

Transforming Geant4 for the Future

Bob Lucas and Rob Roser



Geant4 is Important to HEP

Big computation

- Large ensembles of sequential jobs
- Runs on a worldwide Grid of processors
- \$10M per year of HEP H/W investment

Big data

- HEP has more simulated than collected data
- Geant4 generated most of HEP's O(100) PBytes
- Any useful data curation must also curate the experiment's Geant4 program

Challenges for Geant4's Future

Geant4 is a sequential C++ toolkit

- MC runs are ensembles dispatched to the Grid
- It can take months to simulate a billion particles

CPU capability has plateaued

- Dennard scaling has ended

This constrains progress in HEP

- Precludes more sophisticated physics packages
- Precludes real-time adaptation of data analysis
- Hampers analysis of individual events in isolation

Charge from Dan Hitchcock (ASCR) and Jim Siegrist (HEP)



Department of Energy
Office of Science
Washington, DC 20585
March 7, 2012

Dr. Robert F. Lucas
Information Sciences Institute
University of Southern California
4676 Admiralty Way, Suite 1001
Marina Del Rey, CA 90292

Dr. Robert M. Roser
Particle Physics Division
Fermi National Laboratory
Wilson and Kirk Roads
Mail Station: 318 (CDF ASSY HALL)
Batavia, IL 60510-5011

Dear Drs. Lucas and Roser:

Thank you for agreeing to organize and co-chair a workshop on issues related to, "Transforming Geant4 for the Future."

The Department of Energy's (DOE) Office of High Energy Physics (HEP) and Office of Advanced Scientific Computing Research (ASCR) are co-sponsoring this workshop to identify opportunities and needs for leveraging the powerful physics capabilities of Geant4 into a robust, sustainable software infrastructure. This workshop will identify applied mathematics, computer science and algorithm development challenges in effectively transitioning Geant4 to new computer architectures. This workshop will examine opportunities for discovery enabled by numerical algorithms and optimization tools likely to emerge from current ASCR investments to meet these challenges. This workshop will explore models of collaborative efforts that include applied mathematicians, computer scientists, algorithm developers as well as Geant4 users to optimize productivity and the scientific advances through modeling and simulation. The workshop has the opportunity to influence future HEP and ASCR investment decisions.

The goals of this workshop are to:

- Identify and review the current status, successes and limitations of the Geant4 software toolkit including the international scene;
- Determine the challenges that lie ahead in transforming Geant4 into a software that runs effectively in new architectures;



- Consider the opportunities related to existing algorithms, optimization tools and physics models;
- Ascertain research areas in applied mathematics, computer science, algorithm development and simulation strategies needed to leverage the powerful physics capabilities of Geant4 into a robust, sustainable software infrastructure;
- Create the foundation for information exchanges and collaborations among ASCR and HEP supported researchers, ASCR computing facilities and Geant4 user community;
- Understand the research that may currently be in progress at the international level and identify directions that would not duplicate existing projects;
- Explore possibilities for transformative advances that could ensue through the unique characteristics of the HEP-ASCR collaborations.

The workshop should focus on areas of research and collaboration to position Geant4 to exploit emerging computer architectures, while providing for a strong, diverse and potentially changing user community. We anticipate that the workshop will develop findings in the context of HEP and ASCR missions and the collaborative exchanges between the two communities will seed fruitful directions that enhance the impact of the workshop. We anticipate that you will establish a program committee to organize the workshop and that the workshop will consist of plenary and breakout sessions.

The workshop should be held in the Washington, DC, metropolitan area in early to mid-May 2012 time frame. We request that a written report representing the results of the workshop be prepared by you as workshop chairs, with inputs from panel leads, and other assigned writers. The report should specifically address all workshop goals. We would like a draft version of the Executive Summary, containing an overview of the major findings of the workshop, within 45 days after the workshop. The final report will be used by HEP and ASCR to shape out-year program plans and to inform the Office of Science long-range budget planning process.

Dr. Ceren Susut (Ceren.Susut-Bennett@science.doe.gov), ASCR, and Dr. Lal Chatterjee (Lal.Chatterjee@science.doe.gov), HEP, will be your primary DDE contacts for this workshop and will provide any support needed to organize and conduct a successful workshop.

3

This workshop is an important step toward developing and executing the strategic vision for porting Geant4 in the future through a partnership between HEP and ASCR. Thank you again for agreeing to contribute to this effort.

