HSF Community Overview

Graeme Stewart, for HSF Coordination

Software and Computing Roundtable, 2021-01-12

Future Facilities and Challenges

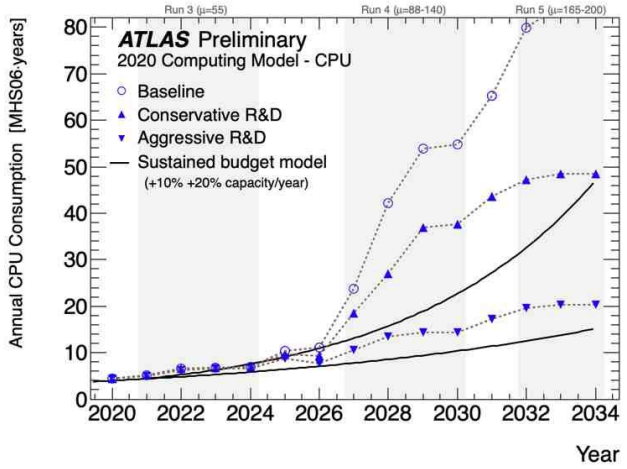

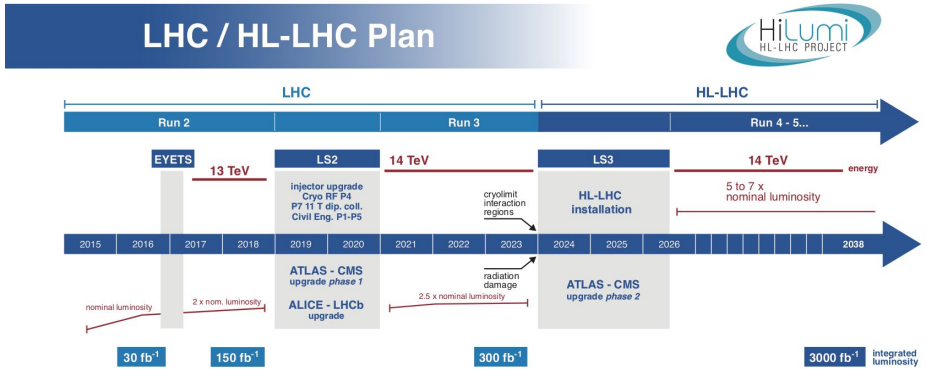

- LHC Run 3 starting very soon, with new ALICE and LHCb detectors
- HL-LHC (runs 4+) bring a additional huge challenge to software and computing
- Outside the LHC there is Belle II, DUNE, plus EIC, FAIR
- Next generation of accelerators now also under intense discussion (ILC, FCC, CEPC)

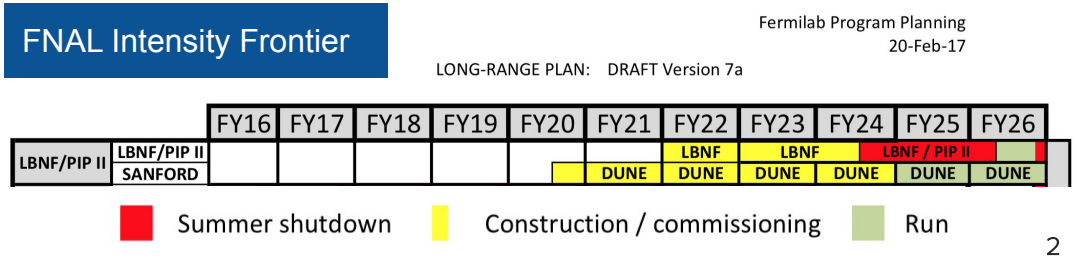

HEP Software and Computing

- High Energy Physics has a vast investment in software
  - Estimated to be around 50M lines of C++
  - Which would cost more than 500M$ to develop commercially
- It is a critical part of our physics production pipeline, from triggering all the way to analysis and final plots as well as simulation
- LHC experiments use about 1M CPU cores every hour of every day, we have around 1000PB of data with 100PB of data transfers per year (10-100Gb links)
  - We are in the exabyte era already
- This is a *huge* and *ongoing* cost in hardware and human effort
- With significant challenges ahead of us to support our ongoing physics programme

