U.S. DEPARTMENT OF ENERGY | Office of Science

CCS 2021, Tsukuba, Oct 8, 2021

Development and application of real-time time-dependent density functional theory (RT-TDDFT) code, INQ, optimized for hybrid CPU-GPU HPC systems

Tadashi Ogitsu

Lawrence Livermore National Laboratory

Work was performed under the auspices of the U.S. DOE by Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

Lawrence Livermore National Laboratory

LLNL is one of seventeen DOE national laboratories and is located about one hour

DOE Leadership Class HPC systems at LLNL

June 2021 TOP500 list

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	**Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442,010.0	537,212.0	29,899
2	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148,600.0	200,794.9	10,096
3	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438

2021

Lassen/Sierra (2018)

El Capitan (2023)

Nonperturbative Studies of Functional Materials Under Nonequilibrium condition (NPNEQ)

Aaron Lindenberg

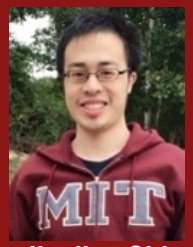

Jiaojian Shi

Ultrafast Experiments

Tadashi Ogitsu

Alfredo Correa

Xavier Andrade

Software Development

INQ code developing team

Das Pemmaraju

vacant (PD)

TDDFT Simulations

TD-TB to extend time & length scale

Liang Tan

David Prendergast

Sangeeta Rajpurohit (PD)

Interpretation Experiments, Spectroscopy Simulations

4

Project Goal and Approach

Develop open source RT-TDDFT software, *INQ*, optimized for the DOE Leadership Class HPC systems

- Programmabiliity (users = scientists)
- Transferability
- Performance

Validate the software using the state-of-art ultrafast experimental capability at SLAC

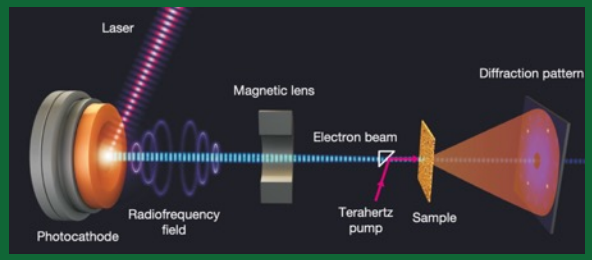

Comprehensive software dissemination activities

- Conferences
- Workshops
- Facility Users Meetings
- Summer School
- Online tutorials

Facilitate Advancement of Ultrafast Sciences

Why (RT-TD)DFT on GPUs

Current situation

- Most of DOE HPC systems are GPU-based
- Around 15% of DOE HPC resources are used by DFT simulations
- Current DFT codes: monolithic design poses challenge in GPU adaptation
- Limited success with high arithmetic intensity: GW, hybrid
  - No successful case for plain DFT (GGA/LDA)

Lassen/Sierra (2018)

El Capitan (2023)

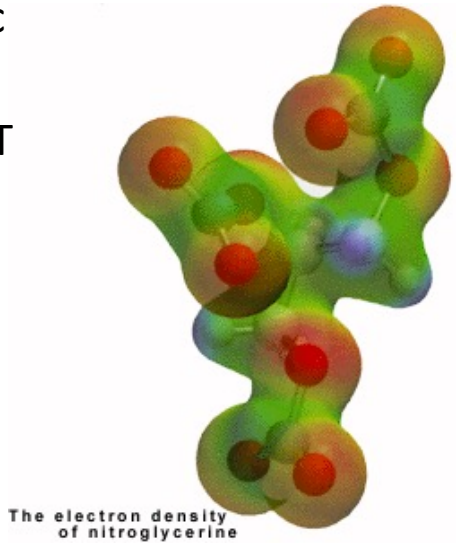

The electron density of nitroglycerine

Walter Kohn

CPU + GPU HPC architecture (schematics)

