SEM4HPC, HPDC 17

READEX Tool Suite for Energy-efficiency Tuning of HPC Applications

Anamika Chowdhury

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

European Commission
Horizon 2020
European Union funding for Research & Innovation

Project Overview

- READEX:

  Runtime **E**xploitation of **A**pplication **D**ynamism for **E**nergy-efficient e**X**ascale Computing

- Starting Date:

  1 September 2015

- Duration:

  3 years

- Funding

  European Commission Horizon 2020 grant agreement 671657

- Collaboration with 6 other institutions all over Europe

Objectives

- Tuning HPC applications dynamically for energy efficiency.
- Improve energy efficiency by influencing tuning parameters
- Switching between configurations
  - Exploit dynamic characteristics
- Develop tool aided auto-tuning methodology.
  - Design-time Analysis
  - Runtime Application Tuning
- Detect at design-time, exploit at runtime.

The READEX Tool Suite

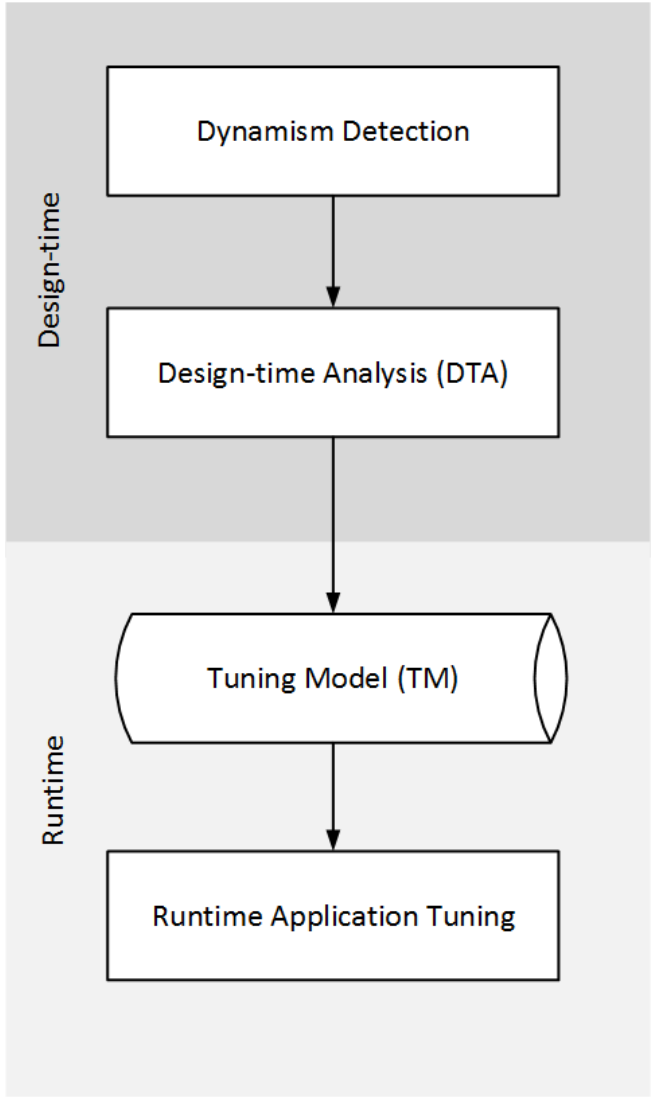

Terminology: Phase Region and Phase

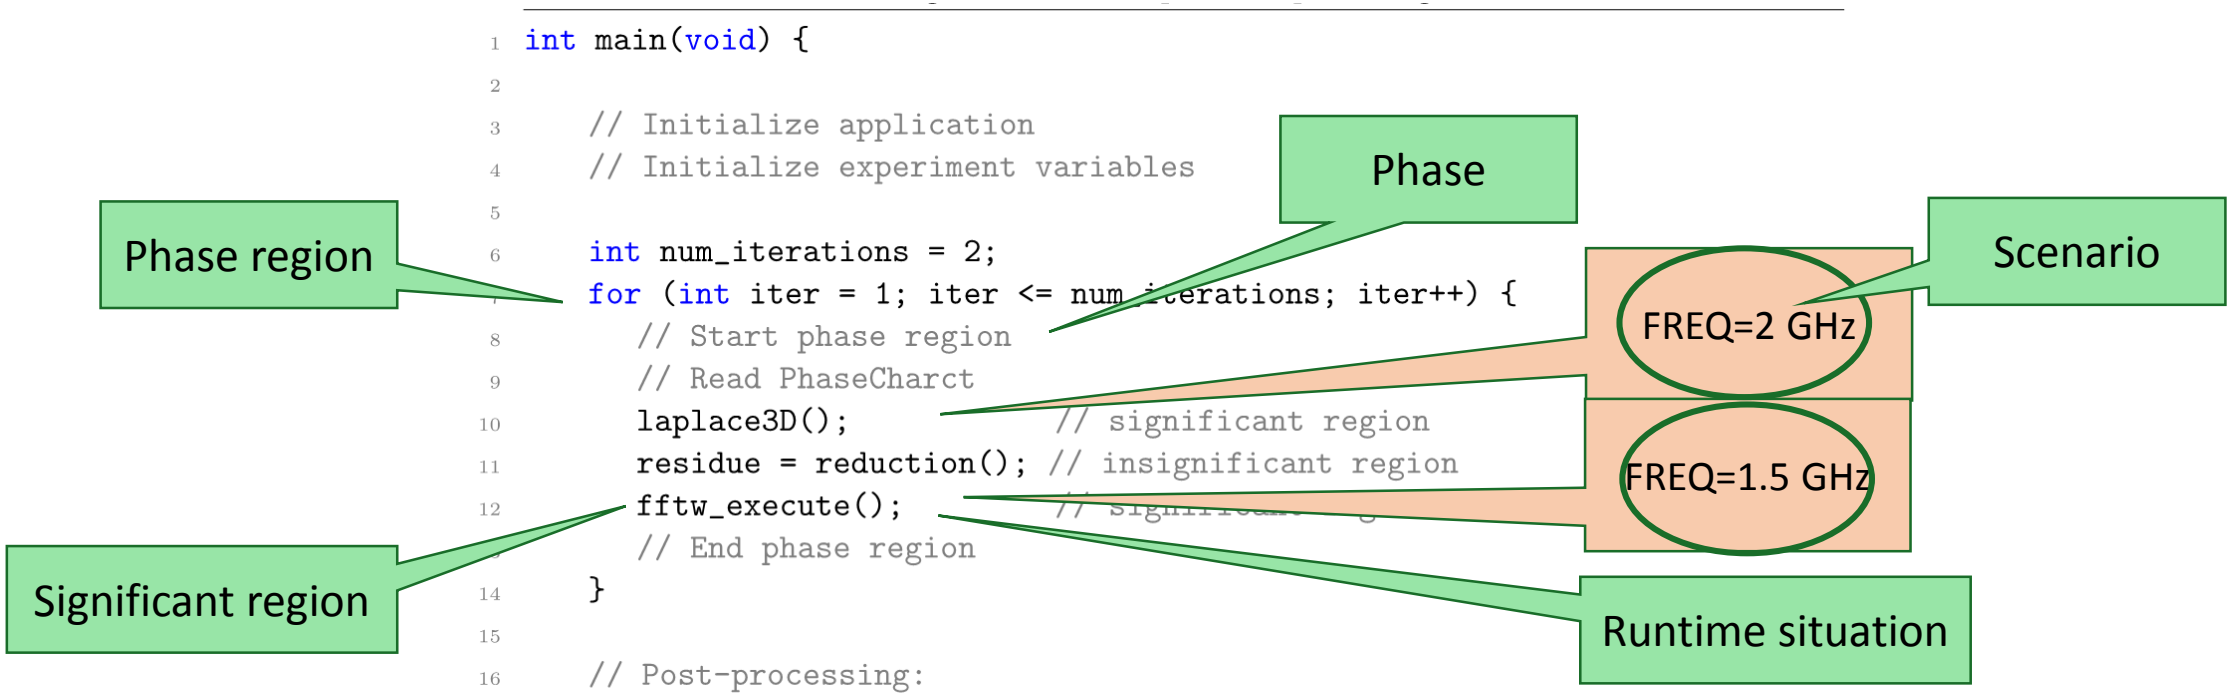

Score-P

- Scalable Performance Measurement Infrastructure for Parallel Codes
  - Common instrumentation and measurement infrastructure

