

Title: Performance portability of GW and many-body perturbation theory codes: the case of Yambo on GPUs

Author: **Nicola Spallanzani** - National Research Council CNR Nano - Modena, Italy

Abstract:

The YAMBO code [1,2] has been installed and tested on a large number of HPC architectures including homogeneous multi-core systems. Recently, an extensive activity of porting of Yambo on heterogeneous architectures, GPUs in particular, has been put in place. This has addressed the kernels computing dipoles, Coulomb cutoff, Hartree-Fock, linear response, GW and Bethe Salpeter equation (BSE). Technically, the porting has been achieved by taking advantage of CUDA Fortran, a programming model which provides a native support for NVIDIA architectures, and exploiting both cuf-kernels directives and CUDA libraries such as cublas, cufft, and cusolver. In this talk will be presented the results obtained. The GPU-aware version of Yambo has just been released to the user community with v4.5.0.

[1] A. Marini et al, Comp. Phys. Comm. (2009)

[2] D. Sangalli et al, J. Phys. Cond. Mat. (2019)