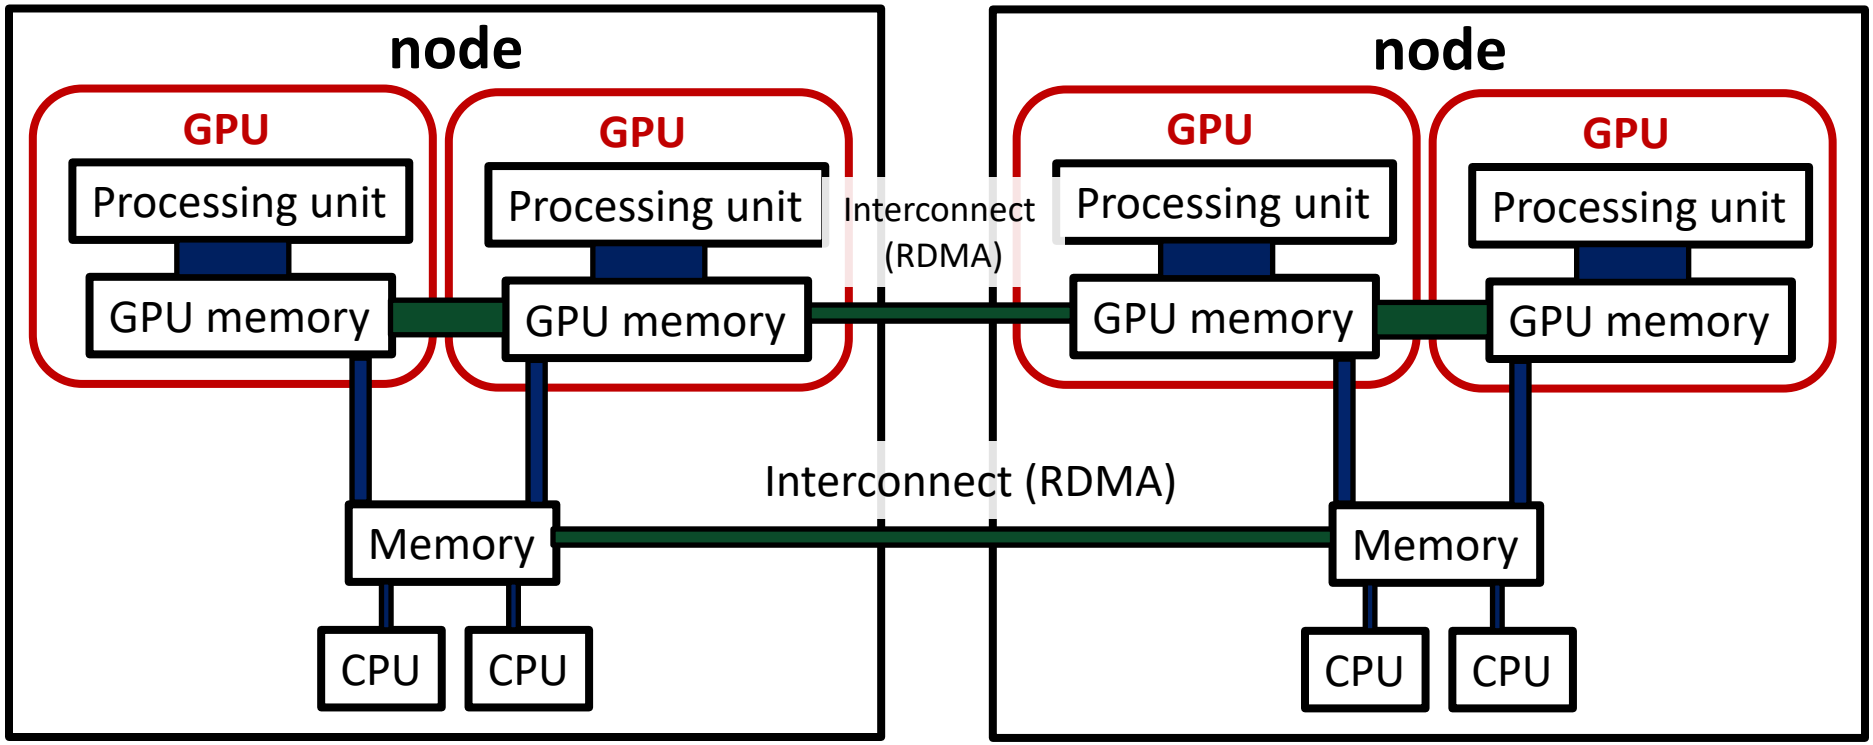

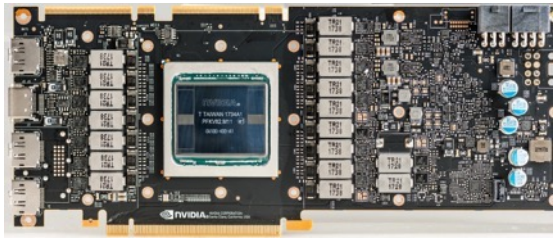

- Computing power: GPU >> CPU
- CPU-GPU copy: order of μ sec overhead
- GPU-GPU copy: intranode significantly faster than internode

GPU programing paradigm: offloading (or not)

- **If predominant portion of computing time is spent by one function, offloading the execution of function to GPU will significantly reduce computing time**

- **Planewave pseudopotential DFT algorithm has three distinct subroutines that necessitates frequent offloading**
  - GPU acceleration is offset by CPU-GPU copy overhead