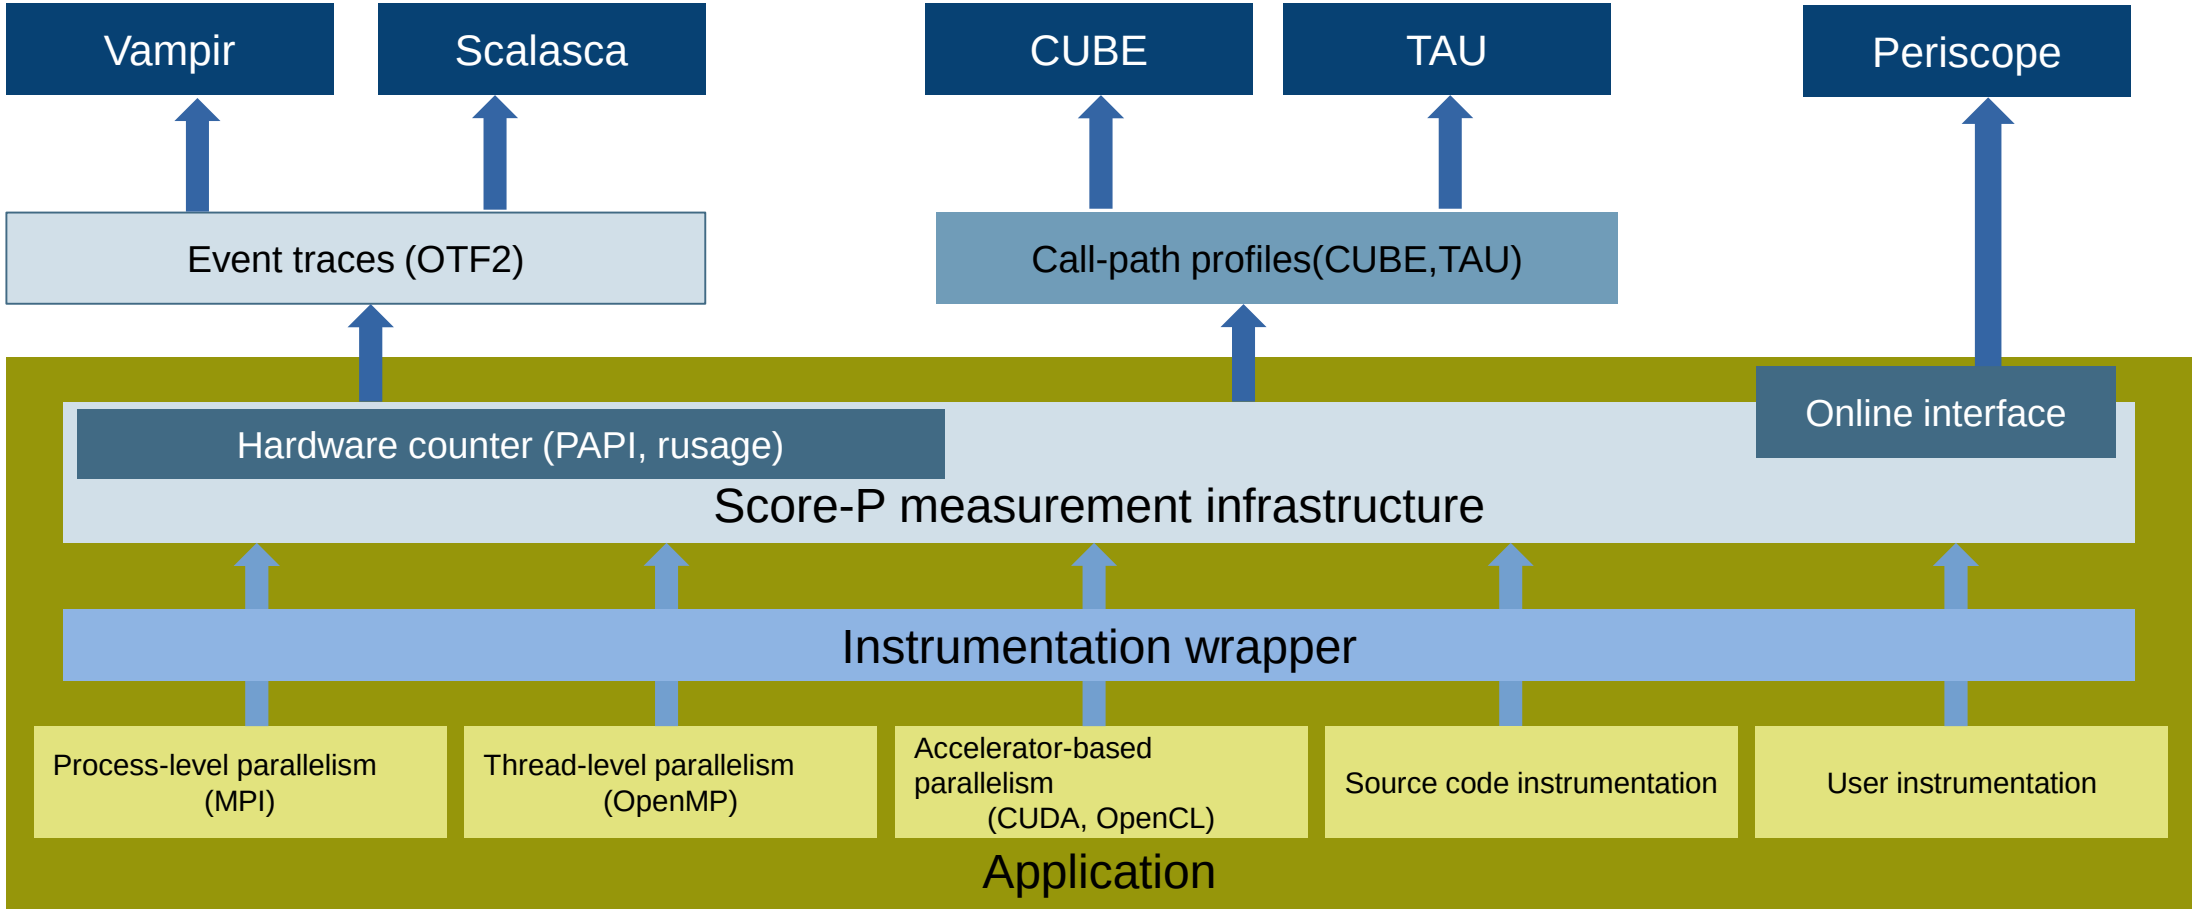

Periscope Tuning Framework (PTF)

