1.1 Nuclear physics

STAR's tenth year of data taking has brought new levels of data challenges, with the most recent year's data matching the integrated data of the previous decade. Now operating at the Petabyte scale, the data mining and production has reached its maximum potential. Over a period of 10 years of running, the RHIC/STAR program has seen a data growing by two orders of magnitudes and despite the increase in the data processing challenge, data production have still been achieved while data analysis as well as science productivity remained strong. In 2010, the RHIC program and Brookhaven National Laboratory earn recognition as number 1 for Hadron collider research[1].

To effectively face the data challenge, all raw simulations had previously been migrated to Grid-based operations. This year, the migration has been expanded, with a noticeable shift toward the use of Cloud resources wherever possible. While Cloud resources had been an interest to STAR as early as 2007, our previous years' reports noted multiple tests and a first trial usage of Cloud resources (Nimbus) in 2008/2009 at the approach of a major conference, absorbing additional workload stemming from a last minute request. This mode of operation has continued as the Cloud approach is increasingly allowing STAR to run what our collaboration has not been able to perform on Grid resources due to its technical limitations (harvesting of resources on the fly has been debated in length by STAR as an unreachable ideal an for experiment equipped with complex software stacks). Especially, Grid usage remains restricted to either opportunistic use of resources for event generator-based production (self-contained program easily assembled) or non-opportunistic / dedicated site usage with a pre-installed software stack maintained by a local crew allowing running STAR's complex workflows. Cloud resources, coupled with virtualization technology, permit relatively easy deployment of the full STAR software stack within the VM, allowing large simulation requests to be accommodated. Even more relevant for STAR's future, recent tests successfully demonstrated that larger scale real data reconstruction is easily feasible. Cloud activities and development remain (with some exceptions) outside the scope and program of work of the Open Science Grid; one massive simulation exercise was partly supported by the ExTENCI satellite project.

STAR had planned to also run and further test the Glow resources after an initially successful reported usage via a Condor/VM mechanism. However, several alternative resources and approaches offered themselves. The use of the Clemson model especially appeared to allow for faster convergence and deliverables of a needed simulation production in support of the Spin program component of RHIC/STAR. Within a sustained scale of 1,000 jobs (peaking at 1,500 jobs) for three weeks, STAR clearly demonstrated that a full fledge Monte-Carlo simulation followed by a full detector response simulation and track reconstruction was not only possible on Cloud but of a large benefit to our user community. With over 12 billion PYTHIA events generated, this production represented the largest PYTHIA event sample ever generated in our community. The usage of Cloud resources in this case expanded the resources capacity for STAR by 25% (comparing to the resources available at BNL/RCF) and, for a typical student's work, allowed a year long science time wait to be delivered in a few weeks. Typically, a given user at the RCF would be able to claim 50 job slots (the facility being shared by many users) while in this exploitation of Cloud resources, all 1,000 slots were uniquely dedicated to a given task and one student. The sample represented a four order of magnitude increase in statistics comparing to

1 http://sciencewatch.com/ana/st/hadron/institutions/

other studies made in STAR with a near total elimination of statistical uncertainties which would have reduced the the significance of model interpretations. The results were presented at the Spin 2010 conference where unambiguous agreement between our data and the simulation was shown. It is noteworthy that the resources were gathered in an opportunistic manner as seen in Figure 3. We would like to acknowledge the help from our colleagues from Clemson, partly funded by the ExTENCI project.

