



UNIVERSITÀ DI PARMA

# GPU accelerated Quantum ESPRESSO

Pietro Bonfà

Department of Mathematical, Physical and Computer Sciences, University of Parma, Italy

14 Jan. 2020 - Jülich



EUROPEAN CENTER OF EXCELLENCE - A H2020 E-INFRASTRUCTURE

H2020-INFRAEDI-2018-1 G.A. 824143

# Topics

---

What's in this presentation:

- ✓ Evolution of **GPU acceleration** in QE: strategies and motivation.
- ✓ **Domain Specific Libraries**, that you may want to check for your own codes.
- ✓ A few details about the programming model used for QE.
- ✗ How to run QE on GPU accelerated platforms
- ✗ How to develop QE on GPU accelerated platforms

# Acknowledgements





<https://gitlab.com/QEF/q-e-gpu>




Collaboration and support from: J. Romero, M. Marić, M. Fatica, E. Phillips (NVIDIA)  
F. Spiga (ARM), A. Chandran (FZJ), I. Girotto (ICTP), Y. Luo (ANL), T. Kurth (NVIDIA), B. Cook (NERSC),  
F. Ferrari, P. Giannozzi (Univ. Udine), A. Ferretti (CNR-S3 Modena), P. Delugas, S. De Gironcoli (SISSA).

# MaX CoE

Mission: materials science codes ready for exascale computing.



**SIRIUS**



EUROPEAN CENTER OF EXCELLENCE - A H2020 E-INFRASTRUCTURE

# What is QE

---

## QUANTUM ESPRESSO

is an integrated suite of Open-Source computer codes for electronic-structure calculations and materials modeling at the nanoscale. It is based on density-functional theory, plane waves, and pseudopotentials.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PW91, PBE, B88-P86, BLYP) to meta-GGA, exact exchange (HF) and hybrid functionals (PBE0, B3LYP, HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdw-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

## Structural Optimization:

- GDIIIS with quasi-Newton BFGS preconditioning;
- Damped dynamics.

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;
- Electron-phonon interactions;
- Third-order anharmonic phonon lifetimes, using the [D3Q](#) package;
- Infrared and (non-resonant) Raman cross-sections;
- EPR and NMR chemical shifts, using the [QE-GIPAW](#) package.
- Phonons for 2D heterostructures ([reference](#))

## Spectroscopic properties:

- $K$ -,  $L1$  and  $L2,3$ -edge X-ray Absorption Spectra (XSpectra package);
- Time-Dependent Density Functional Perturbation Theory (TurboTDDFT package);
- Electronic excitations with Many-Body Perturbation Theory, using the [YAMBO](#) package;
- Electronic excitations with Many-Body Perturbation Theory (GWL package).

## Quantum Transport:

- Ballistic Transport (PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIER90](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PW91, PBE, B88-P86, BLYP) to meta-GGA, exact exchange (HF) and hybrid functionals (PBE0, B3LYP, HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdw-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

package);  
AMBO package;  
ige).

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Quantum Transport:

- Ballistic Transport ( PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIER90](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PW91, B3LYP,

## Structural Optimization:

- GDIIS with quasi-Newton BFGS preconditioning;
- Damped dynamics.

- Damped dynamics.

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;
- Electron-phonon interactions;
- Third-order anharmonic phonon lifetimes, using the [D3Q](#) package;
- Infrared and (non-resonant) Raman cross-sections;
- EPR and NMR chemical shifts, using the [QE-GIPAW](#) package.
- Phonons for 2D heterostructures ([reference](#))

## Spectroscopic properties:

- $K$ -,  $L1$  and  $L2,3$ -edge X-ray Absorption Spectra (XSpectra package);
- Time-Dependent Density Functional Perturbation Theory (TurboTDDFT package);
- Electronic excitations with Many-Body Perturbation Theory, using the [YAMBO](#) package;
- Electronic excitations with Many-Body Perturbation Theory (GWL package).

## Quantum Transport:

- Ballistic Transport (PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIER90](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.



# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PW91, PBE, B88-P86, BLYP) to meta-GGA, exact exchange (HF) and hybrid functionals (PBE0, B3LYP, HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdw-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;
- Electron-phonon interactions;
- Third-order anharmonic phonon lifetimes, using the [D3Q](#) package;
- Infrared and (non-resonant) Raman cross-sections;
- EPR and NMR chemical shifts, using the [QE-GIPAW](#) package.
- Phonons for 2D heterostructures ([reference](#))

## Spectroscopic properties:

- $K$ -,  $L1$  and  $L2,3$ -edge X-ray Absorption Spectra (XSpectra package);
- Time-Dependent Density Functional Perturbation Theory (TurboTDDFT package);
- Electronic excitations with Many-Body Perturbation Theory, using the [YAMBO](#) package;
- Electronic excitations with Many-Body Perturbation Theory (GWL package).

## Quantum Transport:

- Ballistic Transport (PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIER90](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PW91, PBE, B88-P86, BLYP) to meta-GGA, exact exchange (HF) and hybrid functionals (PBE0, B3LYP, HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdw-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

## Structural Optimization:

- GDIIIS with quasi-Newton BFGS preconditioning;
- Damped dynamics.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;
- Electron-phonon interactions;
- Third-order anharmonic phonon lifetimes, using the [D3Q](#) package;
- Infrared and (non-resonant) Raman cross-sections;
- EPR and NMR chemical shifts, using the [QE-GIPAW](#) package.
- Phonons for 2D heterostructures ([reference](#))

## Spectroscopic properties:

- $K$ -,  $L1$  and  $L2,3$ -edge X-ray Absorption Spectra (XSpectra package);
- Time-Dependent Density Functional Perturbation Theory (TurboTDDFT package);
- Electronic excitations with Many-Body Perturbation Theory, using the [YAMBO](#) package;
- Electronic excitations with Many-Body Perturbation Theory (GWL package).

## Quantum Transport:

- Ballistic Transport (PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIER90](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PBE, B88-P86, BLYP) to meta-GGA, exact exchange (HF) and hybrid functionals (PBE0, HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdw-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

## Structural Optimization:

- GDIIIS with quasi-Newton BFGS preconditioning;
- Damped dynamics.

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;
- Electron-phonon interactions;
- Third-order anharmonic phonon lifetimes, using the [D3Q](#) package;
- Infrared and (non-resonant) Raman cross-sections;
- EPR and NMR chemical shifts, using the [QE-GIPAW](#) package.
- Phonons for 2D heterostructures ([reference](#))

- Ballistic Transport (PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIERTOOLS](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to general gradient approximation (GGA), PBE, B88-P86, BLYP to meta-GGA, exact exchange (HF) and hybrid functionals (HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdW-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

## Structural Optimization:

- GDIIIS with quasi-Newton BFGS preconditioning;
- Damped dynamics.

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;

## Spectroscopic properties:

- $K$ -,  $L1$  and  $L2,3$ -edge X-ray Absorption Spectra (XSpectra package);
- Time-Dependent Density Functional Perturbation Theory (TurboTDDFT package);
- Electronic excitations with Many-Body Perturbation Theory, using the [YAMBO](#) package;
- Electronic excitations with Many-Body Perturbation Theory (GWL package).

## Quantum transport:

- Ballistic Transport (PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIERTOOLS](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PW91, PBE, B88-P86, BLYP) to meta-GGA, exact exchange (HF) and hybrid functionals (PBE0, B3LYP, HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdw-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

## Structural Optimization:

- GDIIIS with quasi-Newton BFGS preconditioning;
- Damped dynamics.

