VALIDATION OF GEANT4 BERTINI CASCADE NUCLIDE PRODUCTION USING PARALLEL ROOT FACILITY

A. Heikkinen and T. Lindén, Helsinki Institute of Physics, Helsinki, Finland

Abstract

We present an investigation to validate Geant4 [1] Bertini cascade nuclide production by proton induced reactions on various target elements [2]. The production of residual nuclides is calculated in the framework of an intranuclear cascade, pre-equilibrium, fission, and evaporation model [3].

A 132 CPU Opteron Linux cluster running the NPACI Rocks Cluster Distribution [4] based on Red Hat Enterprise Linux has been used to compute proton induced isotope production cross-section results for the Bertini cascade. We have used the new features of the Parallel ROOT Facility (PROOF), distributed with ROOT version 5, to analyse the isotope cross-section data. Automatic class generation for PROOF event data-analysis has been used on the Rocks cluster [5, 6]. Performance results for the cluster as measured with the ProofBench package are also presented [7].

INTRODUCTION

Cross sections for nuclide production by proton induced nuclear reactions at low and intermediate energies are of importance in many fields of science. They are needed in the design and operation of accelerator based experiments, space technology, and cosmic ray physics. They are essential for radiation protection studies, for radioactive waste transmutation and for simulation of radionuclide production in nuclear medicine [8].

Careful validation of hadronic shower models in Geant4 is needed to ensure a successful analysis phase of the future LHC experiments [9]. As part of a considerable coordinated effort put into physics validation, we study here isotope production from semi-empirical and theoretical Bertini models. A full Monte Carlo experiment, requiring significant machine resources, is needed to study rare isotope production channels.

CLUSTER HARDWARE

The 132 CPU *Ametisti* cluster at our campus (see Fig. 1) is shared between the the Department of Physical Sciences, the Department of Chemistry and Helsinki Institute of Physics. It consists of a HP Proliant DL-585 frontend, which has a shared 2 TB HP SA-6402+MSA-30 SCSI disk system. HP Procurve Gb/s and fast ethernet switches have been used for the three networks on the cluster. The two separate Gb/s networks for communication and NFS-traffic have been implemented to enhance the performance of the shared NFS disk system. The fast ethernet network is used for remote management, allowing remote booting of the nodes and the frontend. The cluster nodes are managed remotely trough a single CPU HP Proliant DL-145 G1 administration node. Table 1 describes the cluster nodes in detail.

Figure 1: The 132 CPU Opteron cluster at the University of Helsinki is used for physics-, chemistry- and HEP-applications. The *Ametisti* cluster is part of the Finnish M-grid, which currently has 1118 AMD Opteron CPUs connected with the Advanced Resource Connector (ARC) middleware [10].

Table 1: Components of the cluster compute nodes:

Nodes	1U HP Proliant DL-145 G1
CPU	dual 1.8–2.2 GHz AMD Opteron
RAM	2–4 GB ECC registered DDR
Disk	2*(80 or 160) GB 7200 rpm PATA
Networks	dual Gb/s, management 100 Mb/s
Display card	PCI ATI Rage XL 8MB
PCI	one 133 MHz 64 bit slot
USB	CDROM / flash memory stick, if needed

CLUSTER SOFTWARE ENVIRONMENT

We use a significantly patched version 3.2 of NPACI Rocks Cluster Distribution as the cluster operating- and management system on *Ametisti*. The batch queue system has been upgraded to Sun N1 Grid Engine (SGE) release 6 update 4. As cluster application software Geant4 v4.7.0.p01 and ROOT v5.08/00b have been used.

Cluster operating- and management system Rocks

NPACI Rocks Cluster Distribution is a RPM based cluster management software for scientific computation based on the source code of Red Hat Enterprise Linux, supporting cluster installation, configuration, monitoring and maintenance. Several Rocks based clusters have made it to the Top500 list [11], for a current list of the installed Rocks clusters, see the Rocks website [12].

Version 3.2. used in this work with kernel 2.4.21-27.0.4 is based on Red Hat Enterprise Linux 3.0. The monitoring of the *Ametisti* cluster is performed using the web-based Rocks Ganglia tool (see Fig. 2).