  - Simple offloading is not suitable for DFT

Our strategy
- **All data resides in GPU memory**
- **Computations are performed only on GPU (to avoid GPU-CPU copy overhead)**

The physics in inq

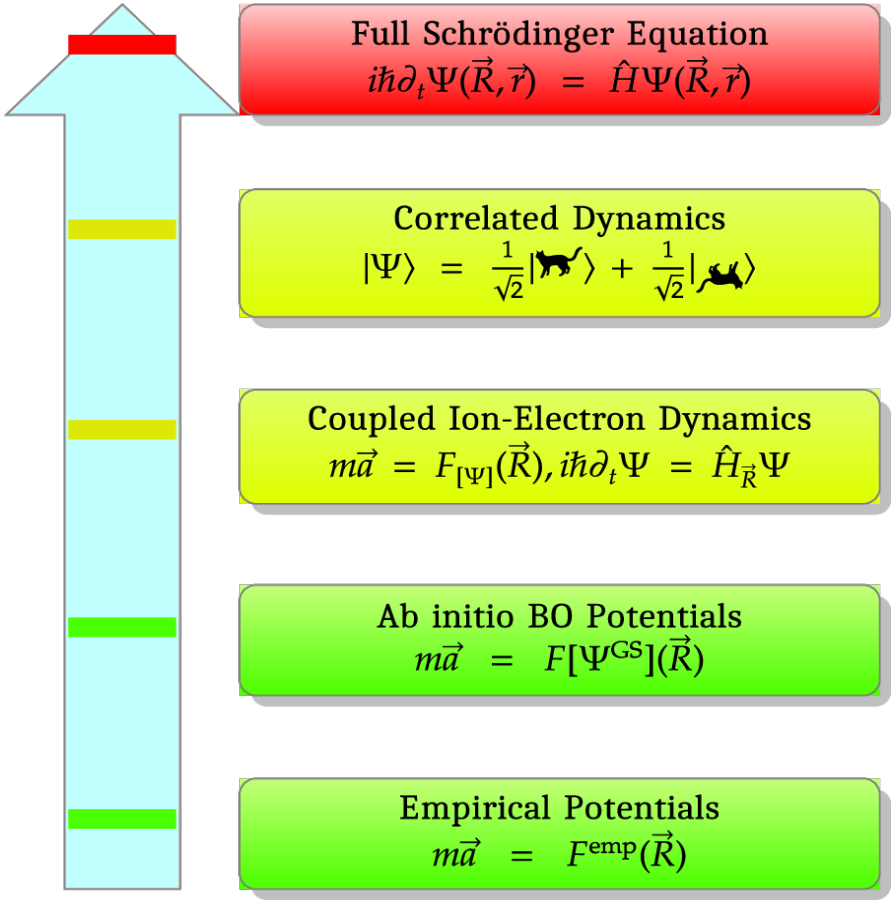

- From time independent quantum mechanics (DFT) to time dependent one (RT-TDDFT)
  - Non-equilibrium (spin-)electron-ion coupled dynamics
  - Materials response to opto-electro-magnetic pulse stimuli
- Platform to explore beyond Ehrenfest dynamics method
  - Programable INQ library for scientists
  - GPU performance ensured by INQ team