Automatic application analysis & tuning

- Tune performance and energy (statically)
- Plug-in-based architecture
- Evaluate alternatives online
- Scalable and distributed framework

Support variety of parallel paradigms

- MPI, OpenMP, OpenCL, Parallel pattern

Developed in the AutoTune EU-FP7 project

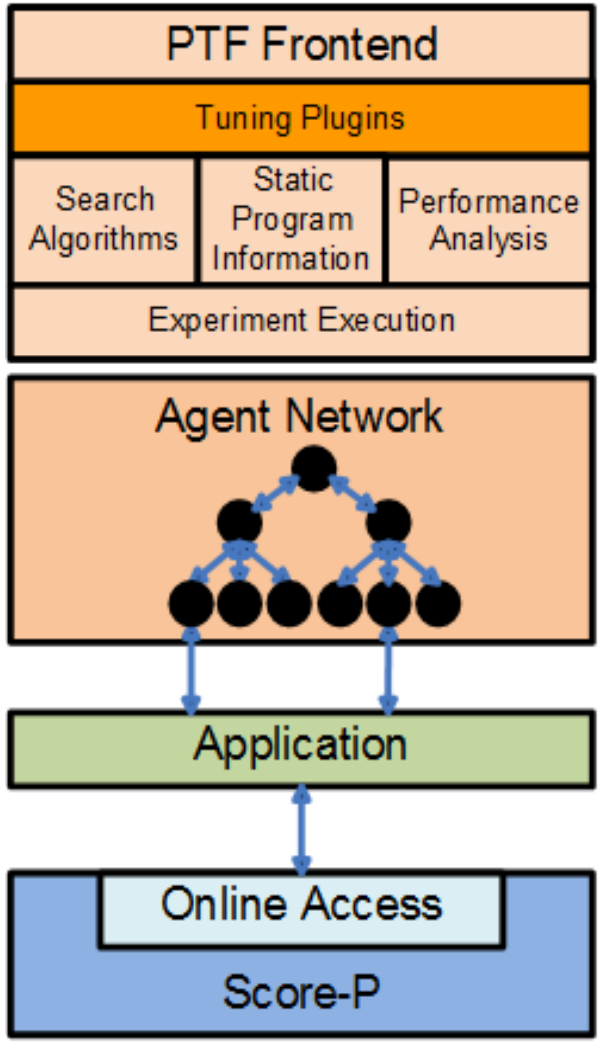

Design Time Analysis

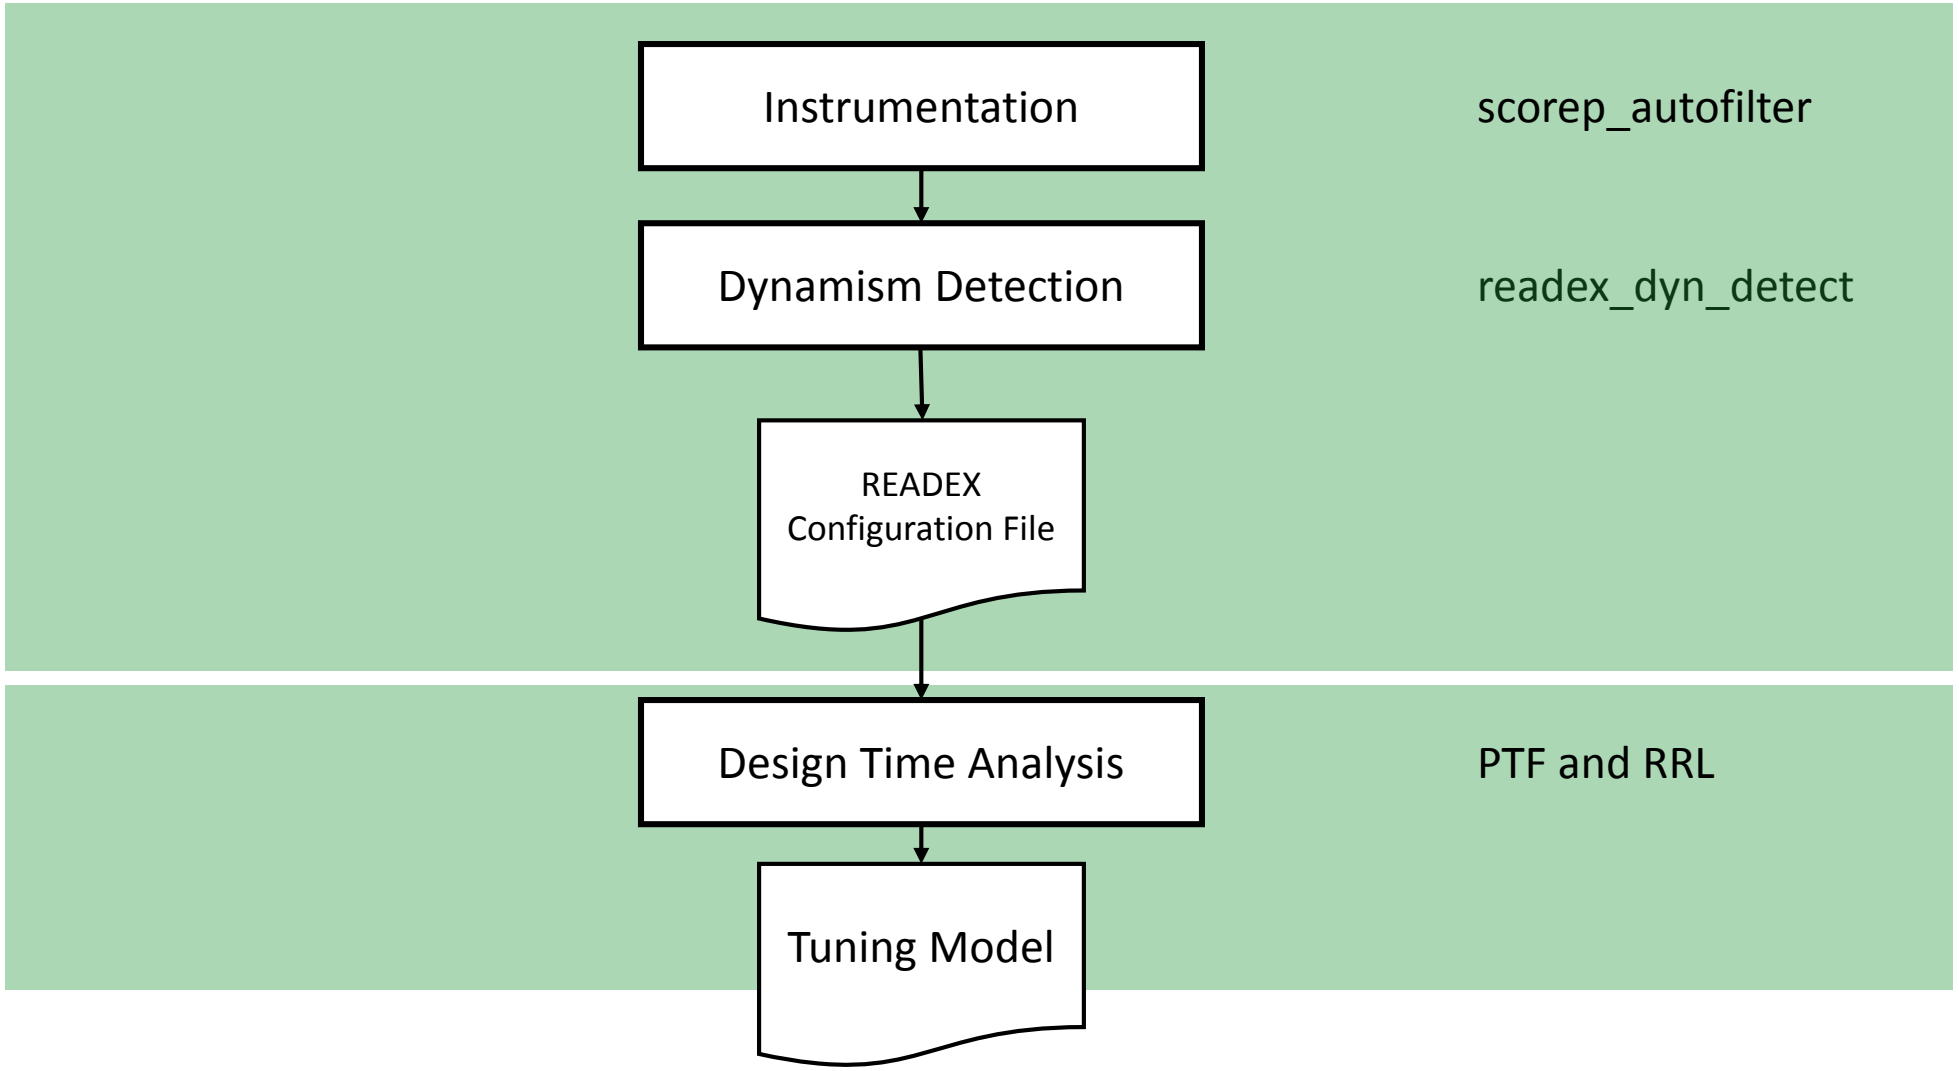

READEX Tuning Plugin

- Tuning plugin supporting
  - Core and uncore frequencies, numthreads parameters
  - Configurable search space via READEX Configuration File
  - Several objective functions: energy, CPUenergy, EDP, EDP2, time
  - Several search strategies: exhaustive, individual, random, genetic
- Approach
  - Experiment with default configuration
  - Experiments for selected configurations
  - Configuration set for phase region
  - Energy and time measured for all runtime situations
  - Identification of static best for phase and specific best configurations for rts‘s