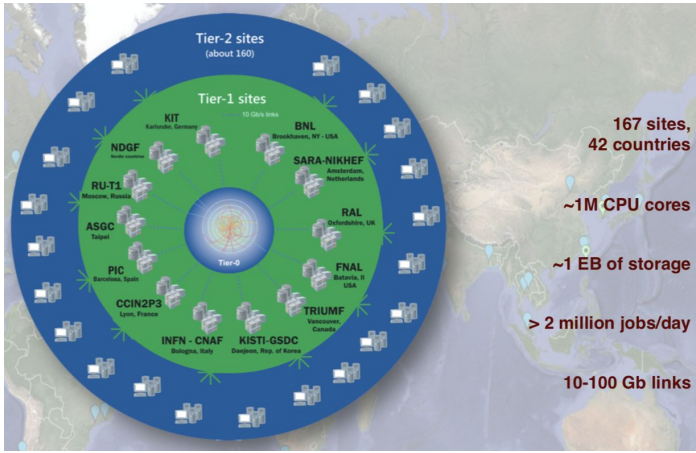

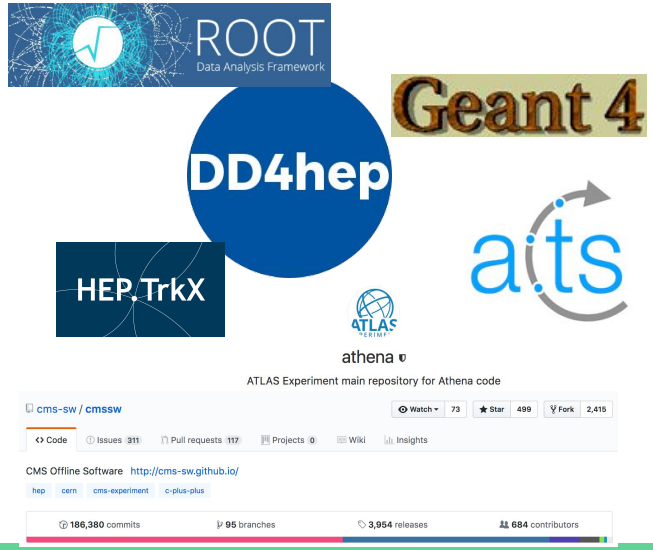

Technology Evolution

- Moore's Law continues to deliver increases in transistor density
  - But, doubling time is lengthening
- Clock speed scaling failed around 2006
  - No longer possible to ramp the clock speed as process size shrinks
  - Leak currents become important source of power consumption
- So we are basically stuck at ~3GHz clocks from the underlying Wm^{-2} limit
  - This is the *Power Wall*
  - Limits the capabilities of serial processing
- Memory access times are now ~100s of clock cycles

