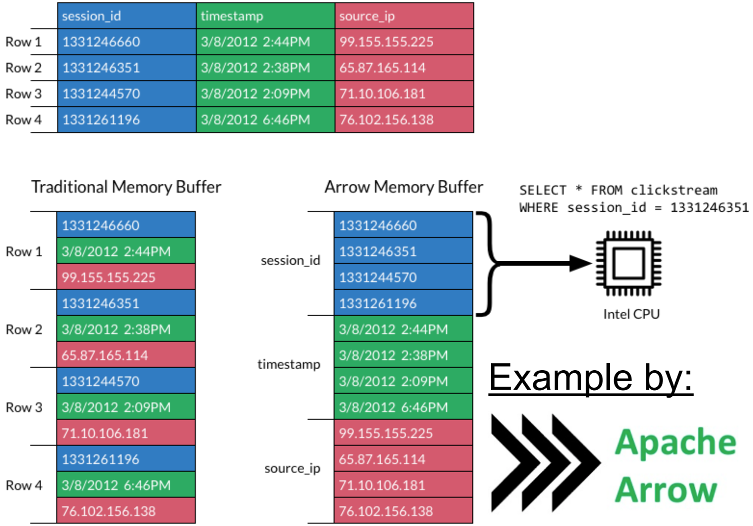

Example by:

