AIDA++ Software LoI Summary

Graeme A Stewart and Frank Gaede

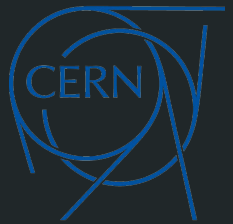

AIDA++ Open Meeting 2019-09-05

Challenges in the Software Domain

- Moore's Law continues to deliver increases in transistor density, but slowing
- Clock speed scaling failed around 2006
  - So we are basically stuck at ~3GHz clocks from the underlying Wm^{-2} limit
- Need **concurrent processing**
- Architectures become increasingly **heterogeneous**
  - GPUs, TPUs, FPGAs
- Memory access times are ~100s of clock cycles
  - Poor data layouts are catastrophic for software performance

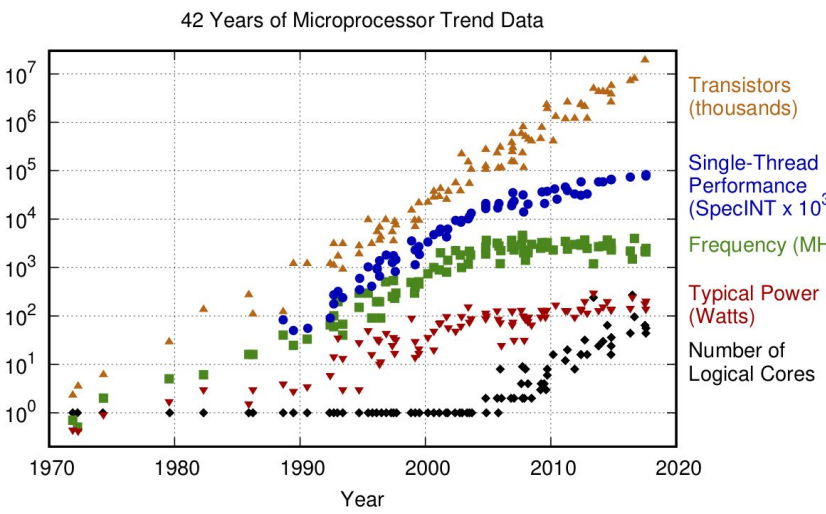

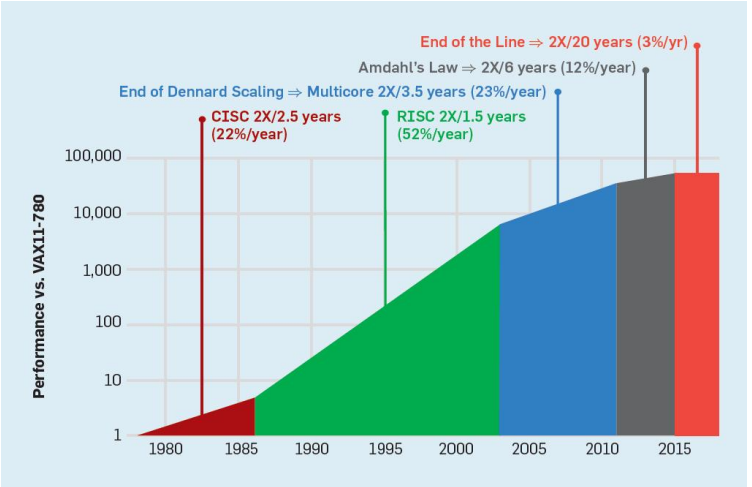

LHC and HL-LHC Challenges

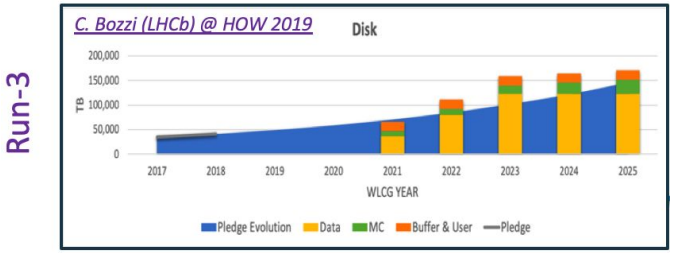

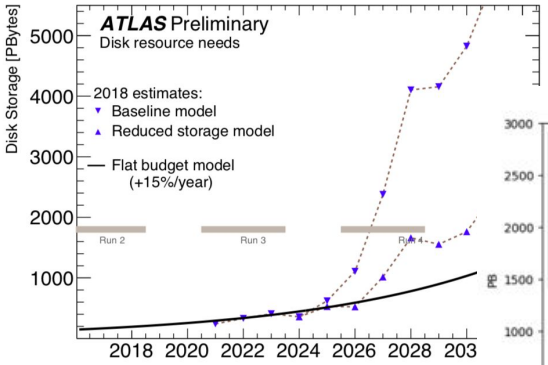

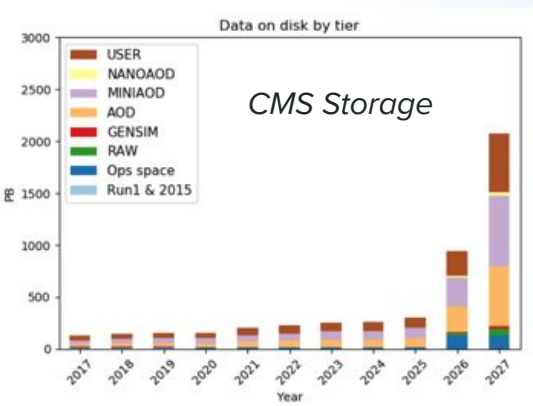

- ALICE and LHCb will have a very large increase in rate for LHC Run-3
  - This puts pressure on both CPU resources and storage
- Move to model of data reduction and *software triggers*
  - Maximise physics within available resources
- HL-LHC factor 4 in instantaneous luminosity for ATLAS and CMS (7.5 x 10^{34})
- Trigger rates of 7.5-10kHz
  - Challenge of *rate* x *complexity*
- Current plots already represent significant improvements over the estimates in the HSF Community White Paper from 2017

AIDA - Challenges for Future Detectors

- Designing and optimizing detectors for future accelerators requires efficient and powerful software tools
- Flexible and modular to study large number of variations in detail
  - For example DD4hep addresses this
- More lightweight design than what is used by large collaborations
  - Less manpower and people only involved part time
  - Ease of use is important
- Need to develop novel algorithms for novel detector designs
