

HIGH-PERFORMANCE MONTE CARLO TOOLS

Michael Mascagni

In April, the Workshop on High-Performance Monte Carlo Tools met at Mississippi's Stennis Space Center to discuss tools that enable Monte Carlo computations on parallel and distributed systems. At this two-day event, speakers presented new work in the areas of Monte Carlo algorithms, parallel Monte Carlo applications, parallel and distributed computing tools, random-number generation, and recent trends in high-performance parallel and distributed computing. The workshop's organizers have recently developed the SPRNG library for parallel random-number generation. In conducting this workshop, their aim was to place this new high-performance Monte Carlo tool in the context of current and future user needs, applications requirements, recent theoretical results, and future computing trends.

Parallel random numbers and Monte Carlo applications

To begin the program, David Ceperley of the National Center for Supercomputing Applications (NCSA) discussed issues related to the testing of parallel random numbers. Next, Simonetta Pagnutti of Italy's ENEA-Bologna described recent results on the theoretical analysis of correlation found in common parallel random-number generators and on the possibility of controlling correlation effects during a Monte Carlo computation. Following that talk, Karl Entacher of Austria's University of Salzburg discussed the geometrical quality of random numbers and another theoretical method for studying the quality of random num-

bers. Miron Livny of the University of Wisconsin then presented his distributed-computing package, Condor, which could be a powerful tool for distributed Monte Carlo, especially when outfitted with SPRNG.

The next cluster of talks focused on different applications areas in Monte Carlo. The speakers were asked to present their application area in detail and

◆

*The workshop's
organizers have
recently developed the
SPRNG library for
parallel random-
number generation.*

◆

think about any special random-number requirements for their application. Todd Urbatsch of the Los Alamos National Laboratory spoke about thermal radiative transfer calculations using time-implicit Monte Carlo on parallel machines. In addition to the algorithmic challenges, this application furnishes some of the most stringent requirements for random numbers. Next, Cornell's Mal Kalos spoke about new Monte Carlo algorithms for solving the fermionic Schrödinger equation. In this talk, he also advocated the use of controllably bad random-number generators (generators constructed with known correla-

tions) to help understand quality in Monte Carlo computations. Columbia University's Pavlos Vranas presented the algorithmics of lattice quantum chromodynamics and described the construction of a 400-gigaflop lattice QCD machine out of commodity digital-signal processors and custom ASICs, a considerable undertaking. James Given of the US National Institute of Standards and Technology concluded this section by talking about a new class of diffusion algorithms for the point solution of elliptic partial-differential equations. After describing the new algorithms, he discussed a wide variety of physical quantities that could be calculated effectively with these new methods

Quasirandom numbers and the future of HPC

The workshop's second day began with three talks on quasirandom numbers. These are numbers that attempt to be as evenly distributed as possible, but unlike the pseudorandom numbers described on the previous day, do not need to pass statistical tests. Los Alamos's Tony Warnock spoke about new results that directly affect the use of quasirandom numbers on parallel machines: error bounds for combining results obtained from different quasirandom streams. Even though quasirandom numbers have quicker convergence than pseudorandom numbers in applications, it turns out that N quasirandom streams combine to reduce the error as $O(N^{1/2})$, just like combining pseudorandom number streams! The next speaker, John Halton from the University of North Carolina, Chapel

Hill, proved that result. Giray Okten of the University of Alaska-Fairbanks then spoke on some new constructions of quasirandom sequences that combine traditional quasirandom sequences with pseudorandom sequences. This series of talks addressed some of the mathematical difficulties behind providing parallel streams of quasirandom numbers. During the extensive discussion on this topic, new results on how to use full-period pseudorandom numbers to provide many related quasirandom streams were presented. This seems like a promising approach that needs more study.

Next, Greg Astfalk, chief scientist of Hewlett Packard's High-Performance System's Division, shared his views on the future of high-performance computing. He made two especially telling points:

- Vector architectural features (such as flat memory, low latency, extremely high sustainable-memory bandwidth, multiple pipes to memory,

hardware support for gather/scatter, efficient single-word access, vector instructions—all the other parts of the vector architecture that matter more than the vector instructions themselves) are dead.

- Almost all vendors seem to be converging on clusters of symmetric multiprocessing systems for their high-end products. These clusters will be based on the same SMPs that make up their server and high-end workstation markets: success will ride on maximally leveraging commodity technology.

The SPRNG tool

The final two speakers spoke about SPRNG and its suitability for this community. NCSA's Ashok Srinivasan described SPRNG in detail, while I presented SPRNG in terms of the design decisions made in its implementation. This led into a broad discussion on whether SPRNG was the right tool for the present and what things an improved

SPRNG should offer in the future. Ashok and I concluded that SPRNG is well-designed for current applications, but that new generators should be designed to better meet the demands of ASCI-class applications. In addition, applications-based testing, coupled with a Web-accessible test site, would serve both the theoretical and applications community. Quasirandom numbers should be made available for parallel and distributed computing, but first a better mathematical understanding of key issues will clearly be required.

Workshop sponsors included the University of Southern Mississippi's Program in Scientific Computing, the Center of Higher Learning and the Programming Environment and Training Program at Stennis Space Center, the NCSA, the University of Illinois at Urbana-Champaign, and Darpa.

Michael Mascagni is an associate professor of mathematics at the University of Southern Mississippi; michael.mascagni@usm.edu.

QUANTUM CATALYSIS

Donald G. Trublar

The recent National Meeting of the American Chemical Society in Dallas included a three-day (March 30–April 1) international symposium on Transition State Modeling in Computational Catalysis, sponsored by the Division of Computers in Chemistry. Forty papers by computational chem-

ists and physicists from 13 countries covered the whole gamut of computational issues involved in current models of catalytic reactions. All kinds of catalysis were represented: acid-base, homogeneous organic and organometallic, metal surfaces, metal oxides and zeolites, and enzymes.

Catalytic reactions are widely used for polymerization and for producing fine chemicals and pharmaceutical molecules. The computational methodologies applied to the wide variety of catalytic processes share a number of common features, providing yet another example of the phenomenon

where
vides
cross
react
port
place
antic
men

Und
reac

Con
grea

actic
but

Mos
cal

lysts
mor

tion
ress

this
ofte

ly, cl
actic

indu
and

men
drug

ing
equa

situ
vanc

C
gine

mol
und

char
on

and,
char

stat
gati

den
bel

pro
con

that