

FUTURE TRENDS IN NUCLEAR PHYSICS COMPUTING

This is the live meeting notes document for the [Future Trends in Nuclear Physics Computing Workshop](#) held on September 28 – 30, 2022, hosted by the [Center for Frontiers in Nuclear Science](#) (CFNS) at Stony Brook, and accessible by zoom. This workshop follows previous editions in [2020](#), [2017](#) and [2016](#).

The design and development of software & computing constitute a cornerstone for the future success of the Nuclear Physics program. The workshop will serve as a forum to discuss priorities for design and development as input for a community white paper to inform the next Long Range Plan for Nuclear Science.

We meet in a time window chosen to be as inclusive as possible for participants. Substantial discussion time is included in the agenda, and session conveners will keep speakers to time in order to preserve the discussion time. This google doc can be used in advance to give the discussions structure and focus, as well as during the workshop itself to moderate and record the discussion and gather input from participants, and after the workshop as the basis for summarizing and report writing. Editing is on, and all participants are encouraged to contribute in all phases.

The workshop will conclude with a discussion of the planned Community White Paper for NSAC Long Range Plan. This workshop's written report will serve as an input to the white paper. If you're interested in contributing in writing the report, please contact the organizers. The organizers wish to thank all participants and particularly the speakers for sharing their time and expertise. We extend tremendous thanks to CFNS for making this workshop possible and helping us at every stage in the organization.

[Zoom details](#)

Introductory remarks

- Welcome and thank you all for joining us, and thank you [Center for Frontiers in Nuclear Science](#) (CFNS) at Stony Brook for hosting us!
 - More virtual than we intended, but that's the world we're (still) in!
- This workshop follows previous editions in [2020](#), [2017](#) and [2016](#)
 - We've made the content of this workshop largely complementary to the coverage of the 2020 workshop, which produced a report [summary](#) that will also inform the white paper we'll discuss on Friday

- In this workshop we'll discuss
 - where we are as a community: NP as a field, NP computing, analysis, AI/ML, exascale, data management
 - making analysis easier: user-centered design, differential workflows, metadata, systematics
 - careers and diversity, equity, inclusion: NSF perspective, demographic studies, panel discussion
 - scaling up computing: streaming, distributed heterogeneous computing, scaling storage, scaling towards EIC
- All of it informing the Community White Paper for NSAC Long Range Plan we will prepare by the end of the year, discussion on Friday
- Speakers please keep to time, conveners will make every effort to ensure this
- We aim for lots of discussion, generally an informal workshop, which works best if we keep time under control
- All invited are encouraged to make use of the live notes to pose a question during the talks (preferable to typing in chat), and/or raise your zoom hand
 - all contributors to live notes recording the discussion are welcome
- we will be recording the workshop, we trust all speakers approve

Sep 28: Where are we as a community?

Conveners (Morning): Paul Laycock (BNL), Torre Wenaus (BNL)

Where are we as a field? - Thomas Papenbrock

- Questions
 - Torre - what exactly are the emulators
 - a (greatly) simplified model involving diagonalization of a small matrix (not emulator in the sense of custom hardware), a fast and efficient approximation of an otherwise computationally expensive calculation
 - Markus - computational scale of computing nuclei from EFT Hamiltonians?
 - maxing out 12hr per job limit on LCFs
 - Torre - scale in computing terms of AI/ML applications?
 - some involve expensive calculations, growing exponentially with mass number of the nucleus, also done on LCFs
- Discussion

Outcome of previous workshops - Markus Diefenthaler

- [2020 workshop summary](#)
- Questions
 - (Torre - did we put the 2020 'summary' on arXiv? We talked about that. No we didn't and probably should.)

