Eric Phipps

Session 4

Bio: Eric joined Sandia National Laboratories in September 2002 and is currently a principal member of the technical staff in the Scalable Algorithms Department. Eric's research focuses on developing new capabilities for predictive simulation and analysis in Sandia's large-scale parallel application codes using techniques based on automatic differentiation and template-based generic programming. His work has recently emphasized developing tools and techniques relevant to emerging extreme-scale computer architectures. He is the lead developer for several related software packages in Trilinos including the Sacado automatic differentiation, Stokhos embedded uncertainty quantification, and LOCA continuation/bifurcation analysis packages. Finally, he leads the development of the GenTen software for tensor-based analysis on emerging extreme-scale architectures.

Title: Automatic Differentiation as an Enabling Technology for Simulation, Analysis, and Scientific Machine Learning

Abstract: Derivative computation is essential in simulation and analysis, from Jacobians for implicit time-stepping to force calculations in molecular dynamics, to parameter sensitivities for optimization and uncertainty quantification. Moreover, recent scientific machine learning (ML) successes have intensified the need for differentiable programming techniques whereby sub-grid/sub-scale ML models are trained on high-resolution indirect data by coupling to traditional physics simulations. Historically, needed derivative computations have been created by hand, which is well-known to be both time-consuming and error prone, often resulting in compromises that reduce complexity and therefore accuracy. Automatic differentiation (AD) is a family of techniques that transform computer code for a given computation into code for its derivative and is a proven alternative. While it is a decades old idea originally created for scientific simulation, it has only recently gained wide-spread recognition through its integration into well-known ML toolkits for training deep learning models (where it is commonly referred to as back-prop). In this talk, we briefly review AD technologies and their history, and then study the relevance of AD to simulation, analysis, and scientific ML by examining several case studies spanning nearly two decades of use at Sandia National Laboratories. We will discuss the approach for AD taken at Sandia, describe successes and challenges of this approach, and conclude with potential future research directions to overcome these challenges.