Pre-Computation of Configurations

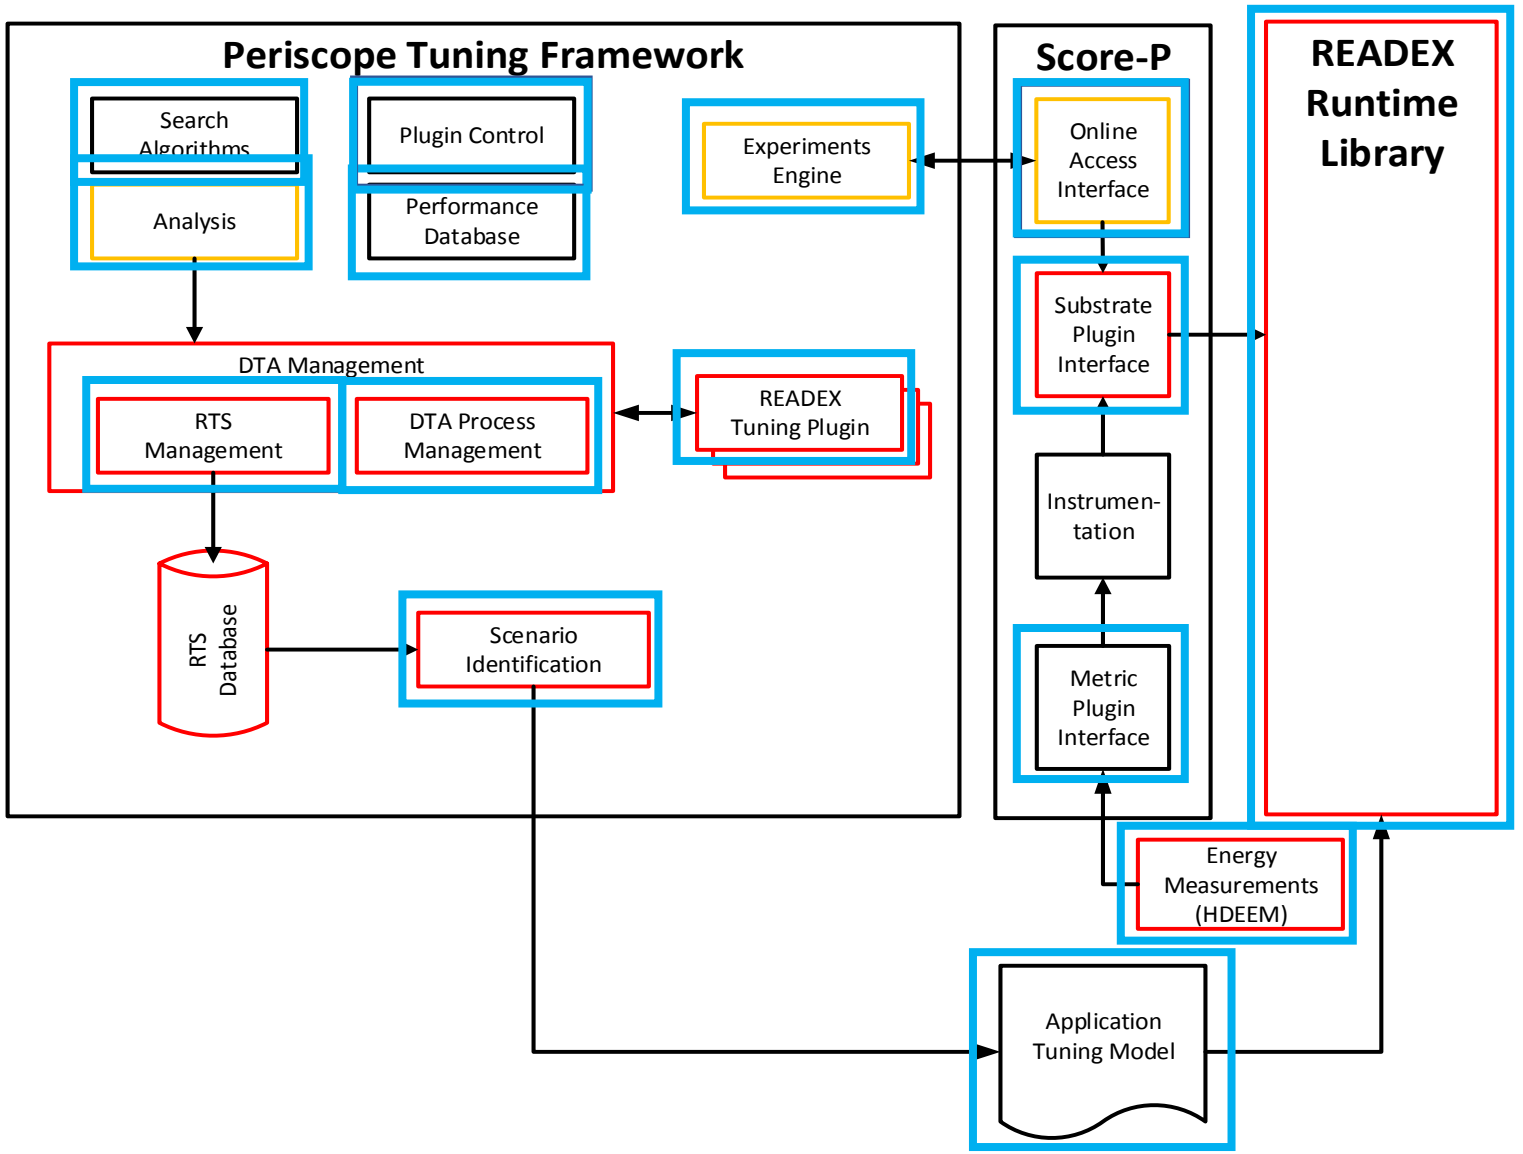

READEX Runtime Library (RRL)