- Paul, Joe - Timeline for white paper? A: End of the year to have an impact on NSAC. A ‘do-ocracy’ of contributors. While keeping the editorial team small enough to complete it on time.
- Discussion

Status of analysis - The python perspective - Jim Pivarski

- Questions, discussion combined with Lorenzo’s talk

Status of analysis - The ROOT perspective - Lorenzo Moneta

- Questions
 - Torre - RNTuple, why is I/O so much faster? A: Layout the same on storage and in memory, very fast as essentially ‘nothing to do’ on read (this is uncompressed files)
 - Stefano - I didn't know about RNTuples and they have async I/O. Which ROOT tag release are they available? A: Best to use 6.26.
 - Torre - profile of GPUs these days? A: Used in many places in ROOT as Lorenzo described. GPU resources fairly limited still. Would a lot more resources be used? Chicken and egg problem. Used primarily for AI/ML.
 - Paul - is there any cross-coordination between the scikit-hep and ROOT teams, not least to make sure that there are no critical areas deserving more attention? A: Yes the teams talk at e.g. forums and workshops like this and the HSF Analysis Ecosystem workshop. The experiments are also an important source of requirements to make priorities clear. Members of the scikit-hep team also attend ROOT team meetings when appropriate.
- Discussion
 - Maybe this is heresy, but,..., some collaborations are using Java and have a lot of investment in Java (i.e. CLAS12). It would be great if there was direct access from Java to ROOT. Is that even possible? A: This still seems difficult. Several efforts mentioned that are dated.

Conveners (Afternoon): Nathan Baltzell (JLab), Ofer Rind (BNL)

AI/ML For Nuclear Physics - Tanja Horn

- Questions
 - Ofer: what resources are these run on? On leadership class facilities? Any limitations from access to resources? A: Maybe only theory applications running at LCF (Michigan/ANL). Yes, always need for more resources, AI-specific.
 - David: AI to become part of physics curriculum? A: The data science “language” is starting to be taught. Finding shared, collaborative interests/projects between data scientists and physics is an important part of it.

- Markus: 5 years ago AI/ML talks had to come from non-physicists, that's changed, what changes and challenges do you expect in the next 5? A: Hopefully will have fully integrated AI in our workflows. Education will help, improving communication between data science and domain-specific fields. Smaller, nuclear physics experiments may be attractive for application of AI/ML due to their scale, #configuration, availability.
- Discussion

Exascale: Status, Guidelines for Applications - Andrew Siegel

- Questions
 - Ofer: What else can be done to foster utilizing these resources in a truly HPC way? A: It's a big overhead/investment from the researchers. It has been getting easier, hopefully it will keep going.
 - Torre: Any trend towards higher-level, portable, libraries? A: Yes. Kokkos and RAJA, for example. Kokkos in particular has been very popular. OpenMP-offload more of a challenge.
 - Torre: On the slide 14 assessment, Celeritas is not there? A: ExaSMR, a nuclear engineering project with particle transport, had the beginnings of Celeritas embedded. Now Celeritas has spun off to a project (outside ECP) of its own.
 - Torre: Is an exascale machine built on a bunch of ARM processors conceivable? A: Power consumption is an issue, the Japanese one only reached about half exascale. There was a US proposal for this, but it's not clear the power envelope is reachable, or that it would be an improvement on GPUs in getting out the full performance with tractable programming.
 - Ofer: ECP project ends next year, what's next? A: Hopefully good use of the current exascale machines by those that are ready, producing scientific results.
- Discussion

Rucio-SENSE integration - Justas Balcas

- Questions
 - Torre: Can you say more about ATLAS using this? A: Initial discussion already happened with UChicago on how to use SENSE in ATLAS workflows. More discussions with all CERN experiments coming in October
 - Ofer: How much is ESnet specific, e.g., versus applicable in EU? A: Mostly still US-based, but a couple sites in Europe already.
 - Markus: Regarding EIC, currently using xrootd, looking into RUCIO. What extra is needed to take advantage of SENSE? A: Still some site-local things to do, but working on everything being moved to ESNet. Prioritization of bandwidth still a detail, see next.
 - Ofer: How are priorities defined? A: Configurable by sites/paths, very configurable. The current approach is relative priorities. Does get complicated when many sites involved.
- Discussion