Sincerely,

Dr. Daniel A. Hitchcock
Associate Director of Science for the
Office of Advanced Scientific Computing

Dr. James Siegrist
Associate Director of Science for the
Office of High Energy Physics

Summary of Charge from Dan and Jim

Bob Lucas and Rob Roser to co-chair workshop
Ceren Susut and Lali Chatterjee are DOE contacts

Goals:

- **Review status, successes and limits of Geant4**
- **Determine challenges posed by new architectures**
- **Consider opportunities in algorithms and optimization**
- **Ascertain research for robust, sustainable code**
- **Create foundation among ASCR and HEP investigators**
- **Understand and not duplicate international efforts**
- **Explore transformative advances via HEP-ASCR collaboration**

Invitation to Participants

March 16, 2012

Dear colleague,

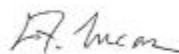
At the behest of Dr. Daniel Hitchcock (DOE SC ASCR) and Dr. James Siegrist (DOE SC HEP), we are organizing a workshop to study the research challenges posed by the rapid changes taking place in the computing technology that underpins much of the research in High Energy Physics, and we are inviting you to join us. Specifically, we would like to focus on the GEANT4 code and how it can be adapted to forthcoming many-core microprocessors in order to increase its overall performance and allow the incorporation of additional physics that will increase its overall precision. GEANT4 in turn is used as an important tool in the design of high energy physics detectors and their subsequent data handling. It is a key component of an experiment's overall physics analysis toolkit. As the field of HEP progresses it will need to understand how to handle the exponentially increasing volume of experimental data that must be collected, stored, simulated and analyzed to enable further discovery of the fundamental workings of nature.

We believe that the challenges of modeling and the handling of data generated by High Energy Physics experiments will be best addressed by multidisciplinary teams including applied mathematicians, computer scientists, and of course, physicists. The workshop will therefore not only address questions of what research is needed to accelerate the pace of scientific discovery in an era of many-core processors and new storage technology, but how best to organize teams whose expertise spans the research domains of both ASCR and HEP. These collaborations must extend beyond the United States as High Energy Physics experiments are increasingly international in character, as is the development of tools such as GEANT4.

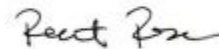
The ultimate goal of this workshop is to draft a report for ASCR and HEP as to how they can organize future research investments so as to maximize the productivity of any collaborative efforts that they may launch along these lines. We expect that this will result in a broad research agenda spanning applied mathematics, performance engineering, scientific data management, and the analysis of such data. We request your participation in both the workshop itself as well as the drafting of the report.

Our workshop will be held in the Washington, DC, metropolitan area, May 8th and 9th, 2012. Precise details of the venue and the agenda will follow shortly. Please let us know if you will be able to join us. Note, due to space limitations at the event location, pre-registration is limited to the invited experts.

Sincerely,



Bob Lucas
Information Sciences Institute
University of Southern California
rlucas@isi.edu



Rob Roser
Particle Physics Division
Fermi National Accelerator Laboratory
roser@fnal.gov

