Improving Energy Efficiency of Scientific Data Compression with Decision Trees

ENERGY 2020

Jun.-Prof. Dr. Michael Kuhn
michael.kuhn@ovgu.de

2020-09-28 – 2020-10-01

Parallel Computing and I/O
Institute for Intelligent Cooperating Systems
Faculty of Computer Science
Otto von Guericke University Magdeburg
http://parcio.ovgu.de

Motivation

- Scientific applications and large-scale experiments generate a deluge of data
- Gap between the computational power and other hardware is widening [2]
  - More data can be computed than can be stored efficiently
  - Additional investments for storage hardware are necessary
  - New hardware requires additional power and incurs additional costs for energy
- Data-intensive fields have increasing costs for storage and energy
  - German Climate Computing Center: each petabyte of disk storage costs roughly 100,000 € plus 3,680 € annually for electricity
    - Almost 200,000 € per year in total for its 54 PiB storage system

Motivation...

- Data reduction is a solution to minimize energy consumption
  - Reducing the amount of storage hardware required to store the data
  - Can consume significant amounts of energy, negating its benefits
- Energy efficiency in supercomputers has been investigated extensively
  - Impact of data reduction on energy efficiency remains largely unexplored
- Data reduction is of great interest in scientific software
  - Algorithms have to be appropriate for their data and must be tuned
  - Decreasing runtime performance has to be avoided
  - Choosing a compression algorithm is a technical decision
    - Decision depends on the data and the software/hardware environments

Motivation...

- Ultimate goal: automatize the decision making process for users
  - Poor manual choices can lead to low compression ratios (CR), decreased performance and increased energy consumption
  - Focus on scientific applications in the context of HPC with huge amounts of data
- We have already analyzed the energy impact of data reduction in [1]
  - Extended for intelligently selecting algorithms and settings
  - Emphasis on the algorithms’ energy consumption
  - Users are able to specify additional criteria to tune behavior

Energy-Efficient Data Compression

- Extended Scientific Compression Library [6]
  - SCIL is a meta-compressor
  - Should pick a suitable chain of algorithms satisfying the user's requirements
  - Extended by a decision component that can use different criteria for selection
  - Support energy-efficient data reduction by using machine learning approaches
- Main goal: provide the most appropriate data reduction strategy for a given data set

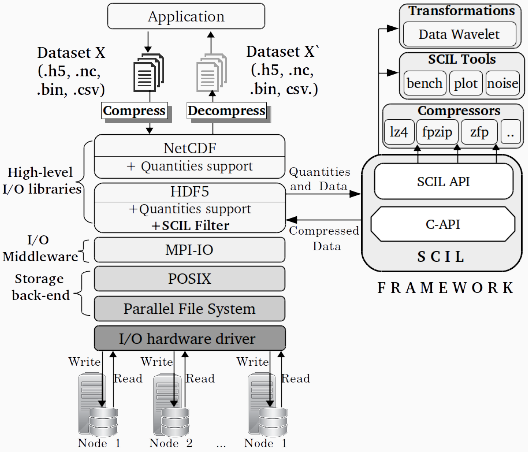

Decision Component

- Decision component takes into account information about the data's structure
  - Machine learning techniques infer best suited algorithms and settings
  - A training step is currently done separately from application runs
    - Output data set is split up into individual variables and then analyzed
    - Collect compression ratios, processor utilization, energy consumption etc.
    - Use wide range of data sets to have a sufficiently large pool of training data
- Decision component currently makes use of decision trees
  - Planned to be extended with other techniques
  - Data is split into a training set and a test set to prevent overfitting

Decision Component...

- Trees are created using the `DecisionTreeClassifier` from scikit-learn
  - Tree representation is then parsed into a file that is usable in SCIL
- Decision component's behavior can be separated into two distinct modes
  1. Data structure is known
     - Can select the optimal compression algorithm and settings
     - Mainly useful for production runs of known applications
     - E. g., train for specific application and choose the best compressor for each subsequent run
  2. Data structure is unknown
     - Use machine learning techniques to infer algorithm and settings
     - Information about storage size, number of elements, data dimensions, data type etc.
     - Possible to tune decision for energy efficiency, compression ratio or performance