- Runtime Application Tuning performed by the READEX Runtime Library.
- Tuning requests during Design-time Analysis are sent to RRL.
- A lightweight library
  - Dynamic switching between different configurations at runtime.
  - Implemented as a substrate plugin of Score-P.
- Developed by Technische Universität Dresden (TUD)

Runtime Scenario Detection and Switching Decision during Production Run

- During Runtime Application Tuning
- Scenario classification

- Switching decision component
- Manipulation of tuning parameters

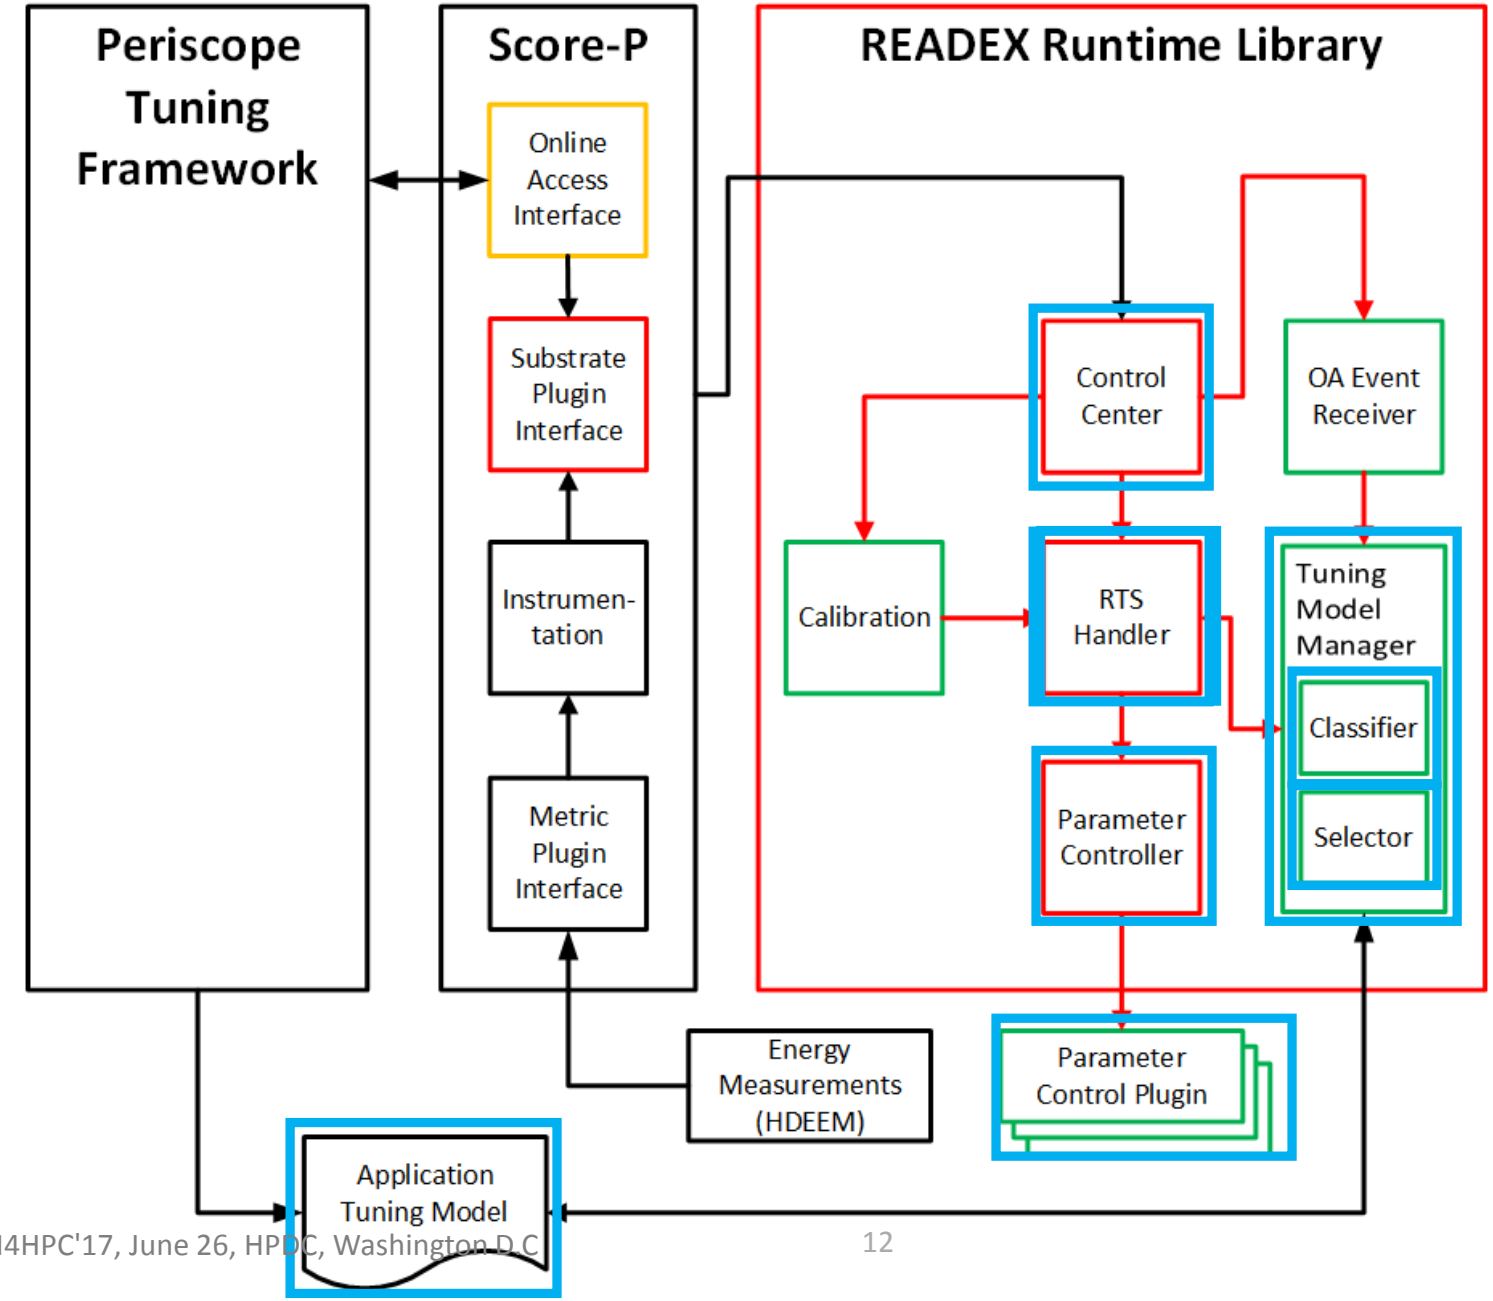

Evaluation of DTA (NPB BT-MZ)

- 16 experiments, exhaustive search
- Significant regions found during dynamism detection
  - exch_qbc, x_solve, y_solve, z_solve

Significant regions (16 experiments)	Energy for worst static configuration (1, 1.6)	Energy for best static configuration (4, 1.6)	READEX Energy	# of thread	Core frequency
exch_qbc	3245	6649	2760	1	2.4
x_solve	74219	41341	39962	4	2.0
y_solve	73536	39497	39497	4	1.6
z_solve	76393	40699	40386	4	2.0
SUM	227393	128186	122605		
savings		43.60%	4.40%		
Energy for phase	376722	284223			
savings		24.60%			

Evaluation of DTA (LULESH)

- 64 experiments, exhaustive search

Significant regions	Energy for worst (16, 1.2)	Energy for best (1, 2.4)	READEX Energy	# of thread	Core frequency
CalcCourantConstraintForElems	6143	5375	5084	14	2.0
CalcKinematicsForElems	23816	15546	15546	1	2.4
