Partners

- Technische Universität Dresden
- Gesellschaft für Numerische Simulation mbH
- Intel Exascale Labs Paris
- IT4Innovations, VSB - Technical University of Ostrava
- National University of Ireland Galway
- Norwegian University of Science and Technology
- Technische Universität München

TECHNISCHE UNIVERSITÄT DRESDEN

gns

OÉ Gaillimh NUI Galway

NTNU

IT4Innovations national supercomputing center

intel

Contact and Funding

www readex.eu

@readex_eu

READEX is funded by the European Union's Horizon 2020 research and innovation programme „FET-Proactive – towards exascale high performance computing" under grant agreement No 671657.

Project term:
2015/09 - 2018/08

Project Coordinator:
Prof. Dr. Wolfgang E. Nagel
wolfgang.nagel@tu-dresden.de

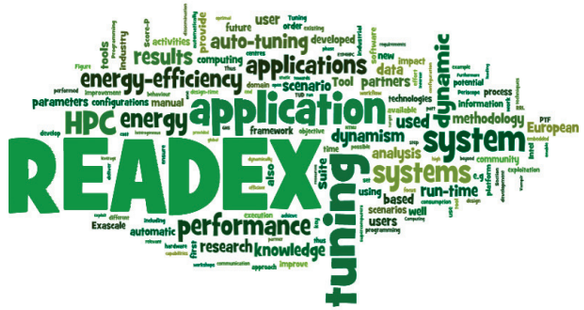

READEX

Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

Energy-Efficient Exascale Computing

www.readex.eu

European Commission | Horizon 2020 European Union funding for Research & Innovation

Overview

The importance of energy efficiency is constantly increasing in High Performance Computing (HPC). While systems can be adapted to individual applications in order to reduce energy consumption, manual tuning of platform parameters is a tedious and often neglected task.

The READEX projects automates this by developing a tools-aided methodology for dynamic auto-tuning that combines technologies from two ends of the computing spectrum: system scenario methodology from the embedded world and auto-tuning from the field of HPC.

READEX Tools-Aided Methodology

The READEX methodology has been designed for exploiting the dynamic behaviour of software. At design time different runtime situations (RTS) are detected and optimized system configurations are determined. RTSs with the same configuration are grouped into scenarios, forming the tuning model. At runtime, the tuning model is used to switch system configurations dynamically.

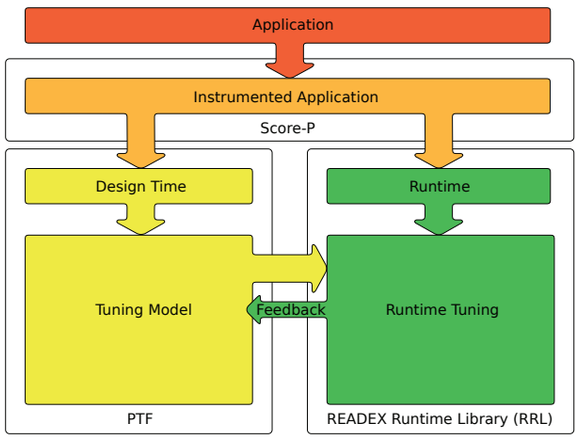

Design Time Analysis

Design time analysis (DTA) is carried out with the Periscope Tuning Framework (PTF). It uses a multi-agent based approach to identify RTSs and to determine optimized system configurations. These are settings for tuning parameters, e.g., core and uncore frequencies. It also provides means for the specification of domain knowledge (DK) to improve the automatic tuning results. Part of the DK is the specification of application tuning parameters, which allows users to offload computation to accelerated devices. The result of DTA is a tuning model that guides runtime tuning.

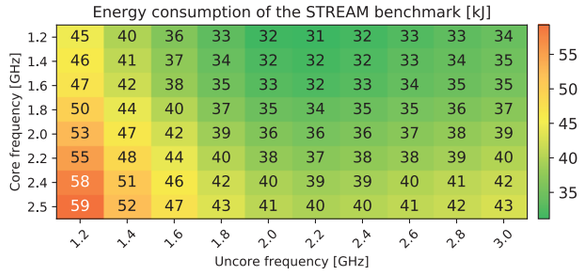

Heatmap of the energy consumption of a STREAM benchmark for different core and uncore frequencies. The data array does not fit in the processor's L3 processor cache.

Runtime-Tuning

During production runs of the user's application, the READEX Runtime Library takes control. It is designed to apply the different configuration in a lightweight manor.

Moreover the READEX Runtime Library will be able to adapt to a changing application behaviour. The latter is to be implemented by state-of-the-art machine learning mechanisms, which are currently under development.

Impact and Validation

In order to validate the impact of the READEX project, several real world applications are employed. In a co-design approach, selected applications are being hand-tuned and both the improvements in energy efficiency and the effort spent will be compared with the automatic tuning approach.

Initial Results on Tuning Potential

ESPRESO is a highly efficient parallel solver, which contains several FETI based algorithms. It is based on a communication layer on top of MPI. The layer was developed specifically for FETI solvers and uses communication-hiding and -avoiding techniques to achieve scalability.

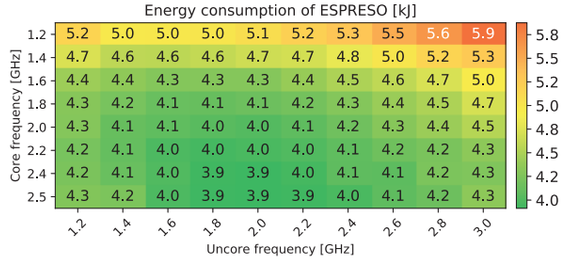

Heatmap of the energy consumption of Espreso for different core and uncore frequencies (static tuning).