Apache Arrow

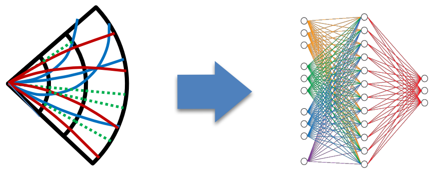

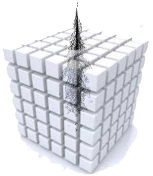

GeantV

Rethink Software & Computing Models

Rethink Software & Computing Models

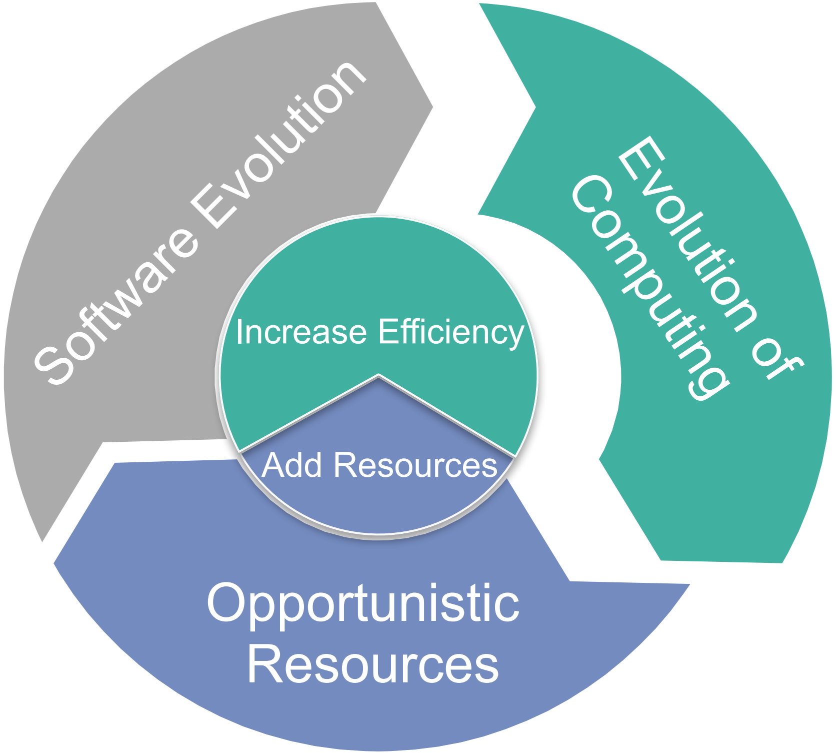

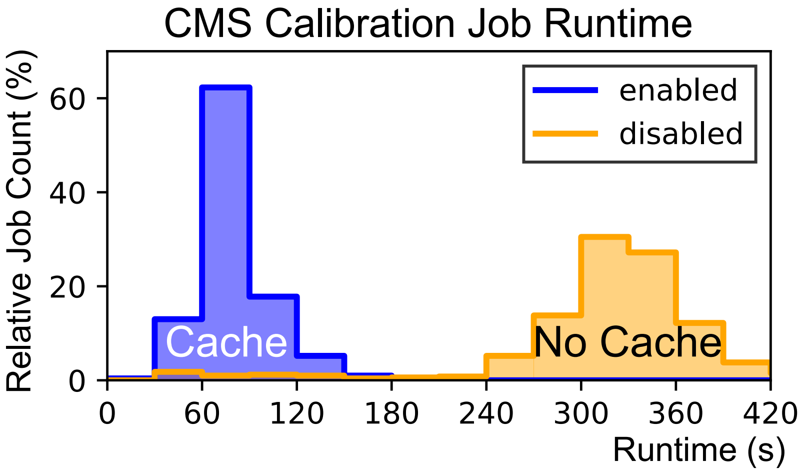

Utilize modern technologies:

- Caching on SSDs, etc.
- NVRAM
- GPUs, etc.

Computing Model Evolution

- WLCG today - more than 150 sites
  (not really economic compared to industry)
- Storage services are most expensive
  (Procurement and operation!)
- WLCG plans storage consolidation
  Reduce data centers: $\mathcal{O}(100) \Rightarrow \mathcal{O}(10)$
- More heterogenous compute resources
  ML fat nodes, opportunistic resources, etc.
- Remote data processing capability becomes more important
- Fast network links and data caches on volatile storage are mandatory

Caching technology is one key component!

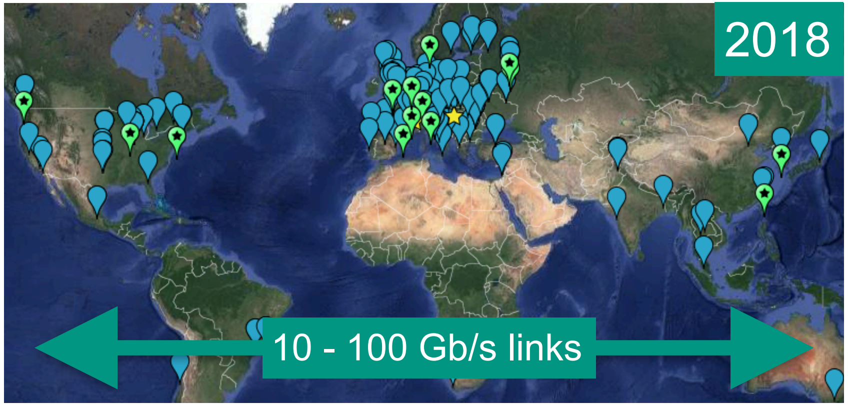

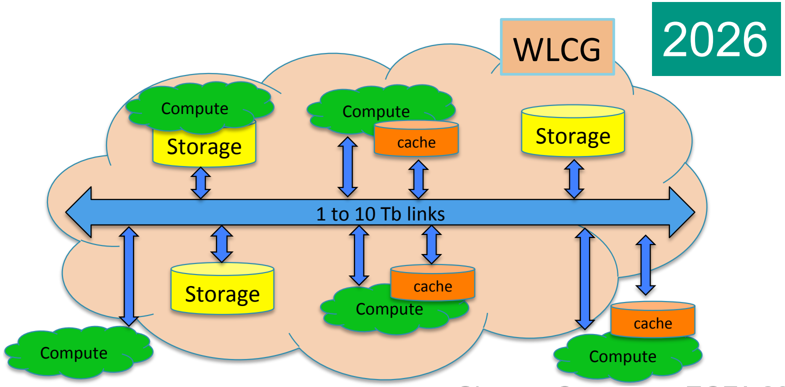

Simone Campana, ECFA 2016

High Performance through Coordinated Caching

- Prototype developed in interdisciplinary doctoral thesis (Physics and Computer Science)
- Established new approach for user data analysis via distributed, coordinated caching on local SSDs
- Introduced data locality to HTCondor batch system
- Performance gain of factor 3-4 on typical recurrent end user data analysis payloads (prototype system)
- Third party funding for a production system was recently granted (hardware)