Open Discussion: Future Trends in Nuclear Physics Computing

- David: Trying to include and hiring other-domain experts directly in our projects has been successful, but still need more. Moving our projects to exascale would be huge. Career development with dual-expertise, computing/software and physics, is important.
 - Torre: We don't hire CS domain experts directly in our physics groups (including my software group), due to institutional structure.
 - Brad: I can confirm this is also a problem within the JLab Halls / Physics Division – definitely something that needs to be addressed/encouraged at the upper management level.
 - Paul: The RUCIO team at CERN is a good example, computer scientists hired to work directly with physicists(?).
 - Torre: Probably worked because RUCIO was conceived from the beginning this way, with open-source ecosystem, devops.
 - Ofer+Torre: IRIS-HEP has been a benefit and another type of associated strategy
- David: Any exascale projects seeped into your production standard running yet? Exploring it for application? We're overwhelmed with moving our code to GPUs, e.g., the possibility that the next generation we have to do it again.
 - Torre: Exploring, yes. The HEP-CCE project in particular has looked at this. Focus so far in portable parallel processing on GPUs (finding the right libraries to work for us, haven't found a perfect fit yet, Andrew's talk was encouraging) and I/O. Nothing close to production yet. The cost-benefit analysis can be difficult/discouraging. The realistic prospect for us using exascale machines for now is AI/ML, and we're working towards that.

Sep 29 am: Making analysis easier

Convener: Brad Sawatzky (JLab), Torre Wenaus (BNL)

User-Centered Design – Wouter Deconinck

- Questions
 - Jan: Sorry, I had to listen in to a different meeting, and could only see the slides. One comment: We need to have a pipeline for people who can write the low-level stuff. We don't have enough of those! While it's good to have nice high level user-centric interfaces, we must it make as easy as possible to "look behind the curtain".
 - WD: User-centered design is not about the user interfaces (or at least not only or primarily), but about the involvement of what users want in how we design the software, including at the low level.
 - Jan: Users might not be programmers though. Users might want a closed curtain.

- Torre: On the point “Physicists have lost their edge in the big data analytics job market”, question to Wouter and others, particularly university-based, how do we gain prominence in the AI/ML job market? Do we lose out because of AI/ML hesitancy in our community arising from “black box” concerns?
 - WD: Physicists still have an advantage in model development, being aware of uncertainties associated with the model, and how to manage those.
 - BDS: Agreed, we can/should focus on that and training our people to use AI/ML as a good tool (with strengths and weaknesses) instead of a black-box.
- Lukas: Management of software frameworks also requires training. Are there resources to support this kind of training?
 - WD: Suggests looking at a number of the Agile certification training that are on the market. DOE may support that kind of thing too. Personal experience has shown it has been difficult to build this into student training programs (so far).
 - Literature Suggestion: [“The Agile Faculty”](#)
- Graeme: I am wondering what difference it makes to the approach that our user/developer boundary is not so well defined. That can be a disadvantage, the model doesn’t apply so well; but it also can be a strength, as we can transition users to contributors and then developers more easily?
 - WD: We can move in the direction of considering software tasks unrelated to one’s dissertation topic as a service task for the experiment, to get closer to an agile team where each team member can address each task. The thin line from user to developer leads to better transitions in the user -> bug reporter -> contributor -> developer -> project leader sustainability chain, so that’s a clear advantage to nurture.
- Brad: Wouter, do you have any lessons-learned or recommendations to help get institutions to shift from the ‘old’ software development models to something newer/better?
 - WD: For entire teams: organize a team training exercise (e.g. [CubeSatScrum](#) which was developed for/with NASA Langley for agile/scrum in engineering contexts, not software); invite an actual agile practitioner to talk (I can recommend and introduce [Guy Beaver](#) who’s in Yorktown); send one or more team members to receive agile certification training (e.g. starting with ICAgile’s [Agile Fundamentals](#)); encourage regular meetings of entire team including users (1 or 2 week intervals), these should go beyond developer meetings where implementation details are discussed; ship products at shorter intervals (1 or 2 month intervals; this does not mean having an ‘always working’ main branch: this is intended as a target for feature inclusion).
-

Differential Workflows – Lukas Heinrich

- Questions
 - SJ: G4 simulations (for example) do not provide gradients. How do we work around this?