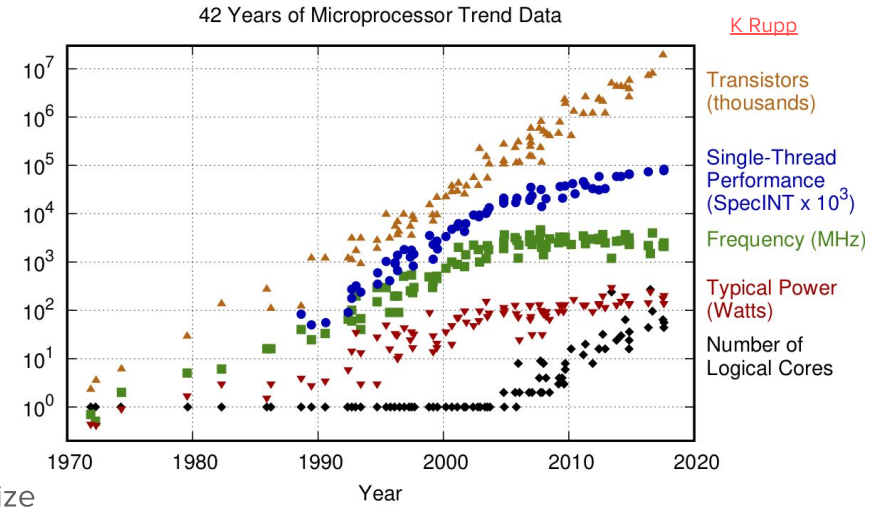

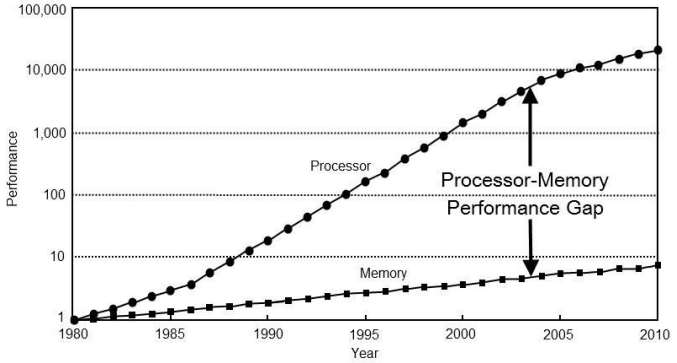

Decreasing Returns over Time

- ACM Conclusion: diversity of new architectures will only grow
  - Best known example is of GPUs
- Some promising new CPU architectures
  - A64FX, M1, AMD, ...
- Computational *power efficiency* is a driver from mobiles to HPCs

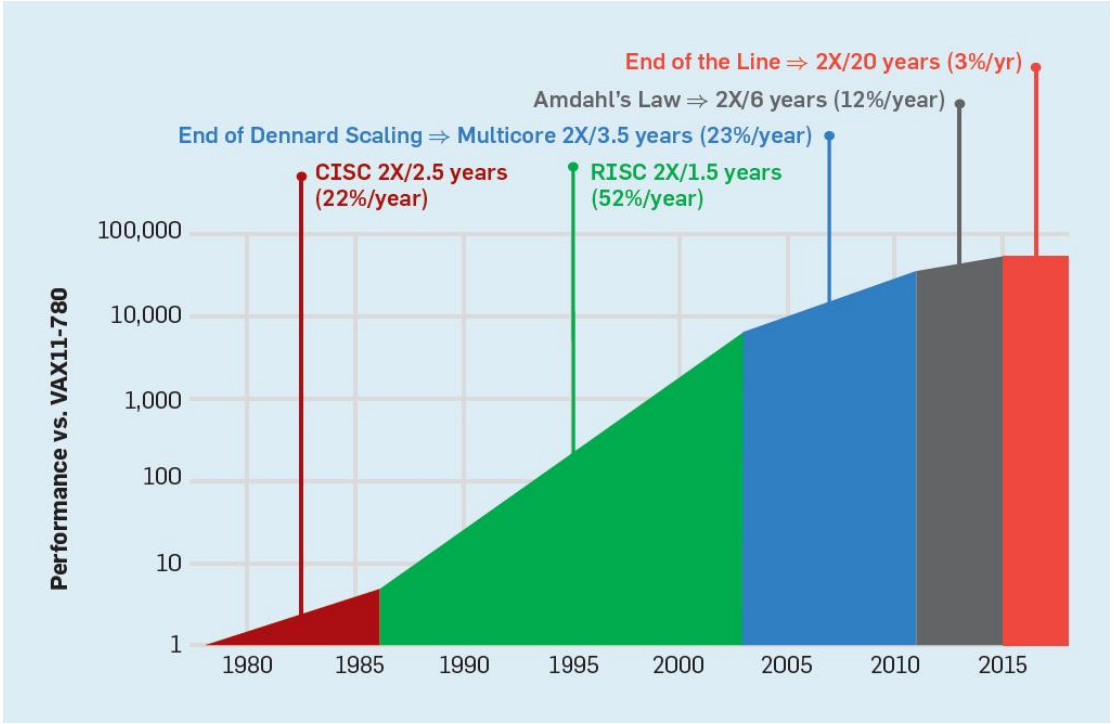

GPUs dedicate far more transistors to arithmetic

HEP Software Foundation (HSF)

- The LHC experiments, Belle II and DUNE face the same challenges
  - HEP software must evolve to meet these challenges
  - Need to exploit all the expertise available, inside and outside our community, for parallelisation and algorithmic improvements
  - New approaches needed to overcome limitations in today's code
- Reduce duplicated efforts and share knowledge
  - Each experiment had its own solution for almost everything (framework, reconstruction algorithms, ...)
  - New experiments should not be starting from scratch, but building on best-of-breed
- The goal of the HSF is to facilitate coordination and common efforts in software and computing across HEP in general
  - *Our philosophy is bottom up, a.k.a. do-ocracy*

Community White Paper, HL-LHC Review and Software Advocacy...