Participants

DOE	8
Labs	31
Academe	9
Total	48

David	Asner	Pacific Northwest National Laboratory
Amber	Boehnlein	SLAC National Accelerator Laboratory
Richard	Brower	Boston University
Paolo	Calafiura	Lawrence Berkeley National Laboratory
Philippe	Canal	Fermi National Accelerator Laboratory
Lali	Chatterjee	DOE Office of High Energy Physics
Gene	Cooperman	Northeastern University
Terence	Critchlow	Pacific Northwest National Laboratory
Pedro	Diniz	University of Southern California
V. Daniel	Elvira	Fermi National Accelerator Laboratory
Michael	Ernst	Brookhaven National Laboratory
Robert	Fowler	University of North Carolina
Salman	Habib	Argonne National Laboratory
Andrew	Hanushevsky	SLAC National Accelerator Laboratory
Jim	Kowalkowski	Fermi National Accelerator Laboratory
David	Lange	Lawrence Livermore National Laboratory
Randall	Laviolette	DOE Office of Advanced Scientific Computing Research
Thomas	LeCompte	Argonne National Laboratory
Steven	Lee	DOE Office of Advanced Scientific Computing Research
Qing	Liu	Oak Ridge National Laboratory
Bob	Lucas	University of Southern California
David	Malon	Argonne National Laboratory
Gabriel	Marin	Oak Ridge National Laboratory
John	Mellor-Crummey	Rice University
Richard	Mount	SLAC National Accelerator Laboratory
Esmond	Ng	Lawrence Berkeley National Laboratory
Boyana	Norris	Argonne National Laboratory
Lucy	Nowell	DOE Office of Advanced Scientific Computing Research
Bruce	Palmer	Pacific Northwest National Laboratory
Karen	Pao	DOE Office of Advanced Scientific Computing Research
Marc	Paterno	Fermi National Accelerator Laboratory
Joseph	Perl	SLAC National Accelerator Laboratory
Allan	Porterfield	University of North Carolina
Lawrence	Price	DOE Office of High Energy Physics
Michael	Procario	DOE Office of High Energy Physics
Kenneth	Roche	Pacific Northwest National Laboratory
Rob	Roser	Fermi National Accelerator Laboratory
Paul	Ruth	University of North Carolina
Allen	Sanderson	University of Utah
Elizabeth	Sexton-Kennedy	Fermi National Accelerator Laboratory
Panagiotis	Spentzouris	Fermi National Accelerator Laboratory
Ceren	Susut	DOE Office of Advanced Scientific Computing Research
Timothy	Tautges	Argonne National Lab
Craig	Tull	Lawrence Berkeley National Laboratory
Brian	Van Straalen	Lawrence Berkeley National Laboratory
Torre	Wenaus	Brookhaven National Laboratory
Dennis	Wright	SLAC National Accelerator Laboratory
John	Wu	Lawrence Berkeley National Laboratory

Logistics

Dates: May 8-9, 2012

Location: Rockville Hilton, Rockville, MD

Run by Donna Nevels

- **Oak Ridge Institute for Science and Education**

Agenda

Tentative Agenda

"Transforming Geant4 for the Future"

May 8

8:00-9:00 am	Registration Open Continental Breakfast	
9:00-9:30 am	Welcome and Goals ASCR and HEP	Dan Hitchcock, ASCR Jim Siegrist, HEP
	Conference Chairs	Bob Lucas, USC Rob Roser, FNAL
9:30-10:00 am	Geant4 overview	Amber Boehnlein, SLAC
10:00-10:15 am	Geant4 Collaboration and History	Asai Makoto, SLAC
10:15-10:30 am	Break	
10:30-11:00 am	Physics uses of Geant4	Tom LeCompte, ANL
11:00-11:30 am	Trends in multi-core architecture and optimization opportunities	Rob Fowler, UNC
11:30-12:00 pm	Exploiting concurrency in Geant4	Jim Kowalkowski, FNAL
12:00-12:30 pm	Scientific data management and analysis challenges	Rob Ross, ANL
12:30-12:45 pm	Charge to Workshop Participants	Bob Lucas, USC Rob Roser, FNAL
12:45-1:45 pm	Lunch on your own	
1:45-3:15 pm	Parallel Sessions Multi-core Optimization Scientific Data Handling and Analysis	
3:15-3:30 pm	Break	
3:30-5:00 pm	Resume sessions	
5:00 – 5:30 pm	Report of Parallel Session progress	
5:30 pm	Adjourn for the day	

May 9

8:00-8:30 am	Continental Breakfast
8:30-10:30 am	Resume breakout discussions
10:30-11:00 am	Break
11:00-12:00 pm	Plenary reports from discussions
12:00-12:15 pm	Closing remarks and path forward
12:15-1:30 pm	Workshop adjourn Working lunch for organizers and chairs
1:30 pm	Report preparation

Challenges Discussed

Low-level Code Enhancements

Multi-threading

S/W Reengineering

Domain Specific Optimization

Early Adoption of ASCR Research

CAD Interface for Geometry

Parallel I/O

Data Access

Distributed Data Management

Workflow and Provenance

Visualization and Analysis

Report


**Rob Fowler from UNC
created a shared SVN
repository**

**All participants were
invited to contribute**

Finalized in Sept. 2012

Transforming Geant4 for the Future

Report from the Workshop on
Transforming Geant4 for the Future
September 2012



Summary

Increasing the capability of Geant4 is important to HEP

ASCR and HEP should investigate this together

- **Optimize today's Geant4 for immediate impact**
- **Refactor Geant4 for future computing systems**
- **Address challenges from the PBytes of data generated**

Must coordinate with worldwide Geant4 collaboration