- LH: see slide 28 for some suggestions, can replace true simulation with a differentiable “surrogate”; attempt to train on gradients
- Torre: Are differentiable programming approaches as amenable to effectively using GPUs as more traditional ML?
 - A: it's more challenging. GPUs, TPUs are optimized for ops like matrix multiplication found in traditional ML. The more ‘custom physics’ we add in differentiable programming approaches, the more we move away from following such patterns. So an area of work now is to try to develop differentiable physics algorithms that do map onto GPUs/TPUs. Another possibility is programmable chips like what was supposed to appear (but won't) on Aurora, ANL's forthcoming new HPC.

Metadata – Paul Laycock

- Questions
 - Markus: to what extent have you looked into what users really want? How much have end users contributed to the discussion? A: refer to the metadata paper, slide 9. Fairly exhaustive. Users want to see ‘posix’ as the read interface even if there's a RDBMS behind. e.g. conditions data on cvmfs.
 - David: on calibrations, somewhat different in nature from conditions, time when the data is processed comes in, not just validity time. A: could talk all day about this :-) do handle this with a CDB solution (see slide 17 for example). Different processing version maps to different global tag. The map makes it efficient to see what changed.
 - Chris: files as intermediary between e.g. a temp reading and getting it into the DB? wind up with a lot of files? A: design doesn't dictate a granularity. Could load up a run's worth of measurements into one calib object. Have to design for the appropriate granularity.
 - This is why the API = production system view is important, it sees the problem as “give me a run's worth of data”, or “give me a job's worth of data”. Think about the access use case rather than the underlying conditions data which may vary with a much higher granularity.
 - Chris: in online environment, architecture doesn't work so well? Want to promptly read a just-recorded value and be confident you got the valid one? A: design is intended for offline, not really for online slow controls DB. But, built into the design is guarantee that the data you read is currently valid (IOV is open ended)
 - ATLAS has an online monitoring global tag, Belle II uses a “snapshot”, so it does depend on what people need. In general, for the detector expert use case you will use your SCADA system to look at slow controls data in as much granularity as you like, you just want to look at your detector, zone in on a channel and a small time window, etc. For an “online monitoring” use case, you define monitoring plots for the whole experiment that aggregate data. You need to build up statistics and there is therefore some latency. So you have a little time, it's not instant. Time in which the SCADA system produces (or is aggregated by another tool) a lower granularity input for this use case, filling monitoring histograms. You may need a higher granularity

for your online monitoring use case, that's fine, it's used in the "Monitoring Global Tag" which is used by one job running on dedicated resources. Tougher is the high-level trigger use case, where you have no latency, and typically need to have looser conditions as calibrations are out-of-date. Here you can still use a global tag. All that said, these use cases require thought and (payload) design.

- Torre: "Data access solutions should ideally look and feel simple." Does 'simple' translate to local (e.g. metadata contained in or traveling with data files)? Or centralized with easy local extraction (c.f. conditions DB)? Or...
 - why not metadata management for analysis using CDB approach? rather than IOV use analysis-appropriate index with appropriate granularity, from lumi block up to a tag corresponding to a selected event sample
 - can we cure people of aversion to databases? (agree we don't want people writing SQL queries, need a friendly API)
 - A: yes, this is what Belle II is doing with their CDB, they use global tags (which support user defined global tags) as the user defined index, using the friendly API on slide 16
-

Systematics – Nicholas Smith

- Questions
 - Markus: also enjoy [Barlow "Systematic Errors: Fact and Fictions"](#). recommendations on modeling systematic uncertainties? combining systematics from various sources and emerging with one uncertainty estimation. A: to combine, have to agree on a parameterization. Make sure you encode your control regions in the likelihood function. Don't have to propagate them if they're all encoded in the likelihood.
 - Torre: move to common tools? A: move to common tools just starting, beginning to take it seriously. cf. Coffea+RDF slide 27.
 - Markus: How best for NP to plug into this community? A: bridging does need to be done, we should keep in touch
 - Paul: cf. slide 19, how good a fit is it to NP/EIC? There's potential for a lot of leverage if it's a good fit. Markus: maps very well.
 - Nick: Is binned analysis used in NP/EIC? Markus: Yes and no. Trend to move away from binning to minimize systematic uncertainties and to maximize information that is being shared with theory and phenomenology. Projects on joint experimental-theoretical research, moving away from unfolding and towards folding where theoretical assumptions and models can be easily varied and updated.