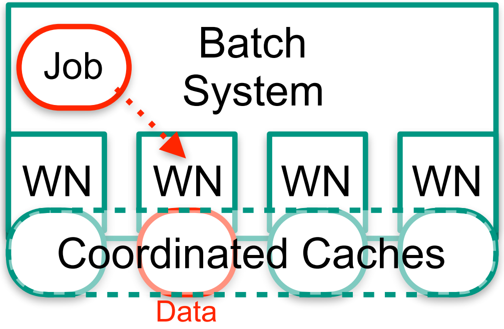

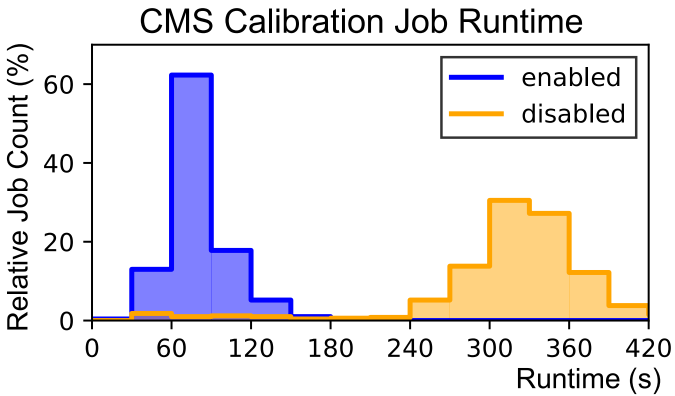

[JOP608MF, JOP664MF, IARIA16MF, JOP762MF, JOP898MF, KITPhD16]

GridKa High-Throughput Analysis Extension

- Knowledge transfer from R&D to the GridKa production system
- Include high throughput data analysis (HTDA) nodes into the GridKa batch farm
- Nodes can serve traditional WLCG jobs as well as HTDA user jobs profiting from caches
- Increase performance of typical recurrent end user data analysis payloads at GridKa
- Caching solution also well suited for recurrent ML trainings exploring different algorithms
- Evaluate performance impact of using pre-placement strategies (Machine Learning?)

