Data-Intensive High-Performance Computing

Michael Kuhn, Anastasiia Novikova, Jannek Squar
michael.kuhn@informatik.uni-hamburg.de

Scientific Computing
Department of Informatics
Universität Hamburg
https://wr.informatik.uni-hamburg.de

2018-06-26

Universität Hamburg
DER FORSCHUNG | DER LEHRE | DER BILDUNG

About Us: Scientific Computing

The group Scientific Computing conducts research and development on high performance storage systems. We develop HPC concepts and apply them to simulation software with a focus on earth system models.

We are an Intel Parallel Computing Center for Lustre.
(“Enhanced Adaptive Compression in Lustre”)

High Performance Storage Systems

- Data reduction
  - **SCIL: Application-specific, lossy compression**
  - Lustre: File system, adaptive lossless compression
  - BigStorage: Energy-efficient data reduction
- Performance
  - I/O monitoring and optimization (15:15)
- I/O interfaces
  - **JULEA: Flexible storage framework for HPC** (15:35)
  - ESDM: Earth system data, HDF5

SCIL

- Scientific Compression Interface Library (SCIL)
  - Design of quantities, tools, compression chain
  - HDF5 plug-in for transparent use
- Domain-specific compression (ratio > 10 : 1)
  - Investigate metrics allowing to define accuracy per variable
  - Design user-interfaces for specifying accuracy
  - Develop a methodology for identifying the required accuracy
  - Implement compression schemes exploiting this knowledge
- Evaluation on synthetic and scientific data (ECHAM, Isabel)
- Part of WP2 of the AIMES project
  - `https://wr.informatik.uni-hamburg.de/research/projects/aimes/start`

SCIL

Accuracy quantities:

absolute tolerance: compressed can become true value ± absolute tolerance

relative tolerance: percentage the compressed value can deviate from true value

relative error finest tolerance: value definining the absolute tolerable error for relative compression for values around 0

significant digits: number of significant decimal digits

significant bits: number of significant decimals in bits

Performance quantities:

compression speed: in MiB/s or GiB/s, or relative to network or storage speed

decompression speed: in MiB/s or GiB/s, or relative to network or storage speed

Supplementary quantities:

fill value: a value that scientists use to mark special data points

SCIL

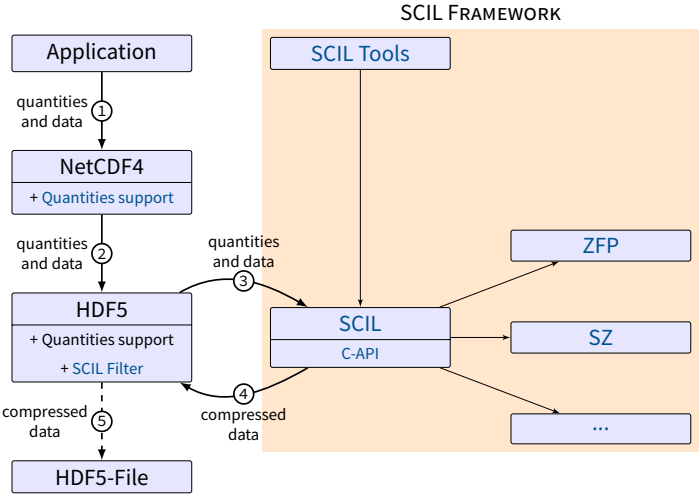
SCIL Framework
Application
quantities and data
1
NetCDF4
+ Quantities support
quantities and data
2
HDF5
+ Quantities support
+ SCIL Filter
compressed data
5
HDF5-File
quantities and data
3
compressed data
4
SCIL Tools
SCIL
C-API
ZFP
SZ
...

JULEA

- JULEA provides a flexible storage framework
  - Contains necessary building blocks for storage systems
  - Facilitates rapid prototyping and evaluation
  - Many projects implement basic functionality from scratch
- Runs completely in user space and has few dependencies
  - Easy to debug and develop
  - Possible to use on clusters without root access
- Feedback and contributions are always welcome
  - First projects using it are running already
  - `https://github.com/wr-hamburg/julea`

JULEA...

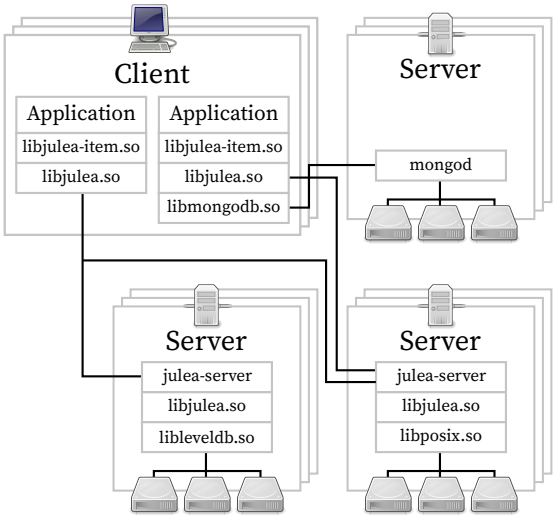
Client
Application
libjulea-item.so
libjulea.so
Application
libjulea-item.so
libjulea.so
libmongodb.so
Server
mongod
Server
julea-server
libjulea.so
libleveldb.so
Server
julea-server
libjulea.so
libposix.so

HPC Concepts and Infrastructure

- Domain-Specific Languages
  - GGDML: Definition and manipulation of grids
  - FortranTestGenerator: Unit tests with capture and replay
- Energy efficiency
  - ArduPower: Open Source, based on Arduino
- Compiler optimizations
  - **DasTool: Translation from OpenMP to MPI-3 RMA** (14:50)
  - MPI-Checker: Static analysis for MPI
- Spack: Software management for supercomputers

DasTool

- Tool to automatize parallelization code modification
  - Probably not as effective as handwritten code ...
  - ... but no knowledge of MPI required by user
- Workflow:
  1. Exchange compiler calls by tool wrapper
  2. Rebuild application
  3. Adapt problem size to utilize additional nodes
  4. Run application as usual

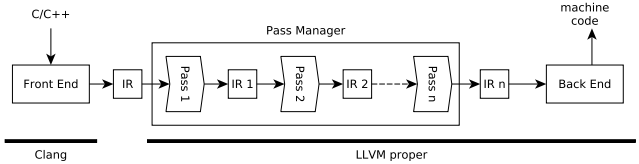

Figure: Pass integration, based on [1]

HPC Education

- PeCoH: Performance engineering and HPC certification (14:30)
- Lectures about high-performance computing and high-performance I/O
- Seminars and software labs in cooperation with geoscience

- Student Cluster Competition
  - Fifth year in a row
  - Visit us at booth B-1366
  - Follow us on Twitter: @UHH_ISC_SCC
  - Vote for us: `https://uhh.de/scc18`

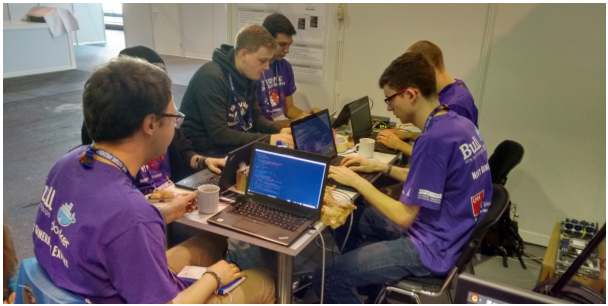

References I

[1] Adrian Sampson. LLVM for Grad Students.
http://www.cs.cornell.edu/~asampson/blog/llvm.html.