

CO-DESIGN IN EPI

13 MAY 2020

DIRK PLEITER (JÜLICH SUPERCOMPUTING CENTRE)

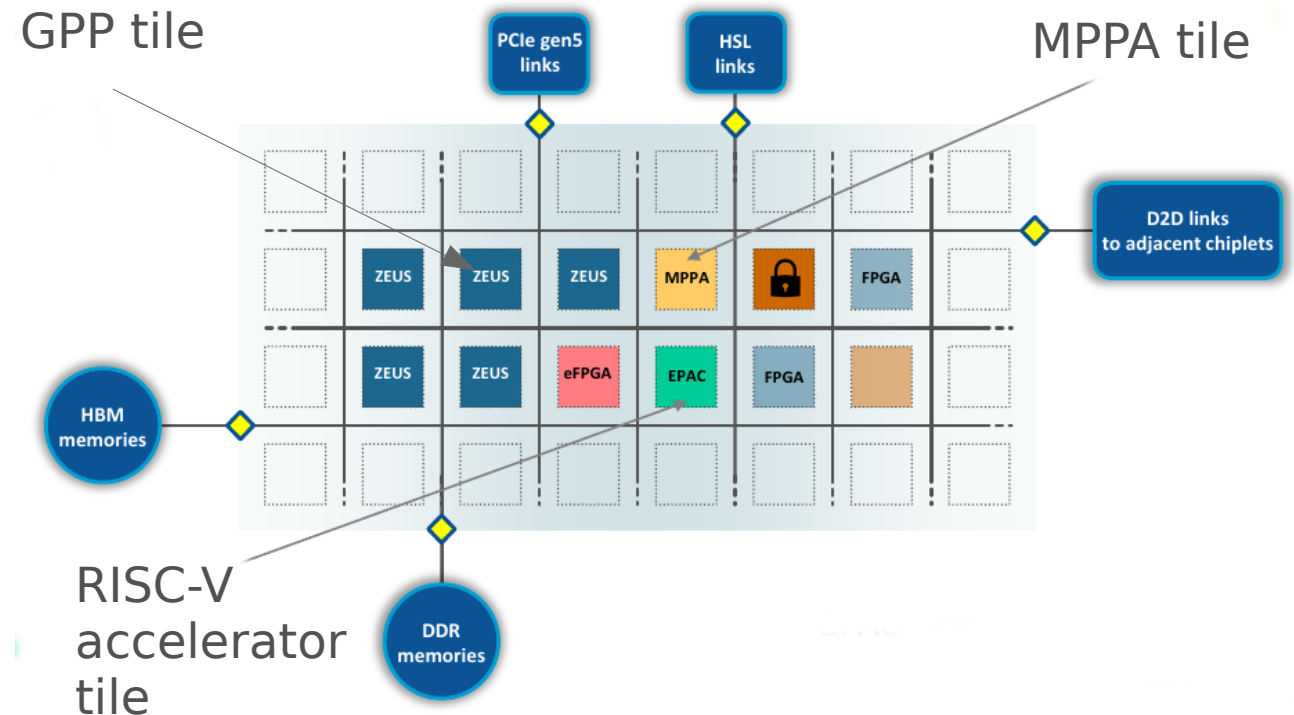


EPI MISSION

- European independence in High Performance Computing Processor Technologies
- EU Exascale machine based on EU processor by 2023
- Based on solid, long-term economic model, Go beyond HPC market
- Address the needs of European industry
 - Example: car manufacturing market