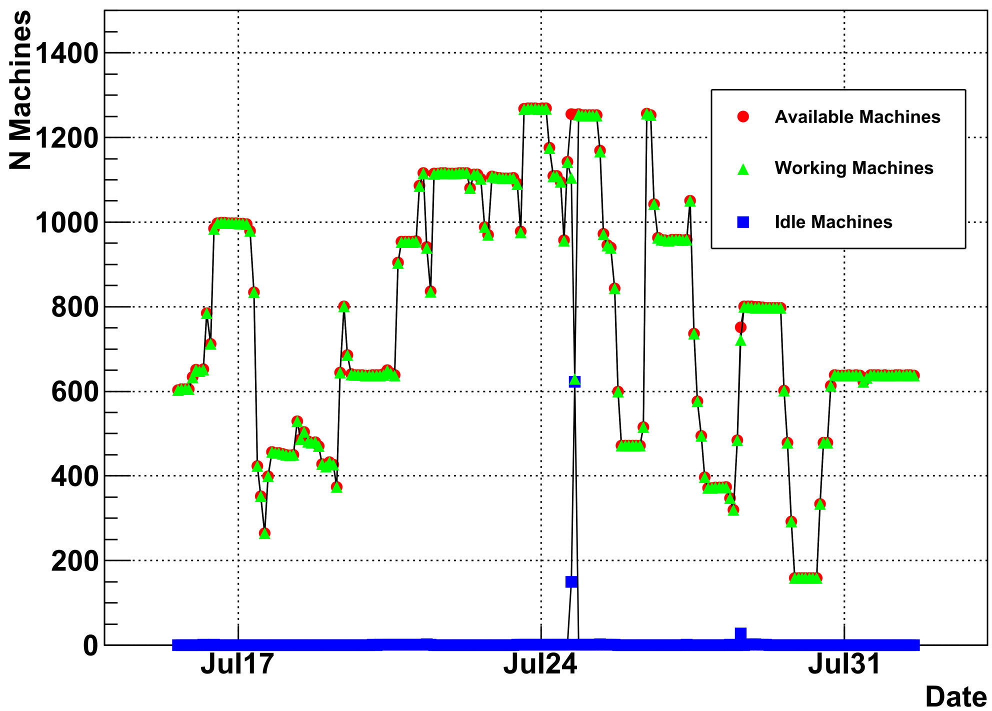

Figure 3: Graph of the number of available machines to STAR (in red), working machines (in green) and idle nodes (in blue) within an opportunistic resource gathering at Clemson University. Within this period, the overlap of the red and green curve demonstrates the submission mechanism allows for immediate harvesting of resources as they become available.

An overview of STAR's Cloud efforts and usage has been presented at the OSG all hand-meeting in March 2010 (see "*Status of STAR's use of Virtualization and Clouds*") and at the International Symposium on Grid Computing 2010 ("*STAR's Cloud/VM Adventures*"). Further overview of activities was given at the Atlas data challenge workshop held at BNL that same month and finally, a summary presentation was given the CHEP 2010 conference in Taiwan in October ("*When STAR Meets the Clouds – Virtualization & Grid Experience*").. Based on usage trend and progress with Cloud usage and scalability, we project that 2011 will see workflow of the order of 10 to 100k jobs sustained as routine operation (see Figure 3).

