

READEX – Runtime Exploitation of Application Dynamism for Energy-efficient eXascale computing

EETHPC @ ISC'18

Andreas Gocht – Technische Universität Dresden

Overview: Project Partners

- Grant agreement No 671657
- Officially started September 1st, 2015
- Technische Universität Dresden/ZIH (Coordinator)
- Norwegian University of Science and Technology
- Technische Universität München
- IT4Innovations, VSB-Technical University of Ostrava
- NUI Galway, Irish Centre for High-End Computing
- Intel France
- Gesellschaft für numerische Simulation mbH



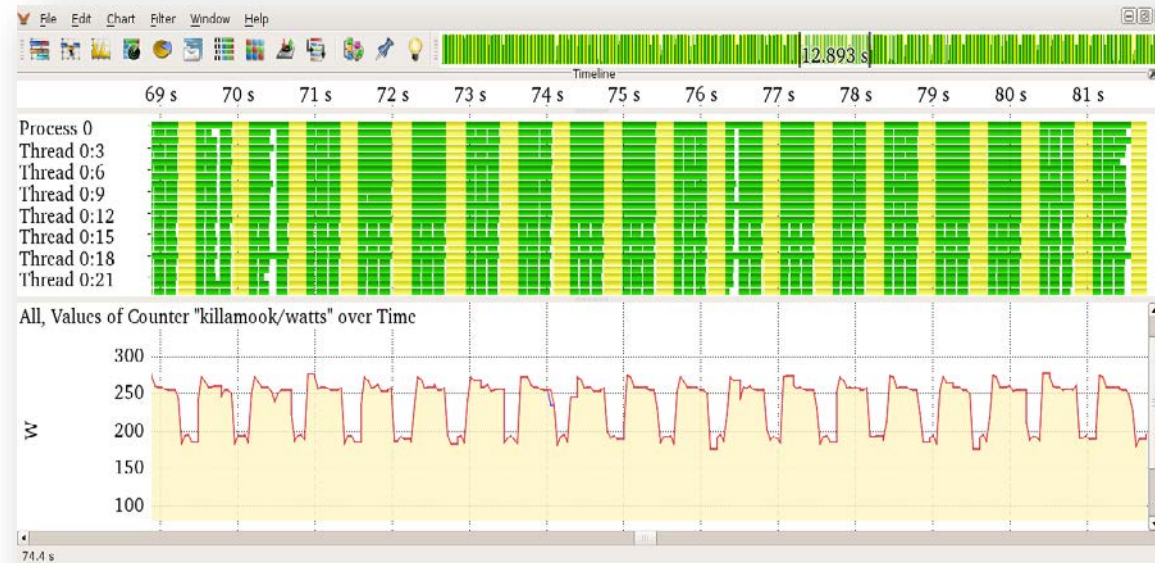
IT4Innovations
national01\$#&0
supercomputing
center@#01%101



