QUAKES

A Parallel Implementation of the Lattice Solid Model

Steffen Abe, David Place and Peter Mora

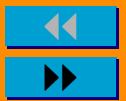


Overview

- Motivation
- Program design
- Benchmarks
- Conclusions




Back
Close

Motivation

- Larger models for increased realism
  - 3D
  - Particle shapes
  - Range of particle sizes
- Single CPU peak performance 1-5 GFlops, doubling every 18 months
- Parallel Computers: tens of GFlops to several TFlops, doubling every 12 months




Basic design decisions

- Object oriented (C++)
- Message passing architecture (MPI)
  - easy to maintain and extend
  - portable (IRIX, Linux, Compaq True64, Windows)
- Split into front-end (GUI, Parser, . . .) and backend (computation)
  - run backend on HPC
  - front-end on Desktop PC/Workstation or high end visualization system
  - attach and remove front-ends dynamically to running backend




Parallelization

- Modified Master/Slave model
- Geometric partition of the model
- Shared boundary particles
- Force computation completely local
- Position of shared particles exchanged once per time step




Back
Close

Benchmarks

- Measures of parallel performance
  - Parallel runtime t_P
  - Speedup $S = t_P/t_S$ (t_S serial runtime)
  - Parallel efficiency $E = S/n$ (n number of CPUs)
- Benchmarks
  - regular triangular lattice of particles
  - various problem sizes
  - constant total problem size
  - constant problem size per CPU