- Early HSF goal to decribe a *global vision for software and computing* for the HL-LHC era and HEP in the 2020s
- **Community White Paper** published in *Computing and Software for Big Science*, https://doi.org/10.1007/s41781-018-0018-8 (and on arXiv)
  - Community engagement: 310 authors from 124 institutes, 14 chapters
- Last year we 'updated' some of the CWP chapters, specifically focused on HL-LHC
  - Analysis, Reconstruction, Detector Simulation and Event Generation
  - HL-LHC Computing Review: Common Tools and Community Software
- This *review process continues* with the LHCC (Large Hadron Collider Committee) which has on ongoing review of HL-LHC preparations
- HSF also had a significant engagement with European Strategy Update and, now, with Snowmass process

4. Other essential scientific activities for particle physics

Computing and software infrastructure

- There is a need for strong community-wide coordination for computing and software R&D activities, and for the development of common coordinating structures that will promote coherence in these activities, long-term planning and effective means of exploiting synergies with other disciplines and industry
- A significant role for artificial intelligence is emerging in detector design, detector operation, online data processing and data analysis
- Computing and software are profound R&D topics in their own right and are essential to sustain and enhance particle physics research capabilities
- More experts need to be trained to address the essential needs, especially with the increased data volume and complexity in the upcoming HL-LHC era, and will also help in experiments in adjacent fields.

d) Large-scale data-intensive software and computing infrastructures are an essential ingredient to particle physics research programmes. The community faces major challenges in this area, notably with a view to the HL-LHC. As a result, the software and computing models used in particle physics research must evolve to meet the future needs of the field. *The community must vigorously pursue common, coordinated R&D efforts in collaboration with other fields of science and industry to develop software and computing infrastructures that exploit recent advances in information technology and data science. Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation.*

HSF Workshops

- Our workshops bring together the community to discuss areas of significant interest and, particularly, to encourage cross-talk between people in different areas
  - Commonly organised with WLCG (Worldwide LHC Computing Grid - computing focused, with experiments, facilities)
- In 2020, due to the pandemic, we had to postpone our intended face-to-face workshop in Lund (lessons learned for virtual events)
- May 2020 - New Architectures, Portability, and Sustainability
  - Held jointly with WLCG
- November 2020 - HSF Software Focused Workshop
  - Training, Event Generators, Detector Simulation, General R&D Session + NP Trends Report
- In 2021 we may focus on smaller ‘one shot’ events to combat Zoom-fatigue
  - So we are very interested to hear other people’s views on this too

HSF Working Groups

- Main engine of the HSF's activities are the Working Groups, who organise around a particular area of interest to HEP
  - These can be topical and cross cutting
- Data Analysis
- Detector Simulation
- Frameworks
- Physics Generators
- PyHEP - Python in HEP
- Reconstruction and Software Triggers
- Software Developer Tools and Packaging
- Training

Software Tools, Packaging and Licensing

- Copyright and Licensing
  - Long neglected inside collaborations
  - Essential to be able to
    - Open source our software
    - Combine with other open source projects
- Copyright
  - Keep as low a number as practicable
  - E.g. © CERN for the benefit of collaboration X
- License
  - Favour liberal licenses for industry collaboration: LGPL, Apache, MIT
  - Avoid GPL for libraries you want non-GPL projects to use

- Software Tools
  - Active group promoting best practice for correctness and performance
- Packaging
  - We don't build our experiment software in isolation
  - Need a software stack, incorporating many components from the open source world and HEP community
  - Preference for tools that are not home grown and have a wider support base
  - Spack actively being prototyped (link to Key4hep project in EP R&D)

[HSF-TN-2016-01; HSF-TN-2016-03; HSF-TN-2020-01]