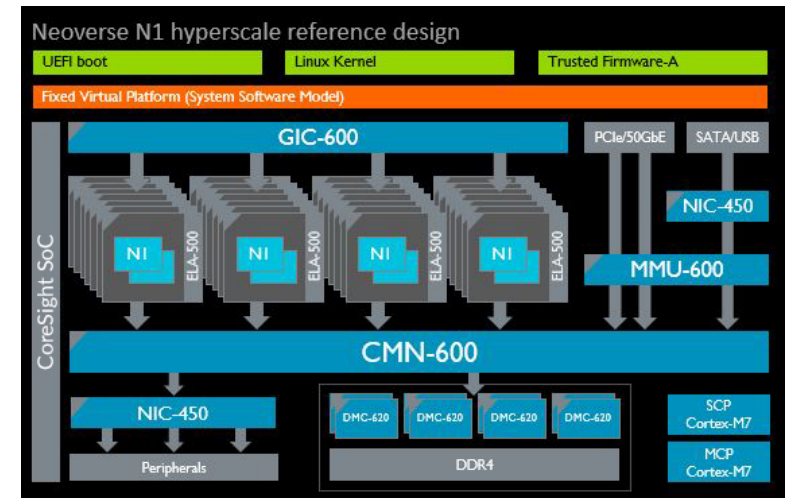
COMMON PLATFORM APPROACH

- Integration of different technological components
 - E.g. Tiles with Arm cores, RISC-V accelerators, MPPA, eFPGA, ...
- Global approach for power management and security
- Modular approach allows (in principle) flexible customization




GPP STREAM

- Based on established processor core technology
 - First generation based on Arm's Zeus core
- Zeus is part of Arm's Neoverse processor family
 - High-performance processors for AI, Cloud, HPC and edge
 - Aim for processor designs with large number of cores
 - Armv8-A ISA



