

New developments of the QUANTUM ESPRESSO code: a combined CUF-OpenACC approach

<u>Ivan Carnimeo</u>

Scuola Internazionale Superiore di Studi Avanzati (SISSA) – Trieste, Italy

QUANTUM ESPRESSO[™] (QE) 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.



Pwscf, the core of QE, computes the electronic structure of molecules and materials by solving the Kohn-Sham (KS) equations, within the **Density Functional Theory** framework:

$$\hat{H}^{KS}\psi_{ik}(\mathbf{r}) = \varepsilon_{ik}\psi_{ik}(\mathbf{r})$$

using Plane-Wave and Pseudopotential methods:

$$\psi_{ik}(\mathbf{r}) = \sum_{\mathbf{G}}^{N_{PW}} C_{\mathbf{G},ik} \frac{e^{i\mathbf{Gr}}}{\sqrt{\Omega}}$$







QUANTUM ESPRESSO workflow



UANTUMESPRESSO







All the most important features of QUANTUM ESPRESSO[™] have been ported to GPU-accelerated architectures in the last years using CUDA Fortran

	QE Version						
	6.4	6.5a1	6.5a2	6.7	6.8	develop	
Davidson							
CG							
ParO							
PPCG							
Energy							
Forces							
Stress							
Magnetism							
USPP						OpenACC	
PAW							
DFT+U							C
XC					OpenACC	OpenACC	
<u>VdW(</u> D3)					OpenACC	OpenACC	
CP						OpenACC	

The most recent developments have been done using OpenACC QUANTUM ESPRESSO on GPU







The CUDA Fortran version of the code has a significantly better performance than the CPU version, but also comes with some drawbacks that we want to tackle using **OpenACC**



SUBROUTINE using_vkb(intento)

```
!USE uspp, ONLY : vkb, vkb d
    implicit none
    INTEGER, INTENT(IN) :: intento
    IF (size(vkb)==0) return
    IF (vkb_ood) THEN
        IF (intento < 2) vkb = vkb_d</pre>
        vkb ood = .false.
    ENDIF
    IF (intento > \emptyset) vkb d ood = .true.
END SUBROUTINE using vkb
SUBROUTINE using vkb d(intento)
    !USE uspp, ONLY : vkb, vkb_d
    implicit none
    INTEGER, INTENT(IN) :: intento
    IF (size(vkb)==0) return
    IF (vkb d ood) THEN
        IF (intento < 2) vkb_d = vkb</pre>
        vkb d ood = .false.
    ENDIF
    IF (intento > 0) vkb_ood = .true.
    CALL upf error('using vkb d', 'no GPU support', 1)
END SUBROUTINE using vkb d
```

If we synchronize host and device variables every time one of the copies changes, the computational burden of the calculations explodes.

N II A N T II M E S P R E S S N

We use flags that keep track whether the last change on a given variable has occurred on the host or device copy

Very efficient but redundant code!!!

P. Bonfa'





In many cases we have to duplicate memory allocations creating both host and device copies of the variables.

This lead to a frequent duplication of subroutines and code in general, that has a strong impact on the maintenance of the code, and makes it more difficult to add new features



OpenACC allows to reduce code duplication

IF (nkb > 0) CALL <mark>init_us</mark>_2(ngk(ik), igk_k(1,ik), xk(1,ik), vkb, .true.)

(e.g. by declaring vkb as present inside init_us_2) but we have to be very careful with synchronizations as the data transfer burden quickly explodes







In this part of the code we have many complex nested loops

do iat=na_smax,na_e

do jat=2, n
 do kat=1, n
 if((jat.ge.iat).or.(kat.ge.jat)) cycle
 linij=lin(iat,jat)
 ijvec1=xyz(1,jat)-xyz(1,iat)
 ijvec2=xyz(2,jat)-xyz(2,iat)
 ijvec3=xyz(3,jat)-xyz(3,iat)

do ktaux=repmin1,repmax1

do ktauy=repmin2,repmax2 do ktauy=repmin3,repmax3 ktau1=ktaux*lat(1,1)+ktauy*lat(1,2)+ktauz*lat(1,3) ktau2=ktaux*lat(2,1)+ktauy*lat(2,2)+ktauz*lat(2,3) ktau3=ktaux*lat(3,1)+ktauy*lat(3,2)+ktauz*lat(3,3) rik2=(ikvec1+ktau1)*(ikvec1+ktau1)+(ikvec2+ktau2)*(ikvec2+ktau2)+(ikvec3+ktau3)*(ikvec3+ktau3) if (rik2.gt.abcthr)cycle

dumvec1=jkvec1+ktau1-jtau1
dumvec2=jkvec2+ktau2-jtau2
dumvec3=ikvec3+ktau3-jtau3





For complex and nested loops **OpenACC** provides more flexible optimization directives than CUF



VdW (D3)







! MODULE exch_lda
!! LDA exchange functionals. !
CONTAINS !
<pre>! SUBROUTINE slater(rs, ex, vx) !\$acc routine (slater) seq</pre>
!! Slater exchange with alpha=2/3
USE kind_l, ONLY: DP
IMPLICIT NONE

- **XClib**, library of functionals with large number of *kernels* and *acc routines*
- Vxc, potential calculation with calls to XClib and FTTs
- XClib and Vxc in OpenACC (no code duplication) with CUDA FFTs (already present)