  - example: Pandora particle flow algorithm for highly granular calorimeters
- Ideally share framework and tools with several new concepts
  - Use synergies and address limited manpower
- Need to also support test beams for these new detectors

AIDA++ LoI Overview for Software

LoIs Received	Total FTE Years	EU FTE Years	Total Costs M€	EU Costs M €
21	141	47	10.4	3.5

- *The Good* - lots of great ideas for software projects to work on
- *The Bad* - exceeds any reasonable budget allocation by factors
- *The Ugly* - dividing an expected budget amongst all projects would leave no one with enough money, so we have to select only some and likely descope as well

LoI Summary Overview

- Advanced software for optimising detectors at (future) accelerators
  - Matching focus of previous AIDA round
  - Tracking, particle flow, EDM, frameworks, geometry
- Proposals also from the online domain
  - Closer to DAQ and real-time reconstruction and analysis
  - FPGAs feature strongly to manage rates
- Some proposals for software that more closely supports sensor development
  - Modeling of Si devices, wireless links, specific detector reconstruction algorithms

Use of machine learning is quite a common theme, arising in a lot of proposals

Criteria and Comments

- Good selection criteria:
  - generality - multi-experiment
  - focus - solvable problems
  - impact - physics reach or resource use
  - innovation - new ideas
  - collaboration - building links between institutions and with the non-academic world
- It has also been seen that it takes quite some time for software to gain a proper foothold in the community
  - Projects going well after one round of funding are much more likely to reach threshold and a critical mass when given continued support
- In the following slides....
  - Comments from Frank and Graeme are indicated in blue

