


CO-DESIGN IN EPI

16 MAY 2019 – EUROHPC SUMMIT WEEK
ESTELA SUAREZ & DIRK PLEITER (JSC)



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**
 - 1). analytical models, high level
 - 2). simulation based (e.g. gem5 simulation engine)
 - 3). reference platform (e.g. Marvell ThunderX2)

CO-DESIGN PROCESS



APPLICATION SELECTION CRITERIA

C1	Relevant (now or in 5-years) for markets addressed by EPI	HPC, HPDA/AI, automotive
C2	Its requirements covers architectural components	CPU perf., I/O, Mem-Cap., Mem-BW, Mem-Lat., vector units, dedicated accel., virt.
C3	Represent a family/class of applications	Highly scalable tightly-coupled, embarrassingly parallel, ML, data analytics
C4	Close relation to code developers	
C5	Licence allows development of mini-apps/derived benchmarks (preferable OpenSource)	
C6	Reference data available from other platforms	x86, commercial ARM, GPGPU, Power
C7	Application uses/covers a software component / programming model relevant for an EPI market	MPI, OpenMP, OpenAcc, PGAS, Berkeley Socket, JVM
C8	Application features relatively simple kernel	
C9	High societal impact	
C10	Part of an existing benchmark suite , or widely known	
C11	Mini-app or kernels already available	

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							
		HPC	HPDA/ AI	auto- motive	rele- vant NOW	In 5 years	CPU perfor- mance	I/O	mem. capac.	mem. BW	mem. Lat.	vector units	Accel- erators (Stencil/ Tensor- r/..)	Virtuali- zation	Highly scalable tighty- coupled	embarrassin- gly parallel	machine learning	data analytics	x86	COTS ARM	GPU	Powe- r	MPI	Open MP	Open ACC	PGAS	Berke- ley socket										
	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	0	1	0	0				
Biology / Medicine	ALYA	1	0	0	0	0	1	1	0	1	1	1	1	0	0	0	1	1	1	1	1	1	0	1	0	1	1	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	1	0	1	1	0	1	1	0	0	0	1	0	0			
Earth sciences / Climate	EC-EARTH	1	0	0	1	1	1	1	0.1	1	1	0	0	0	0	1	0	0	1	0.8	1	0	0	0	1	1	0	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	1	1	1		
	NEMO	1	1	0	1	1	0.5	0.2	0	0.3	0	0	0	0	1	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	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	1	1	1	1	1	1	1			
HEP & Fusion	RTM	1	0	0	1	1	1	0	0	1	0	0.6	1	0	1	0.25	0	0	1	1	1	1	1	0	0	0	0	0	1	0	0.5	0	1	0			
	Grid	1	0	0	0	1	1	1	0	0	1	0	1	1	0	0	0	0	1	1	1	1	1	0	0	0	1	0	0	0	1	0	0	0			
	tmdQCD	1	0	0	1	1	1	0	0	1	0	0	0	1	0	0	0	0	1	1	0	1	0	1	0	0	0	1	0	0	1	0	0	0			
	GYSELA5D	1	0	0	1	1	0.6	0	0	0.3	0.1	0	0	0	1	0	0	0	0	1	1	0	1	1	0	0	0	0	0	1	0	0	0				
Material Sciences	ABINIT	1	0	0	1	1	0.9	0	0	0.2	0	0.6	0	0	1	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	1	1		
	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	0	0	0	0	0	0	0	1	1			
CFD	BigDFT	1	0	0	1	1	0.5	0.8	0.5	0.5	0.2	1	0	0.2	0.1	0	1	0	0	0.3	0.1	1	0.5	0	0.5	1	0	0	0	0	0	0.5	0	0			
	FFnICS	1	0	1	1	0.5	0.8	0.5	0.5	0.2	1	0	0	0.2	0.1	0	1	0	0	0	0.7	1	1	0.5	0	0.5	1	1	1	0	0	0.8	1	1			
	NEKS000	1	0	1	1	1	1	1	0.5	1	1	0	0.7	0.8	0	1	0	0	0	0.7	1	1	0.5	0	0.5	1	1	1	0	0	0.8	1	1				
	OpenFoam	1	0	0	0	0.6	1	1	0.1	0.5	1	0	0	0	0	0	1	0	0	0	1	1	1	0	0	0	0	0	0	1	0	0	0				
	walBerla	1	0	0	0	1	1	1	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	0			
	NAS BT MZ	1	0	0	0	1	1	1	0	1	1	0	1	0	0	1	0	0	0	0	1	1	1	0	1	1	0	0	0	0	1	0	1	1			
Ray tracing / Image / Media	NAS SP MZ	1	0	0	0	1	1	1	0	1	1	0	1	0	0	1	0	0	0	1	1	1	1	1	0	0	0	0	0	1	0	1	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	1	0	0	1	1	0	0	0	1	1	0			
Big Data Analytics	Medical imaging	1	0	0	1	1	1	1	1	0	1	1	0	0.5	0.5	0	0	0	0.2	1	1	0	1	0	0	1	0	0	0	0	0	0	0	0			
	Image processing//Transcoding	0	0	1	0	1	0	1	0.5	0.3	1	0	1	0.1	0	0	0	0	1	1	1	0	1	0	0	1	0	0	0	1	0	0	0	1			
Machine Learning	Spark	0	1	0	1	0	0.5	0	1	0.5	0	0	0	0	0	0	1	1	1	0	0	0	1	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	0	1	0	0	0	0	1	0	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	0			
	Storm	0	1	0	1	0	0.8	0	0.5	0	0	0.5	0	0	0	0	1	0	0	0	1	0	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	0	0	1	1	0	0	0	0	0	0.5	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	1	0	0	0	0	1	0.25	0	0	0	0	0	0	0	0	0	0	0	
Deep Learning	Random Forests	1	1	0.75	0.75	1	1	0	1	1	0.75	0.5	0.75	0	0.5	0.5	1	1	1	0	0	0	0	0	0	0.75	0.25	0	0	0	0	0	0	0	0	0	0
	Deep500	0	1	0	0	0	1	0	0	0	0	0	1	0	0	0	0	1	0	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Cloud	DNN kernels	0	1	0	1	1	1	0	0.5	0.8	0	0	0	1	0	0	0	1	1	0	0	0	0	0	0	0	0.5	0	0	0	0	1	0.1	0.8	0		
	VMcontainers	0	0	0	0	1	1	0.8	0	0.8	0	0	0	1	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
DBMS (Data Base)	SQL	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	0	0	0	0	1	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	0	0	0	0	0	0.5	0.8
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	0	1	0	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	1	0	0	1	1	1	1	1	0	0	0	1	0	0	0	0	1	1	0.5	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	1	0	0	0	0	0	0	0	0	0	0.5	1	0.25	0.5	
	Hydrodynamics	1	0	0	0	1																															