Event Generators

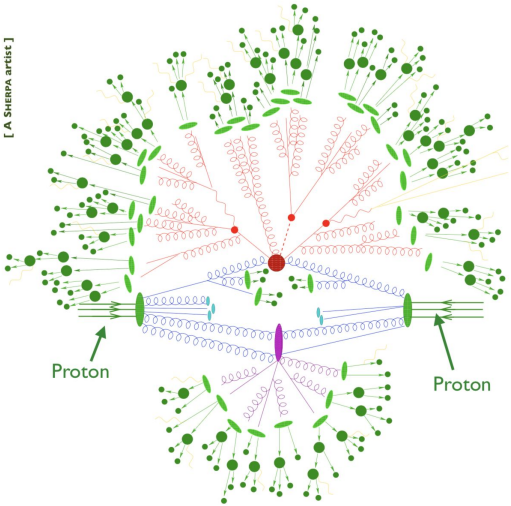

- Base of all simulation
  - In Run-1 leading order generators and little contribution to overall CPU budgets
- Increasing importance for LHC precision measurements
  - ATLAS and CMS use will use higher order generators like Madgraph and Sherpa
  - Technical and physics challenges arise particularly from negative event weights
- HSF Working Group formed after the 2018 computing for event generators workshop
  - Active in a number of areas, such as understanding costs and the physics impact of different event generation choices
  - As well as raising the issue of generators more widely (LHCC talk and related paper), highlighting proper career incentives
  - Involved in porting efforts for running event generation on GPUs (Madgraph making good progress)
- Technical work on leading order generators and hadronisation, filtering and shower step is important
  - Non-thread safe generators (evtgen)
  - Technical improvements in, e.g., Pythia and LHAPDF (Dmitri Konstantinov, EP-SFT)
- Dedicated session in November Workshop highlighted progress

Detector Simulation

- A major consumer of LHC grid resources today
  - Experiments with higher data rates will need to more simulation
  - Faster simulation, with minimal loss of accuracy, is the goal
- Machine learning lends itself to problems like this
  - Calorimeter simulations usually targeted
  - Variational Autoencoders (VAEs) and Generative Adversarial Networks (GANs), etc.
  - Key point is when is it good enough for physics
- Increasingly GPUs are seen as the real target (some architecture convergence)
  - US Celeritas Project
  - EP-SFT AdePT Project (also with UK involvement)
  - This is a hard problem, but also a relatively green-field area
    - Set realistic expectations

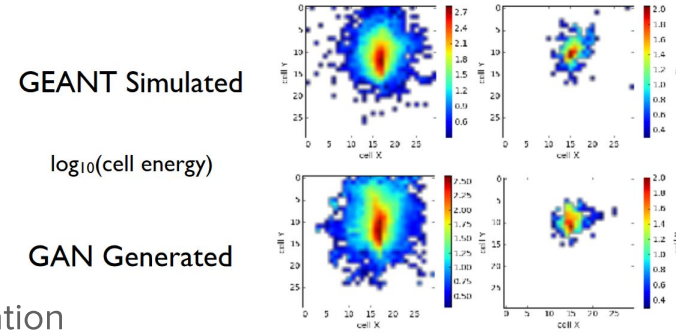

LHCb ECAL Simulated events generated with GAN (F. Ratnikov)

Good summary of recent work at Simulation Session of last workshop

Reconstruction and Software Triggers

- Hardware triggers no longer sufficient for modern experiments
  - More and more initial reconstruction needs to happen in software
  - LHCb and ALICE strategies for Run3 embody this par excellence
- Close to the machine, need to deal with tremendous rates and get sufficient discrimination
  - Pressure to break with legacy code is high
  - Lots of developments rewriting code for GPUs
    - Physics can get better!
  - Lessons learned: keep data model simple, bulk data, be asynchronous, minimise data transfers
  - High quality reconstruction close to the machine is *Real Time Analysis*
- This work is driving more and more interest in GPUs in HEP
  - Choice of LHCb to use Allen for HLT1 is a boost for this R&D line and a general retooling of HEP software
  - Gaudi integration would be of interest to ATLAS, FCC, etc.