- Fine-grained power measurement is necessary
  - Used the ArduPower wattmeter [4], v2 presented at ENERGY 2020
    - Designed to measure different components inside computing systems
    - Provides 16 channels with a variable sampling rate of 480–5,880 Hz
- Main metrics: compression ratio, runtime and consumed energy of each algorithm
- Used three data sets from different scientific domains
  1. ECOHAM: 17 GB, ecosystem model for the North Sea [5] (climate science)
  2. PETRA III: 14 GB, tomography experiments from P06 beamline [3] (photon science)
  3. ECHAM: 4 GB, atmospheric model [7] (climate science)

- Comparable compression ratios with different energy consumptions
  - Compare mafisc and zstd for ECOHAM
- Necessary to select the compression algorithm intelligently
  - Avoid wasting performance/energy
- Detailed analysis available in [1]

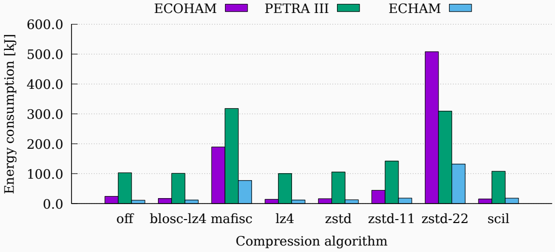

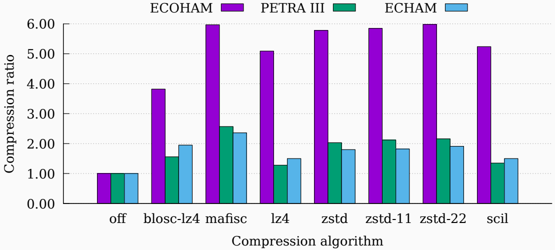

Evaluation

- Run ECOHAM using our SCIL HDF5 plugin to evaluate our approach[1]
- Trained the decision component using two different sets of training data
  1. ECOHAM data
     - Represents case with known application that has been run before
     - Still only 75 % of ECOHAM's output data is used for training
     - Uncertainty corresponds to updated application or changed output structure
  2. ECHAM data
     - Represents case with new application
     - Decision component has to use information gathered from other applications
     - Try to map decisions that make sense for other data sets to the current data

[1] Code and data are available at `https://github.com/wr-hamburg/energy2020-compression`

Evaluation…

- Two different optimization targets, which correspond to different use cases
  1. Optimized for minimal energy consumption per compression ratio
     - Allows shrinking the data with the least amount of energy
  2. Optimized for maximal compression ratio per time
     - Performance is not degraded excessively

- Decision tree using ECOHAM training data
  - Maximal compression ratio per time
  - Multitude of metrics are taken into account
  - Array dimensions, data size, number of elements, size of each dimension and information about data types

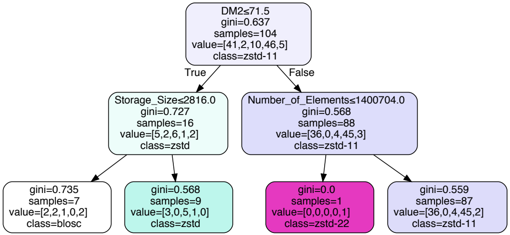

Evaluation. . .

- Run ECOHAM with important compressors
  - Static mode and decision trees
  - Decision trees have access to all compressors
- Four decision trees have been used
  1. ecoham-1: ECOHAM data, minimal energy consumption per CR
  2. ecoham-2: ECOHAM data, maximal compression ratio per time
  3. echam-1: ECHAM data, minimal energy consumption per CR
  4. echam-2: ECHAM data, maximal compression ratio per time

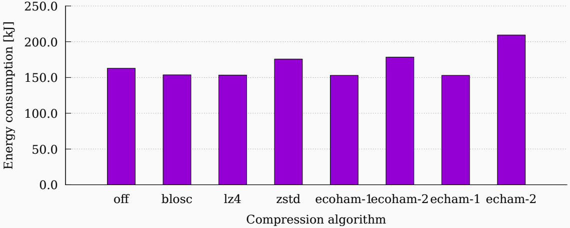

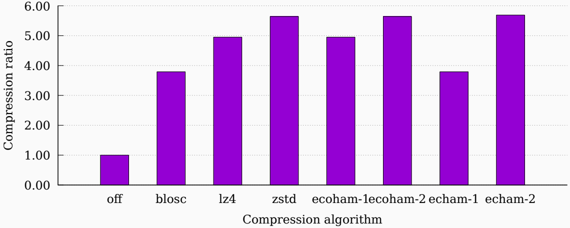

Evaluation...