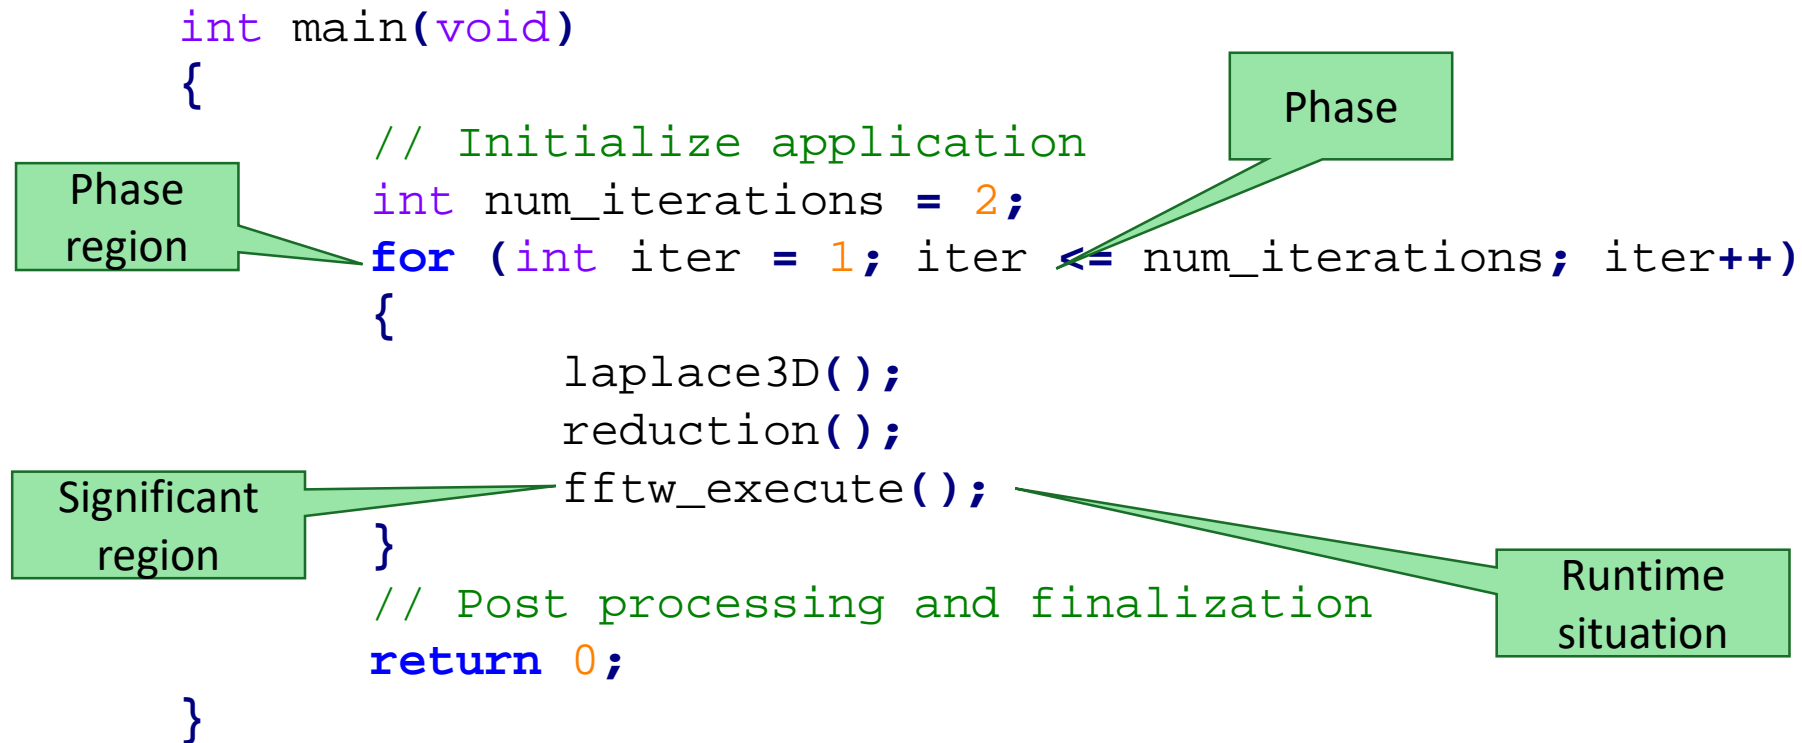
Project Motivation

Applications exhibit dynamic behaviour

- Changing resource requirements
- Computational characteristics
- Changing load on processors over time



Terminology: Region and Region Instance



Terminology: Tuning Parameter

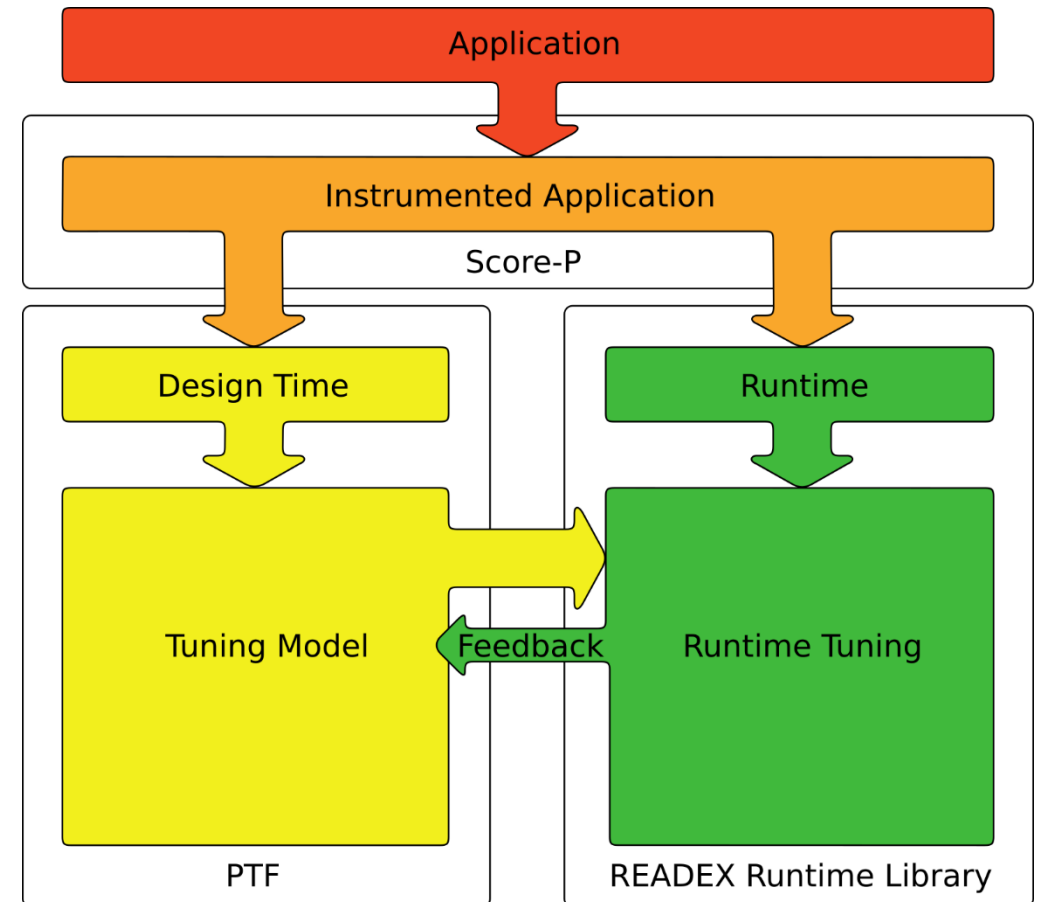
```
int main(void)
{
    // Initialize application
    int num_iterations = 2;
    for (int iter = 1; iter <= num_iterations; iter++)
    {
        laplace3D();
        reduction();
        fftw_execute();
    }
    // Post processing and finalization
    return 0;
}
```

