

MPII - performance tests

C. Kleinwort

DESY/HH tracker alignment chat 19.10.10

Overview

- ★ Bottleneck for parallelization
 - Reading data from disk (PEREAD)
- ★ Mixed Precision Storage
 - Check results vs full (double precision) storage
- ★ Test case: bows&kinks with CRAFT10+MinBias
 - 197k parameters, 4.6M tracks, 7 threads
 - Global matrix 19GB (mixed), 48GB (full)

Reading binary files

- ★ LOOP1, LOOP2

- Each once to detect global parameters/pairs

- ★ LOOPN, iterated


- At least twice to calculate change in function value

- ★ Test case

- 120 binary files with ~ 18GB
- read 2+17 times for ⑥ MP iterations
- ~ 340GB in total

PSI, X5540, 2.83 GHz, 24GB


(F. Meier)



~2h reading
binary files




UHH: X5520, 2.23GHz, 96GB

Courtesy Univ. HH, Si det. lab.



Stability of results

Reference:
double precision, single thread



Δp [μm]

MP rev 68

- ★ Bug fix

- ▶ for uncompressed sparsity info (since rev 65)

- ★ Support for zlib

- ▶ read directly from gzipped binary Cfiles
- ▶ gzipped binaries about factor 2 smaller
- ▶ decompression on the fly: 75-100 MB/s
- ▶ beneficial if decompression faster than disk i/o
- ▶ better: parallelization of decompression
 - ✦ needs parallelization of loop over Cfiles

Summary

★ Reading binary files

- ▶ slow ($\sim 50\text{MB/s}$), limits MP performance (wall time)
- ▶ large improvement (>30) from caching (by OS)
 - ✦ needs a lot more memory (several 10GB)
- ▶ small improvement (≤ 2) with gzipped files possible
 - ✦ needs parallelization of 'Cfiles loop'

★ Stability of results: RMS $\sim 0.01\ \mu\text{m}$

- ▶ vs mixed precision storage and parallelization