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;
- Electron-phonon interactions;
- Third-order anharmonic phonon lifetimes, using the [D3Q](#) package;
- Infrared and (non-resonant) Raman cross-sections;
- EPR and NMR chemical shifts, using the [QE-GIPAW](#) package.
- Phonons for 2D heterostructures ([reference](#))

## Quantum Transport:

- Ballistic Transport ( PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIER90](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEK](#) code.

# What is QE

## Ground-state calculations:

- Self-consistent total energies, forces, stresses;
- Kohn-Sham orbitals;
- Separable norm-conserving and ultrasoft (Vanderbilt) pseudo-potentials, PAW (Projector Augmented Waves);
- Several exchange-correlation functionals: from LDA to generalized-gradient corrections (PW91, PBE, B88-P86, BLYP) to meta-GGA, exact exchange (HF) and hybrid functionals (PBE0, B3LYP, HSE);
- VdW corrections (DFT-D) or nonlocal VdW functionals (vdw-DF);
- Hubbard U (DFT+U);
- Berry's phase polarization;
- Spin-orbit coupling and noncollinear magnetism.

## Structural Optimization:

- GDIIIS with quasi-Newton BFGS preconditioning;
- Damped dynamics.

## Transition states and minimum energy paths

- Nudged Elastic Band method;
- Meta-Dynamics, using the [PLUMED](#) plug-in.

## Ab-initio molecular dynamics

- Car-Parrinello Molecular Dynamics (CP package);
- Born-Oppenheimer Molecular Dynamics (PWscf package).

## Response properties (density-functional perturbation theory):

- Phonon frequencies and eigenvectors at any wavevector;
- Full phonon dispersions; inter-atomic force constants in real space;
- Translational and rotational acoustic sum rules;
- Effective charges and dielectric tensors;
- Electron-phonon interactions;
- Third-order anharmonic phonon lifetimes, using the [D3Q](#) package;
- Infrared and (non-resonant) Raman cross-sections;
- EPR and NMR chemical shifts, using the [QE-GIPAW](#) package.
- Phonons for 2D heterostructures ([reference](#))

## Spectroscopic properties:

- $K$ -,  $L1$  and  $L2,3$ -edge X-ray Absorption Spectra (XSpectra package);
- Time-Dependent Density Functional Perturbation Theory (TurboTDDFT package);
- Electronic excitations with Many-Body Perturbation Theory, using the [YAMBO](#) package;
- Electronic excitations with Many-Body Perturbation Theory (GWL package).

## Quantum Transport:

- Ballistic Transport (PWCOND package);
- Coherent Transport from Maximally Localized Wannier Functions, using the [WanT](#) code;
- Maximally-localized Wannier functions and transport properties, using the [WANNIER90](#) code.
- Kubo-Greenwood electrical conductivity using the [KGEC](#) code.

...and several post processing tools by many research groups!

# What is QE

---

- QUANTUM ESPRESSO is an initiative coordinated by the QUANTUM ESPRESSO Foundation, with the participation of SISSA, CINECA, ICTP, EPFL and *many partners in Europe and worldwide*.
- QUANTUM ESPRESSO is not a single application for quantum simulations; it is rather a distribution of packages performing different tasks and meant to be interoperable.
- Free as in GPLv2 and open development.


# What is QE

- Runs from standalone workstation to massively parallel systems.

```
$ ./configure && make all
```

- Large scientific user base, vehicle for new methods, new theories and new science.

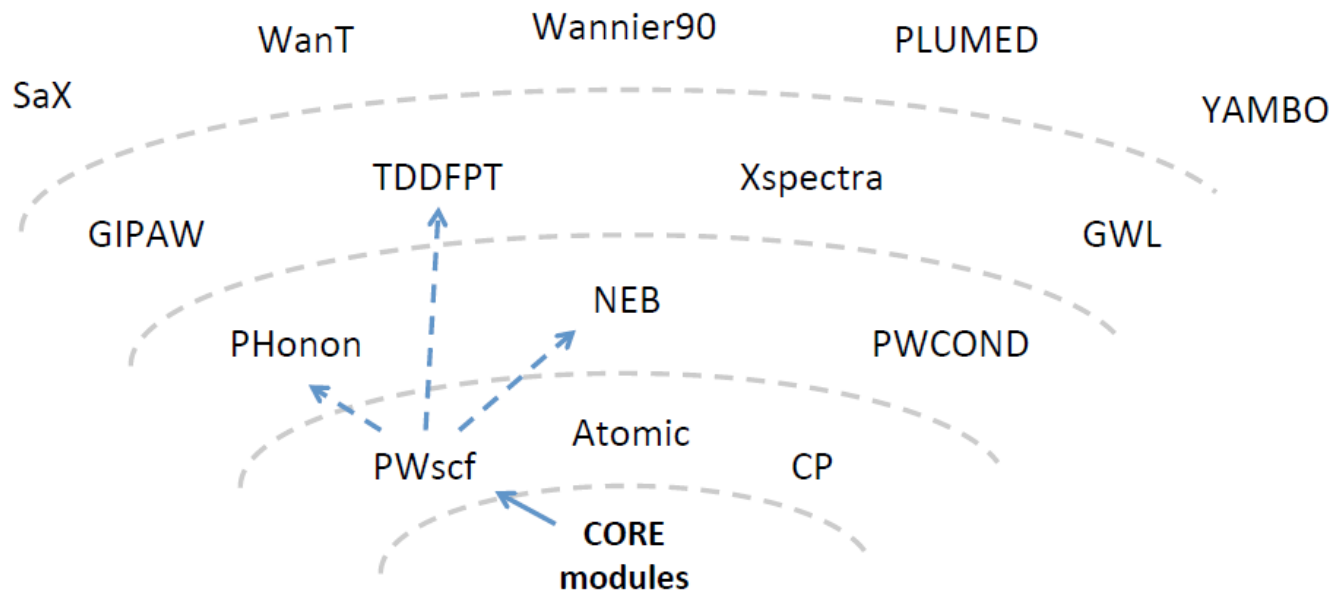
- 600k lines of Fortran
- V6.5 -> 10k downloads
- >50 contributors
- 1600+ registered users
- ...



• Simplify transition of new science to HPC environment.



# What is QE



# What is QE

Some of the time consuming workloads of many packages are already encapsulated in a number of libraries, namely

**LAXLib**

**FFTXlib**

**KS\_Solvers**

**ELPA**


L	A	P	A	C	K
L	-A	P	-A	C	-K
L	A	P	A	-C	-K
L	-A	P	-A	-C	K
L	A	-P	-A	C	K
L	-A	-P	A	C	-K



FFTW, MKL, ESSL, ...