Tuning
Parameter
FREQ=2 GHz

Tuning
Parameter
FREQ=1.5 GHz

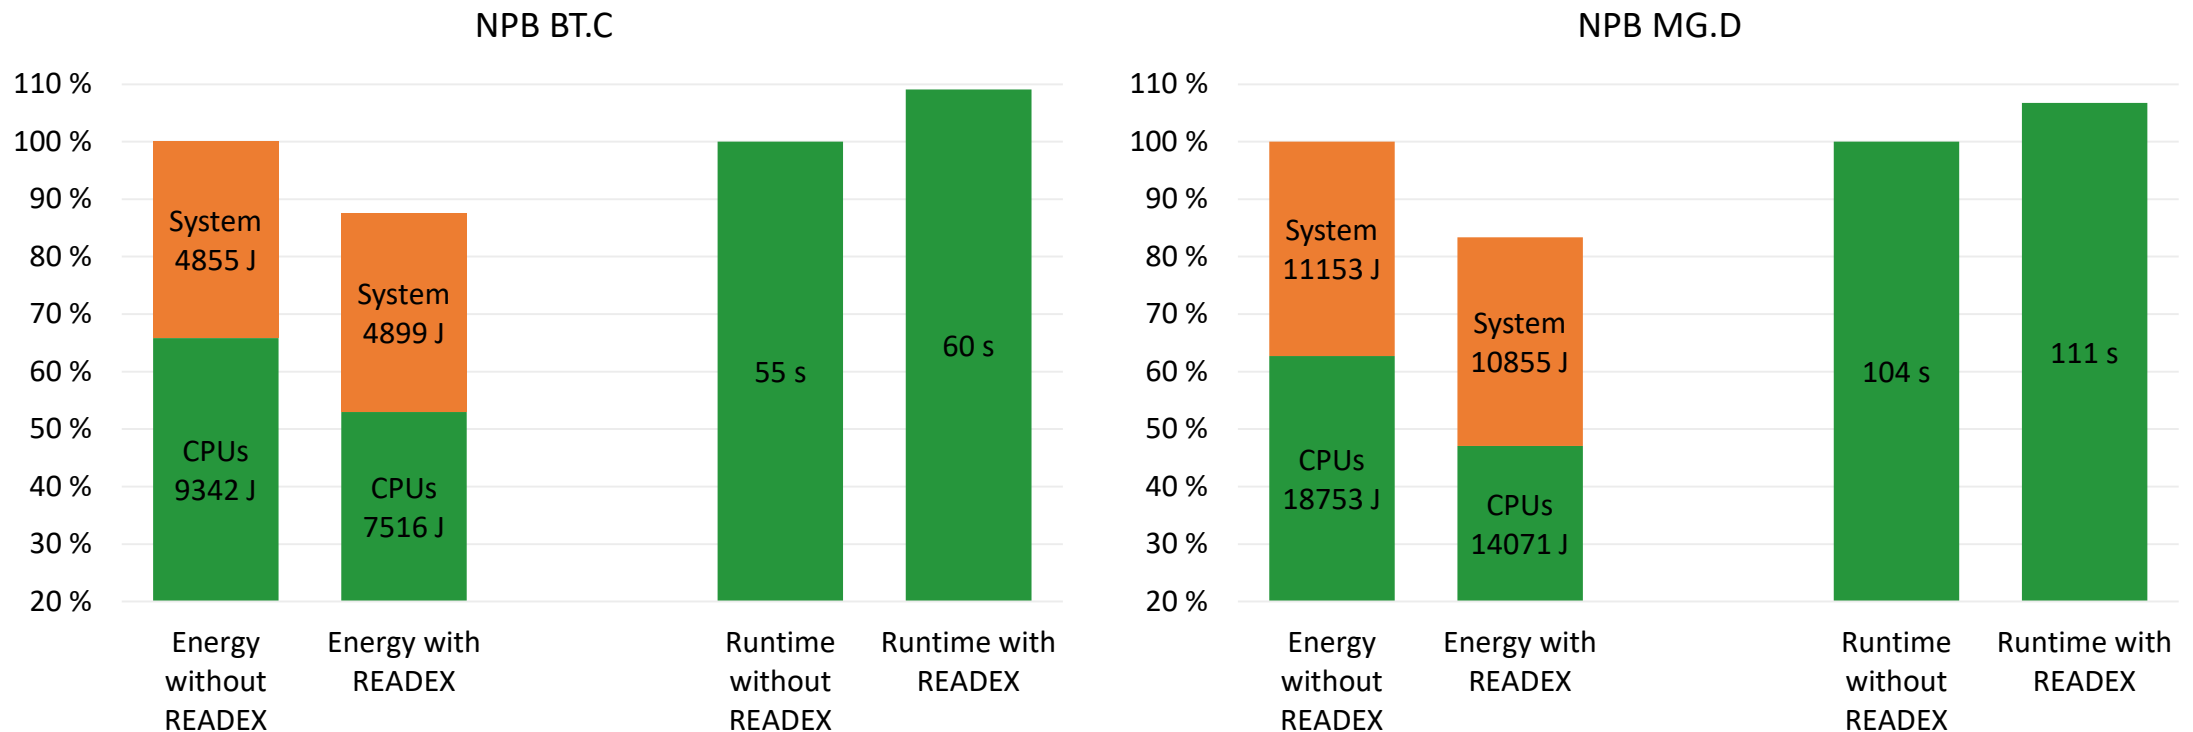
Overview: Workflow

1. Instrument application
 - Score-P provides different kinds of instrumentation
2. Detect dynamism
 - Check whether runtime situations could benefit from tuning
3. Detect energy saving potential and configurations (DTA)
 - Use tuning plugin and power measurement infrastructure to search for optimal configuration
 - Create tuning model
4. Runtime application tuning (RAT)
 - Apply tuning model, use optimal configuration