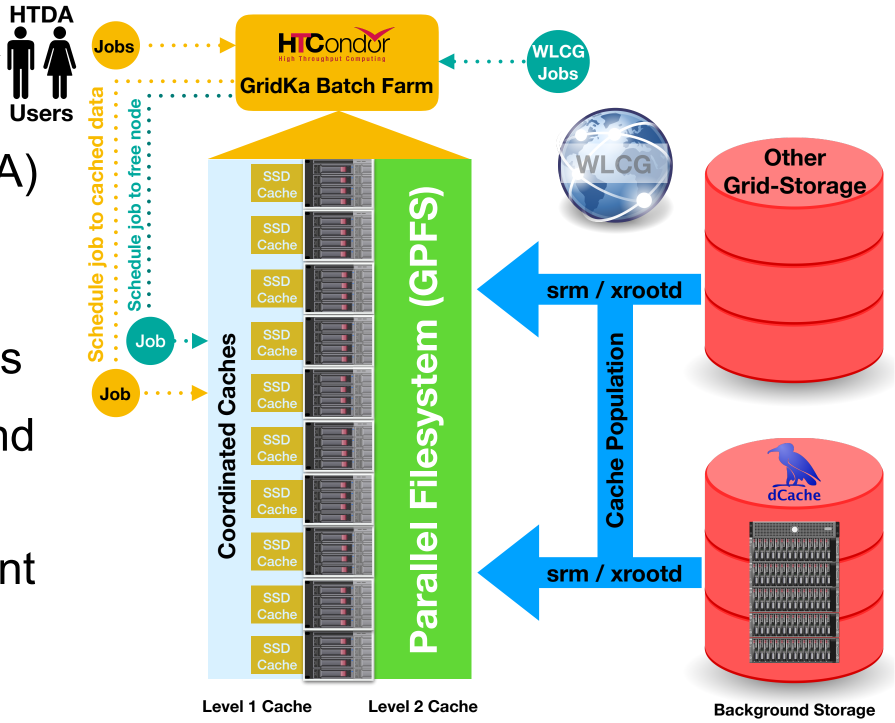

Rethink Software & Computing Models

Rethink Software & Computing Models

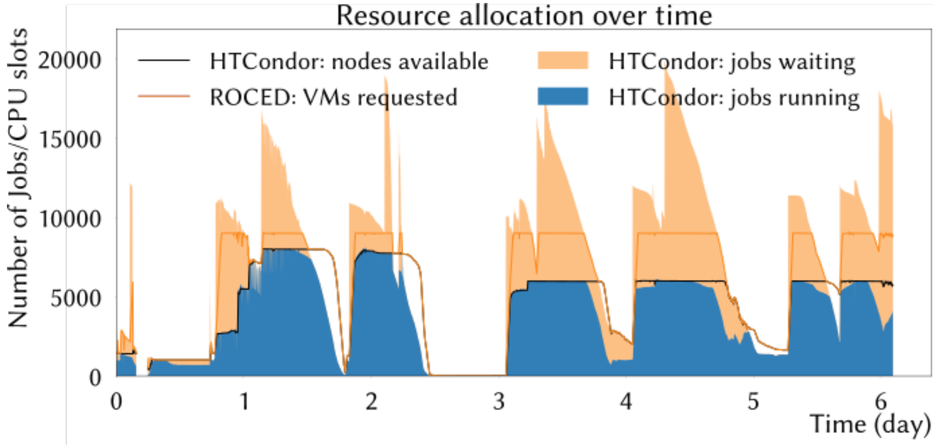

- HPC centers
- Volunteer computing
- Public and private clouds

Success Story - Opportunistic “Tier 1“ for a Day

- Dynamically shared HPC Centre at Freiburg (three diverse communities)
- Virtualization is key component to:
  - Allow dynamic resource partitioning
  - Meet OS & software requirements
- ROCED cloud scheduler developed at KIT
  - On-demand resource provisioning
  - Transparent resource integration
- Suitable for CPU-intense workflows

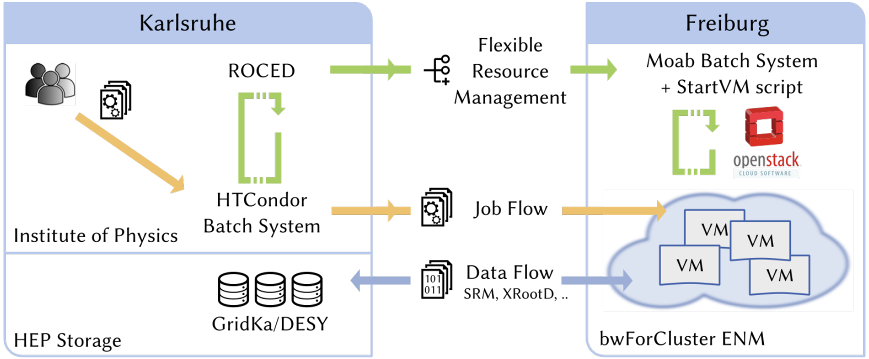

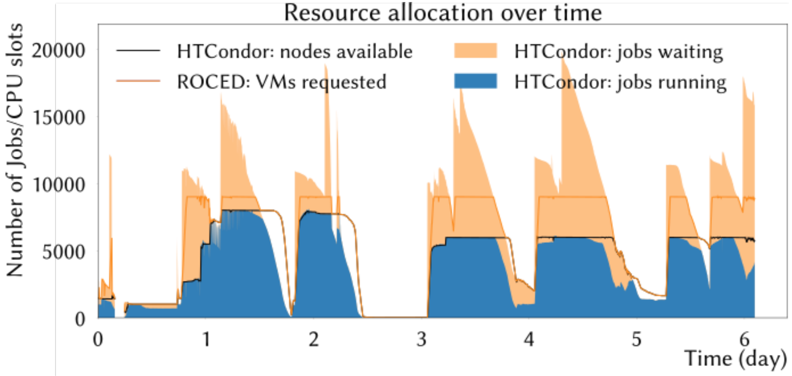

[ACAT17MS, JOP898TH, JOP762TH, JOP664TH]

Dynamic Compute Expansion of GridKa Tier 1

- Transparent on-demand integration of opportunistic resources using ROCED
  - Helix Nebula Science Cloud (based on traditional virtualization)
  - KIT HPC Center (FORHLR II) (based on container technology)
- Automated detection and redirection of suitable CPU-intense workflows
- Evaluate ML for scheduling optimizations [JSSPP18MS]