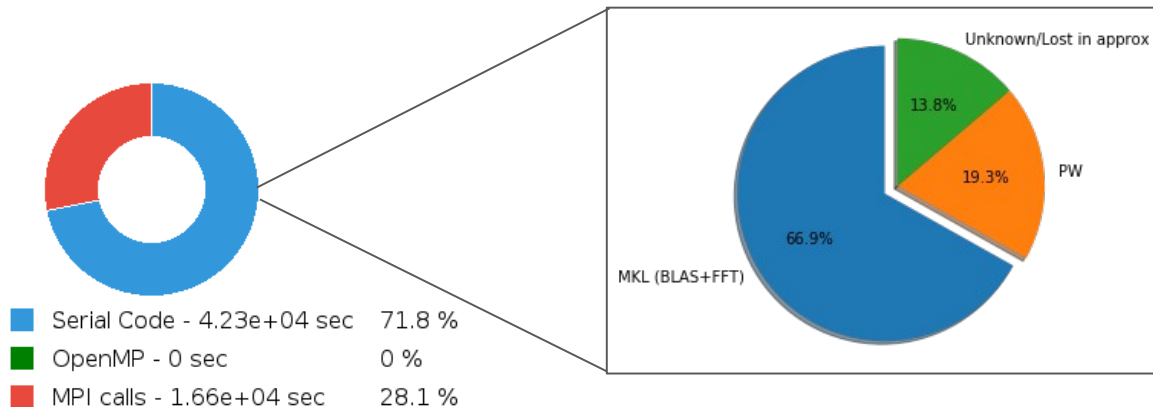


$$|\delta\psi_i\rangle = \frac{1}{D - \epsilon_i}(H - \epsilon_i)|\psi_i\rangle$$



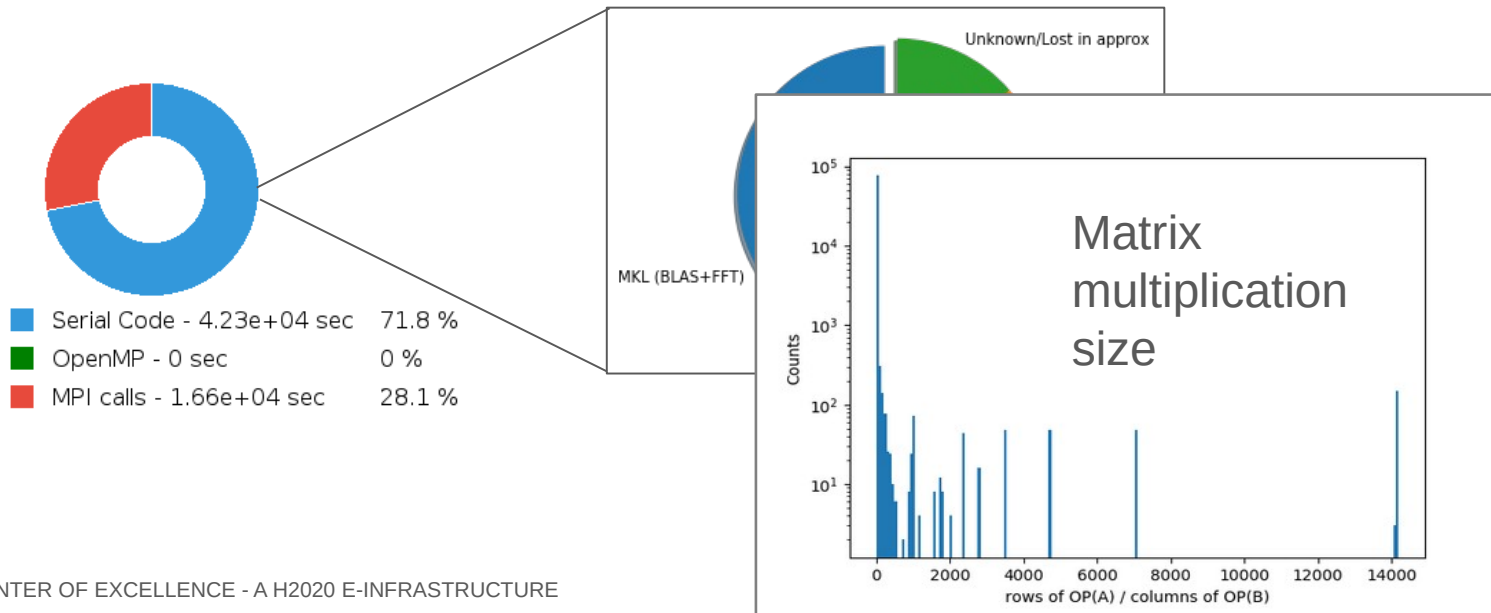
# Profiling

*PWscf* (CPU version) running on a single KNL node with 64 MPI processes  
(best time to solution).



# Profiling

*PWscf* (CPU version) running on a single KNL node with 64 MPI processes  
(best time to solution).





# Profiling

*PWscf* (CPU version) running on a single KNL node with 64 MPI processes  
(best time to solution).

Not a single "portion" of the code takes the majority of the wall-time

Depending on the input, 3D-FFT dominant or LA dominant

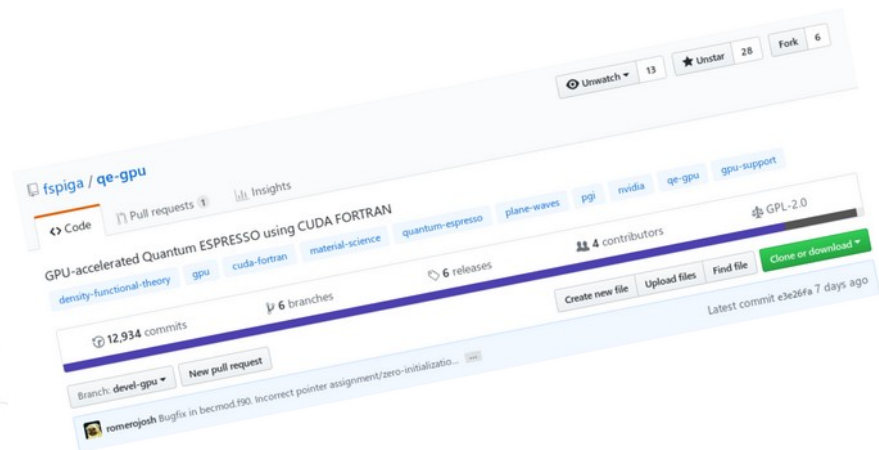
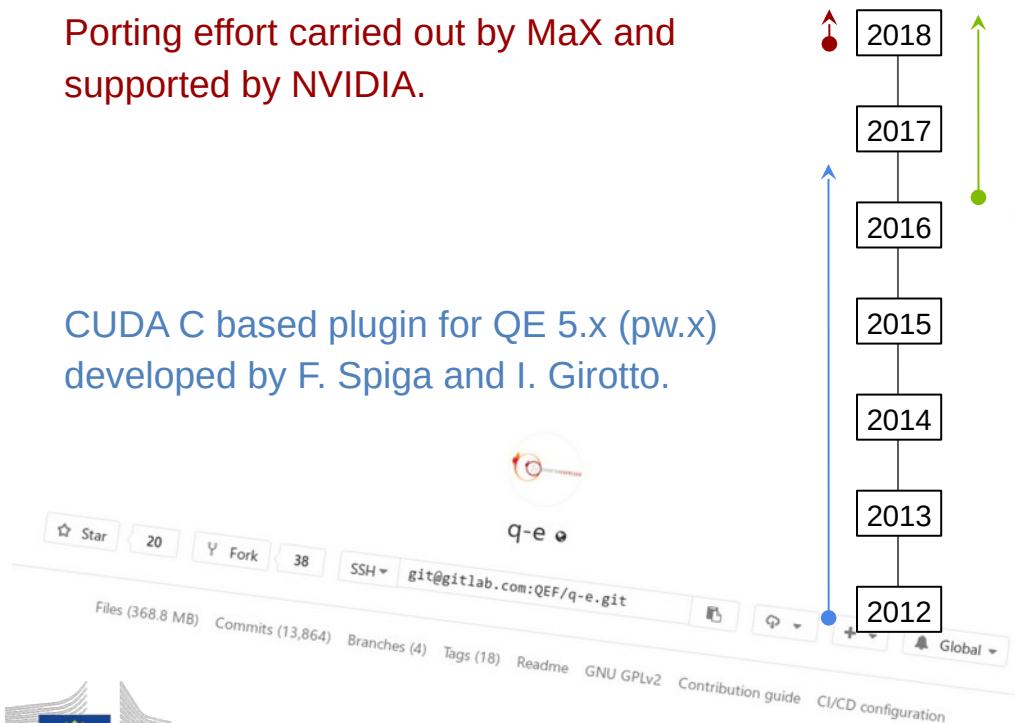


# Past and present QE GPU ports

Porting effort carried out by MaX and supported by NVIDIA.

