Towards EXtreme scale Technologies and Accelerators for euROhpc hw/Sw Supercomputing Applications for exascale

Pasqua D'Ambra

pasqua.dambra@cnr.it

Project PI for CNR and Leader of Task 6.4

Institute for Applied Computing (IAC), Napoli branch
National Research Council of Italy (CNR)

Kick-off meeting, May 5th, 2021

IAC-CNR profile

National Research Council (CNR) is the largest public research institution in Italy. Its scientific structure includes 88 Institutes splitted in 7 scientific Departments with leading expertise in all science sectors

Institute for Applied Computing "Mauro Picone" (IAC) is the more ancient institute of CNR. Currently, it has headquarters in **Roma** and 3 branches in Bari, Firenze and **Napoli**.

The permanent staff has a total of 67 employers, including 50 researchers and technologists (applied mathematicians, physicists, computer scientists) and 17 technicians and administrative officers. Many associate researchers from Universities, Post-Doc research fellows and students (both Master and Ph.D) are involved in the research activities.

Mission of the Institute is to develop **advanced mathematical, statistical and computational methods** for addressing challenging problems with strong relevance to society and industry

www.iac.cnr.it

HPC Research Group @IAC-CNR in Textarossa

Massimo Bernaschi

Pasqua D'Ambra

Fabio Durastante

Salvatore Filippone

Mauro Carrozzo

Dario Pasquini

Gabriele Salvati

HPC Research Group @IAC-CNR in Textarossa

The HPC research Group @IAC-CNR has deep expertise in numerical methods, parallel algorithms and mathematical software tools for High-Performance Scientific Computing

Main current activities focus on the design and development of **highly scalable algorithms and software for sparse linear algebra, leveraging on *de facto* standard programming models and tools for hybrid architectures**

Current research projects:

- 2019-2022: EoCoE-II, Energy oriented Center of Excellence: toward exascale for energy. Role: Leader of WP on Scalable Solvers
- Industrial projects in crypto-analysis and security-related data analysis

Our Role in TEXTAROSSA

Design and develop kernels of a MathLib

initial analysis and co-design in WP1 and full development in WP6

Background Technologies for TEXTAROSSA

BootCMatchG

(Bootstrap AMG based on Compatible Weighted Matching for hybrid CPU/Nvidia GPU)

A software package based on an adaptive Algebraic MultiGrid Method
with a user defined desired convergence rate
for highly accurate solution of sparse linear systems
exploiting hybrid CPU/NVIDIA-GPU nodes
available at www.github.com/bootcmatch (TRL 3.5)

- M. Bernaschi, P. D'Ambra, D. Pasquini, AMG based on compatible weighted matching on GPUs, Parallel Computing. Vol. 92, 2020
- M. Bernaschi, P. D'Ambra, D. Pasquini, BootCMatchG: an Adaptive Algebraic MultiGrid Linear Solver for GPUs. Software Impacts. Vol. 6, 2020 (Invited Paper)