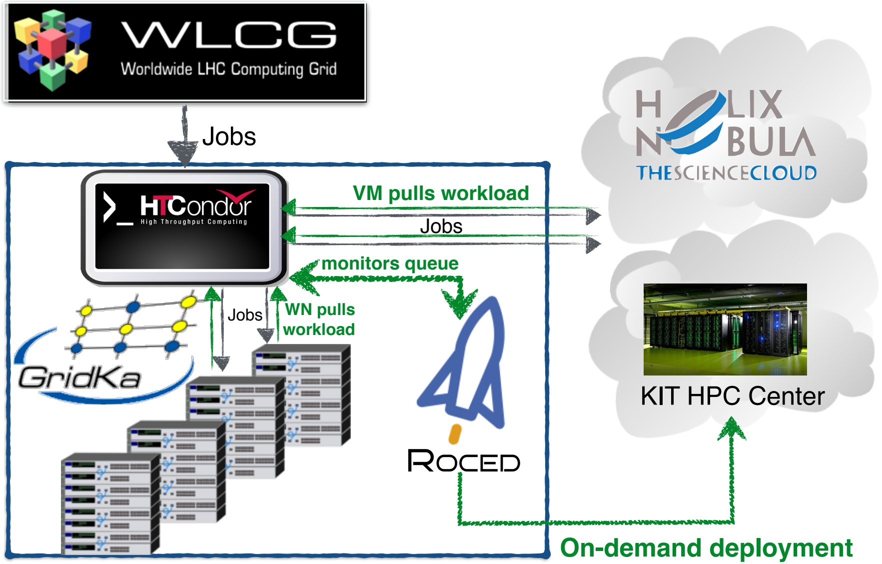

Caching Concepts on Opportunistic Sites

- Opportunistic Resources usually well suited for CPU-intense workflows
- Many opportunistic sites offering fast cloud storage or distributed storage
- Benefit from caching R&D and bring recurrent I/O-intense workflows to the cloud
- Transparent data access also a hot topic in Helix Nebula Science Cloud

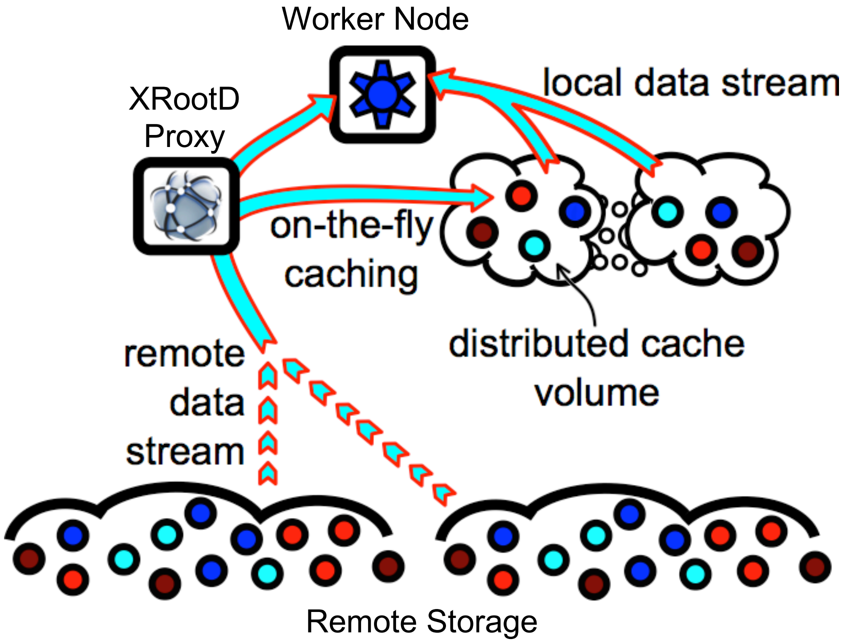

Collaboration in developing a xrootd based caching proxy between KIT and GSI

[ACAT17CH]

Outlook - Cluster of Excellence Proposal

- Algorithm Engineering for the Scalability Challenge (AESC)
- Ongoing proposal in the national Excellence Strategy
- Interdisciplinary research program at KIT
  (Computer science and application domains)
- Make algorithms ready for the big data and many-core era