Reconstruction - Tracking

- Advanced Tracking Tools (#70)
  - Contributing to the increasingly successful Acts tracking toolkit
  - Successor to the AIDA-TT efforts in AIDA-2020
  - Multi-experiment and high impact on resource consumption, particularly for high-pileup
- Development of Machine Learning algorithms for Micro Pattern Gaseous Detectors (#62)
  - Tracking in a specific detector type, main aim to improve physics outputs
  - MPGDs foreseen in a number of future experiments, ML approach is innovative
  - Detector specific; could this be developed as part of the ACTS toolkit (e.g. an alternative dataset)?

Reconstruction - Particle Flow

- Advanced Particle Flow Algorithms for Neutrino and Linear Collider Physics (#73)
  - Continued development of the generic Pandora toolkit, supported in AIDA-2020
  - Adoption by multiple experiments already (e.g. DUNE FD), but extensions are needed (e.g., dual-phase LAr)
  - Software that improves physics significantly, with opportunity to reduce resources; finding a home in many experiments and proposals (N.B UK only partners)
- Development of a new Particle Flow reconstruction algorithm based on Deep Learning techniques for a Dual readout calorimeter with timing (#103)
- Advanced Reconstruction Tools for Imaging Calorimeters (ARTIC) (#145)
  - Further PFA developments, emphasis on machine learning approaches
  - Seek a collaboration in the context of the Pandora developments

Reconstruction - Online and Dual-Readout Fibers

- Embedded event reconstruction in real time with FPGAs (#75)
  - Reconstruction at very high rates (30MHz), but also modular with FPGAs
  - Targeting LHCb VELOPIX reconstruction for Run 3
  - One specific detector for one experiment; timescales seem wrong for AIDA++
- Cluster Counting/Timing: data reduction and pre-processing of drift chamber signals sampled at high rates (#17)
  - Very high rate readouts from drift chambers to be able to do particle ID
  - FPGA data processing to manage data rate and reduce final outputs
  - Quite detector specific and traditional FPGA programming drawbacks; Gas WP better?
- Reconstruction of hadronic final states from heavy-boson decays and search for BSM signatures, using deep learning algorithms (#78)
  - Reconstruction algorithms for dual-readout fibre calorimeters (IDEA detector)
  - Use of machine learning to optimise physics searches
  - Very physics analysis oriented and not clear how useful the approach would be to other detectors

Geometry and Simulation

- Fast simulation (#42)
- Machine learning aided fast simulation (#43)
  - Generalise already known techniques that can reduce simulation time
  - Machine learning is a known area of significant interest already (calorimeter simulation in particular)
    - Integrate a workflow into the whole lifecycle of simulation
  - Nice mix of innovative approach and existing knowledge; significant impact
- Development of a Generic Digitization Package using the DD4hep Toolkit (#1)
  - Building on successful AIDA DD4hep project, which now has wide experiment support
  - Proposal to develop generic hooks for digitisation, bridging the gap between simulation and reconstruction
  - DD4hep proven to be very useful for HEP community. Is digitisation amenable to really a genuinely generic approach?

Detector Design Studies

- Turnkey software stack for detector studies (#76)
  - Produce a coherent suite of software - maximise efficiency, especially for design studies at all future accelerators (CEPC, CLIC, FCC, ILC, ...)
  - Build on best practice from LHC and non-LHC communities, with most widely adopted software
  - Valuable engineering infrastructure task that would also promote commonality inside the WP
- Machine Learning-Optimized Design of a Particle Detector for a Future Collider (#29)
  - Use fast simulation of a complete detector to go from layout to physics performance
  - Coupled with machine learning to drive an optimised design choice
  - Very innovative use of machine learning and ‘whole detector’ optimisation, concerned that fast simulation is not sufficient for this job