”Berry curvature memory through electrically driven stacking transitions,” Xiao, Pemmaraju, Lindenberg et al. *Nat. Phys.* **16**, 1028 (2020)

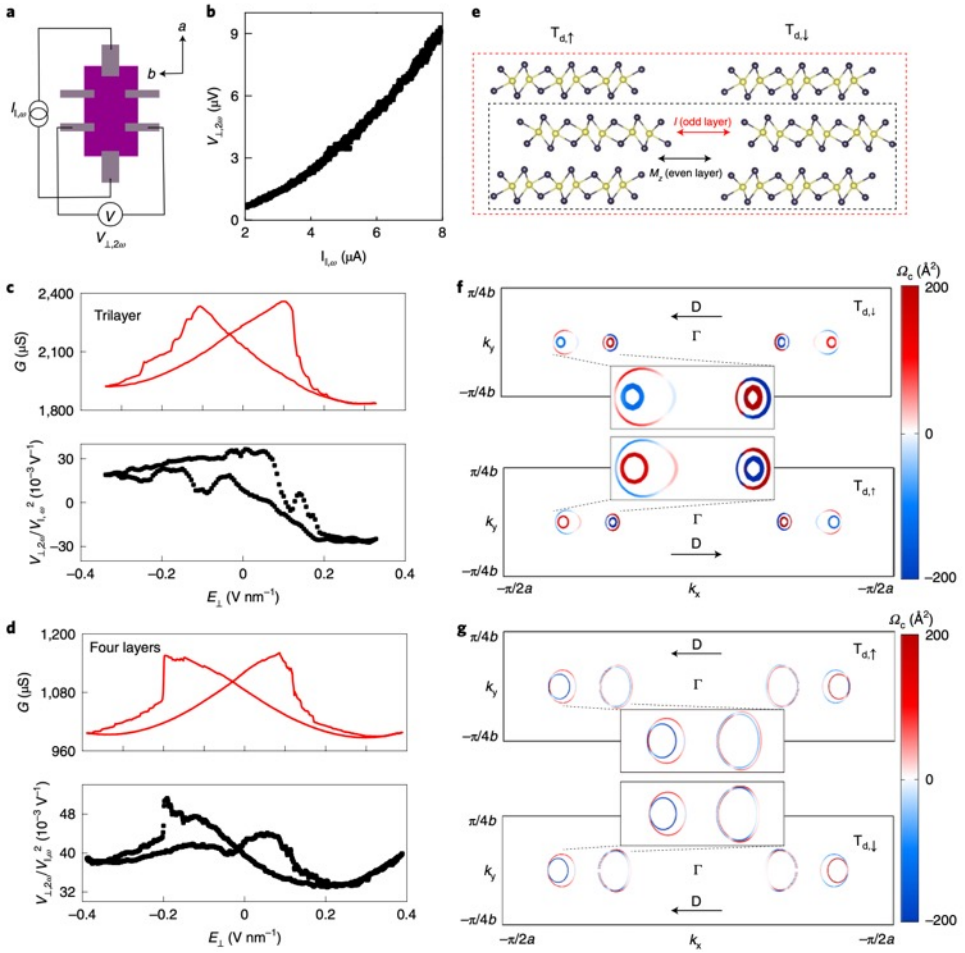

- Electronic stimuli triggers sliding of WTe_2 layers
- Non-linear Hall signals for 3 and 4 layers of WTe_2 show different behaviors, while longitudinal conductance behave similarly (c, d)
- Calculated Berry curvature distribution explains why Berry dipole is reversed for 3 layers WTe_2

