



High Performance Computing: A Tutorial for Scientists

Thomas C. Schulthess, CSCS / ETH Zurich <u>schulthess@cscs.ch</u>

Numerical Approaches to Quantum Many-Body Systems Institute for Pure & Applied Mathematics, UCLA, 01/23/2009

Electronic computing: the beginnings



1943/44: Colossus Mark 1&2 - Britain





1945-51: UNIVAC I

Eckert & Mauchly - "first commercial computer"



1945: John von Neumann report that defines the "von Neuman" architecture



ETH Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Since the dawn of High-performance computing: Supercomputing at Los Alamos National Laboratory

1946: ENIAC 1952: MANIA C I 1957: MANIAC II





Equation of State Calculations by Fast Computing Machines

NICHOLAS METROPOLIS, ARIANNA W. ROSENBLUTH, MARSHALL N. ROSENBLUTH, AND AUGUSTA H. TELLER, Los Alamos Scientific Laboratory, Los Alamos, New Mexico

EDWARD TELLER,* Department of Physics, University of Chicago, Chicago, Illinois (Received March 6, 1953)

A general method, suitable for fast computing machines, for investigating such properties as equations of state for substances consisting of interacting individual molecules is described. The method consists of a modified Monte Carlo integration over configuration space. Results for the two-dimensional rigid-sphere

system have been obtained on the Los Alamos MANIAC and are presented here. These results are compared

to the free volume equation of state and to a four-term virial coefficient expansion.

OF CHEMICAL PHYSICS

1974: Cray 1 - vector architecture

1987: nCUBE 10 (SNL) - MPP architecture 1993: Intel Paragon (SNL)

1993: Cray T₃D

2002: Japanese Earth Simulator - Sputnik shock of HPC

2004: IBM BG/L (LLNL)
2005: Cray Redstorm/XT3 (SNL)
2007: IBM BG/P (ANL)
2008: IBM "Roadrunner"
2008: Cray XT5 (ORNL)

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich





VOLUME 21, NUMBER 6

JUNE, 1953

Outline / goals

- History of scientific computing and HPC
- Overivew of typical HPC architectures
 - Important terminology used later
- Longterm development trends in HPC
 - Understand the brutal facts of HPC today
- Parallel programming models
 - There is more to parallel programming than MPI and OpenMP
- Mapping methods/algorithm onto hardware
 - Importance of computational methematics and computer science
- Example: running QMC/DCA at scale
 - Even simple things can be challenging on 150,000 processors



Scalar / superscalar / vector processors

- Scalar processor: process one data item (integer / floating point number) at a time
- Vector processor: a single instruction operates on many data items simultaneously
- Typical processor today: "pipelined superscalar"
 - Superscalar: simultaneously dispatch multiple instruction to redundant functional units (multiplier or adder)
 - Pipeline: set of processing elements connected in a series
 - Example: 2 multiplies and two add per cycle (4 floating point operations per cycle)

The good news: compiler-level optimization will take care of this!





Distributed vs. shared memory architecture



Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Swiss National Supercomputing Centre

Aspects of performance - typical values in 2009

- Floating point (integer) performance: 2 or 4 per cycle
 - Flop/s = floating point operation per second
 - 2.4 GHz processors: 9.6 GFlop/s
- Memory latency: ~50 ns
- Memory bandwidth: ~10 GB/s
- Network latency ~2-10 µs
- Network bandwidth: ~5 GB/s
- Disk access time ~ ms
- I/O bandwidth ~ MB/s

Cray XT5 node

мD



Outline / goals

- History of scientific computing and HPC
- Overivew of typical HPC architectures
 - Important terminology used later
- Longterm development trends in HPC
 - Understand the brutal facts of HPC today
- Parallel programming models
 - There is more to parallel programming than MPI and OpenMP
- Mapping methods/algorithm onto hardware
 - Importance of computational methematics and computer science
- Example: running QMC/DCA at scale
 - Even simple things can be challenging on 150,000 processors





Computers in the past and today

	1970s (*)	my laptop	improvement
clock (CPU)	6 MHz	2GHz	300 x
Flop/s	6 MFlop/s	8 GFlop/s	10 ³ x
RAM	I 28kB	~2GB	10 ⁵ x
Mem. latency	850ns	~50ns	20 x

(*) Charles Thacker's computer in the 1970s

CSCS Swiss National Supercomputing Centre



Single processor performance is no longer tracking Moore's Law



ETH Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich



Multi-core and heterogeneous processors architectures



Multi-core processors: OpenMP (or just MPI)

NVIDIA G80 GPU: CUDA, cuBLAS





IBM Cell BE: SIMD, threaded prog.



Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Explosion in the number of processing cores



Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Swiss National Supercomputing Centre

Interconnects in the TOP500 systems





Complexity of interconnect



Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Swiss National Supercomputing Centre

Getting ready for Quad Core

- Bytes/flops will decrease
 - XT3 5 GB/sec/2.6 GHZ* 2Flops/clock
 - ▶1 Byte/flop
 - XT4 (dual) 6.25GB/sec/2.6 GHZ* 2Flops/clock/2 processors
 - ► ½ Byte/flop
 - XT4 (quad) 8 GB/sec/2.2GHZ*4Flops/clock/4 processors
 - ►¼ Byte/flop
- Interconnect Bytes/flop will decrease
 - XT3 2 GB/sec/2.6 GHZ* 2Flops/clock
 - ► 1/3 Bytes/flop

XT4 (dual) – 4 GB/sec/2.6 GHZ* 2Flops/clock/2 processors 1/3 Bytes/flop

- XT4 (quad) 4 GB/sec/2.2GHZ*4Flops/clock/4 processors
 - ► 1/10 Byte/flop

HPC in the age of massively parallel processing (MPP) architectures: what does this really mean?

Evolution of the fastest sustained performance in real simulations

```
~1 Exaflop/s
~10<sup>7</sup> processing units
```

1.35 Petaflop/s
Cray XT5
1.5 105 processor cores1.02 Teraflop/s
Cray T3D
1.5 103 processors1.5 Gigaflop/s
Cray YMP
0.8 101 processors

1998

2008

2018

1989



Summary: Brutal fact of modern HPC

- Mind boggling numbers of processing processing units
- Processor complexity (multi-core, heterogeneous)
- Interconnect is a non-trivial part of the HPC system
- Accessing memory is prohibitively expensive compared to the cost of floating point operations
 - 1960s: transistors were expensive, memory access was cheap
 - today: transitors are cheap, memory access is expensive

Key aspect of programming in HPC systems: All about anaging resources



Outline / goals

- History of scientific computing and HPC
- Overivew of typical HPC architectures
 - Important terminology used later
- Longterm development trends in HPC
 - Understand the brutal facts of HPC today
- Parallel programming models
 - There is more to parallel programming than MPI and OpenMP
- Mapping methods/algorithm onto hardware
 - Importance of computational methematics and computer science
- Example: running QMC/DCA at scale
 - Even simple things can be challenging on 150,000 processors





Programming models (I): message passing



- Concurrent sequential processes cooperating on the same task
- Each process has own private space
- Communication is two-sided through send and receive
 - Large overhead!
- Lots of flexibility in decomposing large problems, however, provides only fragmented view of the problem
 - All burden is placed on the programmer to maintain global view
- Examples are message passing libraries like MPI or PVM



Programming models (II): shared memory



- Multiple independent threads operate on same shared address space
- Easy to use since there is only one type of memory access
 - One-sided remote access (low overhead)
- Application view remains integrated (global view)
- Shared memory hardware doesn't scale (local & remote memory access)
- It is difficult to exploit inherent data locality - degradation of performance!
- Examples are OpenMP or Pthreads
 - Compiler directive used with C, Fortran, ...



Programming models (III): data parallel



- Concurrent processing of many data elements in the same manner
- Executing only one process (on many processors)
- Major drawback: does not permit independent branching
 - Not good for problems that are rich in functional parallelism
- Popular examples are C* and HPF





Programming models (IV): distributed shared memory

.

- Also called partitioned global address space (PGAS) model
- Independed threads operate in shared memory space
 - preserve global view of program
- Shared space is locally partitioned among threads
 - allows exploiting data locality
- "Single program multiple data stream" (SPMD) execution
 - independent forking (functional parallelism)
- Popular examples: UPC and co-Array Fortran
- May still not have the same flexibility as Message Passing Model



Distributed shared memory or PGAS: keeping the best from all other models





ETH Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich



Unified Parallel C (UPC)

see http://upc.lbl.gov

```
#include <upc.h>
#define SZ 20
main(){
  static shared int array[SZ];
  shared int *ar ptr;
  static shared int step=10;
  int argument, i;
  ar ptr = array + MYTHREAD;
  upc forall (i=0; i<SZ; i++; i) {</pre>
    argument = step*i;
    *ar_ptr = some_function(argumet);
    ar ptr += TREADS;
  upc barrier;
  if(MYTREAD==0) {
    ar ptr = array;
```

}

}

for (i=0; i<SZ; i++, ar_ptr++) {
 argument = step*i;
 printf ("%d \t %d \n", *ar_ptr, argument)</pre>

- Simple extension to ISO C to implement PGAS model in SPMD mode of programming
- Available on most systems (as extension to gcc or LBL UPC compiler)
- "Simpler" to program than MPI or OpenMP (!)
- Limitation: supports only onedimensional decomposition of shared arrays
 - more complex decomposition are supported by DARPA HPCS languages (Chapel, X10, Fortress)