CUDA C based plugin for QE 5.x (pw.x) developed by F. Spiga and I. Girotto.

Independent CUDA Fortran based port of QE 6.1 (pw.x) developed by F. Spiga and NVIDIA. Provides much better performance, limited features implemented.



# QE-GPU-Plugin

✓ ✓ Self contained

## phiGEMM: a CPU-GPU library for porting Quantum ESPRESSO on hybrid systems

Filippo Spiga, Ivan Girotto

*Irish Centre for High-End Computing (ICHEC), Dublin, Ireland*

20th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP), 2012

DOI: [10.1109/PDP.2012.72](https://doi.org/10.1109/PDP.2012.72)

```

1  !
2  SUBROUTINE addusdens (rho)
3  !
4  !
5  USE realus,          ONLY : addusdens_r
6  USE control_flags,   ONLY : tqr
7  USE noncollin_module, ONLY : nspin_mag
8  USE fft_base,        ONLY : dfftp
9  USE kinds,           ONLY : DP
10 !
11 IMPLICIT NONE
12 !
13 !
14 REAL(kind=dp), intent(inout) :: rho(dfftp%nnr,nspin_mag)
15 !
16 IF ( tqr ) THEN
17   CALL addusdens_r (rho,.true.)
18 ELSE
19   #if defined(__CUDA)
20     CALL addusdens_g_gpu (rho)
21   #else
22     CALL addusdens_g (rho)
23   #endif
24 END IF
25 !
26 RETURN
27 !
28 END SUBROUTINE addusdens

```







```

Modules/mp.f90:#if defined(__CUDA) || defined(__PHIGEMM )
Modules/mp.f90:#if defined(__CUDA) || defined(__PHIGEMM )
PW/src/vloc_psi.f90:#if defined(__CUDA) && !defined(__DISABLE_CUDA_VLOCPSI) && ( !defined(__MPI) || defined(__USE_3D_FFT) )
PW/src/vloc_psi.f90:#if defined(__CUDA) && !defined(__DISABLE_CUDA_VLOCPSI) && ( !defined(__MPI) || defined(__USE_3D_FFT) )
PW/src/rdiaghg.f90:#if defined(__CUDA) && defined(__MAGMA)
PW/src/newd.f90:#if defined(__CUDA) && !defined(__DISABLE_CUDA_NEWD)
PW/src/cdiaghg.f90:#if defined(__CUDA) && defined(__MAGMA)
PW/src/addusdens.f90:#if defined(__CUDA) && !defined(__DISABLE_CUDA_ADDUSDENS)

```

# QE-GPU-Plugin

- ✓✓ Self contained
- ✓ Good performance






# QE-GPU-Plugin

- ✓ ✓ Self contained
- ✓ Good performance

## TESLA K80 DELIVERS 5-10X BOOST IN KEY APPLICATION PERFORMANCE

■ NVIDIA® Tesla® K80 ■ NVIDIA Tesla K20 ■ CPU



CPU Server: Dual Socket E5-2698v3@2.3GHz 3.6GHz Turbo (Haswell EP) HT off, GPU Server Dual Socket E5-2698v3@2.3GHz 3.6GHz Turbo (Haswell EP) HT off, Dual K20/K80 GPU Boost enabled



# Porting strategy

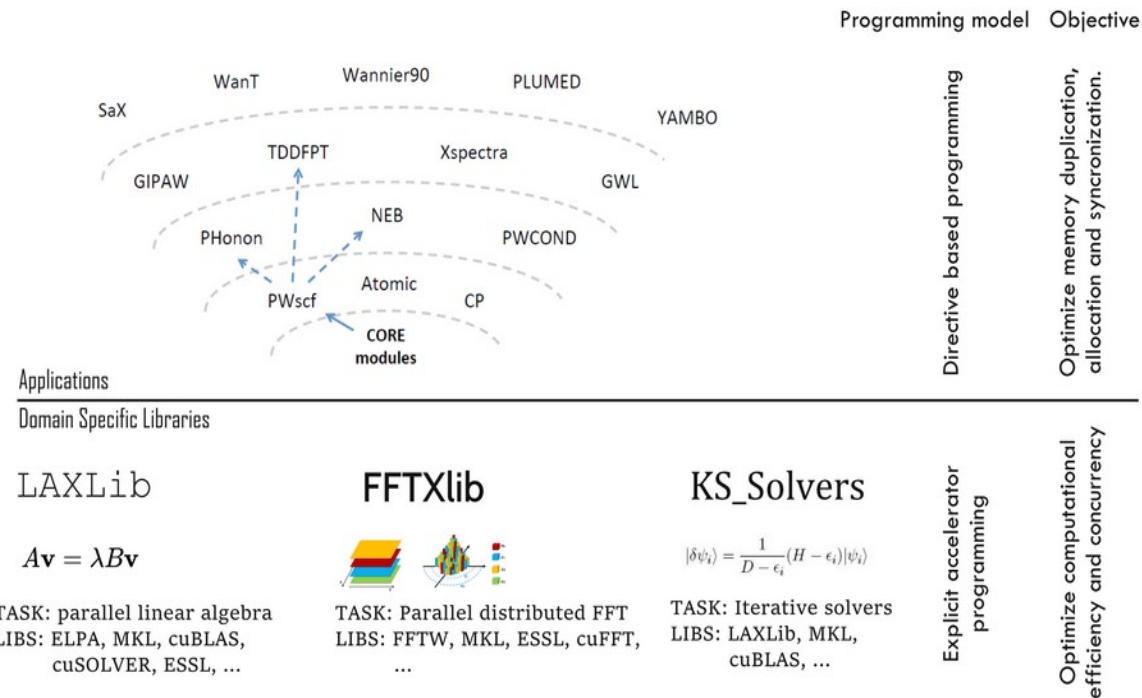
Desiderata:

- Fortran: CUDA Fortran / OpenACC / OpenMP  $\geq 4.5$ .
- Explicit memory management:
  - derived types
  - generic hybrid architecture support.
- Preserve modularity.
- Preserve user experience.



- Prepare full featured and well tested **CUDA-enabled libraries** containing performance critical kernels.
- **Directive based acceleration** of QE's applications.

# Porting strategy



# CUDA Fortran

---

- 😊 Fortran equivalent for CUDA C++
- 😊 Syntax is similar to CUDA, but more concise.
- 😊 Full set of libraries and **interfaces**.
- 😞 Complete syntax only on PGI compilers.
- 😞 Partial implementation on IBM compilers, useless for QE.