Sep 29 pm: Careers, diversity, equity, inclusion

Conveners: Kolja Kauder (BNL), Markus Diefenthaler (JLab)

An NSF Perspective on Careers and DE&I – Bogdan Mihaila

Questions:

- Christine: What happened to submissions by women during the pandemic? Do you have data?
 - General number of proposals hasn't changed during pandemic, except for some slight delays. No data on sub-divisions
 - C: Early reports were of a major drop-off, would like to see that
 - B: That data not normally released, can't be made available
- Anthony: Many losses at high school / middle school. Are there funds to address this?
 - No. Also problematic to work with minors. Has to be done with care. Falls under Outreach
 - A: Understood, I had to receive training to work with minors myself recently
 - B: For grad and undergraduate students:
<https://www.nsf.gov/pubs/2021/nsf21065/nsf21065.jsp>
Also encourage proposals for specific activities
- Kolja: Are there funds to research why the pipeline leaks?
 - Yes, there are, and would be happy to help and direct
- Markus: Focus here is software. Is there support for primary software people who don't publish so much science
 - Cyberinfrastructure is important. This is a career choice, and it moves you from science to CSSI. Funds for CI professionals are finite (and small). OIC has a career and core program for this.
 - But reach out to us one-on-one and we'll help as much as we can

Demographic Studies of Major Conferences in Heavy-Ion Physics – Christine Nattrass

- Questions:
- Jan: Double blind?
 - Tested blinding, improved acceptance from 23% to 48% (compared to also Nobel Laureates, didn't catch the numbers)
CN: references:
Here's the one where a Nobel prize winner and a research scientist submitted a paper together and they solicited lots of reviews, with single blind reviews much more likely to accept the Nobel prize winner's paper
 - https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4190976
 - NASA moves to double blind
 - <https://www.nature.com/articles/d41586-019-02064-y>
 - and its follow up
 - <https://www.nature.com/articles/d41586-021-03538-8>
 - "Since NASA introduced a double-blind review system to reduce bias, more successful proposals are coming from astronomers who haven't been awarded observation time before." That was the enduring effect rather than gender equity.

- Double blind review in journals leads to more papers with female first authors getting accepted
- <https://pubmed.ncbi.nlm.nih.gov/17963996/>
- Kolja: How to double blind in practice, I know my field
 - There's a study that showed actually only 50% success rate when trying to ID!
CN: see <https://pubmed.ncbi.nlm.nih.gov/17963996/>
- Amber:
 - Hosting a conference (CHEP) soon, would like advice to for plenary speakers
 - HI is very resistant to double blind. If I were chair I'd do it by fiat and let people get mad at me.
 - Open call for self nominations could help
 - Conferences were trying to match our numbers - shame works.
 - Torre: I like the self-nomination, we should do that. A: And we will
 - How robust are these skimming tools? Might be interesting topic for CHEP 2023:
 - Not at all. Started with a google doc that grew. Then we had to grow. Pushing inSpire to collect demographic information, could even be semi-anonymous (and should)
- Markus:
 - Had a 25% panel which spoke up when our program wasn't diverse enough and we revised to 25% quota. But you advise against it, can you elaborate?
 - You need to know your population, and it changes. It's a good cross check but the quota shouldn't be never exceeded
- David:
 - I've heard some other disciplines (e.g. Biology) may have better gender equity or even more female than male scientists. Have you seen any similar studies for those and can we learn anything from that?
-