Constant total problem size - runtime

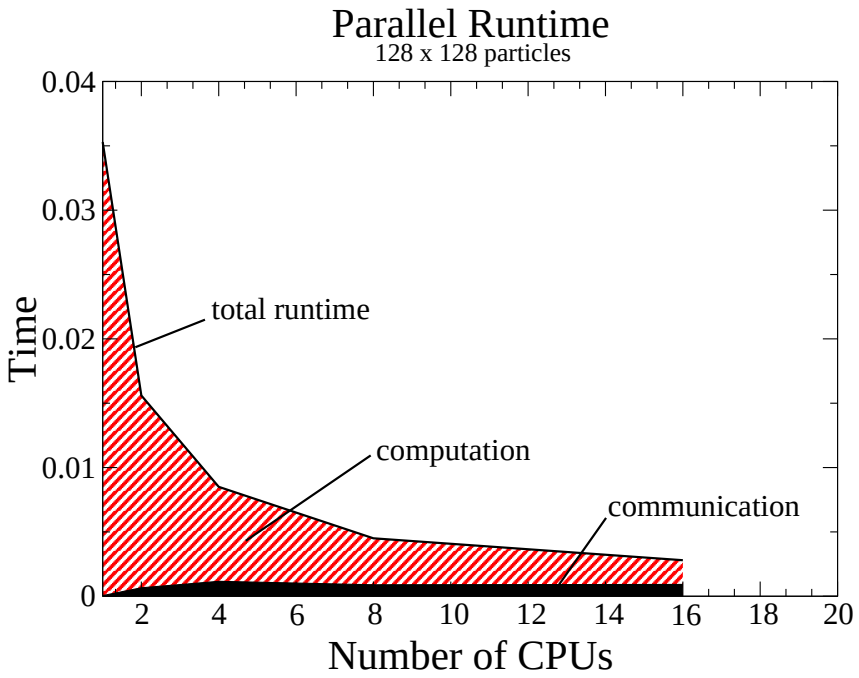

Constant total problem size - speedup

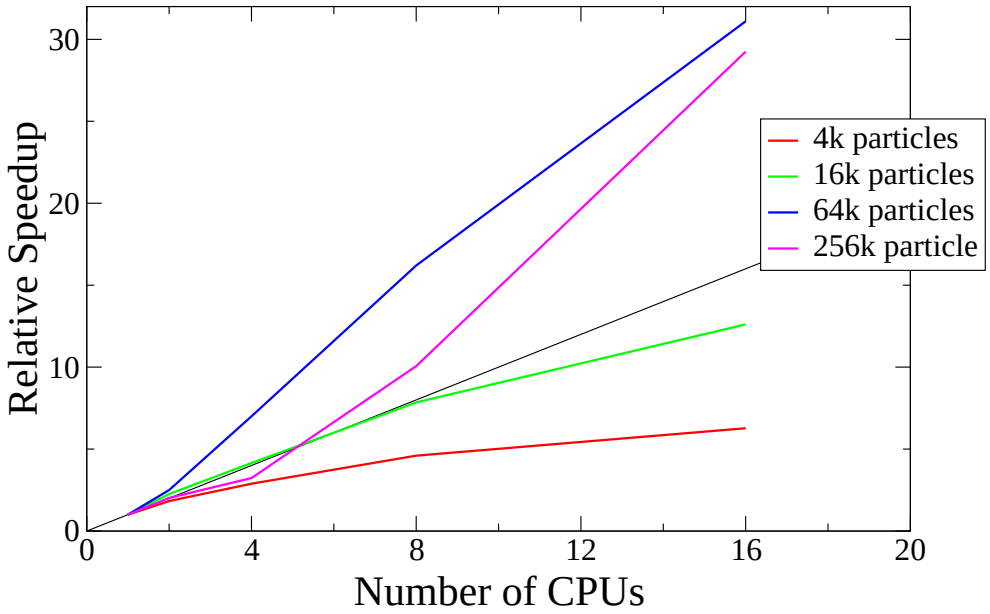

Constant total problem size - parallel efficiency

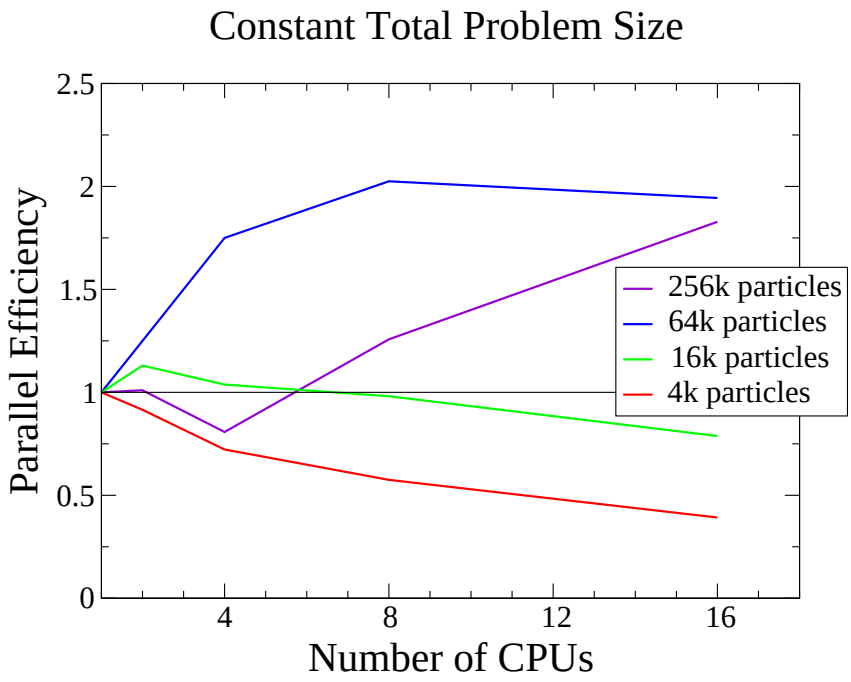

Single CPU performance