# CUDA Fortran

```

1 attributes(global) subroutine increment(a, b)
2   implicit none
3   integer, intent(inout) :: a(:)
4   integer, value :: b
5   integer :: i, n
6
7   i = blockDim%x*(blockIdx%x-1) + threadIdx%x
8   n = size(a)
9   if (i <= n) a(i) = a(i)+b
10
11 end subroutine increment

```

## CUDA kernels in Fortran

### CUDA equivalent syntax

Full support for Fortran  
intrinsic types.

```

1 call vaddkernel <<<(N+31)/32,32 >>> (A,B,C,N)
2
3
4 type(dim3) :: g, b
5 g = dim3((N+31)/32, 1, 1)
6 b = dim3( 32, 1, 1 )
7 call vaddkernel <<< g, b >>> ( A, B, C, N )

```

# CUDA Fortran

Allocation done by the host,  
according to “device” attribute

Just copy (no need for cuda  
APIs for sync. copies).

```

1  real, device, allocatable :: a(:, :)
2  real, allocatable :: b(:)
3  attributes(device) :: b
4
5  real, device, allocatable :: a(:, :), c
6  allocate( a(1:n, 1:m), STAT=ivar )
7  ! CHECK ivar
8  allocate(c)
9  ...
10 deallocate( a, c )
11
12
13
14 module mm
15   real, device, allocatable :: a(:)
16   real, device :: x, y(10)
17   real, constant :: c1, c2(10)
18   integer, device :: n
19   contains
20     attributes(global) subroutine s( b )
21 end module mm

```

# CUDA Fortran

Allocation done by the host,  
according to “device” attribute

Just copy (no need for cuda  
APIs for sync. copies).

```
1 program cuf_memory
2
3 #ifdef USE_CUDA
4   use cudafor
5 #endif
6 implicit none
7
8 ! Define the floating point kind to be single/double_precision
9 integer, parameter :: fp_kind = kind(0.0d0)
10 !integer, parameter :: fp_kind = kind(0.0)
11
12 ! Define
13 real (fp_kind), dimension(:,,:), allocatable :: A, B, C
14 real (fp_kind) :: rand_vals(10,10)
15 #ifdef USE_CUDA
16   attributes(device):: A,B,C
17 #endif
18
19 CALL RANDOM_NUMBER(rand_vals)
20
21 allocate(A(10,10))
22 allocate(B(10,10))
23 allocate(C(10,10))
24
25 A=1._fp_kind
26 B=2._fp_kind
27 C=rand_vals
28
29 deallocate(A,B,C)
30
31 end program cuf_memory
```



# CUDA Fortran

```
#ifdef USE_CUDA
  use cudafor
  use cublas
#endif
```

```
47
48   time_start= wallclock();
49
50   #ifdef USE_CUDA
51   ▶ istat=cudaDeviceSynchronize()
52   #endif
53   ▶ call dgemm('n','n',m1,m1,m1,alpha,A,m1,B,m1,beta,C,m1)
54   #ifdef USE_CUDA
55   istat=cudaDeviceSynchronize()
56   #endif
57
58
59   time_end= wallclock();
```

Fortran interfaces for most (all?!) NVIDIA Cuda Runtime and NVIDIA Libraries.

## Cuf kernels, directive based automatic kernel generation:

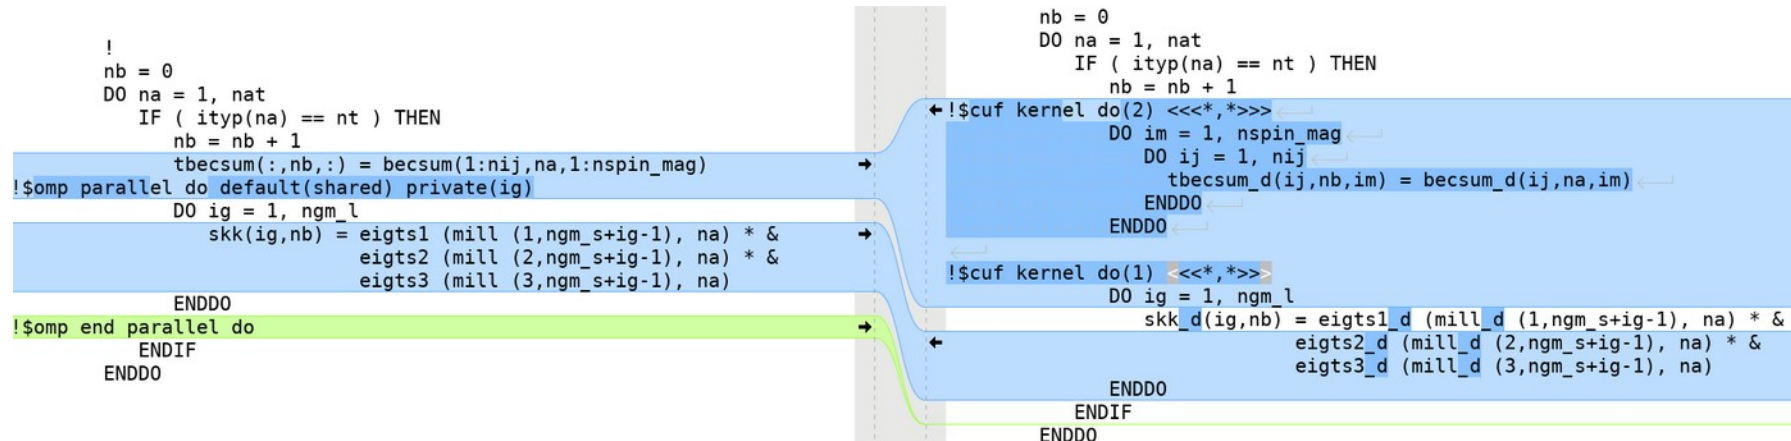
```
1 program incTest
2   use cudafor
3   implicit none
4   integer, parameter :: n = 256
5   integer :: a(n), b
6   integer, device :: a_d(n)
7   a = 1
8   b = 3
9   a_d = a
10  !$cuf kernel do <<<*,*>>>
11  do i = 1, n
12    a_d(i) = a_d(i)+b
13  enddo
14  a = a_d
15  if (all(a == 4)) write(*,*) 'Test Passed'
16 end program incTest
```

## **Cuf kernels**, a few simple rules:

- Scalars are private by default
- Reduction automatically detected (only scalars)
  - Only loop based constructs (possibly nested)