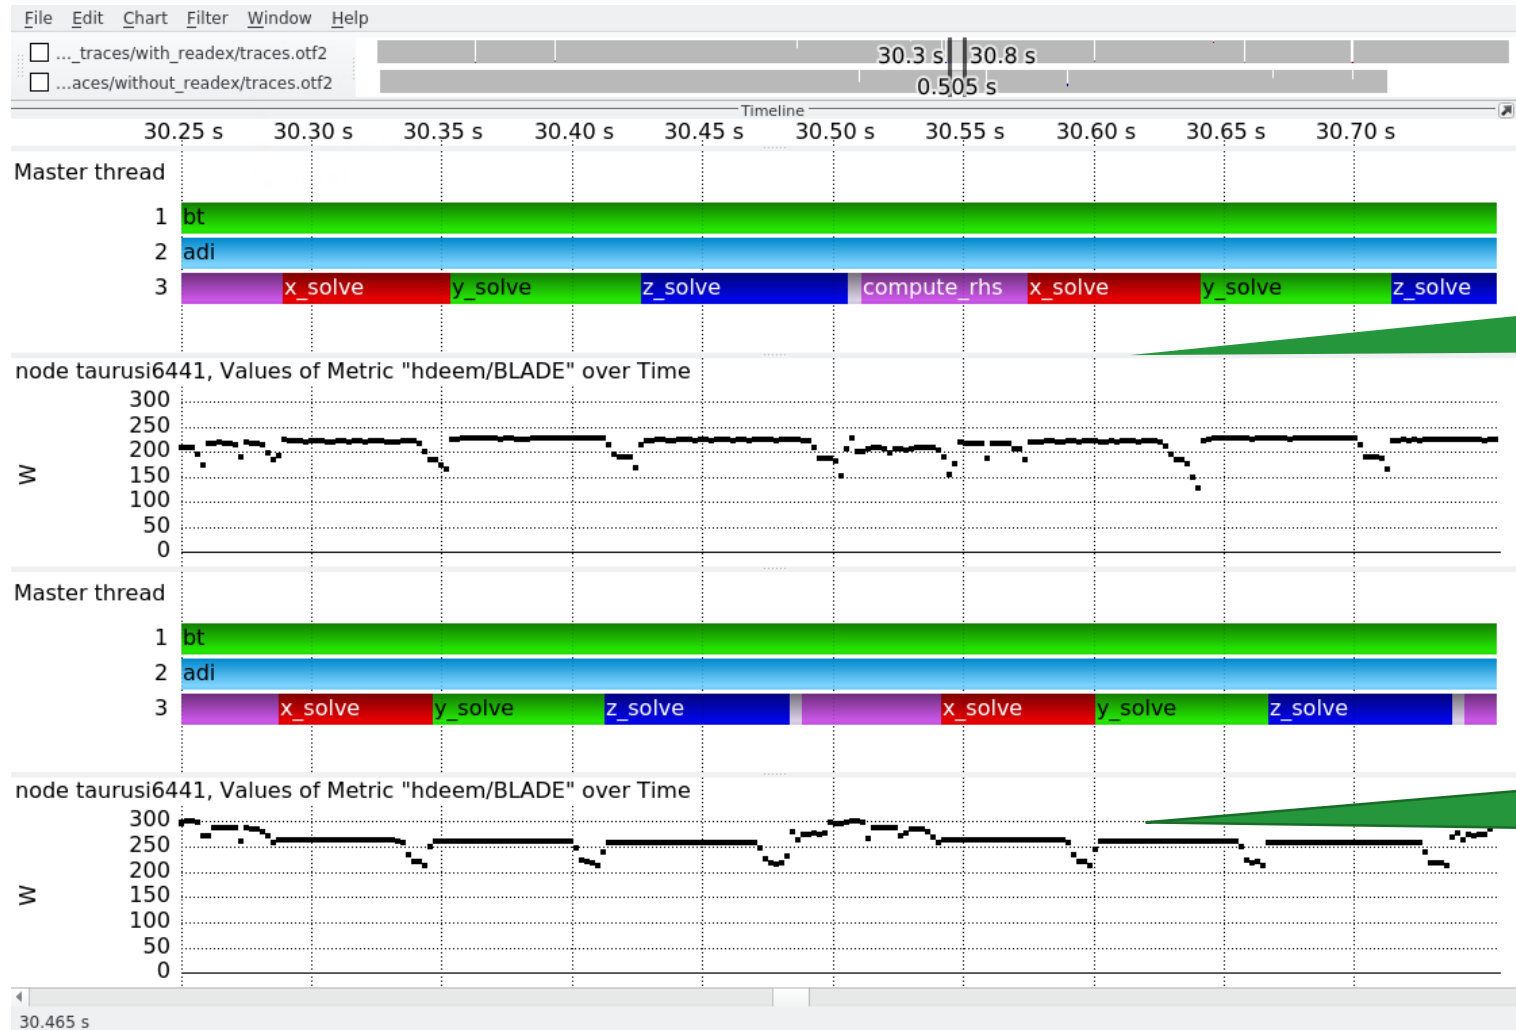


Current Status

- Application instrumentation, dynamism detection, DTA, and RAT are implemented
- Promising results:



Example: NAS-OMP BT.C Benchmark

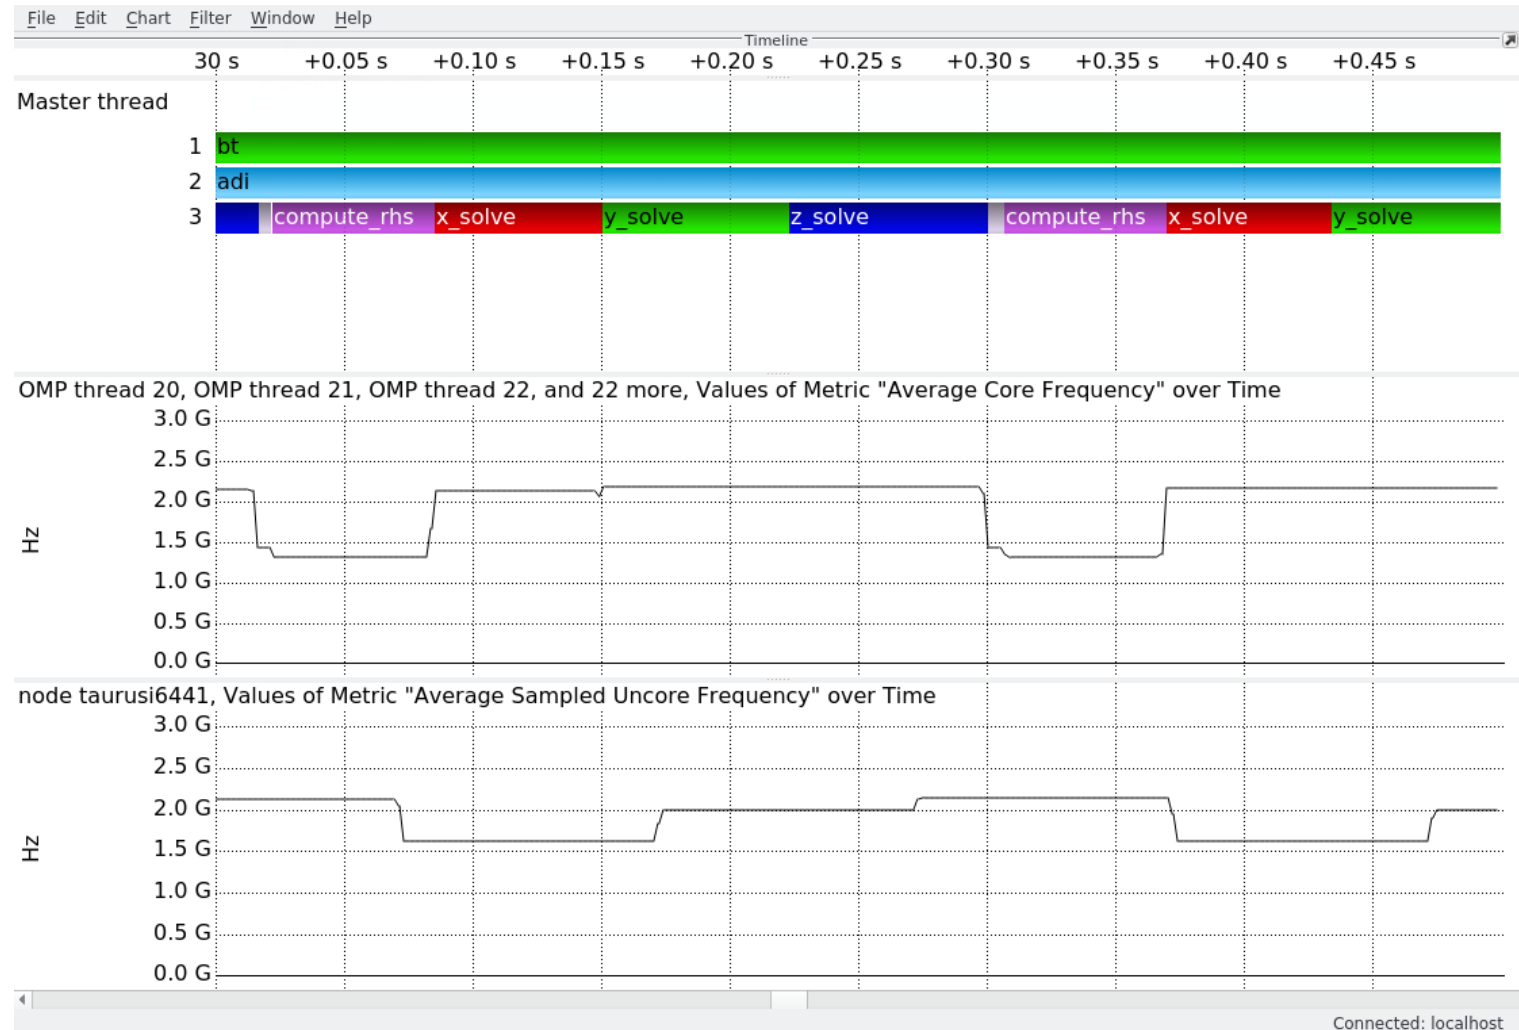


With READEX

Without READEX

Comparison of power consumption and runtime of the BT benchmark

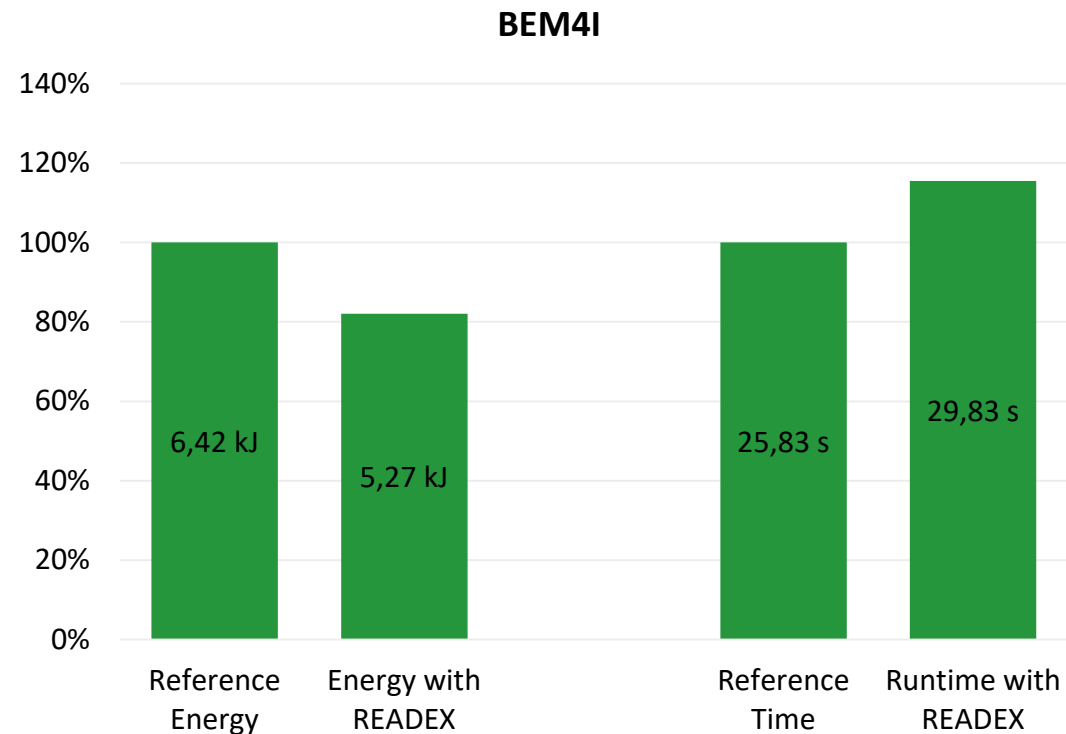
Example: NAS-OMP BT.C Benchmark



Selected tuning parameter: core and uncore frequency of the BT benchmark

Current Status - BEM4I

- BEM4I solver for partial differential equations
- Benchmark solves Dirichlet boundary value problem for 3D Helmholtz equation