BootCMatchG at work on Nvidia Titan V

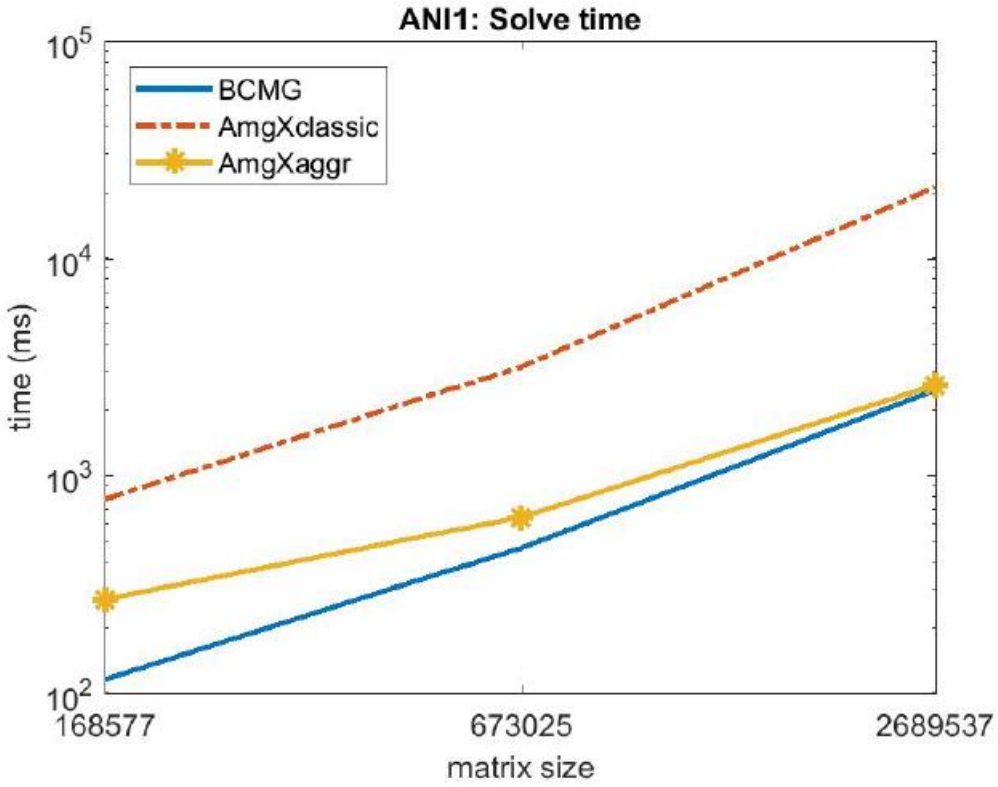

Solving systems up to 2.7 millions of dofs in ~ 1 sec

Our solver outperforms NVIDIA AmgX library implementing similar methods

Background Technologies for TEXTAROSSA

PSCToolkit (Parallel Sparse Computation Toolkit)

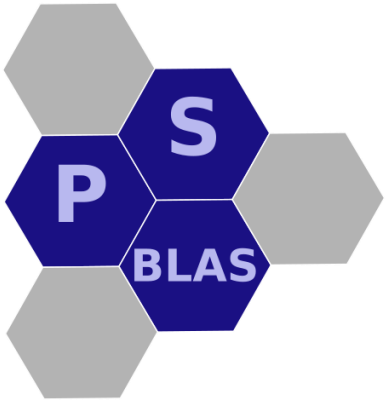

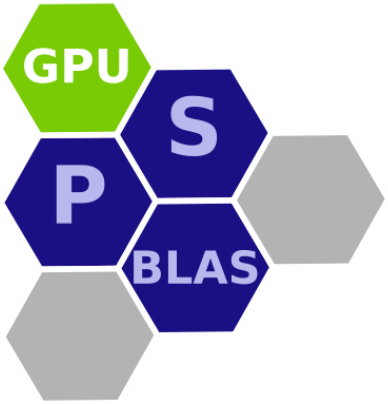

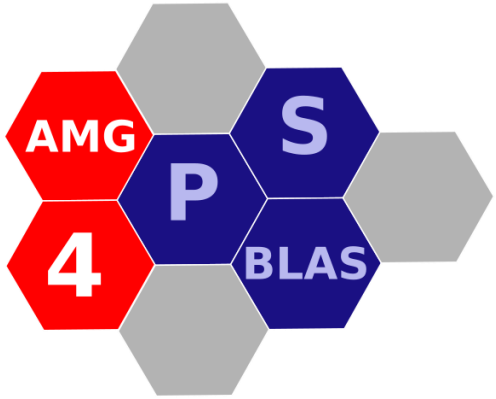

A software framework for scalable solution of sparse linear systems at extreme scale on hybrid CPU/NVIDIA-GPU architectures

available at psctoolkit.github.io (TRL 3.5)

Recently selected as Excellent Innovation by EU Innovation Radar

P. D'Ambra, F. Durastante, S. Filippone, AMG Preconditioners for Linear Solvers towards Extreme Scale, to appear on SIAM Journal on Scientific Computing
Preprint available at [2006.16147] (arxiv.org)

PSCToolkit at work on Piz Daint (Tier0-PRACE)