Heterogeneous Computing

- Optimization of the Event Data Model Toolkit PODIO (#28)
  - Data layouts for modern architectures, based on Plain Old Data; adopted by FCC
  - Follow in project from AIDA-2020 - GPUs, HDF5 backend foreseen
  - Critical part of high throughput efficient computing, code generation is good for heterogeneity; applications to Test Beam and DAQ WP
- Efficient Computing with a Heterogeneous Infrastructure of Dynamic Networked Accelerators (#142)
  - Scheduling work in a heterogeneous landscape is a much more complex issue than on a CPU
  - Development of a heterogeneous scheduler to optimise costs and improve flexibility, measure throughput for real HEP workloads on different architectures
  - Unsolved and important problem right now (though progress in ALFA/O2 domain); can be difficult to provide an output to match the many frameworks in use right now

Detector and Sensor Design Software

- Creating low power and radiation resistant FPGA-based AI accelerators using BondMachine (#48)
  - Hardware and software co-design for a low-power FPGA with high radiation tolerance
  - Specifically implementing a machine learning based algorithm for online tasks
  - Innovative idea; relationship between hardware and (changing) software not totally clear nor the specific DAQ problem being addressed
- Development of a common tool for Monte Carlo simulations of semiconductor detectors (#97)
  - Develop an open-source simulation toolkit adapted for state-of-the-art sensors
  - New physics models to optimise designs before an expensive fabrication process
  - Could save a lot on development costs and improve outcomes; impressive list of partners; better as part of a silicon sensors work package?

Detector Infrastructure and Monitoring

- Layer-to-layer wireless mesh network for event triggering, readout and data reduction at future hadron collider experiments (#96)
  - Simulation of wireless readouts, particularly for high data rate detectors (tracking layers)
  - Incorporate first steps of data processing into detector electronics, e.g. track stubs
  - Build a small ‘telescope’ to demonstrate concepts at a test beam
  - Important future technology for HEP; relationship to existing software unclear and partly a hardware project (link to other WPs)
- Advanced Real-Time Data Processing for Luminosity and Luminous Region Measurement at Colliders (#156)
  - Real time algorithms for precision luminosity measurements during VdM scans, incorporate timing information
  - Use machine learning algorithms to process additional data and improve accuracy
  - Expected gains from machine learning piece are a bit unclear (is real-time needed?)

Computing and Visualisation

- MADE-Service (MAchine and DEep Learning as a Service) (#155)
  - Streamlined service for ML training and inference
  - Open source toolkit for advances ML models
  - Computing problem, some relationship to frameworks (#142); added value over commercial or generic offerings unclear
- Dynamic Augmented Reality for radiation display and visual detector operation (#157)
  - Visualisation of DCS and DAQ information in an augmented reality environment
  - Improve detector commissioning and maintenance by improving information presented to operators and designers; additional outreach and training applications
  - Novel idea for sure; hard to design effectively to really get value; relationship to existing technology (event displays) unclear