- Surface mesh with 12288 triangular elements
- Energy savings: 18%




What we are currently working on

- Application Tuning Parameters (ATPs)
- Runtime Calibration
- Porting to other machines

Application Tuning Parameters (ATPs)

- Provide an interface to allow the user to offload decisions
- Example: Preconditioner in ESPRESO



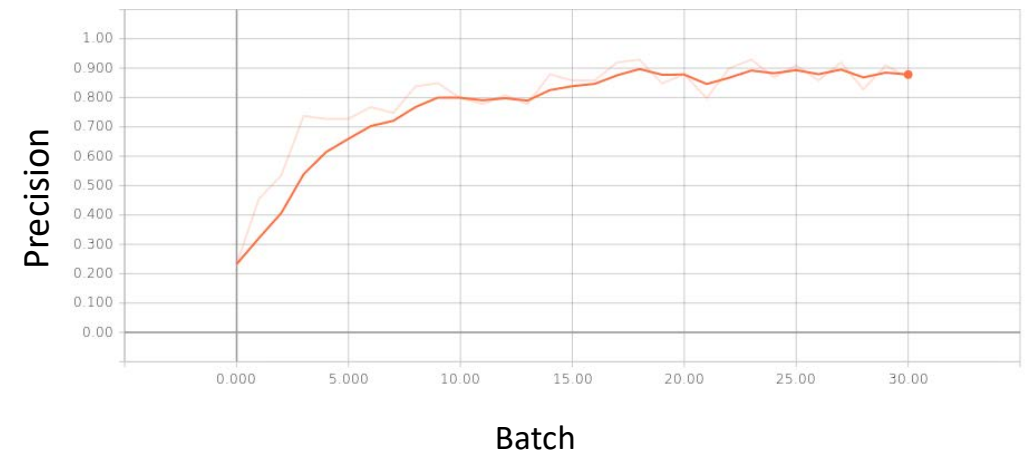
Runtime Calibration

What happens if a region has not been seen during DTA?