CalcMonotonicQGradientsForElems	10833	7317	7116	15	2.4
CalcMonotonicQRegionForElems	14340	13238	12291	5	2.0
CalcVolumeForceForElems	109217	57075	57075	1	2.4
EvalEOSForElems	54791	30402	29837	3	2.4
SUM	**219140**	**128953**	**126949**		
savings		**41,2%**	**1.6%**		
Energy for phase	261980	167978			
savings		**35,9%**			

Evaluation of DTA (LULESH)

- 19 experiments, individual search

Significant regions	Energy for worst (16, 1.2)	Energy for best (1, 2.4)	READEX Energy	# of thread	Core frequency
CalcCourantConstraintForElems	6265	5097	5097	1	2.4
CalcKinematicsForElems	23328	15528	15399	4	2.4
CalcMonotonicQGradientsForElems	10645	7181	7021	15	2.4
CalcMonotonicQRegionForElems	14242	13011	12210	16	2.4
CalcVolumeForceForElems	108368	56151	56151	1	2.4
EvalEOSForElems	42172	27773	27773	1	2.4
SUM	**205020**	**124741**	**123651**		
savings		**39.2%**	**0.9%**		
Energy for phase	247782	162489			
savings		**34.4%**			

Conclusion and Outlook

- Inter-phase dynamism
- Handling multiple input files
- Domain knowledge specification
  - Allows the user to provide domain knowledge as *identifiers.*
    - Application Tuning Parameters
    - Input identifiers
- To know more about the project, go to http://www.readex.eu/
- Demo video can be shown after session
  - Also available on the project website.

Discussion

European Commission
Horizon 2020
European Union funding for Research & Innovation