Supplemental Material - Huey-Wen Lin

- Unfortunately cannot attend
- DE&I studies for lattice QCD
- 1. Is there any gender/race bias in hep-lat primary publication? Machine-Learning Evaluation of Author Ethnicity and Gender <https://arxiv.org/abs/2107.12991>
This was a quick survey of how race and gender may impact publication, a key measure of career achievement. The network of discerning author ethnicity didn't work as well as I would like; I would welcome feedback/suggestions from ML experts in this workshop that can help to refine it.
- 2. Some Statistics on Women in Lattice QCD <https://arxiv.org/pdf/1612.09367.pdf> in which I studied the gender ratio of conference participants and plenary speakers, starting from 2000.
I want to point out that many many times I heard people complain that the reason for having a low fraction of female speakers is because we don't have enough females available. This is a factor but not the sole or even dominant reason. I looked into re-invitation rates (Fig. 5 of this proceeding), and rate at which male speakers give 2 plenary talks during the study periods are 17 times higher than for women.

- Talk:  2021-11-09_hwlin_GSU.pdf
-

Panel Discussion

- Wouter: What are some actively harmful (if unintentional) practices in our software and computing communities that may have a negative impact on attraction and retention of women and other underrepresented groups?
 - Anthony: 1) Financial pressures on students (graduate and undergraduate) are not resolved by vague promises of future earnings; 2) Archaic practices such as the qualifying exam are not reflective of actual research activity and push good people out of the field.
 - Rosi: 1) Finances are indeed a concern; 2) Codes of conduct and community standards can help prevent peers pushing out underrepresented students (toxicity is coming from peers, not (only) the old sexist people, in particular when suddenly in competition for the same opportunities) -> bystander training is useful.
 - Amber: There is a cumulative effect, “death by a 1000 cuts”, that leads to being too tired to go on.
 - Wouter:
- Markus: Do the panelists have any experience or lessons related to “gate keepers”?
 - Rosi: Sometimes code review turns into “this is all awful” which is not a constructive attitude in terms of developing workforce and capacity. Maybe code reviewers could benefit from some training in this regard. Setting up pathways for people to be on-boarded into coding tasks may expand the developer community.
 - Amber: There is a lot of passion over small decisions, which can be off-putting.
 - Zhaoyu:
- Markus: How can we address the passion over small decisions into productive practices?
 - Anthony: In sPHENIX conveners are usually helpful to keeping meetings on track.
 - Amber: We need structures where not always the loudest voice is rewarded. We need to get beyond the ability of someone spending a weekend on something that is good enough getting it accepted as the de facto solution.
 - Rosi: EIC currently requires a lot more work than some people are willing to put in. That requires that we expand our community. Leaving work at the end of the day to pick up your kids should not put you at the risk of being seen as a less valuable member of the community.
 - Torre: Impressed that EIC has made important decisions in a much more structured and shorter time than ATLAS in the past.

Sep 30 am: Scaling up computing

Conveners: Brad Sawatzky (JLab), Ofer Rind (BNL)