# QE-GPU CUDA Fortran

✓ Single programming language: Fortran + CUDA Fortran



# QE-GPU CUDA Fortran

✓ Single programming language: Fortran

```
!
nb = 0
DO na = 1, nat
  IF ( ityp(na) == nt ) THEN
    nb = nb + 1
  !$omp parallel do default(shared) private(ig)
  DO ig = 1, ngm_l
    skk(ig,nb) = eigts1 (mill (1,ngm_s+ig-1), na) * &
    eigts2 (mill (2,ngm_s+ig-1), na) * &
    eigts3 (mill (3,ngm_s+ig-1), na)
  ENDDO
!$omp end parallel do
ENDIF
ENDDO
```

```
IF (.not.okvan) RETURN
CALL start_clock ('addsdens')
!
ALLOCATE (aux (ngm, nspin_mag) )
aux (1,:) = (0.05, 0.00)
!
! With k-point parallelization, distribute G-vectors across processors
! ngm_s = index of first G-vector for this processor
! ngm_e = index of last G-vector for this processor
! ngm_l = local number of G-vectors
CALL divide (inter_pool_comm, ngm, ngm_s, ngm_e)
!
! for the extraordinary unlikely case of more processors than G-vectors
IF ( ngm_l <= 0 ) GO TO 10
!
ALLOCATE (qnod(ngm_l), qgn(ngm_l) )
ALLOCATE (ylnk(ngm_l, lnaeq * lnaeq))
CALL using_becsum(0)
CALL ylnr2 (lnaeq * lnaeq, ngm_l, q(1,ngm_s), qg(ngm_s), ylnk0)
DO ig = 1, ngm_l
  qnod (ig) = sqrt (qg (ngm_s+ig-1) )
ENDDO
DO nt = 1, ntyp
  IF ( upf(nt)%vamp ) THEN
    !
    ! nij = max number of (ih,jh) pairs per atom type nt
    !
    nij = nh(nt)*(nh(nt)+1)/2
    !
    ! count max number of atoms of type nt
    !
    nab = 0
    DO na = 1, nat
      IF ( ityp(na) == nt ) nab = nab + 1
    ENDDO
    !
    ALLOCATE ( skk_d(ngm_l,nab), tbcsum_d(nij,nab,nspin_mag), aux2_d(ngm_l,nij) )
    !
    nab = 0
    DO na = 1, nat
      IF ( ityp(na) == nt ) THEN
        nab = nab + 1
      !
      tbcsum (1,nab) = becsum(1:nij,na,1:nspin_mag)
    ENDDO
    !
    !$cuf kernel do(2) <<<<,nab
    DO ij = 1, nij
      tbcsum_d(ij,nab,ln) = becsum_d(ij,na,ln)
    ENDDO
    !
    !$cuf kernel do(3) <<<<,nab
    DO ig = 1, ngm_l
      skk_d(ig,nb) = eigts1_d (mill_d (1,ngm_s+ig-1), na) * &
      eigts2_d (mill_d (2,ngm_s+ig-1), na) * &
      eigts3_d (mill_d (3,ngm_s+ig-1), na)
    ENDDO
    !
    DO is = 1, nspin_mag
      ! sum over atoms
      CALL cublasDgemv( 'N', 'I', 2*ngm_l, nij, nab, 1.0_dp, skk_d, 2*ngm_l, &
      tbcsum_d(1,1,1), skk_d, 0.0_dp, aux2_d, 2*ngm_l )
      ! sum over ln indices of Q (ln)
      DO ih = 1, nh (nt)
        DO jh = ih, nh (nt)
          !
          CALL qvan2_gpu (ngm_l, ih, jh, nt, qnod_d, qgn_d, ylnk0_d)
          !$cuf kernel do(3) <<<<,nab
          DO ig = 1, ngm_l
            aux_d(ngm_s+ig-1,1) = aux_d(ngm_s+ig-1,1)+aux2_d(ig,1,jh)*qgn_d(ig)
          ENDDO
        ENDDO
      ENDDO
    ENDDO
    !
    DEALLOCATE (aux2_d, tbcsum_d, skk_d)
    !
    ENDDO
    !
    DEALLOCATE (ylnk0_d)
    DEALLOCATE (qgn_d, qnod_d)
    !
    GO CONTINUE
  !
  aux_h = aux_d
  CALL mp_sum(aux_h, inter_pool_comm)
  !
  ! add aux to the charge density in reciprocal space
  !
  rho(1,:) = rho(1,:) + aux_h(1,:)
  !
  DEALLOCATE (aux_h)
  !
  CALL stop_clock ('addsdens')
  RETURN
END SUBROUTINE addsdens_g
```

# QE-GPU CUDA Fortran

---

- ✓ Single programming language: Fortran + CUDA Fortran
- ✓ Very good performance (shown at the end of this presentation).
- ✗ Code duplication:
  - Performance (eg. preserve cache blocking optimizations)
  - Preserve both CPU and GPU version.

```

!$acc kernels loop present(qmod(ngy), ylmk0(ngy, lmaxq * lmaxq), qg, grad, lpx, lpl, ap) &
!$acc num_workers(256) collapse(1) if(on_device)
do ig = 1, ngy
!
!
qg(ig) = (0.d0, 0.d0)
qm = qmod(ig) * dq
px = qm - int(qm)
ux = 1.d0 - px
vx = 2.d0 - px
wx = 3.d0 - px
i0 = INT(qm) + 1
i1 = i0 + 1
i2 = i0 + 2
i3 = i0 + 3
uvx = ux * vx * sixth
pwx = px * wx * 0.5d0
do lm = 1, lpx(ivl, jvl)
lp = lpl(ivl, jvl, lm)
!
! find angular momentum l corresponding to combined index lp
! (l is actually l+1 because this is the way grad is stored, check init_us_1)
!
if (lp == 1) then
l = 1
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 4) then
l = 2
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
elseif (lp <= 9) then
l = 3
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
elseif (lp <= 16) then
l = 4
sig = CMPLX(0.0d0, 1.0d0, kind=DP)
elseif (lp <= 25) then
l = 5
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 36) then
l = 6
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
else
l = 7
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
endif
work = grad(i0, ijv, l, np) * uvx * wx + &
grad(i1, ijv, l, np) * pwx * vx - &
grad(i2, ijv, l, np) * pwx * ux + &
grad(i3, ijv, l, np) * px * uvx
qg(ig) = qg(ig) + sig * CMPLX(ap(lpl(ivl, jvl), ylmk0(ig, lp)) * work, 0.d0, kind=DP)
enddo

```

## OpenACC

```

ig= threadIdx%x+BlockDim*x*(BlockIdx%x-1)
if (ig <= ngy) then
! compute the indices which correspond to ih,jh
dq = 1.0_DP / dq
qg(ig) = 0.d0

```

## CUDAFortran kernel

```

qm = qmod(ig) * dq
px = qm - int(qm)
ux = 1.d0 - px
vx = 2.d0 - px
wx = 3.d0 - px
i0 = INT(qm) + 1
i1 = i0 + 1
i2 = i0 + 2
i3 = i0 + 3
uvx = ux * vx * sixth
pwx = px * wx * 0.5d0

do lm = 1, lpx(ivl, jvl)
lp = lpl(ivl, jvl, lm)
if (lp == 1) then
l = 1
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 4) then
l = 2
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
elseif (lp <= 9) then
l = 3
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
elseif (lp <= 16) then
l = 4
sig = CMPLX(0.0d0, 1.0d0, kind=DP)
elseif (lp <= 25) then
l = 5
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 36) then
l = 6
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
else
l = 7
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
endif
!sig = sig * ap(lpl(ivl, jvl)
work = grad(i0, ijv, l, np) * uvx * wx + &
grad(i1, ijv, l, np) * pwx * vx - &
grad(i2, ijv, l, np) * pwx * ux + &
grad(i3, ijv, l, np) * px * uvx
qg(ig) = qg(ig) + sig * CMPLX(ylmk0(ig, lp) * work * ap(lpl(ivl, jvl), 0.d0, kind=DP)
enddo

```

```

!$acc kernels loop present(qmod(ngy), ylmk0(ngy, lmaxq * lmaxq), qg, grad, lpx, lpl, ap) &
!$acc num_workers(256) collapse(1) if(on_device)
do ig = 1, ngy
!
!
qg(ig) = (0.d0, 0.d0)
qm = qmod(ig) * dq
px = qm - int(qm)
ux = 1.d0 - px
vx = 2.d0 - px
wx = 3.d0 - px
i0 = INT(qm) + 1
i1 = i0 + 1
i2 = i0 + 2
i3 = i0 + 3
uvx = ux * vx * sixth
pwx = px * wx * 0.5d0
do lm = 1, lpx(ivl, jvl)
lp = lpl(ivl, jvl, lm)
!
! find angular momentum l corresponding to combined index lp
! (l is actually l+1 because this is the way grad is stored, check init_us_1)
!
if (lp == 1) then
l = 1
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 4) then
l = 2
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
elseif (lp <= 9) then
l = 3
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
elseif (lp <= 16) then
l = 4
sig = CMPLX(0.0d0, 1.0d0, kind=DP)
elseif (lp <= 25) then
l = 5
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 36) then
l = 6
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
else
l = 7
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
endif
work = grad(i0, i1, l, np) * uvx * wx + &
grad(i1, i2, l, np) * pwx * vx - &
grad(i2, i3, l, np) * pwx * ux + &
grad(i3, i0, l, np) * px * uvx
qg(ig) = qg(ig) + sig * CMPLX(ap(lpl(ivl, jvl), ylmk0(ig, lp)) * work, 0.d0, kind=DP)
enddo

