EUROPEAN PROCESSOR INITIATIVE

FRAMEWORK PARTNERSHIP AGREEMENT IN EUROPEAN LOW-POWER MICROPROCESSOR TECHNOLOGIES

EPI Co-design approach

- Bi-directional and iterative interaction process between:
  - o application experts and
  - o hardware (HW) and system-software (SW) developers
- Multi-level suite of benchmarks
  - o from very low-level synthetic benchmarks to high-level applications
- Methodology with multi-level models & simulators
  - o analytical models, high level
  - o simulation based (e.g. gem5 simulation engine)
  - o reference platform (e.g. Marvell ThunderX2)
- Node-level co-design parameters (e.g.):
  - o GPP: SVE length, number of SVE pipelines per core
  - o Accelerator: vector registers length, ratio accelerator-vs-GPP cores
  - o Memory (size and BW): Cache, HBM, DDR

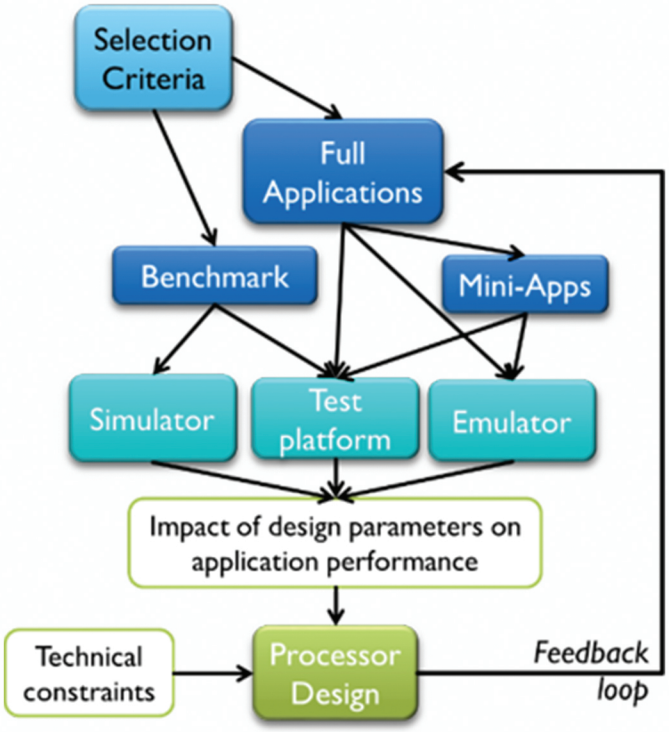

The success of the processor technologies developed in EPI depends to a large extent on how well they fulfil the needs of the end-users: scientific and industrial application developers. Co-design is understood as a bi-directional and iterative interaction process between application owners and hardware (HW) and system-software (SW) developers, in which benchmarks and applications are employed to identify the impact of design decisions onto the performance of the codes.

www.european-processor-initiative.eu
European Processor Initiative
@EuProcessor
European Processor Initiative

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647

EUROPEAN PROCESSOR INITIATIVE

FRAMEWORK PARTNERSHIP AGREEMENT IN EUROPEAN LOW-POWER MICROPROCESSOR TECHNOLOGIES

EPI benchmark suite

The EPI benchmark suite is based on a multi-level set of benchmarks with increasing code complexity: from very low-level test codes (simple flow of instructions to test functional units), passing by synthetic benchmarks, and reaching up to mini-apps (simplified versions of an application preserving its main characteristics but with much fewer lines of code), or even full application codes (for performance evaluation at later stages of the project). All disclosable elements of the EPI benchmark suite are available in a public code repository: https://gitlab.version.fz-juelich.de/epi-wp1-public

Covered application fields

The selection criteria applied to choose the codes in the EPI benchmark suite ensures that all components of the EPI design can be tested, and that fields with large weight on the current and predicted user portfolio on HPC systems are covered. The selection includes not only "traditional" HPC applications but also emerging fields such as deep learning (e.g. convolutional neural networks) and high performance data analytics:

- Biophysics, Biology, Medicine
- Earth Sciences, Climate
- High energy physics, fusion
- Material Sciences
- CFD, hydrodynamics
- PDE
- Image / Media
- Automotive
- Cryptography
- HPDA
- Machine Learning, Deep learning
- Cloud, Data Base
- Reference benchmarks (HPL, HPCG, Stream, DGEMM...)

The multi-level benchmark suite contains synthetic benchmarks, mini-applications, and full-fledged codes. This shall allow a wide range of studies, from detailed design analysis and simulations in the early stages of the project, up to platform evaluation and comparison when EPI evaluation vehicles are available. In this context it is important to highlight that the benchmark suite is by no means "frozen" or "closed". The current selection of applications is to be considered as a pool from which the most suitable workloads are selected, depending on the study that needs to be performed. If specific needs are identified later on in the project, which cannot be addressed with the current codes, new codes might be added.

Interaction with architects

A constructive and efficient interaction between WP1 and the hardware and software architects is crucial in order to enable real co-design, i.e. identifying and giving answer to the most crucial design questions based on the analysis of application requirements. This interaction has been established through regular face-to-face technical meetings. In them, a methodology has been developed based on running a small set of benchmarks and use cases with the simulation tools available to the project: Gem5, SESAM, and MUSA. This has allowed to give indications on the preferred balance between various design parameters, such as number of cores per chip, size of caches, length of vector units, or memory capacity.

Once the first test systems are available, the EPI benchmark suite will be used to evaluate and compare performance and usability with other technologies, providing further co-design input for next generation EPI processors chips. The results of these co-design activities will have an impact on the final design of the EPI processor technologies, ensuring that they fulfil the requirements of their targeted markets.

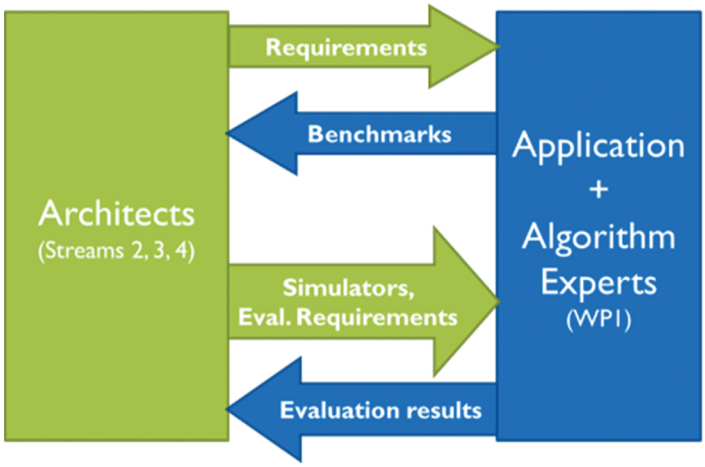

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 826647