Budget Requests

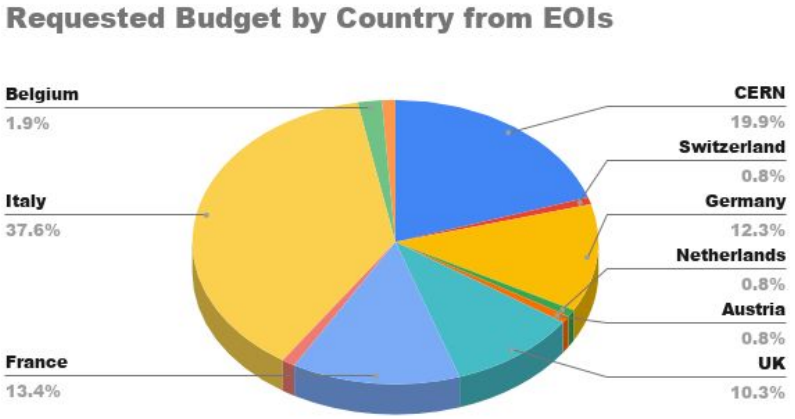

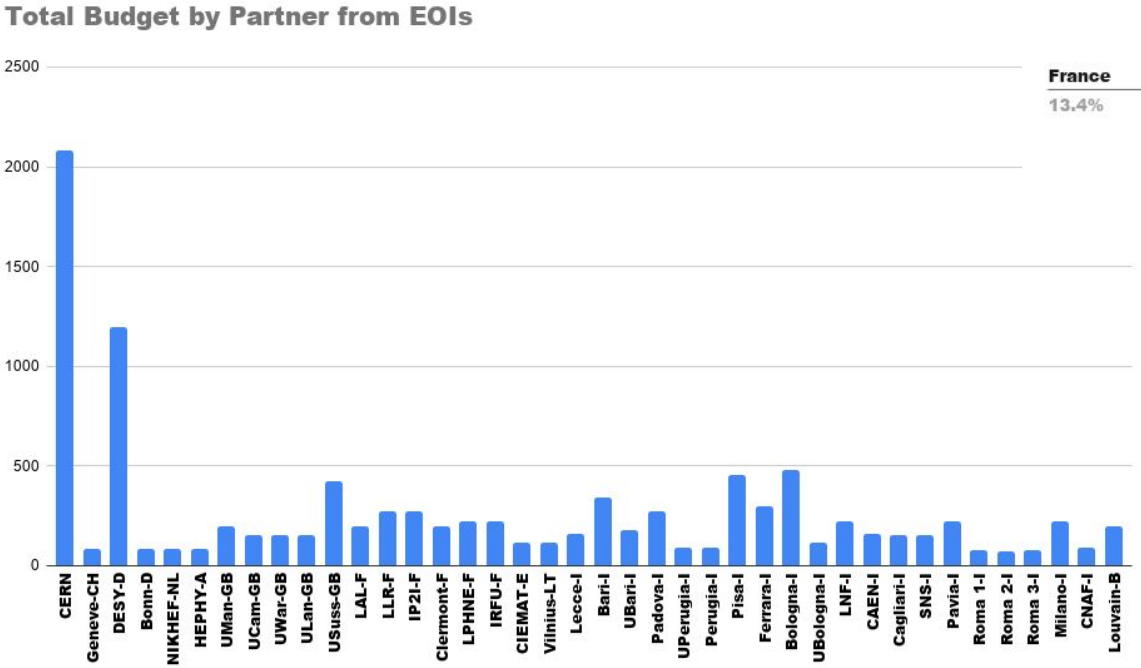

- Very uneven at the current EoI level
- Very large number of potential partners

Final Thoughts

- Many interesting EoIs showing a wide application of software to many parts of the HEP lifecycle
- Proposed budget probably overcommitted by a factor of x3-4
  - 40 partners, too large by a similar factor
- We will need to descope significantly
- Guiding principle should be what is most useful for the community, addressing the key goal of AIDA projects
  - Optimising detectors for future HEP experiments (including upgrades)
  - Prefer generic over specific EoIs, with more than one user or application
- We would like a consistent work package with synergy between tasks
  - Including synergies with other work packages

Backup

AIDA, AIDA-2020 WP3 Advanced Software

Task	Task leaders
Task 3.1 Scientific coordination	Witold Pokorski (CERN) Frank Gaede (DESY)
Task 3.2 Detector Description for HEP (DD4hep) and Unified Solids (USolids) extensions	Witold Pokorski (CERN)
Task 3.3 Alignment and conditions data (test beam)	Chris Parkes (UNIMAN)
Task 3.4 Event Data Model (EDM) toolkit and framework extensions	Benedikt Hegner (CERN) Michel Jouvin (CNRS)
Task 3.5 DDG4 (Detector Description Geant 4): Geant4 based simulation toolkit	Markus Frank (CERN)
Task 3.6 Advanced Tracking Tools	Frank Gaede (DESY)
Task 3.7 Advanced particle flow algorithms	John Marshall (UCAM)

