NAME: Yves Roblin, Michael Spata, Anthony Cuffe

AFFILIATION: Jefferson Lab

EMAIL: roblin@jlab.org, spata@jlab.org, cuffe@jlab.org

EXPERTISE: Accelerator Science, simulation and operations

KEY DOE NETWORK/TRANSPORT CHALLENGE

Our Key DOE Network/transport challenges are near real-time integration of data with simulation and with designing multi-scale simulations that run across multiple different computing platforms.

The nation's accelerator facilities are complex machines that require substantial network bandwidth throughout their life cycle. Accelerators are networking intensive from the initial design phase, during intense commissioning phases up to full power and to the actual data taking for the physics experiments. The current approaches for design and commissioning are fragmented leading us to rely on experience and intuition to reach a good outcome. We strongly believe accelerator and network/transport experts working together are necessary to tackle these challenges. Being able to support such an approach would allow for cost savings during the initial design phase and commissioning by uncovering optimizations that may arise when considering the system as a whole and would advance the state of the art in several areas of computational science.

Building an Electron Ion Collider was recently strongly endorsed by the Nuclear Science Advisor Committee (http://science.energy.gov/~/media/np/nsac/pdf/2015LRP/2015_LRPNS_091815.pdf) Accelerator physics simulations for designing such a machine must include effects that encompass a wide dynamic range, from the macro beam parameters all the way to collective effects on a per particle basis. Currently, the approach is to layer the simulations at each level and consider them separately. At the deepest level are particle in cell calculations that are now implemented on graphical unit processor farms (GPU) while the top level simulations involve adjusting lattice optics for cost or performance reasons and typically run on traditional parallel computing farms. Simulations to address specific parts of the accelerator complex, which is typically an aggregate of many machines are done by different teams of accelerator scientists. Clearly, R&D in networking and transport would be required to progress from a collection of independent calculations to true multi-scale simulations.

During the commissioning phase of accelerator operations, the ramp to full power, generally requiring operating at the limits of the accelerator's parameters often proves extremely time-consuming and challenging. A recent example has been the relatively slow commissioning of the upgraded Large Hadron Collider. To diagnose these performance limitations one has to be able to synchronously sample many systems at high data rate to capture and resolve the root cause. Simulations of these same large complex systems in parallel could help recover the facility to full power in a minimal time. With limited operational budgets and fewer operating weeks, it is vitally important to develop efficient methods of troubleshooting these systems failures to maximize the physics output of the accelerators.

We believe that there are opportunities for exploring the interconnections between all of these and therefore derive the performance constraints that it places on the networks at each level. Ultimately the goal would be to have all these data sources seamlessly feed the next level and enable for a global simulation of the entire system.