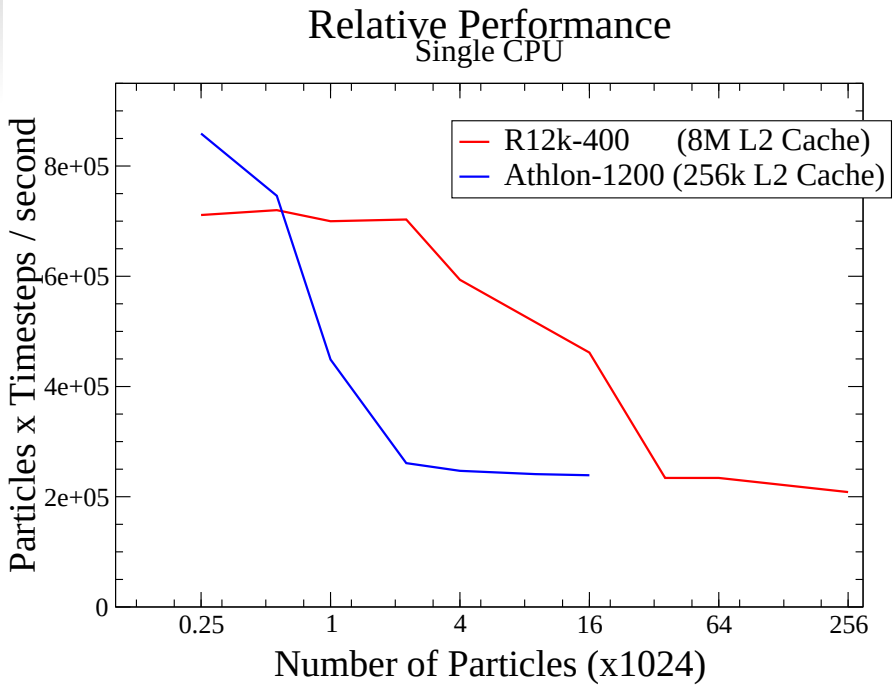

Constant problem size per CPU - parallel efficiency

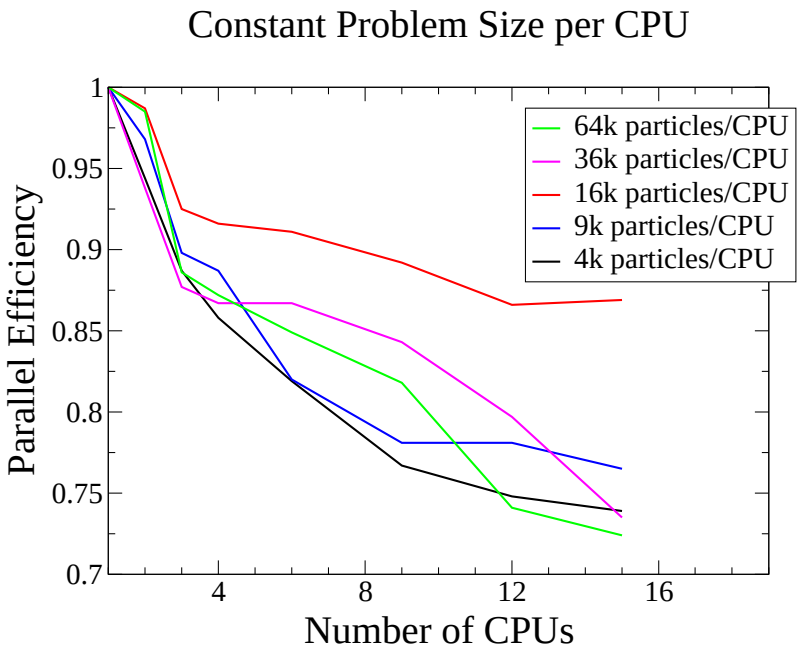

Constant problem size per CPU - runtime

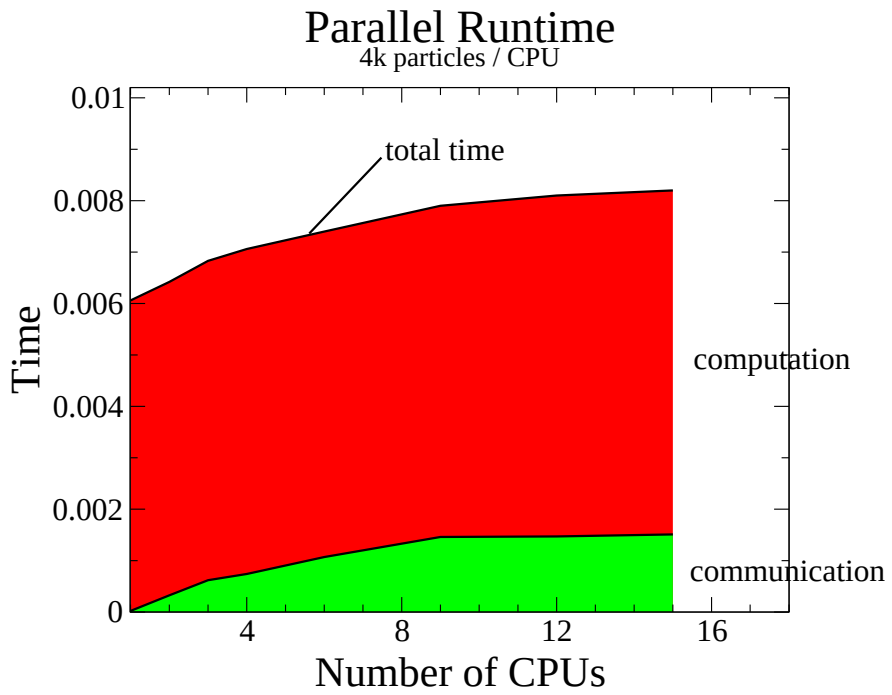

Constant problem size per CPU

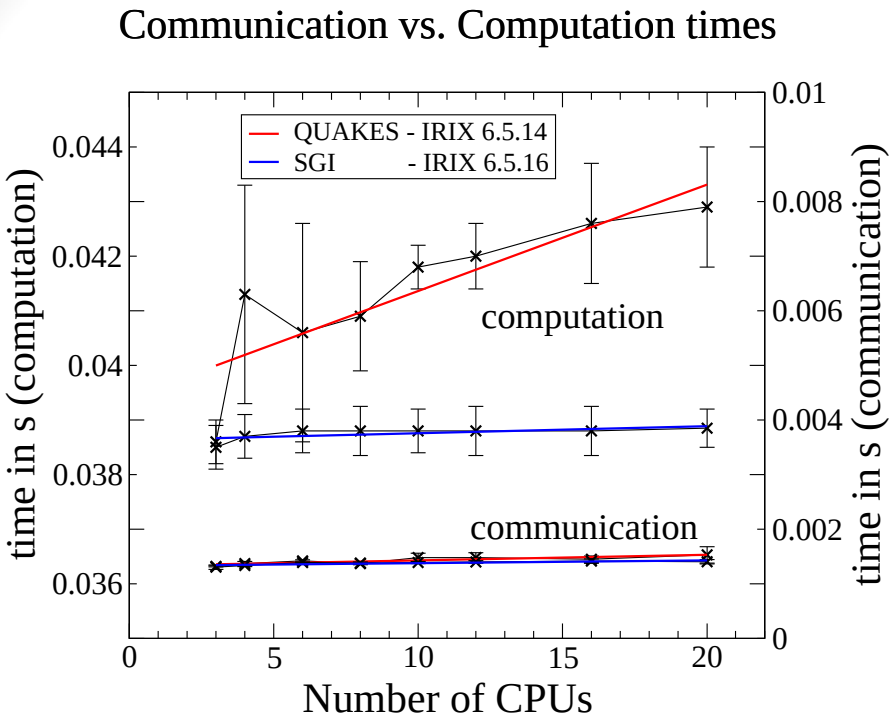

“Large” Benchmark

- constant problem size per CPU
- 256 $\times$ 256, wave propagation
- testing primarily force computation and integration
- SGI Benchmarks
  - run on 128 CPU SGI 3800 (SGI)
  - using 1-120 CPUs
- Compaq Benchmarks
  - run on 128 CPU AlphaServer SC40 (VPAC)
  - using 1-25 CPUs (loaded system)




