READEX

Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

European Commission

www.readex.eu

RUNTIME EXPLOITATION OF APPLICATION DYNAMISM FOR ENERGY-EFFICIENT EXASCALE COMPUTING

OVERVIEW

- Exploit dynamic behaviour of HPC applications to achieve improved energy-efficiency and performance
- Develop a tools-aided scenario based dynamic auto-tuning methodology
- Bring together experts from embedded systems and HPC

READEX TOOLS-AIDED METHODOLOGY

Design Time Analysis

- Detect different Runtime Situations (RTS)
- Determine optimal system configurations
- Group the different optimal configurations into scenarios
- Save the collected information into the tuning model

Runtime Tuning

- Load the tuning model
- Switch the configuration according to the tuning model

DESIGN TIME ANALYSIS

- Lead by the Periscope Tuning Framework (PTF)
  - Supports different tuning strategies
  - Multiple objectives
  - Supports a variety of tuning knobs
- Executed by the READEX Runtime Library
  - RTS support
  - Extendable due to plugin infrastructure
- Optional ATP Library
  - Adds Application Tuning Parameters (ATPs)
  - Allows the tuning of different code paths or offloading decisions to external devices like Xeon Phi or Nvidia Tesla

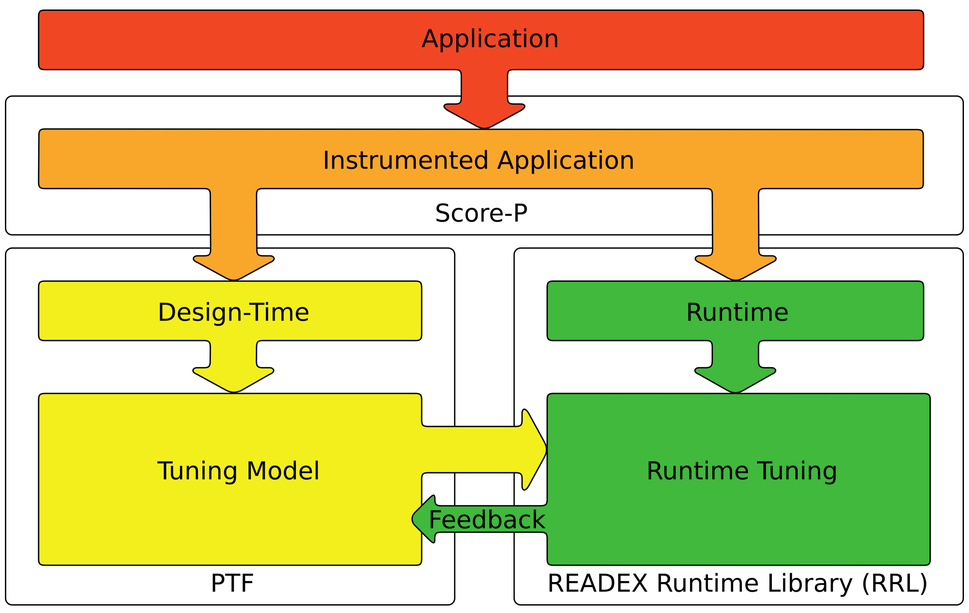

RUNTIME TUNING

- Lead by the READEX Runtime Library
  - Lightweight
  - Loads the tuning model generated at design-time
  - Sets the optimal configuration according to the tuning model
- Optional calibration mechanism
  - Calibrates regions which are not seen during design-time
  - Adjusts the tuning model if the application behaviour changes during runtime
  - Uses state-of-the-art machine learning techniques

VALIDATION

- Industry-grade simulation codes Indeed (FEM), Elmer (FEM), and OpenFOAM (CFD)
- CORAL and Proxy Apps benchmark suites
- ESPRESO library
  - Highly efficient parallel solver
  - Contains several FETI based algorithms
  - Based on a communication layer on top of MPI, which uses communication-hiding and -avoiding techniques
- On Bullx DLC B720 Intel Xeon 2650v3 (@TU Dresden)

INITIAL RESULTS ON TUNING POTENTIAL

1st Step: Static Tuning

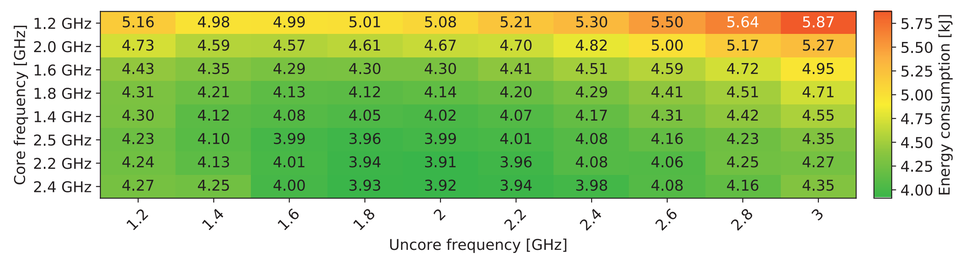

Energy consumption heatmap of the ESPRESO library for different core and uncore frequencies

2nd Step: Dynamic Tuning

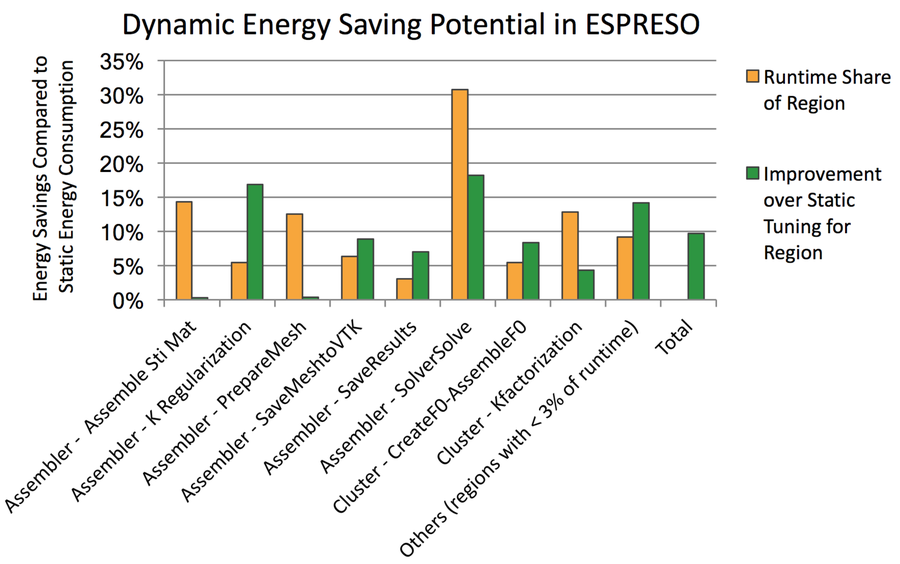

FUNDING AND PARTNERS

- Funded by the European Union's Horizon 2020 research and innovation programme „FET-Proactive – towards exascale high performance computing".
- Grant agreement No 671657
- TU Dresden, TU Munich, Norwegian University of Science and Technology, National University of Ireland Galway, IT4Innovations, Intel Exascale Labs Paris, Gesellschaft für Numerische Simulation mbH

FURTHER INFORMATION

readex.eu
researchgate.net/project/READEX

readex_eu

Robert Schöne: robert.schoene@tu-dresden.de
Wolfgang E. Nagel: wolfgang.nagel@tu-dresden.de