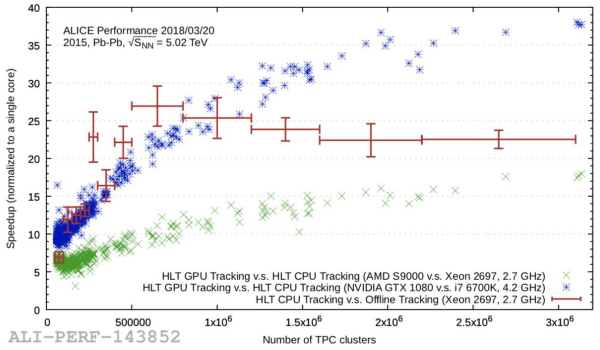

arXiv:1905.05515

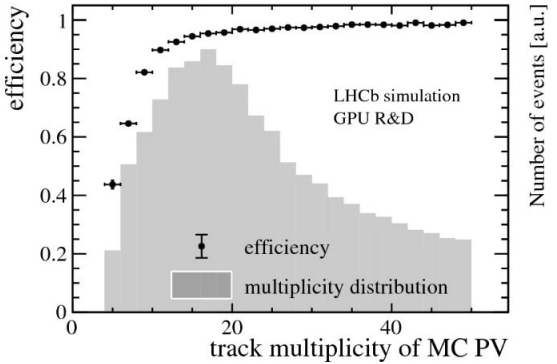

Allen: A High-Level Trigger on GPUs for LHCb, doi:10.1007/s41781-020-00039-7

Analysis

- Scaling for analysis level data also a huge challenge for all LHC experiments
- Efficient use of analysis data can come with combining many analyses as carriages in a train like model (pioneered by PHENIX then ALICE)
  - Also goes well with techniques like tape carousels (ATLAS scheme for rotating primary AOD data from tape systems into a disk buffer)
  - Interest in analysis clusters, specialised for analysis operations over the generic grid resources
- Reducing volume of data needed helps hugely
  - CMS ~1kB nanoAOD makes a vast difference to analysis efficiency and “papers per petabyte”
- Declarative models (RDataFrame and friends)
  - Say what, not how and let the backend optimise
  - E.g. split and merge, GPU execution

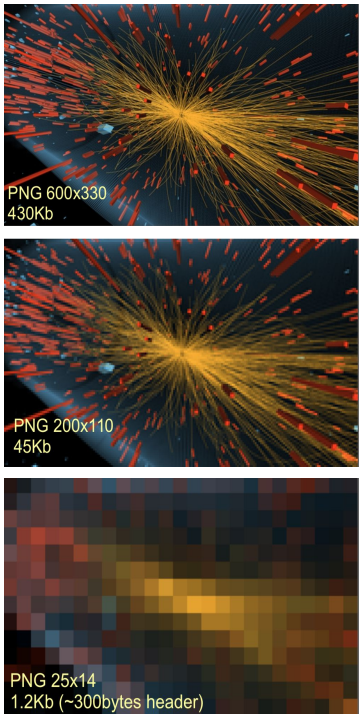

And More!

- Compute Accelerator Forum
  - Co-organised with CERN openlab and SIDIS (Software Institute for Data Intensive Science)
  - Foster discussion and teaching
  - Fundamental aspects of software engineering for compute accelerators and heterogeneous computing platforms
  - Monthly meetings
- Next meeting tomorrow:
  - SYCL (codeplay)
  - Fast Deep Learning Inference for Particle Physics (HLS4ML)
- Google Summer of Code and Google Season of Docs
  - HSF acts as an umbrella organisation for HEP
  - We have also hosted some astro projects in the past
  - Nuclear Physics projects would be very welcome
- HSF Activity Areas can provide a meeting point for other activities
  - Differentiable computing
  - Licensing
  - Event displays (Phoenix)

HSF Outlook

- Particle physics is in inherently international effort, with an excellent tradition of cooperation in many different domains
  - Detector R&D, Experiments, WLCG, Common Software
- HEP Software Foundation tries to foster much more a shared vision
  - This encourages diverse R&D!
  - Helps garner success in attracting funding to this area
  - Recognised links to other main players (WLCG, LHCC, IRIS-HEP, EPPSU, Experiments)
  - Nuclear Physics is very much a close cousin in science and software and computing
- Very happy to be a new partner in the Round Table series
  - We hope this fosters exchange of ideas and best practice as well as work on areas of common interest

Getting Involved with HSF...

- Join the HSF Forum, hsf-forum@gmail.com
  - Few messages a week with updates, jobs, items of interest
  - Owned by the community - please just post items of relevance
  - Bi-weekly coordination meetings are open to all
- Join a working group, https://hepsoftwarefoundation.org/what_are_WGs.html
  - Follow the group's meetings and discussions
  - Suggest a meeting topic
- Indico Main Page also lists workshops and training events
- Or propose a new activity!