Computing Models that Feature Streaming – Graham Heyes

- Questions

- David Lawrence: How do we balance our own investment of effort vs. what's going on in industry?
 - It's a difficult balance, we have been doing it for years. We start out when no solution exists but sometimes an external solution comes along during the development process and we have to change course.
 - Best course of action is to look at the science up front, break the challenges into components and decide which may fit with existing solutions. If none exist then keep in mind to be careful about your API, in anticipation that you may need to plug in other solutions later
- Jerome Lauret: Similar to the question from David [and possibly naive] would be how the industry solutions today compare or are usable by experiments that have pioneered streaming DAQ (ALICE, etc ...) ?? What are the lessons I learned?
 - LHC experiments are typically very specific in what the goal is and have at their disposal the resources to implement a tailored solution from scratch. I can't personally speak to their own lessons but from my experience in the NP community, with much smaller development teams, the pace of code development makes taking some existing solution very appealing even if it is not a perfect fit.
 - Comment from Jerome: but we could also look at their success (or not as they are finding the first wave of issues) and adapt from there ... leveraging some ideas from the industry.
 - Response from speaker: agreed, and that is the first step in any of these developments. Look at the lessons from the community and industry. The issue here is that streaming data is a hot topic and the industry is moving quickly.
- Torre: How are you going to use your FPGA accelerated transport?
 - Started looking within JLab to transfer data from one Hall to the data center. CH computing isn't provisioned to do this, Farm can, but the challenge is the more complicated path between the Hall systems and the central computing systems. FPGA based 'steering' can help address this potential weakness by re-steering data 'on the fly'
 - Also looking at project at BNL NSLS-II to stream light source data and with LBNL/NERSC to load share across geographically distributed data centers
- Graeme: LHC experiments that are more on the streaming side are ALICE and LHCb, where they did decide to have a fixed sized farm for their use case, because the load was rather well understood (appropriate buffers were vital). Other comment: early data reduction is vital. Would you think NP will be different?
 - This has been something that has always been true about LHC experiments. Their challenges are seemingly mountainous but can be reduced because they are so well defined. Important to understand your detector to filter data early in the process. Dynamic load balancing is more important at labs like JLab where compute

workloads vary wildly on the weeks to months timescale. Also at light sources where there are often hundreds of beamlines.

- Brad: In your view, what are the critical path technologies that need to be developed to make SRO work at scale? Are there particular technologies or SW bottlenecks that need attention?
 - It is fairly clear how to make this all happen with conventional CPUs. Work on using heterogeneous hardware needs attention. How to efficiently terminate a data stream in the memory space of a GPU or some other more exotic device. Future compute power should be sufficient but the challenge will be moving the data where it needs to be quickly, efficiently, and reliably.

Utilizing Distributed Heterogeneous Computing – Tadashi Maeno

- Questions
 - Brad Sawatzky: How would you rank the priority of the items to be addressed in the near term?
 - Latency is most important currently, but also difficult to pinpoint a general workflow description. Nothing currently that good on the market.
 - Graeme Stewart: There are some generic workflow management packages on the market with extensive documentation and help ([snakemake](#), [luigi](#)) which are popular with users... could those be plugged in on the backend of PANDA?
 - Had a look at luigi as well as airflow and prefect before. Got the impression at that time that they were coupled with their workload schedulers and it was difficult to replace the layer with PANDA. They have plugin structures to integrate resources, so one possibility is to treat PANDA as a resource provider, but that would not be a right solution since PANDA itself is a workload scheduler. Didn't look at snakemake closer yet.

Adapting and Scaling Storage for NHEP – Tejas Rao

- Questions
 - Brad: What are your thoughts on proprietary vs open sources solutions? (Long term / archival considerations, cost, vs performance, etc?)
 - Trying to use OS solutions due to cost, but vendor support is helpful with proprietary solutions (example, data migration capabilities supported with proprietary). Need to keep an eye on what is popular in industry (helps with long term viability/support).
 - Doug: What is the performance loss for using object storage vs Posix?
 - Random workloads are comparable; perhaps 50% hit on streaming workloads
 - Markus: What are thoughts on CEPH?
 - CEPH does support multiple interfaces. Most institutions favor the object storage interface. Large institutions with 100PB+ data volumes are using it in this way. POSIX interface is not as robust due to complexity of the POSIX requirements. Performance/reliability have improved a lot over the last several years. Needs

development to better support SSD/Flash media (system is still tuned for spinning media).

- CEPH also has an interface for erasure-coding (under development). Looks promising
- RedHat is supporting CEPH
- Chris P: comments that CEPH is working well for them
- Tony W: On slide 5: notes year 3 data rate 100 Gb/sec, but sPHENIX is 300 GB/sec. It looks like the challenge may be raw storage, not bandwidth to disk.
 - Keeping in mind that future projections are uncertain, yes, that looks to be true.
- Doug: Do you expect changes to workflow required as underlying storage method changes (ie. towards object storage)
 - Short answer is yes, tagging objects with metadata and moving queries to a DB will likely be required. User software will need to be adapted to this approach.