INQ: a new code written from scratch for GPUs

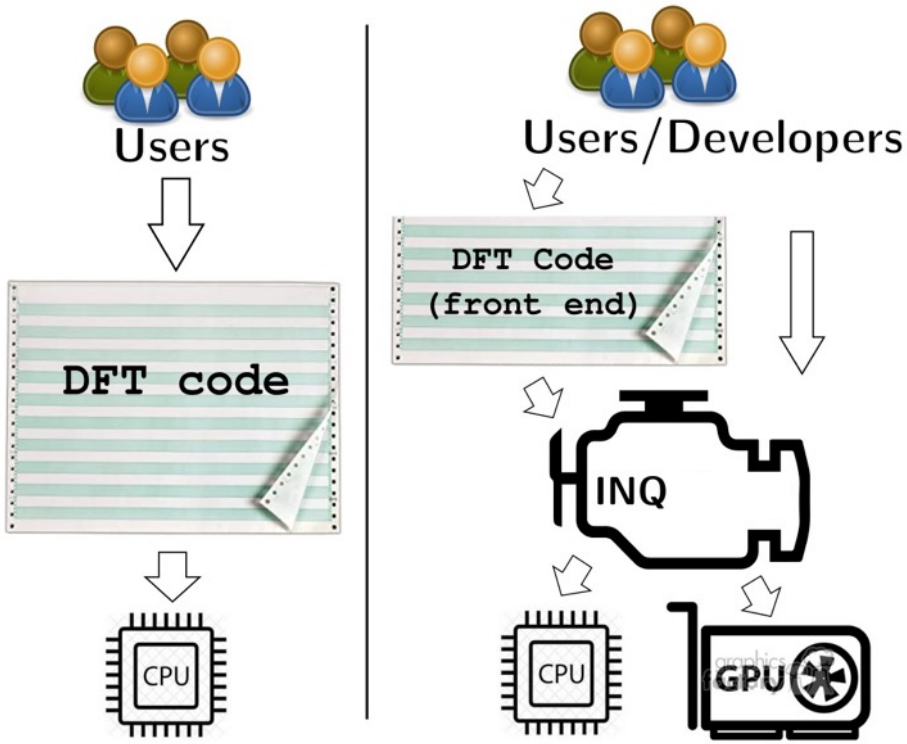

- **Complexity of GPU programing is hidden within INQ library**
- **Users can use INQ as a standalone (RT-TD)DFT code**
- **Other codes can use INQ library to run on the GPU**