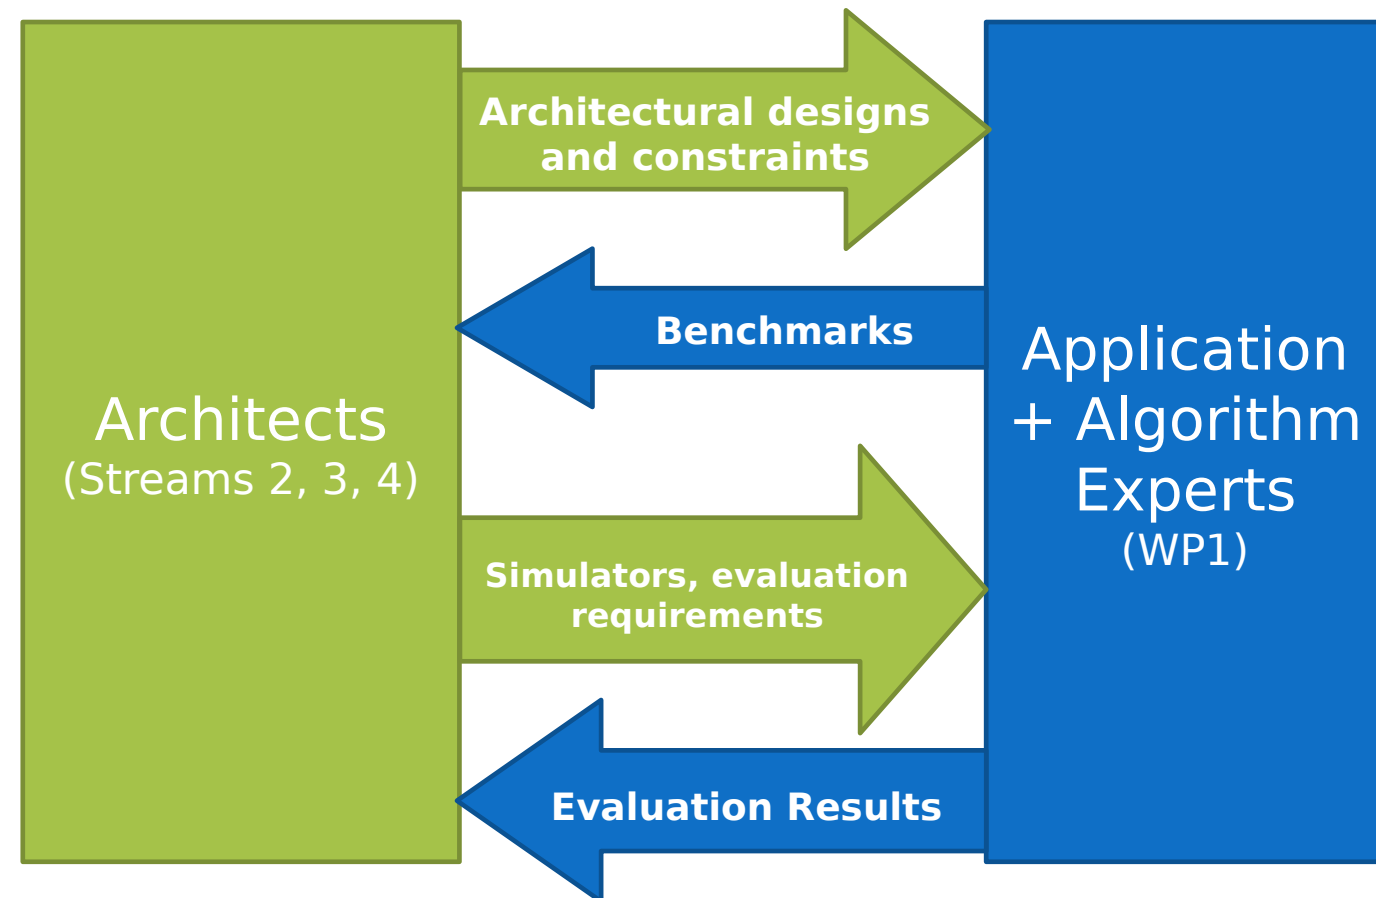
ACCELERATOR STREAM

- Vector processing unit
 - Based on RISC-V vector ISA
- Domain specific accelerators
 - Stencil/tensor accelerator (STX)
 - Variable precision processor



PROCESSOR-LEVEL CO-DESIGN IN EPI

- **Bi-directional and iterative** interaction process between:
 - application experts and
 - hardware (HW) and system-software (SW) developers
- Goal:
 - Identify **application's requirements**
 - **Feed these into design** of EPI's HW+SW technologies



CO-DESIGN SCOPE IN EPI

- Focus on **giving quality feedback to HW/SW designers**
 - Not enough effort for full application porting
 - Co-design between Applications and HW/SW
 - Co-design between HW and SW happens in other WPs
- **Multi-level suite of benchmarks**
 - From very low-synthetic benchmarks to high-level applications
- Methodology with **multi-level models & simulators**
 - Analytical models, high level
 - Simulation based (e.g. gem5 simulation engine)
 - Reference platform (e.g. Marvell ThunderX2, HiSilicon Hi1616)

FIELDS PRESENT IN CURRENT APPLICATION SELECTION

- Biophysics
- Biology/Medicine
- Earth Sciences/Climate
- HEP & Fusion
- Material Sciences
- CFD
- Hydrodynamics
- PDE
- Image / Media
- Automotive
- Cryptography
- HPDA
- Machine Learning
- Deep learning
- Cloud
- Data Base
- Reference benchmarks (HPL, HPCG, Stream, DGEMM...)