AIDA++ Software LoI Summary

Number	Name	Contact	Institutues	Full Costs	FTE Months
1	**Development of a Generic Digitization Package using the DD4hep Toolkit**	Markus Frank	CERN, *BNIP*	300	36
17	**Cluster Counting/Timing: data reduction and pre-processing of drift chamber signals sampled at high rates**	Francesco Grancagnolo	INFN Lecce, Bari, *BNIP, CAEN*	486	54
28	**Optimization of the Event Data Model Toolkit PODIO**	Frank Gaede, Graeme Stewart	DESY, CERN	600	84
29	**Machine Learning-Optimized Design of a Particle Detector for a Future Collider**	Tommaso Dorigo	INFN Padova, UC Louvain, U Clermont Auvergne, *HSE*	600	208
42	**Fast simulation**	Anna Zaborowska	CERN, DESY, UMan	600	108
43	**Machine learning aided fast simulation**	Anna Zaborowska	CERN, LAL, DESY	600	108
48	**Creating low power and radiation resistant FPGA-based AI accelerators using BondMachine**	Daniele Spiga	INFN Perugia, CNAF, Pisa, U Perugia	360	60
62	**Development of Machine Learning algorithms for Micro Pattern Gaseous Detectors**	Gianluigi Cibinetto	INFN Ferrara, Bologna, LNF	450	50
70	**Advanced Tracking Tools**	Hadrien Grasland	LAL, CERN, DESY	600	84
73	**Advanced Particle Flow Algorithms for Neutrino and Linear Collider Physics**	Steve Green	U Cam, U Warwick, U Lancaster	456	108
75	**Embedded event reconstruction in real time with FPGAs**	Giovanni Punzi	INFN Pisa, Ferrara, Cagliari, SNS	600	100
76	**Turnkey software stack for detector studies**	André Sailer	CERN, DESY, INFN, *IHEP*	642	72

AIDA++ Software LoI Summary

Number	Name	Contact	Institutues	Full Costs	FTE Months
78	**Reconstruction of hadronic final states from heavy-boson decays and search for BSM signatures, using deep learning algorithms**	Roberto Ferrari	INFN Bologna, Milano, Pavia, Roma, U Sussex, *Yonsei U, Kyungpook U, Iowa SU, Korea U*	600	100
96	**Layer-to-layer wireless mesh network for event triggering, readout and data reduction at future hadron collider experiments**	Francesco Cresciolo	LPHNE, IRFU CEA	450	70
97	**Development of a common tool for Monte Carlo simulations of semiconductor detectors**	Dominik Dannheim	CERN, U Bonn, DESY, U Geneve, NIKHEF, HEPHY	500	75
103	**Development of a new Particle Flow reconstruction algorithm based on Deep Learning techniques for a Dual readout calorimeter with timing.**	Biagio Di Micco	INFN LNF, Milano, Padova, Pavia, Roma I, Roma Tre; CERN; U Sussex	600	100
142	**Efficient Computing with a Heterogeneous Infrastructure of Dynamic Networked Accelerators**	Lucia Silvestris	INFN Bari, CERN, U Bari	540	80
145	**Advanced Reconstruction Tools for Imaging Calorimeters (ARTIC)**	Vincent Boudry	IP2I Lyon, LLR	550	84
155	**MADE-Service (MAchine and DEep Learning as a Service)**	Daniele Bonacorsi	U Bologna, INFN Bologna, U Vilnius, CIEMAT	450	48
156	**Advanced Real-Time Data Processing for Luminosity and Luminous Region Measurement at Colliders**	Georg Auzinger	CERN	328	36
157	**Dynamic Augmented Reality for radiation display and visual detector operation**	Giulio Aielli	INFN Roma Tor Vergata, Bologna	140	30