Code Structure

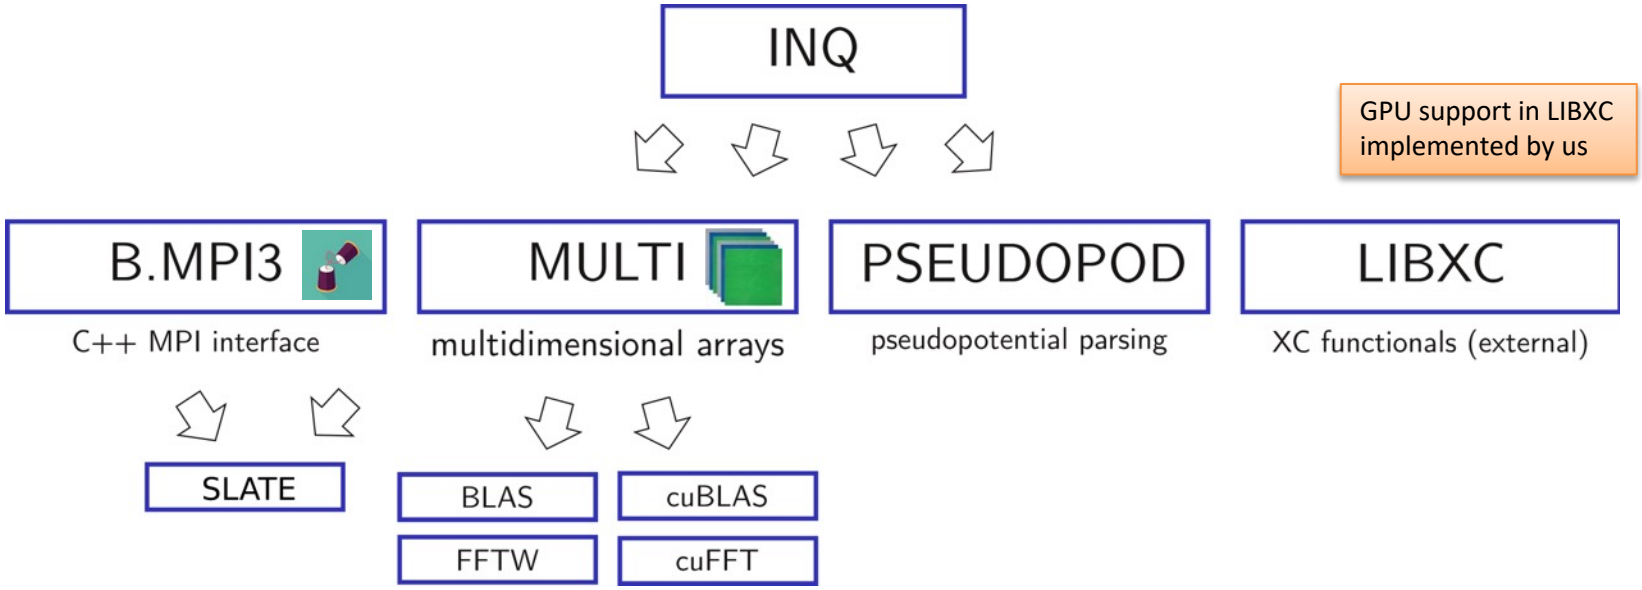

GPU support in LIBXC implemented by us

GPU programming with INQ

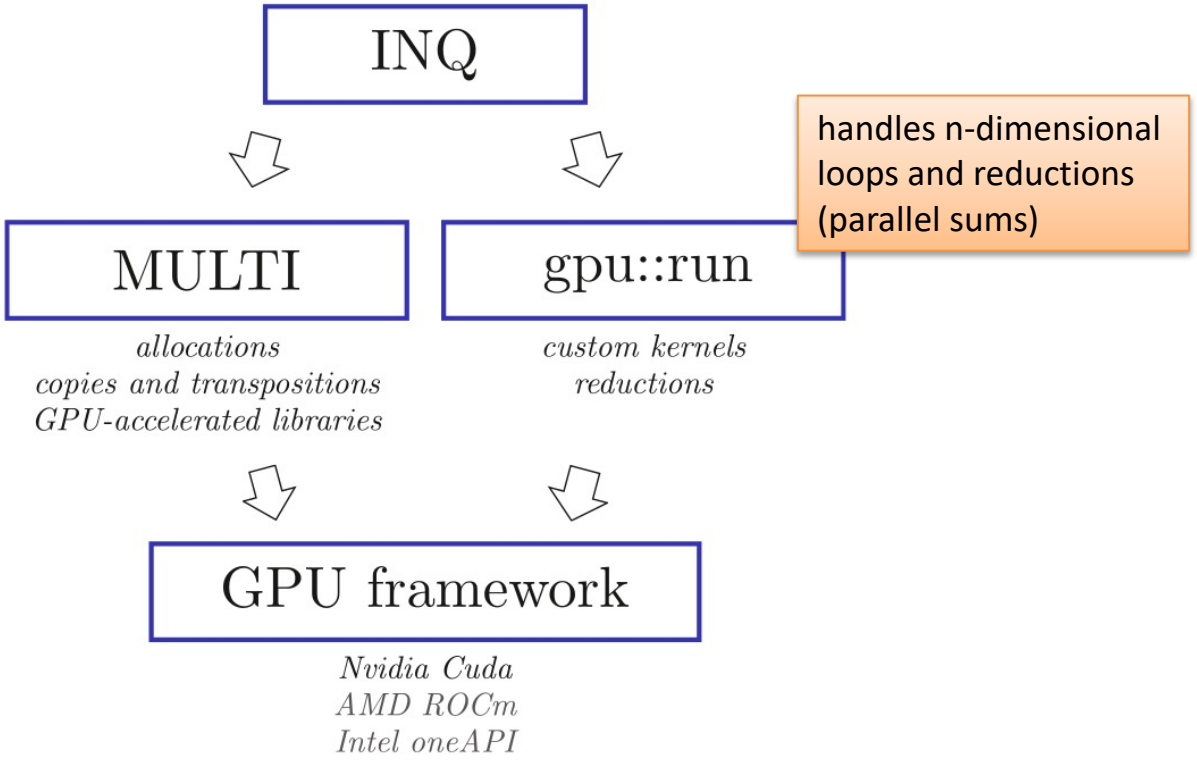

Only CUDA (NVIDIA) is supported now, however, adaptation of INQ library to AMD and Intel GPUs will be performed by INQ team. No change in main routine will be necessary