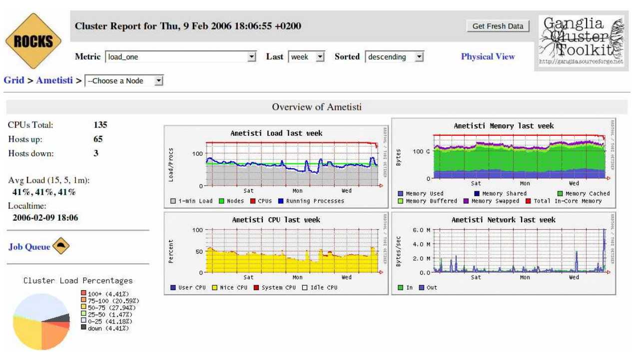

Figure 2: Rocks Ganglia monitoring page for *Ametisti*. The plot shows that sometimes the batch queues have unused CPUs, that could be used for interactive PROOF sessions.

Batch queue system SGE

Rocks comes equipped with the SGE batch queue system. A test queue for very short jobs or installation configuration development has been set up on *Ametisti*. This queue has been used in the static PROOF configuration. Larger jobs can be handled by adjusting the number of CPUs in the queue accordingly. The PROOF team has worked on a dynamic configuration file that can interact with different batch queue systems, but there is not yet any implementation for SGE.

Data analysis system ROOT/PROOF

PROOF is an extension to ROOT which allows the parallel analysis of large ROOT trees distributed over many computers. PROOF can be used on a set of inhomogenous workstations or in a homogenous cluster environment. It uses a master slave architecture with possible layers of submasters and load balancing through pull mode task distribution. It has been designed to be used on systems ranging from single or multi core workstations to large clusters with $\mathcal{O}(100)$ CPUs [7] or even collections of clusters connected with grid middleware. The best performance is obtained when the data set to be analyzed is distributed over the local disks of the cluster nodes.

We installed ROOT v5.08/00b on the frontend node and it was distributed with NFS to all compute nodes. The PROOF configuration was implemented in the Rocks XML configuration tree [9]. The NFS performance is not a problem, because of the dedicated Gb/s NFS-network, and because only a small part of the cluster is usually available for PROOF usage.

RESULTS

We have validated nuclide production by proton-induced reactions on various elements. In order to evaluate the code, we have generated $\sim 10^8$ single interactions at proton energies of 20 MeV to 8.1 GeV. The experimental data are compared with semi-empirical and theoretical models available in a Bertini cascade hybrid model based on intranuclear cascade, preequilibrium, fission, and evaporation reactions (see Fig. 3). In many cases the performance was found to be comparable to codes such as HETC, CEM, LAHET, or CASCADE [2].

The number of different isotopes the Geant4 Bertini code can produce is very high, and the results are limited by statistics for the rarely produced isotopes (see Fig. 4). In addition to nuclide production cross sections used for validation, our Monte Carlo experiment has provided information on the energy and direction of the projectiles.

The computations in this work were done on the *Ametisti* cluster described previously, building on the experiences of a small test setup reported in [9].

We used the *Run_Simple_Test.C* -script available in ProofBench package [7] to verify the setup of our PROOF cluster and compare the relative PROOF performance of *Ametisti* with *Mill*, which is a 32-bit 2.133 GHz AMD Athlon Rocks cluster with 64 CPUs. Processing eight million events with four nodes (eight slaves) with one file per slave (14 GB of data in total) with PROOF was 1.8 times faster on the 64-bit 1.8 GHz Opteron *Ametisti*-nodes than on the 32-bit 2.133 GHz *Mill*-nodes. The corresponding ratio for the packet processing CPU time per slave was 2.1.

CONCLUSIONS

The PROOF package develops at a rapid place with many useful features added during the year 2005, like the graphical user interface to PROOF clusters, using stateless connections (for details see Ref. [14]). The authentication has developed and the default SSH passphraseless authentication used on Rocks works in spite of the warning messages given, so getting it to work was easier than on our

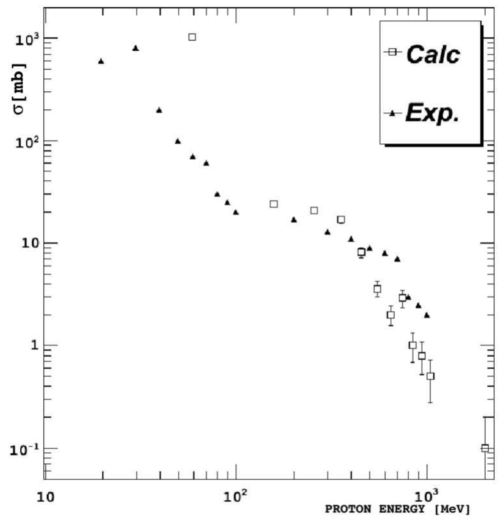

Figure 3: Example of Geant4 Bertini cascade code simulation. Proton induced isotope production is studied for the case of 39-Y-89(p,X)40-Zr-88 and validated against experimental data [13].