CURRENT SELECTION

Field	App candidate	Selection criteria																																
		C1 (relevance in EPI market)					C2 (architecture feature it tackles/stresses)								C3 (app. family)				C4	C5	C6				C7 (Ref. data)				C8	C9	C10	C11		
		HPC	HPDA/ AI	auto- motive	relevant NOW	In 5 years	CPU perform ance	I/O	mem. capac.	mem. BW	mem. Lat.	vector units	Acceler ators (Stencil/ Tensor r/..)	Virtual ization	Highly scalable tightly- coupled	embarrassin gly parallel	machine learning	data analytics			x86	COTS ARM	GPU	Power	MPI	Open MP	Open ACC	PGAS	Berke ley socke t					
	Total	39	14	9.3	41.3	43.4	40	6.7	22.85	33.65	17.4	25.75	18.4	1	30.25	9.7	10	9	24.35	33.4	42	25.8	17.45	15.85	30.5	32.6	6	3	0.5	24	26.4	18.55	21.1	
Biophysics	GROMACS	1	0	0	1	1	1	0	0	0	0	1	0.5	0	1	0	0	0	0.3	0.5	1	1	1	0.6	1	1	0	0	0	0.2	1	1	0.5	
	CP2K	1	0	0	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0.5	1	1	0	1	0	0	0	0	0	0	1	0	0		
Biology / Medicine	ALYA	1	0	0	0	1	1	0	1	0	1	0	0	0	1	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	1	0	0	
	Genome assembly	1	0	0	1	1	1	0	1	1	1	1	1	0	0	0	1	1	1	1	1	0	1	0	1	1	0	0	0	0	1	0	0	
Earth sciences / Climate	NEST	1	0	0	1	1	1	0	1	0	1	0	0	0	1	0	0	0	0	1	1	1	1	0	1	1	0	0	0	0	1	0	0	
	EC-EARTH	1	0	0	1	1	1	1	0.1	1	1	0	0	0	1	0	0	0	1	0.8	1	0	0	0	1	1	0	0	0	0	1	0	0	
	ECMWF DWARF	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
	NEMO	1	1	0	1	1	0.5	0.2	0	0.3	0	0	0	0	1	0	0	1	0	1	0	1	1	0	1	1	1	0	0	0	1	1	1	
	SPECFEM3D	1	1	0	1	1	0.6	0.2	0	0	0	0.2	1	0	1	0	0	0	0	1	1	1	1	1	1	1	1	0	1	1	0	1	0	1
	RTM	1	0	0	1	1	1	0	0	1	0	0.6	1	0	1	0.25	0	0	0	1	1	1	1	0	0	0	0	1	0	0.5	0	1	0	
HEP & Fusion	Grid	1	0	0	0	1	1	0	0	1	0	1	1	0	1	0	0	0	1	1	1	1	0	0	1	1	0	0	0	1	0	0	0	
	tmlQCD	1	0	0	1	1	1	0	0	1	0	0	1	0	1	0	0	0	1	1	1	1	0	0	1	1	0	0	0	1	0	0	0	
Material Sciences	GYSELASD	1	0	0	1	1	0.6	0	0	0.3	0.1	0	0	0	1	0	0	0	0	0	0	1	1	0	1	1	0	0	0	0	1	1	0	0
	ABINIT	1	0	0	1	1	0.9	0	0	0.2	0	0.6	0	0	1	0	0	0	1	1	1	0	1	0	1	0.1	0	0	0	0	0	1	0	
	Quantum ESPRESSO	1	0	0	1	1	0.8	0.2	0.3	0.5	0.5	0.5	1	0	1	0.2	0	0	1	1	1	1	1	1	1	1	0	0	0	1	1	1	1	
	BigDFT	1	0	0	1	1	0.6	0	0	0	0	0	0.4	0	1	0	0	0	0	0	1	1	1	1	0	1	1	0	0	0	1	0	0	
CFD	FEniCS	1	0	1	0.5	0.8	0.5	0.5	0.2	1	0	0.2	0.1	0	1	0	0	0	0.3	0.1	1	0.5	0	0.5	1	0	0	1	0	0	0.5	0	0	
	NEK5000	1	0	1	1	1	1	0.5	1	0	0.7	0.8	0	0	1	0	0	0	0.7	1	1	0.5	0	0.5	1	1	1	0	0	0.8	1	1	1	
	OpenFoam	1	0	0	0.6	1	1	0.1	0.5	1	0	0	0	0	1	0	0	0	1	0.5	1	0.8	0.2	0.5	1	0	0	0	0	1	1	1	0.1	
	waLBerla	1	0	0	1	1	1	0	0	1	0	1	1	0	1	0	0	0	1	1	1	0	1	0	1	1	0	0	0	0	1	0	0	
	NAS BT MZ	1	0	0	1	1	1	0	1	1	0	1	0	0	1	0	0	0	0	1	1	1	0	1	1	1	0	0	0	1	0	1	1	