Outline / goals

- History of scientific computing and HPC
- Overivew of typical HPC architectures
 - Important terminology used later
- Longterm development trends in HPC
 - Understand the brutal facts of HPC today
- Parallel programming models
 - There is more to parallel programming than MPI and OpenMP
- Mapping methods/algorithm onto hardware
 - Importance of computational methematics and computer science
- Example: running QMC/DCA at scale
 - Even simple things can be challenging on 150,000 processors





Sketch of the Dynamical Cluster Approximation



Solve many-body problem on cluster > Essential assumption: Correlations are short ranged

DCA method: self-consistently determine the "effective" medium



Hirsch-Fye Quantum Monte Carole (HF-QMC) for the quantum cluster solver Hirsch & Fye, Phys. Rev. Lett. 56, 2521 (1998)

Partition function & Metropolis Monte Carlo $Z = \int e^{-E[\mathbf{x}]/k_{\rm B}T} d\mathbf{x}$ Acceptance criterion for M-MC move: $\min\{1, e^{E[\mathbf{x}_k] - E[\mathbf{x}_{k+1}]}\}$

Partition function & HF-QMC: $Z \sim \sum_{s_i,l} \det[\mathbf{G}_c(s_i,l)^{-1}]$ $N_c \qquad N_l \approx 10^2$ $N_t = N_c \times N_l \approx 2000$

Acceptance: $\min\{1, \det[\mathbf{G}_{c}(\{s_{i}, l\}_{k})] / \det[\mathbf{G}_{c}(\{s_{i}, l\}_{k+1})]\}$



Update of accepted Green's function:

 $\mathbf{G}_c(\{s_i,l\}_{k+1}) = \mathbf{G}_c(\{s_i,l\}_k) + \mathbf{a}_k \times \mathbf{b}_k$

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich

Swiss National Supercomputing Centre



CSCS

HF-QMC with delayed updates (or Ed updates)

 $\mathbf{G}_{c}(\{s_{i},l\}_{k+1}) = \mathbf{G}_{c}(\{s_{i},l\}_{0}) + [\mathbf{a}_{0}|\mathbf{a}_{1}|...|\mathbf{a}_{k}] \times [\mathbf{b}_{0}|\mathbf{b}_{1}|...|\mathbf{b}_{k}]^{t}$

Complexity for *k* updates remains $O(kN_t^2)$

But we can replace *k* rank-1 updates with one matrix-matrix multiply plus some additional bookkeeping.





Performance improvement with delayed updates



 $N_c = 16$ $N_l = 150$ $N_t = 2400$





Swiss Federal Institute of Technology Zurich

Swiss National Supercomputing Centre



Outline / goals

- History of scientific computing and HPC
- Overivew of typical HPC architectures
 - Important terminology used later
- Longterm development trends in HPC
 - Understand the brutal facts of HPC today
- Parallel programming models
 - There is more to parallel programming than MPI and OpenMP
- Mapping methods/algorithm onto hardware
 - Importance of computational methematics and computer science
- Example: running QMC/DCA at scale
 - Even simple things can be challenging on 150,000 processors







DCA++ code from a concurrency point of view





DCA++: strong scaling on HF-QMC



Eidgenössisch: Swiss Federal Institute of Technology Zurich

Swiss National Supercomputing Centre

Weak scaling on Cray XT4

Eidgenössische Technische Hochschule Zürich

Swiss Federal Institute of Technology Zurich

- HF-QMC: 122 Markov chains on 122 cores
- Weak scaling over disorder configurations



Swiss National Supercomputing Centre

Cray XT5 portion of Jaguar @ NCCS

Peak: 1.382 TF/s Quad-Core AMD Freq.: 2.3 GHz 150,176 cores Memory: 300 TB For more details, go to WWW.nccs.gov

Sustained performance of DCA++ on Cray XT5

Weak scaling with number disorder configurations, each running on 128 Markov chains on 128 cores (16 nodes) - 16 site cluster and 150 time slides

Eidgenössische Technische Hochschule Zürich Swiss Federal Institute of Technology Zurich Number of Cores

New algorithm to enable 1+ PFlop/s sustained performance in simulations of disorder effects in high-*T_c* superconductors

T. A. Maier P. R. C. Kent

T. C. Schulthess

G. Alvarez

M. S. Summers

- E. F. D'Azevedo
- J. S. Meredith
- **M. Eisenbach**
- **D. E. Maxwell**
- J. M. Larkin
- **J. Levesque**

Physics Software Comp. mathematics **Computer Science Computer Center** Hardware vendo