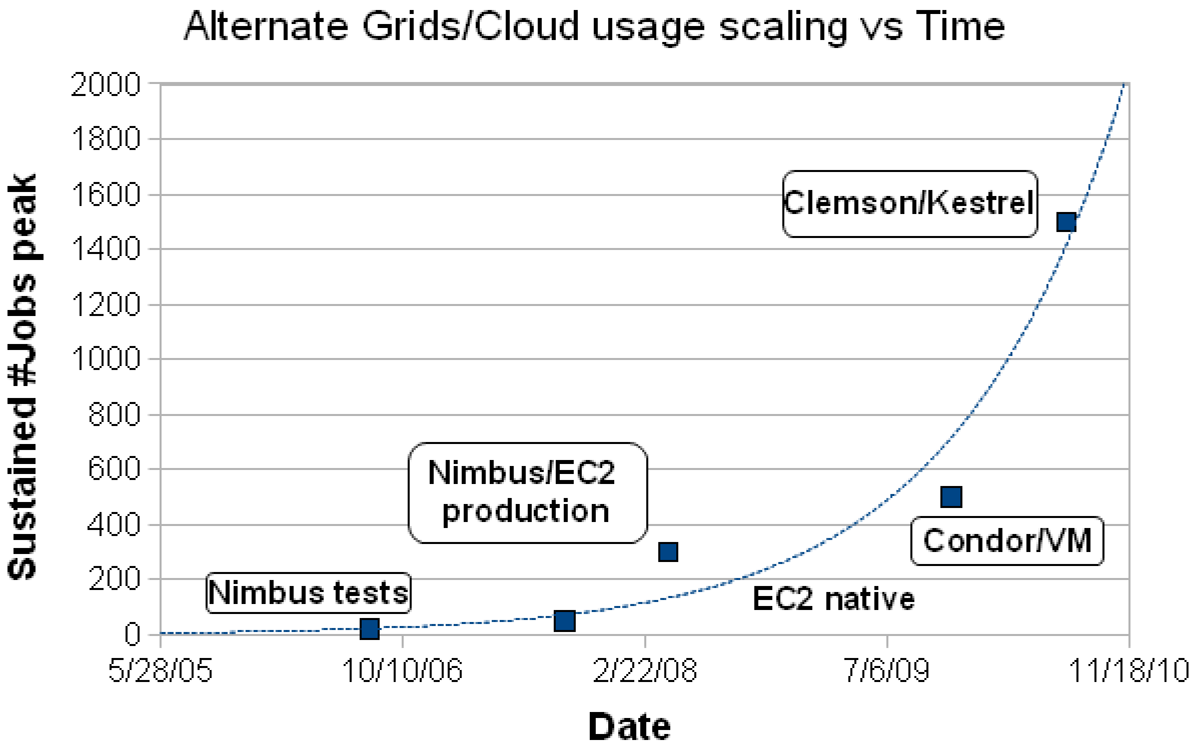

Figure 3: Summary of our Cloud usage as a function of date. As seen, the rapid progression of the exploitation and usage may indicate that a 10,000 job scale in 2011 may be at reach.

From BNL, we steered Grid-based simulation productions (essentially running on our NERSC resources), and STAR has in total produced 4.8 Million events representing a total of 254,200 CPU hours of processing time using the standard OSG/Grid infrastructure. During our usage of the NERSC resources, we re-enabled the SRM data transfer delegation mechanism allowing for a job to terminate and pass to a third party service (SRM) the task of transferring the data back to the Tier0 center, BNL. We had previously used this mechanism but not integrated it into our regular workflow as the network transfers allowed for immediate globus-based file transfer with no significant additional time added to the workflow. However, due to performance issues with our storage cache at BNL (outside of STAR's control and purview), the transfers were recently found, at times, to add a significant overhead to the total job time (41% impact on total job time). The use of a 0.5 TB cache on the NERSC side and the SRM delegation mechanism allowed mitigation of the delay problems. In addition to NERSC, large simulation event generations were performed on the CMS/MIT site for the study of prompt photon cross section and double spin asymmetry. Forty-three million raw PYTHIA events were generated, amongst which 300 thousand events were passed to GEANT as part of cross section / pre-selection speed up (event filtering at generation), a mechanism designed in STAR to cope with large and statistically challenging simulations (cross section-based calculations require however to generate with a non-restrictive phase space and count the events passing our filter and the one being rejected). Additionally, 20 billion PYTHIA events (1 million filtered and kept) were also processed on that facility. The total resource usage was equivalent to about 100,000 hours of CPU hours spanning over a period of two months total.

STAR has also begun to test the resources provided by the Magellan project at NERSC and aims at pushing a fraction of its raw datasets to the Magellan Cloud for immediate processing via an hybrid Cloud/Grid approach (a standard Globus gatekeeper will be used as well as data transfer tools), while the virtual machine capability will be leveraged for provisioning the resources with the most recent STAR software stack. The goal of this exercise is to provide a fast lane processing of data for the Spin working group with processing of events in near real time. While near real-time processing is already practiced in STAR, the run support data production known as "FastOffline" currently uses local BNL/RCF resources and passes over a sample of the data only once. The use of Cloud resources would allow outsourcing yet another workflow in support of the experiment scientific goals. This processing is also planned to be iterative, each pass using more accurate calibration constants. We expect by then to shorten the publication cycle of results from proton+proton 500 GeV Run 11 data by a year. During the Clemson exercise, STAR had designed a scalable database access approach which we will also use for this exercise. In essence, leveraging the capability of our database API, a "snapshot" is created and uploaded to the virtual machine image and a local database service is started. The need for a remote network connection is then abolished (as well as the possibility of thousands of processes overstressing the RHIC/BNL database servers). A fully ready database factory is available for exploitation. Final preparations of the workflow are in discussion, and if successful, this modus-operandi will represent a dramatic shift in the data processing capabilities of STAR. Raw data production will no longer be constrained to dedicated resources but allowed on widely distributed Cloud based resources).