	NAS SP MZ	1	0	0	1	1	1	0	1	1	0	1	0	0	1	0	0	0	0	1	1	1	0	1	1	1	0	0	0	1	0	1	1	
	AVBP	1	0	1	1	1	0.8	0	0	0	0.2	0.1	0	0	1	0	0	0	0	0	1	1	0	0	1	1	1	0	0	0	1	1	0	
Ray tracing / Image / Media	Medical imaging	1	0	0	1	1	1	1	1	1	0	1	1	0	0.5	0.5	0	0	0.2	1	1	0	1	0	0	1	0	0	0	0	1	0	0	
	Image processing/(Trans)coding	0	0	1	0	1	1	0.5	0.3	1	0	1	0.1	0	0	0	0	0	1	1	1	0	1	0	0	1	0	0	0	1	0	0	1	
Big Data Analytics	Spark	0	1	0	1	0	0.5	0	1	0.5	0	0	0	0	0	1	1	1	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	
	Hadoop	0	1	0	1	0	0.5	0	0	0	0	0	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	
	Flink	0	1	0	1	0	0.4	0	0	0.4	0.3	0	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	
	Storm	0	1	0	1	0	0.8	0	0	0.5	0	0	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	
Machine Learning	K-Nearest Neighbours	0	1	0.8	0.7	0.8	1	0	0.8	0.7	0.5	0.6	0	0	0	0	1	1	0	0	0	0	0.5	0.25	0	0	0	0	0	0	0	0	0	
	High Performance Fuzzy Computing	1	1	0.75	0.75	1	1	0	0.75	0.75	0.25	0.75	0.75	0	0.5	0.75	1	1	0	0	0	0	0	1	0.25	0	0	0	0	0	0	0	0	
	Random Forests	1	1	0.75	0.75	1	1	0	1	0.75	0.5	0.75	0	0.5	0.5	1	1	1	0	0	0	0	0	0.75	0.25	0	0	0	0	0	0	0	0	
Deep Learning	Deep500	0	1	0	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	
	DNN kernels	0	1	0	1	1	1	0	0.5	0.8	0	0	1	0	0	1	1	0	1	1	0	0	0	0	0	0.5	0	0	0	1	0.1	0.8	0	
Cloud	VMcontainers	0	0	0	0	0	0.8	0	0.8	0	0	0	0	1	0	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	
DBMS (Data Base)	SQL	0	0	0	1	1	0	1	1	1	0.5	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0	0	1	0	0
	NoSQL	0	0	0	1	1	0	1	1	0.5	1	0	0	0	0	0	0	1	0	1	0	0	0	0	0	0	0	0	0	0.5	0	0.8	0	
Automotive	Apollo Autonomous Driving Framework	0	0	1	1	1	1	0	0	0.5	0.5	1	0	0	0.25	0	1	0	0.1	1	1	1	1	0	0	1	0	0	0	1	1	0.5	0.5	
	Lane tracking	0	0	1	1	1	1	0	0	1	0	1	1	0	0.5	0.5	0	0	1	1	1	1	1	0	0	1	0	0	0	1	1	0	0.5	
Crypto acceleration	OpenSSL	0	0	1	1	1	0.5	0	0.1	0.1	0.1	0	1	0	0	0	0	0	0.25	0.5	1	1	0	0	0	0	0	0	0	0.5	1	0.25	0.5	
Hydrodynamics	Lulesh	1	0	0	0	1	1	0	1	1	1	0	0	0	1	0	0	0	1	1	1	1	0	1	1	1	0	0	0	1	1	1	1	
	HYDRO	1	0	0	0	1	1	0	1	1	1	1	0	0	1	0	0	0	1	1	1	1	0	1	1	1	0	0	0	1	1	0	1	
PDE	minighost	1	0	0	0	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	1	0	1	0	1	0	1	1	0	0	0	0	0	
Reference benchmarks	HPL	1	0	0	1	1	1	0	0.5	0.1	0	1	0	0	0	1	0	0	0	0	0	0	0	0	1	1	0	0	0	1	0	1	1	
	HPCG	1	0	0	1	0.8	0.5	0	0.5	1	0.1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	0	0	0	1	0	1	1	
	Vector SpMV	1	1	0	1	1	0	0	1	1	1	1	0.5	0	1	0	1	0	1	1	1	1	0	1	1	1	0	0	0	1	1	0	1	
	DGEMM	1	1	0	1	1	1	0	1	1	1	1	0.5	0	1	0	1	0	1	1	1	1	0	1	1	1	0	0	0	1	1	0	1	
	Somier	1	0	0	1	1	1	0	1	1	1	1	0	0	1																			