Runtimes

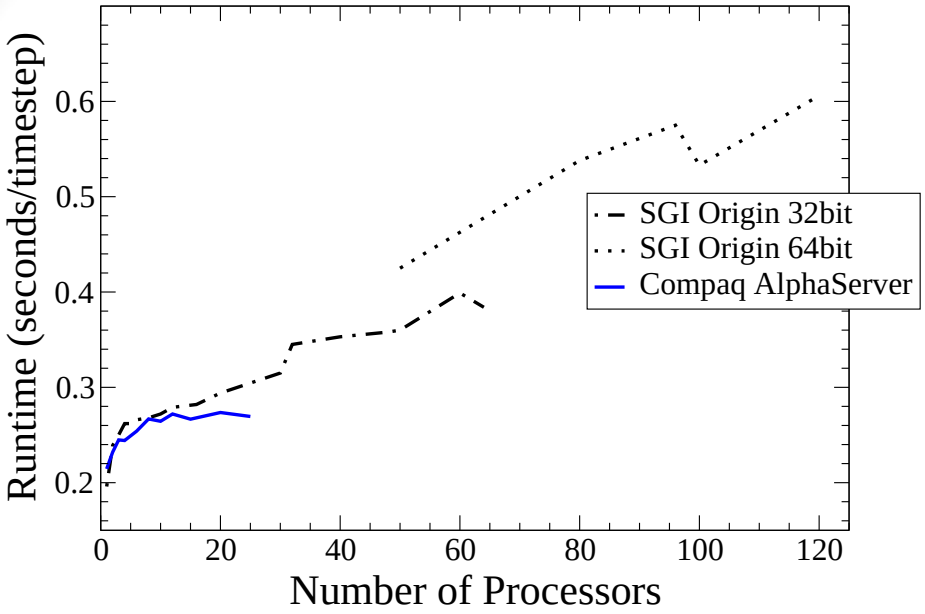
Runtime (seconds/timestep)
0.2
0.3
0.4
0.5
0.6
0
20
40
60
80
100
120
Number of Processors
SGI Origin 32bit
SGI Origin 64bit
Compaq AlphaServer

Speedup

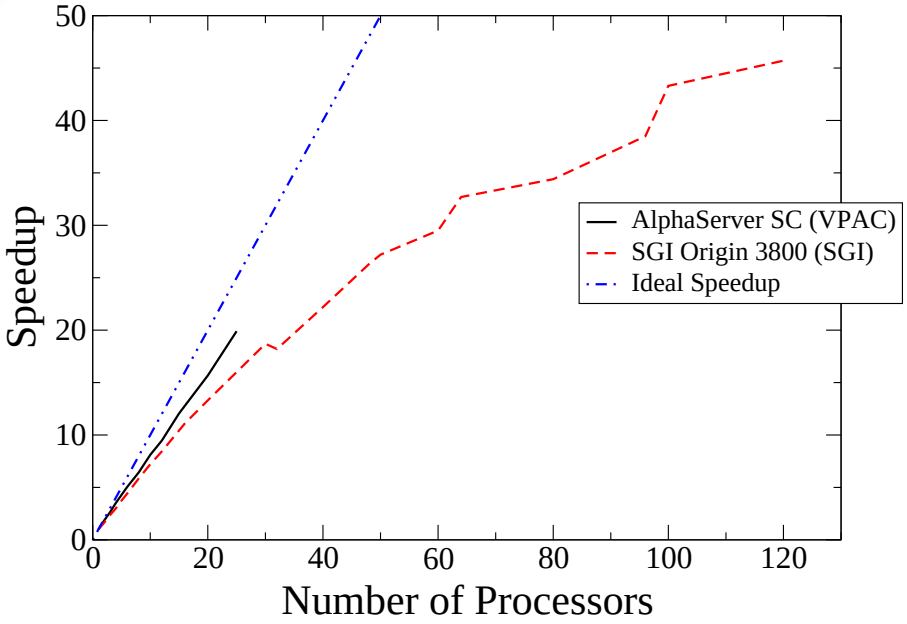
Speedup
Number of Processors
0
10
20
30
40
50
0
20
40
60
80
100
120
AlphaServer SC (VPAC)
SGI Origin 3800 (SGI)
Ideal Speedup

Conclusions

- New implementation of the LSM enables use of large parallel computer systems for the simulation of large models
- Program design (C++,MPI) provides flexibility and portability
- Very good scalability can be achived on different computer architectures