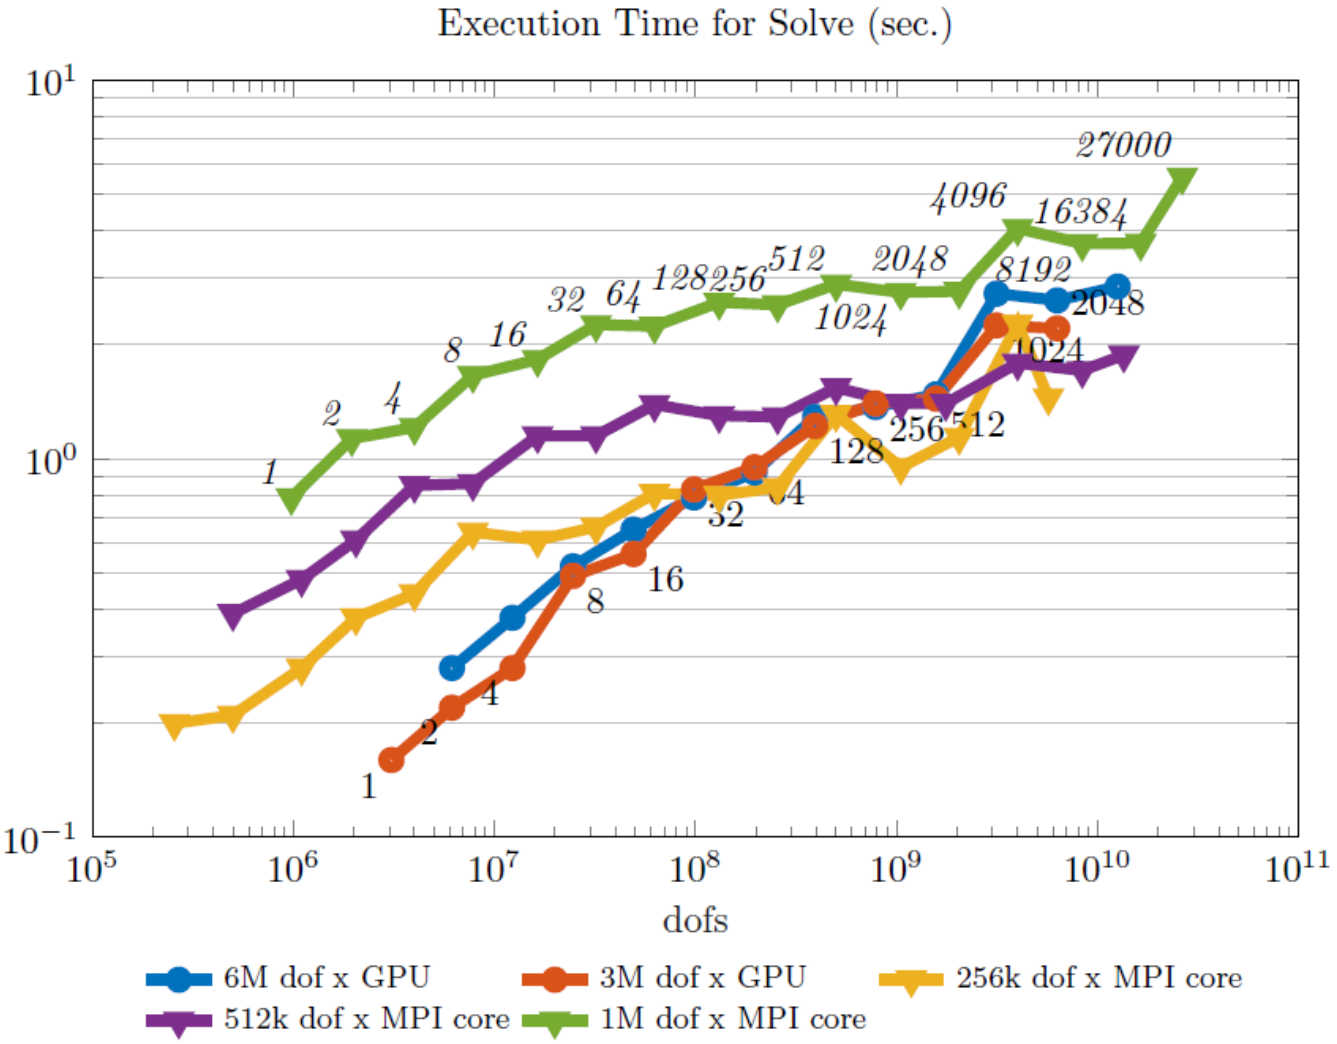

Solving systems up to 10 billions of dofs

in ~4 sec on 27700 CPUs and in ~3 sec on 2048 GPUs

Role and use of background, plus project goals

Optimization and acceleration of algorithms looking for trade-off between energy-to-solution and time-to-solution promoting reproducibility and wide scalability

Focus on **key numerical building blocks for (memory/communication-bound) sparse matrix/graph computations**, including:

- sparse matrix-sparse matrix and sparse matrix-vector products (SPMM & SPMV)
- maximum weight/maximum cardinality (sparse) graph matching (MWMCM)
- Communication Avoiding Krylov-type iterative methods (CA-KM): Preconditioned Conjugate Gradient (PCG) and Preconditioned Generalized Minimal Residual (PGMRES)
- Algebraic Multigrid Preconditioners based on aggregation (AMG)

  - **Target Platform: Multi node systems equipped with accelerators and connected through state-of-the-art network technologies**

Ambition

Design, Develop and Test a first kernel of TEXTAROSSA MathLib

- Empower the proposed energy-efficient HPC framework with efficient kernels of a mathematical library, designed and developed **following a co-design approach**

- Design, develop and test new numerical methods and algorithms to ensure **scalable and reduced energy-power for main compute/data-intensive kernels** of HPC/HPDA/AI applications

KPI: energy efficiency and scalability for dealing sparse linear systems up to 10^{13} – 10^{15} dofs on hybrid architectures

Detailed activity and interaction with other WPs

- **performance analysis of memory/communication-bound kernels in sparse matrix/graph computation for bi-directional co-design** with impact on new HW/SW platforms proposed in the project (develop in WP1, interaction with WP2/WP4/WP5)

- **definition of a set of benchmark modules and test cases for final assessment of project results** (develop in WP1, interaction with WP2/WP4/WP5)

- **use of heterogeneous HW** (target platform GPU with some experiments on FPGA); **re-design of algorithms for possible use of mixed-precision**; **experiments with test bed of programming environments and run-time supports** for best trade-off between load-balancing and communication, **final tests and assessment of results** (develop in WP6, interaction with WP2/WP4/WP5)

- Communication and dissemination activities (WP7)

Questions and answers

Any questions?