```

OpenACC

```

ig= threadIdx%x+BlockDim%x*(BlockDim%x-1)
if (ig <= ngy) then
! compute the indices which correspond to ih,jh
dq = 1.0_DP / dq
qg(ig) = 0.d0

qm = qmod(ig) * dq
px = qm - int(qm)
ux = 1.d0 - px
vx = 2.d0 - px
wx = 3.d0 - px
i0 = INT(qm) + 1
i1 = i0 + 1
i2 = i0 + 2
i3 = i0 + 3
uvx = ux * vx * sixth
pwx = px * wx * 0.5d0

do lm = 1, lpx(ivl, jvl)
lp = lpl(ivl, jvl, lm)
if (lp == 1) then
l = 1
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 4) then
l = 2
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
elseif (lp <= 9) then
l = 3
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
elseif (lp <= 16) then
l = 4
sig = CMPLX(0.0d0, 1.0d0, kind=DP)
elseif (lp <= 25) then
l = 5
sig = CMPLX(1.0d0, 0.0d0, kind=DP)
elseif (lp <= 36) then
l = 6
sig = CMPLX(0.0d0, -1.0d0, kind=DP)
else
l = 7
sig = CMPLX(-1.0d0, 0.0d0, kind=DP)
endif
!sig = sig * ap(lpl(ivl, jvl)
work = grad(i0, i1, l, np) * uvx * wx + &
grad(i1, i2, l, np) * pwx * vx - &
grad(i2, i3, l, np) * pwx * ux + &
grad(i3, i0, l, np) * px * uvx
qg(ig) = qg(ig) + sig * CMPLX(ylmk0(ig, lp)) * work * ap(lpl(ivl, jvl), 0.d0)
enddo

```

CUDA Fortran kernel




# Libraries

## LAXLib



L A P A C K  
L -A P -A C -K  
L A P A -C -K  
L -A P -A -C K  
L A -P -A C K  
L -A -P A C -K

## FFTXlib



FFTW, MKL, ESSL, ...

## devXlib



# QE Libraries

- Full API support:

```
IF( use_tg ) THEN
!
CALL invfft ( 'tgWave', tg_psic_d, dffts )
!
CALL tg_get_group_nr3( dffts, right_nr3 )
!
```

- Unit testing:

```
[@node153 tests]$ mpirun -np 1 ./test_fft_scalar_gpu.x
fortran_tester:      0 error(s) for      32 test(s)
fortran_tester: all tests succeeded
[@node153 tests]$ mpirun -np 4 ./test_fft_scatter_mod_gpu.x
fortran_tester:     51 error(s) for     224 test(s)
fortran_tester: tests failed
```

# QE Libraries

## LAXLib

Solution of **dense eigenvalue problem** of real or complex hermitian matrices.



L	A	P	A	C	K
L	-A	P	-A	C	-K
L	A	P	A	-C	-K
L	-A	P	-A	-C	K
L	A	-P	-A	C	K
L	-A	-P	A	C	-K

Both **serial** and **distributed parallel** implementation.

Extract **all eigenpairs** or a **subset**.

# QE Libraries

## LAXLib

ELPA

L	A	P	A	C	K
L	-A	P	-A	C	-K
L	A	P	A	-C	-K
L	-A	P	-A	-C	K
L	A	-P	-A	C	K
L	-A	-P	A	C	-K

NVIDIA / **Eigensolver\_gpu**

<> Code

! Issues 0

🔗 Pull requests 0

📁 Projects 0

Branch: master ▾

**Eigensolver\_gpu** / lib\_eigsolve /

**CU Solver**

Initially custom code now part of CUDA 10.1.

Available at

<https://gitlab.com/max-centre/components>

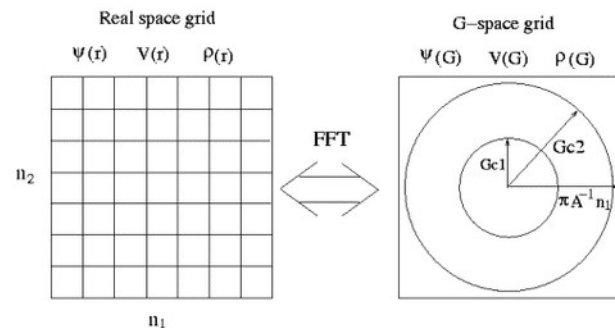
# QE Libraries

## FFTXlib



FFTW, MKL, ESSL, ...

Sparse FFT in reciprocal space.  
Parallel, distributed, accelerated.  
Both pencil and slab decomposition.



Available at <https://gitlab.com/max-centre/components>

## FFTXlib



FFTW, MKL, ESSL, ...

The local potential contribution is computed more efficiently in real space:

$$\psi_{ik}(\mathbf{G}) \xrightarrow{FFT} \psi_{ik}(\mathbf{r})$$


$$[v_{KS}\psi_{ik}](\mathbf{r}) = v_{KS}(\mathbf{r})\psi_{ik}(\mathbf{r})$$

$$[v_{KS}\psi_{ik}](\mathbf{r}) \xrightarrow{FFT} [v_{KS}\psi_{ik}](\mathbf{G})$$

Available at <https://gitlab.com/max-centre/components>

# QE Libraries


- For each band, FFT to real space, multiplication, FFT to reciprocal space.  
→ Many independent small 3D FFTs ( $10^1 \rightarrow 10^3$ )



# QE Libraries

For each band, FFT to real space, multiplication, FFT to reciprocal space.


→ Many independent small 3D FFTs ( $10^1 \rightarrow 10^3$ )






# QE Libraries

- Many small 3D FFTs ( $10^1 \rightarrow 10^3$ )
- Overlap of communication and computation
- Batched work



8 bands



# QE Libraries

- QE allocates many small auxiliary workspaces. This impacts substantially the performances of the accelerated version of the code.
- Optimize memory allocation: GPU memory is limited.