F. Ferrari Ruffino

XC functionals





			v_of_rho [473,573 m	s]					
v_xc [444,624 ms]									
v_xc_lda [63,488 ms]	v_xc_gga [381,107 ms]								
	ff)					ff		ff
			v_of_rho [62,619 ms]						
v_xc [34,631 ms]									
v_xc_lda [14,679 ms]			v_xc_gga [19,939 ms]						
					cuMemFree		с с	cu	c
-	-					1			

- 1) 40% less code w.r.t. CUF
- 2) Same performance as CUF
- 3) No need to rewrite large portions of code when passing from CPU to OpenACC

F. Ferrari Ruffino



CaSiN - PBE

Car-Parrinello molecular dynamics using plane waves and pseudopotentials: most of the simulation time is spent in the main_loop repeated for each step of the dynamics



The code has been partially accelerated using CUDA Fortran but significant parts still run on CPU



P. Delugas

Car-Parrinello MD

- The exch_corr part benefits from the work done in XClib with a 2X speedup without any code change
- Replacing the multithreading with OpenACC directives we observe a reduction from ~100 to few ms
- Currently the overall acceleration of each step is ~ 3X
- Further speedup can be achieved
 - Using accelerated FFTs
 - Improving data locality and syncronization

P. Delugas, F. Ferrari Ruffino, S. Orlandini



cr (197.915.s) main_loop [5.389 s] cr_mid [5.389 s] struct [486.119 ms] moor [292.298 mc] vortmo [1.137 s] moor [292.298 mc] vortmo [1.137 s] moor [292.298 mc] pread64 pread64







Car-Parrinello MD

THE EXASCALE TRANSITION



For those parts of the code where **OpenACC** and **OpenMP** overlap, compilation directives (taken from the **Yambo Code**) are employed to choose one of the two, on the basis of the available architectures

1/1 - + [* □	@XPS-13 32:130	1/1 • + [²	@XPS-13 32:130
1: pdelugas@login01:/m100_scratch/userexternal/pdelugas/q-e/CPV/src 👻		1: pdelugas@login01:/m100_scratch/userexternal/pde	elugas/q-e/CPV/src 👻
! in the root directory of the present distribution, ! or http://www.gnu.org/copyleft/gpl.txt . !		! DEV_OMP do DEV_ACC parallel loop	
#if defined (_OPENACC) #ifndefOPENACC #defineOPENACC #endif #endif		D0 ig = 1, SIZE(vtemp) vtemp(ig)=(0.d0,0.d0) END D0 D0 is=1,nsp DEV_OMP do DEV_ACC parallel loop	
<pre>#if defined (OPENACC) #define DEV_ACC !\$acc #define DEV_OMP !!! #define START_WSHARE DEV_ACC kernels #define FND WSHARE DEV_ACC end kernels</pre>		DO ig=1,dffts%ngm vtemp(ig)=vtemp(ig)+CONJ END DO END DO	lG(rhotmp(ig))*sfac(ig,is)*vps(ig,is
<pre>#else #define DEV_ACC !!! #define DEV_OMP !\$omp #define START_WSHARE DEV_OMP workshare #define END_WSHARE DEV_OMP end workshare #endif</pre>		<pre>DEV_OMP do geduction(+:zpseu) DEV_ACC parallel loop reduction(+:zp D0 ig=1,dffts%ngm zpseu = zpseu + vtemp(ig) END D0 DEV_OMP end parallel</pre>	oseu)
! SUBROUTINE vofrho_x(nfi, rhor, drhor, rhog, drhog, rhos, tlast, eil, ei2, ei3, irb, eigrb, sf	rhoc, tfirst, & ac, tau0, fion)	<pre>! DEV_ACC update self(vtemp(1)) epseu = wz * DBLE(zpseu) !</pre>	
<pre>computes: the one-particle potential v in real spac computes: the one-particle potential v in real spac the total energy etot, the forces fion acting on the ions, the derivative of total energy to cell particular the derivative the derivative the derivative total energy to cell particular the derivative total energy total e</pre>	arameters h	<pre>IF (gstart == 2) epseu = epseu ! CALL mp_sum(epseu, intra_bgrp epseu = epseu * omega !</pre>	ı - DBLE(ˈvtemp(1)) o_comm)

P. Delugas, P. Bonfa', A. Ferretti



- CUF-OpenACC interoperability is one of the key aspects of QE acceleration
- OpenACC logic for handling device variables is helpful for reducing code duplication
- Portability of the code to different hardware architectures is very important
- Other important codes of the QUANTUM ESPRESSO suite will be accelerated (e.g. TDDFPT, PHonon)





QUANTUM ESPRESSO GPU development group

- Paolo Giannozzi, University of Udine
- Stefano Baroni, SISSA
- Pietro Delugas, SISSA
- Pietro Bonfa', University of Parma
- Ivan Carnimeo, SISSA
- Fabrizio Ferrari Ruffino, CNR-IOM
- Oscar Baseggio, SISSA
- Elena De Paoli, CNR-IOM

Nvidia Technical support

- Filippo Spiga, Nvidia
- Louis Stuber, Nvidia

CINECA Technical support

Sergio Orlandini, CINECA