Scaling Up Towards EIC Computing – David Lawrence

• Questions

- Doug: Is EIC computing talking with ES-NET as a stakeholder in their capacity planning activities? ESnet6 (<https://www.es.net/welcome-esnet6/>) is
 - Jerome : the answer is yes :-) ... both from JLab and BNL sides. However, the main point here is that we need to firm-up the plan (somewhat tight to the detector design) and make sure it fits within the bandwidth / capacity planned.
- Doug: (Comment) - The WLCG data grid has evolved into a more flat model and much less hierarchical with data flowing between all the data centers. Tier 1 centers have tape storage requirements and additional higher level of MOU responsibilities.
- Doug: Where is Echelon 0 located? Is located at the accelerator site?
 - Answer - Yes.
- Markus (Comments):
 - Slide 10–15: network bandwidth was discussed in Streaming RO 10 (<https://indico.jlab.org/event/519/contributions/9543/>, slide 17). Conclusion was that it was hard to say what bandwidth would be, but 400 Gb/sec was likely
 - Slides 5, 14: suggests that FPGA may not be desired in order to ensure online and offline software are the same
- Markus Is feedback loop to accelerator and experiment included in planning? Accelerator and experiment operations rely on information from the streaming / online system.
 - Answer - Not yet. But discussion has started.
- Tony: Re: Butterfly model: if the SDCC went down we would automatically migrate the data to JLab. Earlier slide indicates the online buffer is on the order of days. Is the Butterfly model focussed on redundancy or supporting 'real-time' data streaming? I think promoting this model should emphasize redundancy (and disaster recovery) rather than un-scheduled unavailability of an echelon 1.
 - A: Butterfly model is focused on facilitating distributed computing and storage workflows. Fail-over and redundancy come along as additional advantages.

- Torre: on Markus' comment, no FPGAs for offline compatibility of the software. FPGAs may be used to accelerate algorithms in the time critical streaming environment, but via an API that is available offline also. e.g. at BNL there's R&D on using FPGAs in a streaming pipeline for fast ML inference. Can still do the inference offline without the acceleration.
 - A: FPGAs may well be much more ubiquitous in the future and be 'standard'. The issues with developing software for these devices (as noted in earlier) talk do remain challenges though. Consistent APIs need to be developed before this will be practical (ie. in an 'autonomous' heterogeneous computing environment).

Sep 30 pm: Community White Paper for NSAC Long Range Plan

Convener: Jan Bernauer (SBU), Markus Diefenthaler (JLab)

Structure

The white paper will be input for NSAC. We should aim for 10 pages and then likely end up with 20 pages. So each of the sections will be 1-2 pages and will grow with community input.

1. Executive Summary (Markus Diefenthaler, Torre Wenaus)
2. Making Analysis Easier (Nathan Baltzell, Wouter Deconinck)
 - 2.1. Points mentioned during discussion: Working across communities, FAIR data principles, new ideas for simulation and analysis.
3. Meeting Data Intensive Challenges (Bryan Hess, Xin Zhao)
 - 3.1. Pointed mention during discussion: Network focus
4. Near Real-Time Data Processing and Analysis (Marco Battaglieri, Jan Bernauer)
5. Heterogeneous Computing (TBC, Tom Evans)
6. AI/ML (Jin Huang, Malachi Schram)
 - 6.1. Pointed mention during discussion: AI/ML plays a role in many topics we have (analysis, streaming, ...), What we need to do this (access to facilities, ...), Open data [is this for the analysis section?] - analysis level (before the histogram)
7. Software as a Community (Kolja Kauder, Rosi Reed)
8. How Does the NP Vision Relate to ASCR (Stefano Tognini, TBC)

Each section should start with a summary where we are and should focus on where we are going and *what is needed to get there*. We will also mention the EIC Software: Statement of Principles and how this related to the white paper.

Timeline

- ☐ First draft by end of October (small group)
- ☐ Review inside NP community, deadline for comments by end of November.
- ☐ Incorporate comments during first 2 weeks of December
- ☐ Send to NSAC mid December