

Center for Information Services and High Performance Computing (ZIH)

Analysis and Optimization of a Molecular Dynamics Code using PAPI and the Vampir Toolchain

May 2, 2012

Thomas William Zellescher Weg 12 Willers-Bau A 34 +49 351 - 463 32446

Thomas.William@ZIH.TU-Dresden.de



Introduction

- 2 Serial Analysis
- 3 PAPI measurements
- 4 Source Code Analysis
- 5 Source Code Optimization
- 6 Tracing and Visualization













- IU
- ZIH
- FutureGrid
- MD







- Classical molecular dynamics simulation of dense nuclear matter consisting of either fully ionized atoms or free neutrons and protons
- Main targets are studies of the dense matter in white dwarf and neutron stars.
- Interaction potentials between particles treated as classical two-particle central potentials
 - No complicated particle geometry or orientation
- Electrons can be treated as a uniform background charge (not explicitly modeled)





MD Code Details

- Particle-particle-interactions have been implemented in a multitude of ways
- Located in different files PP01, PP02 and PP03
- The code blocks are selectable using preprocessor macros
 - **PP01** is the original implementation with no division into the Ax, Bx or NBS blocks
 - PP02 implements the versions in use by the physicists today
 - 3 different implementations for the Ax block
 - 3 implementations for the Bx block
 - No manual blocking (NBS)
 - PP03 includes new routines
 - 3 Ax blocks
 - 8 Bx blocks
 - Can be blocked using the NBS value





- Two sections of code each have two or more variations
- One section is labelled A and the other B
- Variations are numbered
- MD can be compiled as a serial, OpenMP, MPI or MPI+OpenMP

make MDEF=XRay md_mpi ALG=PP02 BLKA=A0 BLKB=B2 NBS="NBSX=0"





The structure of the algorithm is simple.

- First reads in a parameter file runmd.in
- And an initial particle configuration file md.in
- Program then calculates the initial set of accelerations
- Enters a time-stepping loop, a triply nested "do" loop





Runtime Parameters

!Parameters :								
= ' i o	simulation type!							
=	0.00,	!start time						
=	25.00,	!time step (fm/c)						
=	2,	!groups						
=	50,	!steps per group						
!Measurement :								
=	2,	!groups						
=	2,	!per group						
=	25,	!steps between						
ze =	50,	!temp normal.						
=	50,	!center-of-mass						
		!motion cancel.						
	: = 'io = = = = = t: = = ze =	: = 'ion-mixture', = 0.00, = 25.00, = 2, = 50, nt: = 2, = 2, = 2, = 2, = 50, ze = 50, = 50,						

Figure : Runtime parameters, snippet from runmd.in





The Main Loop

```
do 100 ig = 1, ngroup
       initialize group ig statistics
       do 40 j=1, ntot
          do i=1, nind
              !computes forces
              lupdates x and v
              call newton
          enddo
          call vtot
       enddo
       compute group ig statistics
40
       continue
100 continue
```

Figure : Simplified version of the main loop





- Forces are calculated in the newton module in a pair of nested do-loops
 - Outer loop goes over target particles
 - Inner loop over source particles
- Targets are assigned to MPI processes in a round-robin fashion
- Within each MPI process, the work is shared among OpenMP threads











- Cray XT5m provided by the FutureGrid project
- XT5m is a 2D mesh of nodes
- Each node has two sockets each having four cores
- AMD Opteron 23 "Shanghai" (45mm) running at 2.4 GHz
- 84 compute nodes with a total of 672 cores
- pgi/9.0.4 using the XT PE driver xtpe-shanghai





5k particles measurement takes
27k particles measurement takes
55k particles measurement takes
10 hours





Overview PP01, PP02, and PP03



Figure : Overview of the runtimes for all (143) code-block combinations with an ion-mix input dataset of 55k particles. The naming scheme for the measurements is "source-code-file A-block B-block blockingfactor"





PP02 Code Blocks



Runtime, 55k particle run

TECHNISCHE UNIVERSITÄT DRESDEN

Center for Information Services 8

High Performance Computing

-	
##	do 90 j=i+1,n-1
##	! Block A
## #if defin	ned(A0)
##	r2=0.0d0
movsd	%xmm2, %xmm1
movq	%r12, % rdx
movq	%r15, % rcx
movl	\$8, % eax
.align	16
.LB2_555:	
## lineno:	138





```
##
      do 90 j=i+1,n-1
##
      !----- Block A ------
## #if defined(A0)
##
        r2=0.0d0
##
        do k=1.3
##
           xx(k)=x(k,i)-x(k,j)
 movlpd .BSS2+48(%rip),%xmm2
 movlpd .C2 291(%rip),%xmm0
 mulsd
         %xmm2,%xmm2
 addsd
        %xmm1.%xmm2
 movlpd
         %xmm2,344(%rsp)
 sqrtsd
         %xmm2.%xmm2
 balvom
         %xmm2,448(%rsp)
 mulsd
         md globals 10 +120(%rip),%xmm2
 subsd
         %xmm2,%xmm0
 .p2align
         4..1
```