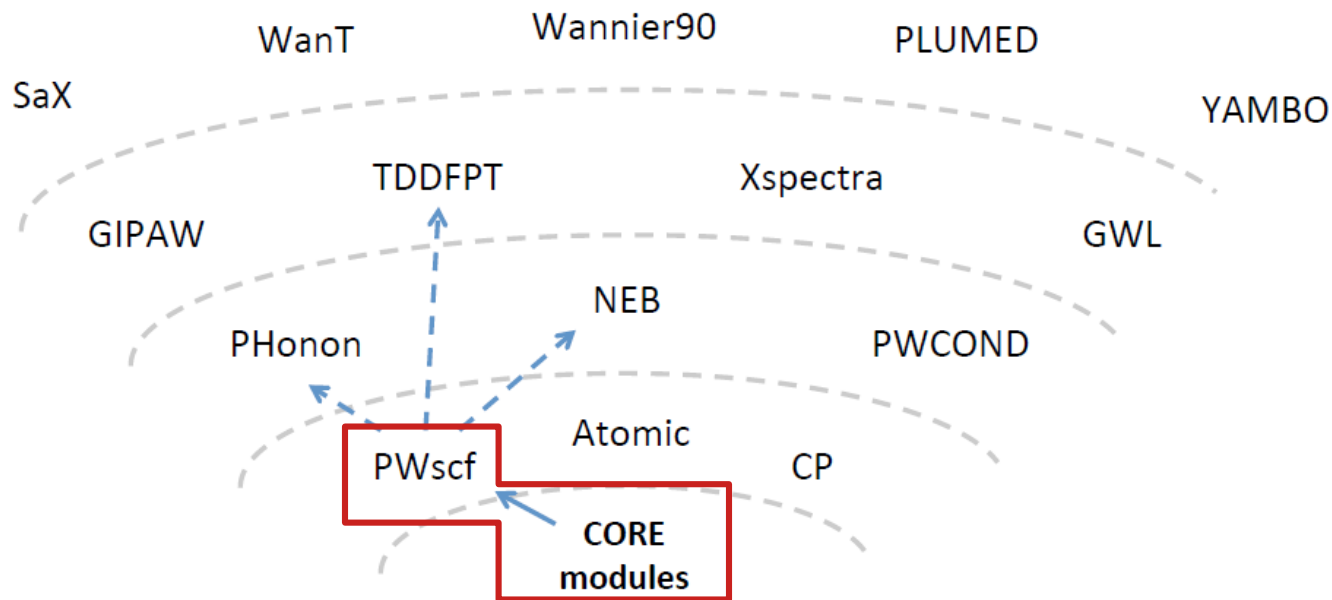
## devXlib

Available at  
<https://gitlab.com/max-centre/components>

```
USE buffer_module, ONLY : gpu_buffer
!  
implicit none  
!  
REAL, POINTER :: work(:)  
gpu_buffer%lock_buffer(work, 10, ierr)  
[...]  
gpu_buffer%release_buffer(work, ierr)
```

# QE Codes

- GPU acceleration currently available for **PWscf**. CP and PHonon planned.




# Current status and evolution

Application: *pw.x*

GPU version	Total Energy (K points)	Forces	Stress	Collinear Magnetism	Non-collinear magnetism	Gamma trick	US PP & PAW	EXX	DFT+U	All other functionalities
v5.4	<b>A</b>	<b>W</b>	<b>W</b>	<b>B</b> (?)	<b>U</b>	<b>A</b>	<b>A</b>	?	<b>W</b> (?)	<b>W</b> (?)
v6.1	<b>A</b>	<b>A</b>	<b>A</b>	<b>A</b>	<b>U</b>	<b>W</b> (*)	<b>A</b>	<b>U</b>	<b>U</b>	<b>U</b> (*)
v6.3	<b>A</b>	<b>W</b>	<b>W</b>	<b>A</b>	<b>A</b>	<b>A</b>	<b>A</b>	<b>W</b>	<b>W</b>	<b>W</b>
V6.5	<b>A</b>	<b>A</b>	<b>W</b>	<b>A</b>	<b>A</b>	<b>A</b>	<b>A</b>	<b>A</b>	<b>A</b>	<b>W</b>

**A**ccelerated, **W**orking, **U**navailable, **B**roken


# Benchmarks



## Marconi @ CINECA

Model: Xeon E5-2697 v4 (BDW) @ 2.30 GHz  
Cores: 2x18 = **36**  
RAM: 128 GB/node

Q3 2016  
1.3 TFLOPs



## Galileo @ CINECA

Model: Xeon E5-2630 v3 (HSW) @ 2.40 GHz  
Cores: 2x8 = **16**  
Accelerators: **2 x K80**  
RAM: 128 GB/node

Q1 2015  
0.6 + 2x2.9 TFLOPs




## Piz Daint XC50 @ CSCS:

Model: Xeon E5-2690 v3 (HSW) @ 2.60 GHz  
Cores: 1x12 = **12**  
Accelerators: **1 x P100**  
RAM: 64 GB/node

Q4 2016  
0.5 + 4.7 TFLOPs

# Benchmarks

Best time to solution obtained with *pw.x* v6.2, with and without GPU support, and with the GPU port of v6.1 done by NVIDIA.



MnSi, bulk, ferromagnetic.  
64 atoms, 365 KS states, USPP.

## **Piz Daint XC50 @ CSCS:**




Processors: 12-cores Intel Haswell 2.60 GHz  
Accelerators: 1 NVIDIA P100  
RAM: 64 GB/node

## **Galileo @ CINECA**

Processors: 2\*8-cores Intel Haswell 2.40 GHz  
Accelerators: 2 NVIDIA K80  
RAM: 128 GB/node

# Benchmarks



Best time to solution obtained with *pw.x* v6.3, with and without GPU acceleration, and with the GPU port of *pw.x* v6.1 done by NVIDIA.



Large test case, 2D,  
26 k-points, 96 atoms,  
326 KS states.

# Benchmarks



Best time to solution obtained with *pw.x* v6.3, with and without GPU acceleration, and with the GPU port of *pw.x* v6.1 done by NVIDIA.






# Benchmarks

Best time to solution obtained with *pw.x* v6.3, with and without GPU acceleration, and with the GPU port of *pw.x* v6.1 done by NVIDIA.



# Benchmarks

Best time to solution obtained with *pw.x* v6.4, with and without GPU acceleration.



Very large benchmark:

AuCONH benchmark: 586 atoms, 1531 KS

EUROPEAN CENTER OF EXCELLENCE - A H2020 E-INFRASTRUCTURE



Processor	Components	
	CPU	GPU
Type	POWER9	V100
Count	9,216	27,648
	2 × 18 × 256	6 × 18 × 256
Peak FLOPS	9.96 PF	215.7 PF
Peak AI FLOPS		3.456 EF

Courtesy of Dr. Ye Luo

# HT Benchmarks




## New Galileo @ CINECA

Model: Xeon E5-2697 v4 (BDW)

@ 2.30 GHz

Cores: 1x18; RAM: 128 GB/node



NVIDIA  
V100



Random structures from COD:



+ QE Input generator  
materialscloud.org

- **Magnetic / Non Magnetic**
- **Metal / Insulator**
- **Efficiency / Accuracy**
- **K-points grid:**  
**Very fine / Fine / Normal / Gamma**

# HT Benchmarks



# Porting other QE applications

- Preserve *all* functionalities
  - Feature testing already available...enough?
  - More feature and unit tests
  - Create verification scheme
- Preserve accelerated function modularity
  - For debugging
  - For code maintainability
  - For simpler development
- Directive based



```
IF (use_gpu) THEN
    call g2_kin_gpu( ik )
ELSE
    call g2_kin( ik )
END IF
```

```
IF( use_tg ) THEN
!
CALL invfft ( 'tgWave', tg_psic_d, dffts )
!
CALL tg_get_group_nr3( dffts, right_nr3 )
!
!$cuf kernel do(1) <<<,>>>
DO j = 1, dffts%nr1x * dffts%nr2x * right_nr3
    tg_psic_d (j) = tg_psic_d (j) * tg_v_d(j)
ENDDO
!
CALL fwfft ( 'tgWave', tg_psic_d, dffts )
!
```

# Lessons learnt

---

- Separation of concerns simplifies subsequent application porting
  - Reduces branches in the code.
  - Cleaner and simpler adoption of accelerated functions in new code.
- Modularization helps (duplicated) data management
  - Porting pushes modularization forward
- CUDA Fortran helped removing boilerplate.
- (Moving to GIT helped a lot)



# Remaining challenges

---

- HW availability limits GPU programming:
  - Policy for contributions?
  - Ratio between CPU and GPU efforts?
- Different hybrid solutions about to appear.
- Is OpenMP the final answer?
- Source code duplication.

# Remaining challenges

---

- HW availability limits GPU programming:
  - Policy for contributions?
  - Ratio between CPU and GPU efforts?
- Different hybrid solutions about to appear.
- Is OpenMP the final answer?
- Source code duplication.

Suggestions?  
Thanks for your attention!