EPI PROCESSOR CO-DESIGN PARAMETERS

- General purpose cores
 - SVE length
 - Number of SVE pipelines per core
- Accelerator cores
 - Ratio of accelerator cores versus general purpose cores
 - Accelerator core design
- NOC and memory sub-system
 - NOC topology and bandwidth
 - Cache sizes and bandwidth
 - HBM size and bandwidth
 - DDR size and bandwidth

Other parameters to be considered for system level co-design


- Network architecture
- I/O architecture

EXAMPLE: DGEMM

L2 cache
exploration with
BLIS microkernel

CREDITS:

- P.Petrakis, V. Papaefstathiou et al. (FORTH): simulation execution an analysis
- B.Brank, S.Nassyr (FZJ): BLIS micro-kernel
- A.Portero (FZJ): Gem5 simulator setup




EXAMPLE: DGEMM

SVE
microkernel
instruction
mix

- To get peak numbers we need to feed the 2 SVE units with 2 FMA instruct. per cycle.
- For a 4-wide core the code must have more than 50% FMA instructions in the mix.
- Peak value at: ~2.3 FMA/cycle and ~1.45 loads/cycle

CREDITS:

- P.Petrakis, V. Papaefstathiou et al. (FORTH): simulation execution an analysis
- B.Brank, S.Nassyr (FZJ): BLIS micro-kernel
- A.Portero (FZJ): Gem5 simulator setup



CO-DESIGN IN JAPANESE EXASCALE PROGRAM

- SPIRE (Strategic Programs for Innovative Research) selected 5 strategic application areas
 - Large fraction of compute resources will be dedicated to SPIRE
 - Selected areas: Life science/drug manufacturing, new materials/energy creation, Global change prediction for disaster prevention/mitigation, manufacturing technology, the origin of matter and the universe
 - 9 social and scientific priority issues identified
- Selected 9 target apps from each area for “co-design”

Program	Target Application	
		Brief description
① GENESIS		MD for proteins
② Genomon		Genome processing (Genome alignment)
③ GAMERA		Earthquake simulator (FEM in unstructured & structured grid)
④ NICAM+LETK		Weather prediction system using Big data (structured grid stencil & ensemble Kalman filter)
⑤ NTChem		molecular electronic (structure calculation)
⑥ FFB		Large Eddy Simulation (unstructured grid)
⑦ RSDFT		an ab-initio program (density functional theory)
⑧ Adventure		Computational Mechanics System for Large Scale Analysis and Design (unstructured grid)
⑨ CCS-QCD		Lattice QCD simulation (structured grid Monte Carlo)

CO-DESIGN IN US EXASCALE PROGRAM


- 10 research areas driven by national labs
 - Nuclear energy, climate, carbon capture and storage, wind energy, combustion, magnetic fusion energy, advanced manufacturing, chemical sciences, precision medicine for cancer, urban system sciences, cosmology, high energy and nuclear physics, accelerator physics, astrophysics, geoscience, metagenomics, nuclear materials, seismic, power grid
- 5 co-design centres for different application motifs
 - CODAR: Co-Design Center for Online Data Analysis and Reduction; COPA: Co-Design Center for Particle Applications; AMReX: Block-Structured AMR Co-Design Center; CEDD: Center for Efficient Exascale Discretizations; ExaGraph: GraphEx Co-Design Center; ExaLearn: Co-Design Center for Exascale Machine Learning Technologies
- ECP proxy applications program <https://proxyapps.exascaleproject.org/>
 - Option to submit more applications
 - Standards for proxy-applications defined

Proxy App
AMG
CANDLE Benchmarks
Ember
ExaMiniMD
Laghos
MACSio
miniAMR
miniQMC
miniVite
NEKbone
PICSARlite
SW4lite
SWFFT
thornado-mini
XSBench

PARTNERS INVOLVED IN EPI CO-DESIGN



Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación



UNIVERSITÀ DI PISA