ECHNISCHE

RESDE









Floating Point Instructions



15/34

FPU idle times

FPU idle time in %





Branch Miss Predictions









```
#if defined(A0)
  r_{2}=0.000
  do k=1.3
      xx(k) = x(k, i) - x(k, j)
      if (xx(k), gt + half(k)) xx(k) = xx(k) - x(k)
      if (xx(k), |t, -ha|f|(k)) xx(k) = xx(k) + x|(k)
      r_{2}=r_{2}+x_{x}(k) * x_{x}(k)
  enddo
#elif defined(A1)
  r_{2}=0.0d_{0}
  do k=1.3
      xx(k) = x(k, i) - x(k, j)
      xx(k) = xx(k) - aint(xx(k) * halfli(k)) * xl(k)
      r2=r2+xx(k)*xx(k)
  enddo
#elif defined(A2)
  xx(:) = x(:,i) - x(:,j)
  xx=xx-aint(xx*halfli)*xl
  r_{2}=xx(1)*xx(1)+xx(2)*xx(2)+xx(3)*xx(3)
#else
```





TECHNISCHE

DRESDEN

```
#if defined(B0)
  r = sart(r2)
  fc = exp(-xmuc*r)*(1./r+xmuc)/r2
  do k=1.3
     fi(k) = fi(k) + zii(i) * fc * xx(k)
     fi(k,i) = fi(k,i) - zii(i) * fc * xx(k)
  enddo
#elif defined(B1)
  r = sqrt(r2)
  fc = exp(-xmuc*r)*(1./r+xmuc)/r2
  fi(:) = fi(:) + zii(i) * fc * xx(:)
  f_{i}(:,i) = f_{i}(:,i) - z_{ii}(i) * f_{c} * x_{x}(:)
#elif defined(B2)
  if(r2.le.rcutoff2) then
     r = sqrt(r2)
     fc = exp(-xmuc*r)*(1./r+xmuc)/r2
     fi(:) = fi(:) + zii(i) * fc * xx(:)
     f_i(:, i) = f_i(:, i) - z_{ii}(i) * f_c * x_x(:)
  endif
#else
```









OpenMP Overhead in Vampir



Code Changes

- omp parallel do for the j-loop is removed
 - omp parallel region around the i-loop
- j-loop is parallelized explicitly
 - Entire ij-loop section in a 3rd loop over threads
- Each iteration of this outer loop is assigned to an OpenMP thread





New MPI/OpenMP Version in Vampir







Parallel Efficiency of the OpenMP Application



High Performance Computing

Parallel Efficiency of the MPI Version



Parallel Efficiency (MPI only)

Center for Information Services 8

High Performance Computing

Parallel Efficiency of the Hybrid Version



Parallel Efficiency (MPI+OpenMP)

TECHNISCHE UNIVERSITÄT DRESDEN













8 core inter-node run of the MPI-only version of MD





8 core inter-node run of the MPI-only version of MD



Table : MPI-only, PP02 A0 B2 run, all processes, accumulated exclusive time per function

	672 cores		8 cores	
function name	time in s	%	time in s	%
accel_ion_mix	13293.4	72.7	11735.0	97.9
MPI_Allreduce	1885.8	10.3	44.1	0.4
sync	939.4	5.1	0.1	0
newton	937.2	5.1	4.9	0.0
MPI_Bcast	807.1	4.4	8.2	0.1
vtot_ion_mix	189.6	1.0	188.2	1.6
MPI_Init	112.8	0.6	0.1	0
MPI_Barrier	107.5	0.6	1.2	0











Conclusion

- The analysis of MD found the best serial version
- Using a cut-off sphere reduces interactions to 19%
- Runtime decrease although the performance counter indicate otherwise
- Moved the OpenMP parallelization from the inner to the outer loop
- Improved the parallel efficiency
- Dual socket, dual chip 16-threads AMD Interlagos system shows a 97.6% parallel efficiency up to 32 cores

The changes to the source code will be included in future versions of the Spec OpenMP benchmark.





Thank you for your attention.







AMD Interlagos OpenMP Efficiency





Kraken - Old Version



- Code has been extended and rewritten multiple times
- Several different implementations of the same semantics
 - Loop reimplemented using the Fortran array syntax
 - Cut-off sphere to restrict interactions to a subset of relatively nearby nucleons/ion
- User must specify which variation in the force calculation routines to use when building the code
 - nucleon
 - pure-ion
 - ion-mixture





672 core run of the MPI-only version of MD





8 core inter-node run of the MPI-only version of MD