EPI WP1 BENCHMARK SUITE

- Partially public:
<https://gitlab.version.fz-juelich.de/epi-wp1-public>
- Public repositories include
 - Vanilla version
 - Ported version (currently SVE only)
- Non-public repositories include
 - Non-public codes
 - Built and test scripts

EPI WP1 GPP SIMULATOR ENVIRONMENT

- Gem5 simulator (JSC+Forth)
 - Relatively accurate, but **very** slow → simulation of a few cores
- SESAM simulator (CEA)
 - Simulator for multi-core system-on-a-chip architecture
 - Simulations at different levels of detail possible
- MUSA simulator (BSC)
 - Multi-level simulator for simulations at different scales

CURRENT CO-DESIGN ACTIVITIES

- Exploration of different GPP configurations
 - Cache configuration
 - Number of cores and clock speeds
 - Network-on-chip configurations and HBM attachment
- Exploitation of accelerators
 - Limitation: Programming models and simulators at an early stage
- Approach: Focus on a small set of synthetic benchmarks and application kernels
 - xgemm, stream, FFT, OpenFlowLB, RTM kernels

NEW ARM SIMD ISA: SCALABLE VECTOR EXTENSION (SVE)

■ Key feature: Vector length agnostic

- Vector length not defined at compile time
- Multiple lengths supported by ISA: 128, 256, ..., 2048 bit

■ Required hardware support for VLA

- Update of predication registers
- Update of loop counters

```
void daxpy(double a, double *restrict x,
          double *restrict y, int n)
{
    int i;
    for (i = 0; i < n; i++)
        y[i] += a * x[i];
}
```

Arm clang 19.3

```
mov     x8, xzr
mov     z0.d, d0
whilelo p1.d, xzr, x9
ptrue   p0.d

.LBB0_2:
ld1d {z1.d}, p1/z, [x0, x8, lsl #3]
ld1d {z2.d}, p1/z, [x1, x8, lsl #3]
fmad z1.d, p0/m, z0.d, z2.d
st1d {z1.d}, p1, [x1, x8, lsl #3]
incd x8
whilelo p1.d, x8, x9
b.mi .LBB0_2

.LBB0_3:
ret
```

SVE: PORTING OPTIONS

■ Auto-vectorisation

- Rely on compiler to identify opportunities for vectorisation
 - Compilers: gcc-8, clang-8, armclang 19.x (or more recent)
- Pro: portable code
- Con: Compiler still often fail to recognise opportunities for vectorisation

■ ARM C Language Extensions (ACLE) for SVE

- Makes features of the ARM architecture directly available in C and C++ programs
https://static.docs.arm.com/100987/0000/acle_sve_100987_0000_00_en.pdf
 - Supported by gcc, armclang (not clang)
- Pro: Control on code generation
- Con: Non-portable code

SVE: COMPILING AND EXPLORING

- Need access to Arm-based platform

- E.g. Juawei prototype at JSC <https://trac.version.fz-juelich.de/armlab/wiki/Public/Account>

- Compilation and emulation

- Allow compiler to generate SVE instructions, e.g. `gcc -march=armv8-a+sve ...`
 - Use Arm Instruction Emulator for execution, e.g. `armie -msve-vector-bits=256 ...`

- Performance analysis

- Initial assessment based on instruction count:
`armie -msve-vector-bits=256 --iclient libinscount.so -- ./myprog.x`
 - Full performance analysis requires simulation

OPTIONS FOR FUTURE ACTIVITIES

- Provisioning of expertise on co-design applications
 - E.g. on workloads
- Porting to Arm SVE (and later accelerators)
 - Access to Arm-based platforms can be provided
 - EPI can help with links to other initiatives supporting Arm enablement
- Performance exploration
 - Access to simulators in very selected cases
 - Beware of required expertise and effort

PARTNERS INVOLVED IN EPI CO-DESIGN



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA



UNIVERSITÀ DI PISA