- Runtime Calibration

Currently two approaches in development:

- Performance Counter based, e.g. using Neural Nets
- „Search“ based using Q-Learning



Training run of a Neural Network

Thank you for your attention

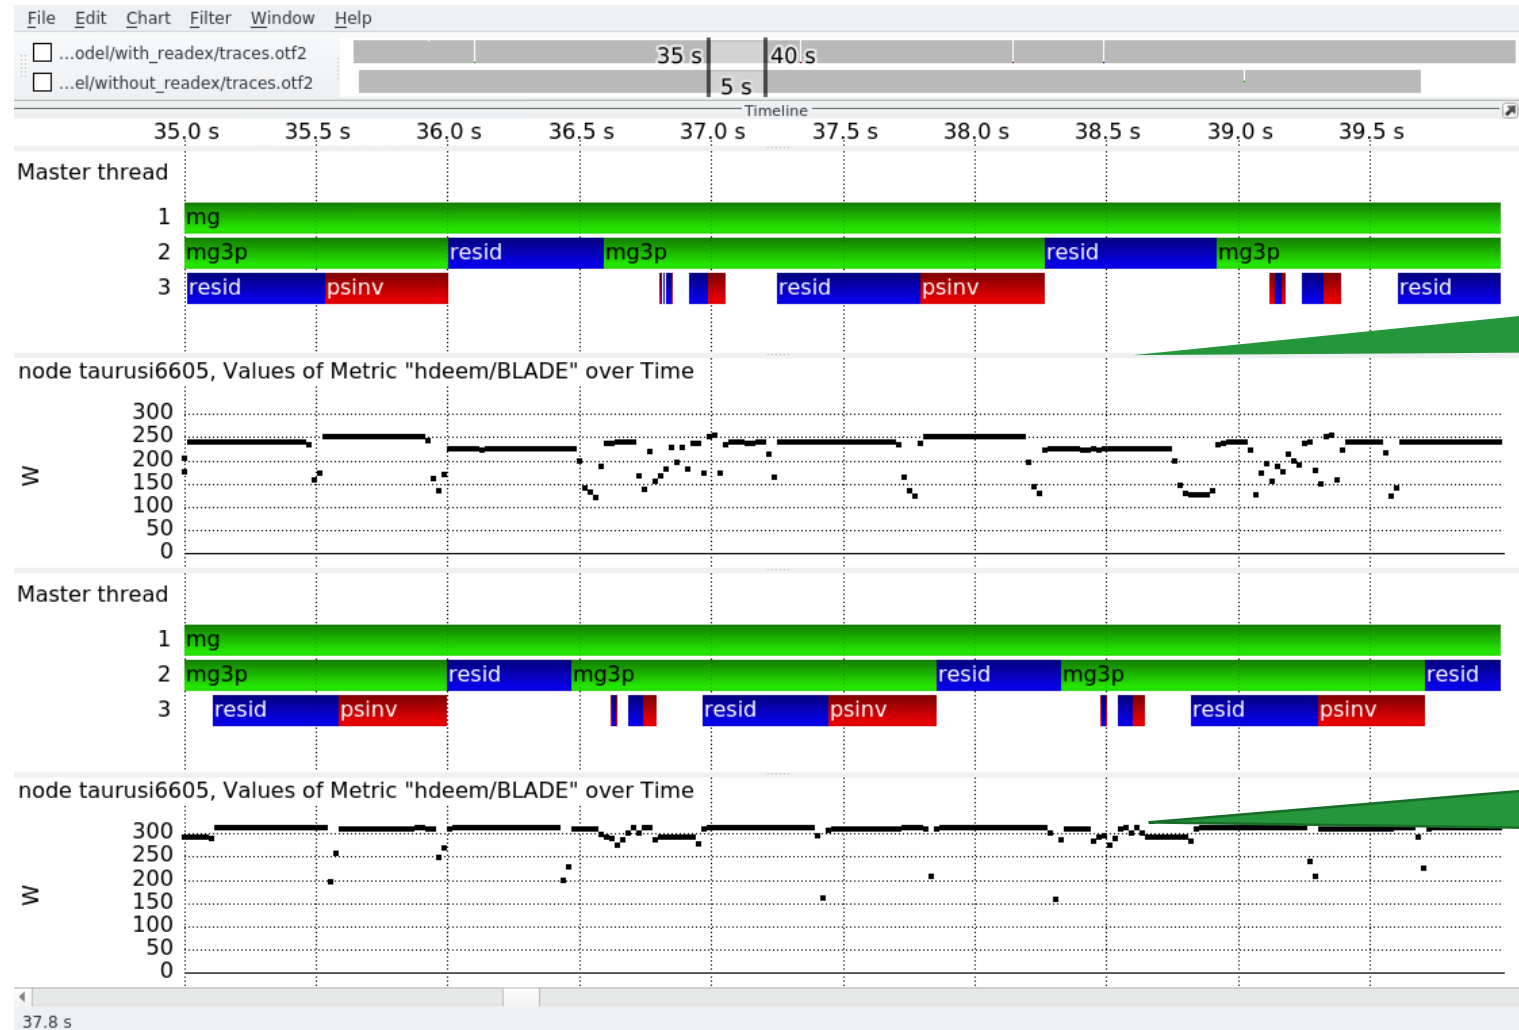
Questions?

Funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 671657.