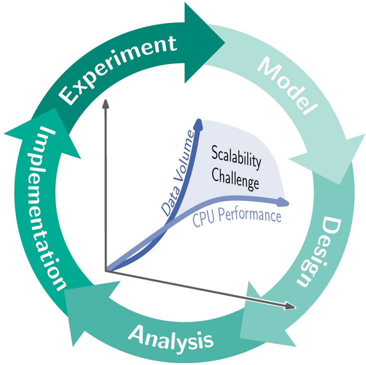

Scheduling Algorithms

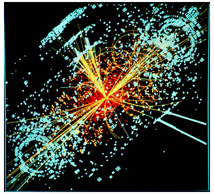

Tracking Algorithms

Caching Concepts

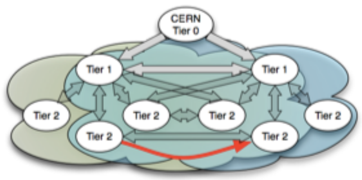

Workflow Management

Manuel Giffels

Outlook - Innovative Digital Technologies for Exploring Universe and Matter

- Joint proposal by HEP, Physics of Hadrons and Nuclei, Astroparticle Physics
- <u>Covered Topics:</u>
  - Development of technologies to utilize heterogeneous computing resources (Integration of Opportunistic Resources, Caching Technologies, Workflow Management)
  - Application and testing of those technologies in heterogenous computing resources
  - Deep Learning - Achieving knowledge through profound data-driven methods (Hardware-related Data Processing, Object Reconstruction, Simulation, Quality of Network Predictions)
  - Event reconstruction: Cost- and energy efficient utilization of computing resources (Alternative Algorithms and Architectures like GPUs)
- Proposal is reviewed in the scope of Digital Agenda programme (BMBF)

<u>Proposal of:</u>

RWTH AACHEN UNIVERSITY, universität bonn, FAU FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG, GOETHE UNIVERSITÄT FRANKFURT AM MAIN, UNI FREIBURG, GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN, UHH Universität Hamburg DER FORSCHUNG | DER LEHRE | DER BILDUNG, KIT Karlsruhe Institute of Technology, JOHANNES GUTENBERG UNIVERSITÄT MAINZ, LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN, BERGISCHE UNIVERSITÄT WUPPERTAL

<u>Associated Partners:</u> CERN, DESY, GridKa, GSI, WWU MÜNSTER

Publications

[JSSPP18MS] M. Soysal et al., Analysis of Job Metadata for Enhanced Wall Time Prediction, to be published (2018)

[ACAT17MS] M. Schnepf et al., Mastering Opportunistic Computing Resources for HEP, to be published (2017)

[ACAT17CH] C. Heidecker et al., Opportunistic Data Locality for HEP Analysis Workflows, to be published (2017)

[JOP898TH] T. Hauth et al., On-demand provisioning of HEP Compute Resources on Clouds Sites and Shared HPC Centers, Journal of Physics **898**, 5 (2017)

[KITPhDEK17] E. Kühn, Online Analysis of Dynamic Streaming Data, KIT Dissertation (2017)

[JOP898MF] M. Fischer et al., Opportunistic Data Locality for End User Data Analysis, Journal of Physics **898**, 5 (2017)

[JOP762TH] T. Hauth et al., Dynamic provisioning of a HEP computing infrastructure on a shared hybrid HPC system, Journal of Physics **762**, 1 (2016)

[JOP762EK] E. Kühn et al., A scalable architecture for online anomaly detection of WLCG batch jobs, Journal of Physics **762**, 1 (2016)

[IEEE16EK] E. Kühn et al., Online Distance Measurement for Tree Data Event Streams, IEEE 681-688, (2016)

[IARIA16MF] M. Fischer et al., Data Locality via Coordinated Caching for Distributed Processing, Cloud Computing 2016, 113-118 (2016)

[KITPhD16] M. Fischer, Coordinated Caching for High Performance Calibration using $Z\rightarrow\mu\mu$ Events of the CMS Experiment, KIT Dissertation (2016)

[JOP762MF] M. Fischer et al., Data Locality via Coordinated Caching for Distributed Processing, Journal of Physics **762**, 1 (2016)

[JOP664EK] E. Kühn et al., Active Job Monitoring in Pilots, Journal of Physics **664**, 5 (2015)

[JOP608EK] E. Kühn et al., Analyzing data flows of WLCG jobs at batch job level, Journal of Physics **608**, 1 (2015)

[ICDM15EK] E. Kühn et al., Clustering Evolving Batch System Jobs for Online Anomaly Detection, IEEE 1534-1535 (2015)

[JOP664MF] M. Fischer et al., High Performance Data Analysis via Coordinated Caches, Journal of Physics **664**, 9 (2015)

[JOP608MF] M. Fischer et al., Tier 3 batch system data locality via managed caches, Journal of Physics **608**,1 (2015)

[JOP664TH] T. Hauth et al., Dynamic provisioning of local and remote compute resources with OpenStack, Journal of Physics **664**, 2 (2015)