- Decision component correctly uses energy-efficient algorithms
    - ecoham-1: ECOHAM's data structure is known, data is reduced effectively
    - echam-1: data structure unknown, still chooses appropriate compressors
- Optimization for maximal compression ratio increases energy consumption
    - echam-2: data structure unknown, choices are not as effective as in previous case

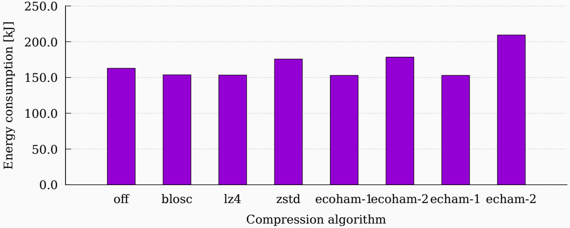

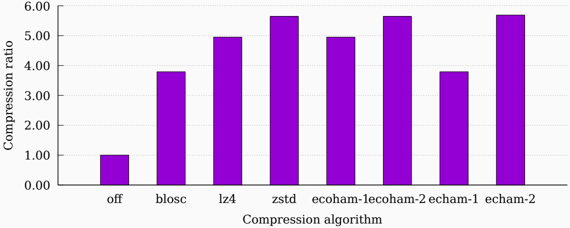

Conclusion

- Amount of data saved by compression heavily depends data structure
  - Preconditioners, algorithms and settings might work well for one data set, but might increase energy consumption for others
- Fine-grained per-variable analyses identified strategies for three data sets
  - Trained the decision component for our real-world evaluation
- Demonstrated that decision component chooses appropriate compressors for both known and unknown applications
  - Can be further tuned for energy efficiency or compression ratio
  - Achieved satisfactory compression ratios without increases in energy consumption
  - Slight increases in energy consumption allow significantly boosting compression ratios

Future Work

- Training currently has to be performed in a separate step
    - Training data collection can be more tightly integrated with production runs
    - Training mode to capture and analyze applications' output data during regular runs
    - Send samples to a training service that analyzes them in more detail
- HDF5's current filter interface is too limiting to fully exploit all possibilities
    - Operates on opaque buffers, impossible to access single data points
    - Use data variance to further tune compressor's behavior
    - Chains of compressors can lead to additional space savings

Questions

Thank you for listening! If you have any questions, please send me an e-mail to

`michael.kuhn@ovgu.de`

References

References

[1] Yevhen Alforov, Thomas Ludwig, Anastasiia Novikova, Michael Kuhn, and Julian M. Kunkel. **Towards Green Scientific Data Compression Through High-Level I/O Interfaces.** In *30th International Symposium on Computer Architecture and High Performance Computing, SBAC-PAD 2018, Lyon, France, September 24-27, 2018*, pages 209–216, 2018.

[2] Renhai Chen, Zili Shao, and Tao Li. **Bridging the I/O performance gap for big data workloads: A new NVDIMM-based approach.** In *49th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2016, Taipei, Taiwan, October 15-19, 2016*, pages 9:1–9:12, 2016.

[3] DESY. **PETRA III.** `http://petra3.desy.de/index_eng.html`, 2015. Retrieved: April, 2020.

[4] Manuel F. Dolz, Mohammad Reza Heidari, Michael Kuhn, Thomas Ludwig, and Germán Fabregat. **ArduPower: A low-cost wattmeter to improve energy efficiency of HPC applications.** In *Sixth International Green and Sustainable Computing Conference, IGSC 2015, Las Vegas, NV, USA, December 14-16, 2015*, pages 1–8, 2015.

[5] Fabian Große et al. **Looking beyond stratification: a model-based analysis of the biological drivers of oxygen depletion in the North Sea.** *Biogeosciences Discussions*, pages 2511–2535, 2015. Retrieved: April, 2020.

[6] Julian M. Kunkel, Anastasiia Novikova, Eugen Betke, and Armin Schaare. **Toward Decoupling the Selection of Compression Algorithms from Quality Constraints.** In *High Performance Computing - ISC High Performance 2017 International Workshops, DRBSD, ExaComm, HCPM, HPC-IODC, IWOPH, IXPUG, P^3MA, VHPC, Visualization at*

Scale, WOPSSS, Frankfurt, Germany, June 18-22, 2017, Revised Selected Papers, pages 3–14, 2017.

[7] Erich Roeckner et al. **The atmospheric general circulation model ECHAM 5.** *Max Planck Institute for Meteorology*, 2003.