Thank you for your attention

Backup

Example: NAS-OMP MG.D Benchmark

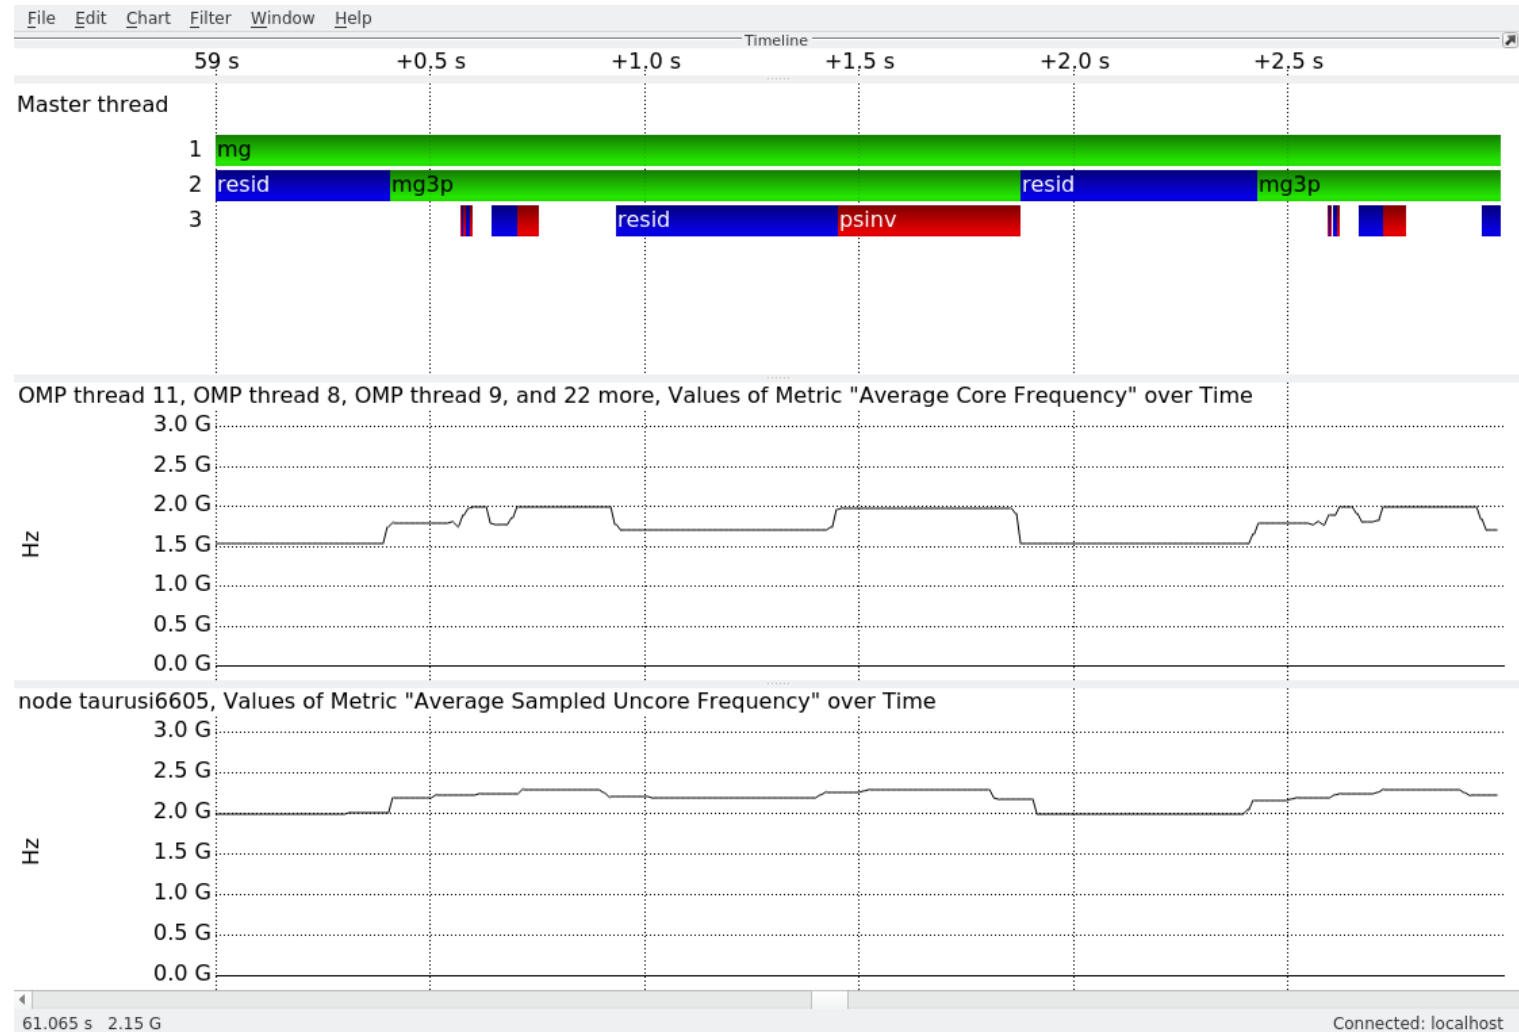


With READEX

Without READEX

Comparison of power consumption and runtime of the MG benchmark

Example: NAS-OMP MG.D Benchmark



Selected tuning parameter: core and uncore frequency of the MG benchmark