The OSG infrastructure has been heavily used to transfer and redistribute our datasets from the Tier0 (BNL) center to our other facilities. Noticeably, the NERSC/PDSF center holds full sets of analysis ready data (known as micro-DST) for the Year 9 data and, on the approach of the Quark Matter 2011 conference, we plan to make available the year 10 data allowing to spread user analysis over multiple facilities (Tier2 centers in STAR typically transfer only subsets of the data, targeting local analysis needs). Up to 7 TB of data can be transferred a day and over 150 TB of data were transferred in 2010 from BNL to PDSF.

As a collaborative effort between BNL and the Prague institution, STAR is in the process of deploying a data placement planer tool in support of its data redistribution and production strategy. The planer is based on reasoning as per {from where / to where} the data has to be {taken / should be moved} to achieve the fastest possible plan, whether the plan is a data placement or a data production and processing turn-around. To get a baseline estimate of the transfer speed limit between BNL and PDSF, we have reassessed the link speed. The expected transfer profile is given by Figure 3. We expect this activity to reach completion by mid-2011.

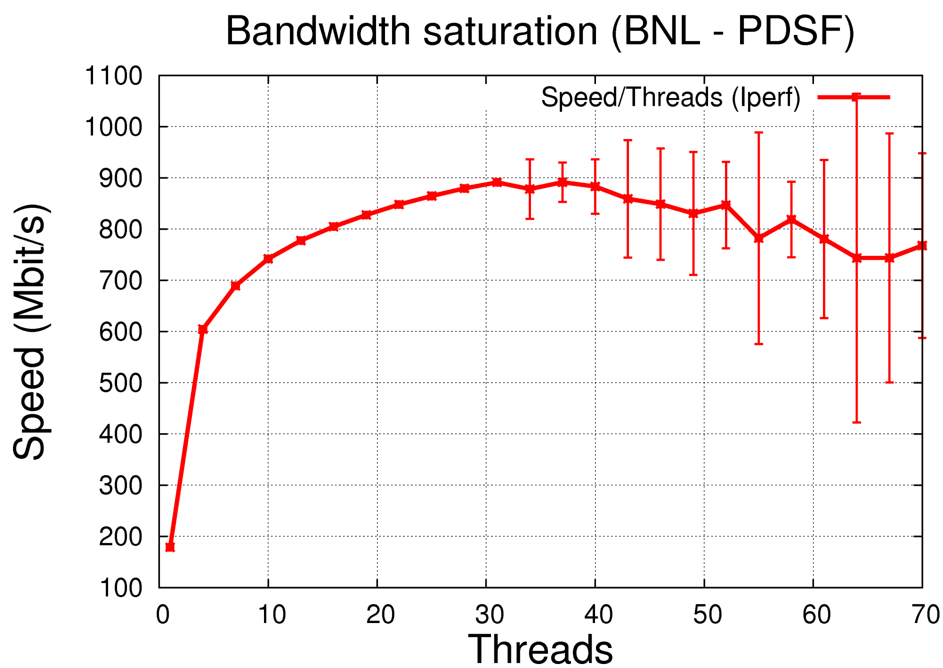

Figure 3: Transfer speed maximum between BNL and NERSC facility. The speed maximum is consistent with a point to point 1 Gb/sec link.

All STAR physics publications acknowledge the resources provided by the OSG.