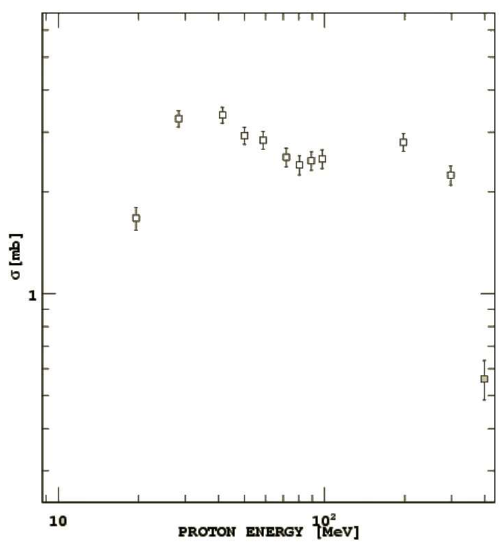

Figure 4: In the cluster environment we are able to study rare Monte Carlo events with a good statistical accuracy. Here proton absorption is demonstrated for the case 8-O-16(p,X)9-F.

test cluster [9]. PROOF is a useful and powerful tool with a lot of potential and it would deserve more extensive documentation.

Utilizing the parallellism of present and future computers is a challenge to data analysis software in HEP. ROOT and PROOF seems to develop in a very promising way to be able to harness the CPU-power needed for analysis of large data samples. We have found these tools very useful in Geant4 physics validation studies.

As a continuation of this work, we are currently parallellizing the Geant4 validation code using the TOP-C library.

ACKNOWLEDGEMENTS

We highly appreciate the PROOF team who has been very helpful in debugging and finding fixes for problems found.

REFERENCES

[1] Geant4 Collaboration, "GEANT4: A Simulation Toolkit", Nuclear Instruments and Methods in Physics Research, NIM A 506 (2003), 250-303.

[2] A. Heikkinen, "Validation of Geant4 Bertini cascade nuclide production", Proceedings of FrontierScience 2005, Milan, Italy, September 12-17, 2005.

[3] A. Heikkinen, N. Stepanov, and J.P. Wellisch, "Bertini intranuclear cascade implementation in Geant4", arXiv: nucl-th/0306008.

[4] P. Papadopoulos, M. Katz, and G. Bruno, "NPACI Rocks: Tools and Techniques for Easily Deploying Manageable Linux Clusters", Concurrency Computat: Pract. Exper. 2002; 00:1-20.

[5] F. Rademakers, M. Goto, P. Canal, R. Brun, "ROOT Status and Future Developments", arXiv: cs.SE/0306078.

[6] M. Balli et al., "Parallel Interactive and Batch HEP-Data Analysis with PROOF", Proceeding of ACAT05, May 22 - 27, 2005 DESY, Zeuthen, Germany.

[7] M. Ballintijn et al., "Super scaling PROOF to very large clusters", Proceedings of Computing in High Energy Physics, CHEP'04, Interlaken, Switzerland, 27 Sep – 01 Oct 2004.

[8] Th. Schiekel et al., "Nuclide production by proton-induced reactions on elements ($6 \leq Z \leq 29$) in the energy range from 200 MeV to 400 MeV", Nucl. Instr. and Meth. in Phys. Res. B 114 (1996) 91-119.

[9] A. Heikkinen and T. Lindén, "Validation of the GEANT4 Bertini Cascade model and data analysis using the Parallel ROOT Facility on a Linux cluster", Proceedings of Computing in High Energy Physics, CHEP'04, Interlaken, Switzerland, 27 Sep – 01 Oct 2004.

[10] www.nordugrid.org/middleware/.

[11] The list of the 500 fastest computers according to the Linpack benchmark, www.top500.org.

[12] The homepage of NPACI Rocks cluster distribution, www.rocksclusters.org.

[13] M. Khankhasayev ed., "Nuclear Methods for Transmutation of Nuclear Waste", Proceedings of the Workshop, Dubna, Russia, 2931 May, 1996.

[14] G. Ganis, "PROOF : status and developments", Root 2005 users workshop.