Scientific Computing

Performance and Efficiency in Climate Models

Sandra Schröder, Michael Kuhn, Nathanael Hübbe, Julian Kunkel,
Timo Minartz, Petra Nerge, Florens Wasserfall, Thomas Ludwig

Scientific Computing
Department of Informatics
University of Hamburg

2012-02-17

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

informatik
die zukunft

Outline

Our Group's Research

- Research Group *Scientific Computing*, University of Hamburg, located at the German Climate Computing Center (DKRZ)
- High Performance Computing, with a focus on climate research
- Communication between climate researchers and computer scientists

Performance and Efficiency

- Performance analysis of climate models is necessary
- Climate researchers' view: interested in results, optimizations might influence their accuracy
- Computer scientists' view: make the best use of the available memory and processor resources
- Many tools are available that can help finding potential bottlenecks

Analysis of a Climate Model

- Climate models are a good example for number crunching applications
  - Challenging in terms of parallelization, I/O and optimization
- Analysis of GETM:
  - **G**eneral **E**stuarine **T**ransport **M**odel
  - Visualizing and analysing communication and I/O patterns with *HDTrace*
  - Analysis of code segments with *gprof*
- Test environment:
  - GETM v2.0
  - Testcase: box_cartesian
  - MPICH2 v1.4
  - NFS

Analysis of a Climate Model - box_cartesian

- Contains a little box with a parabolic bottom topography
- Calculation domain:
  - Equi-distant plane grid
  - 100×30 grid points
  - Grid spacing: $dx = 1000$ m, $dy = 1000$ m
  - Period: 24 h, timestep: 10 s

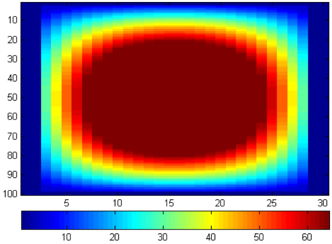

Figure: Topography of box_cartesian.

High-Level Optimizations

- High-Level Analysis
  - Making a rough estimate of the performance and efficiency of the climate model
  - Includes the analysis of I/O patterns and communication within all processes of a parallel program
- Analysis with Sunshot
  - Tool for inspecting the communication between several processes in parallel programs
  - Also supports the visualization of NetCDF functions and underlying I/O system calls

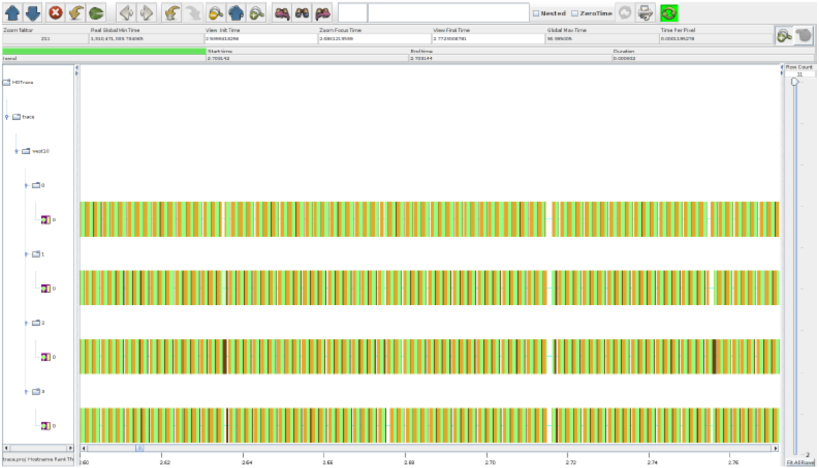

Figure: Timelines of four processes in Sunshot. Dominating functions are synchronisation (`nf90_sync`) and asynchronous communication (`WAIT_ALL`).

High-Level Optimizations

- Large number of `nf90_sync` calls
  - Nearly 96% of the whole I/O
  - Total runtime with synchronization: 115 s
  - Total runtime without synchronization: 35 s
- `MPI_WAITALL` accounts for 85% of the duration of MPI communications
  - Asynchronous communication in MPI
  - We have to
    - Get to know communication patterns in GETM better
    - Discuss *if* and *how* the asynchronous communication can be handled more efficiently

Low-Level Optimizations

```
do j=jmin,jmax
  do i=imin,imax
    C(i,j)=1/4*(A(i  ,j  )/B(i  ,j  )&
               +A(i+1,j  )/B(i+1,j  )&
               +A(i  ,j-1)/B(i  ,j-1)&
               +A(i+1,j-1)/B(i+1,j-1))
  end do
end do
```

- Many sequences for optimizations were found with gprof
- Example: multiple calculations in a loop
- For each point (i, j) all quotients $A(i, j)/B(i, j)$ are calculated four times
- One solution: performing every calculation only once and buffering the results
- Typical time-memory tradeoff

Low-Level Optimizations

```
do j=jmin,jmax
  do i=imin,imax
    D(i,j)=A(i,j)/B(i,j)
  end do
end do

do j=jmin,jmax
  do i=imin,imax
    C(i,j)=1/4*(D(i  ,j  )&
               +D(i+1,j  )&
               +D(i  ,j-1)&
               +D(i+1,j-1))
  end do
end do
```

Summary and Outlook

- Analysis tools are useful to detect performance problems of (parallel) scientific applications
- Performance analysis of the climate model GETM with Sunshot and gprof
- Optimizations presented here are only a fraction of all the potential ones

- Meetings with GETM developers for information exchange between computer scientists and climate researchers
- To do: modify GETM to use ADIOS
  - ADaptable IO System
  - Simplified API for I/O operations
  - How can this affect I/O performance in GETM?