INQ is a programing platform for scientists

Traditional Paradigm

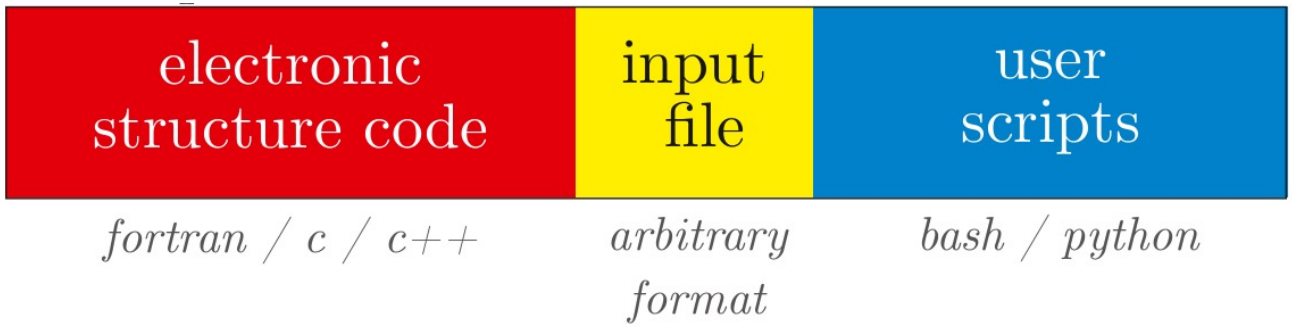

INQ Paradigm

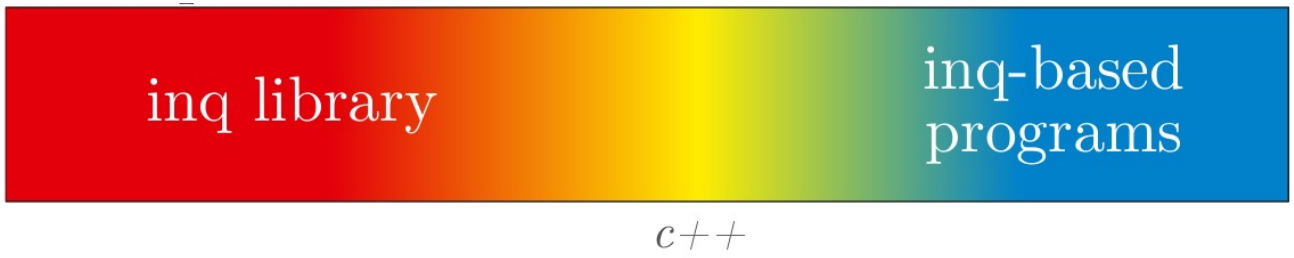

The users write a custom (RT-TD)DFT program using C++ and INQ library

Example: how INQ modules could be used

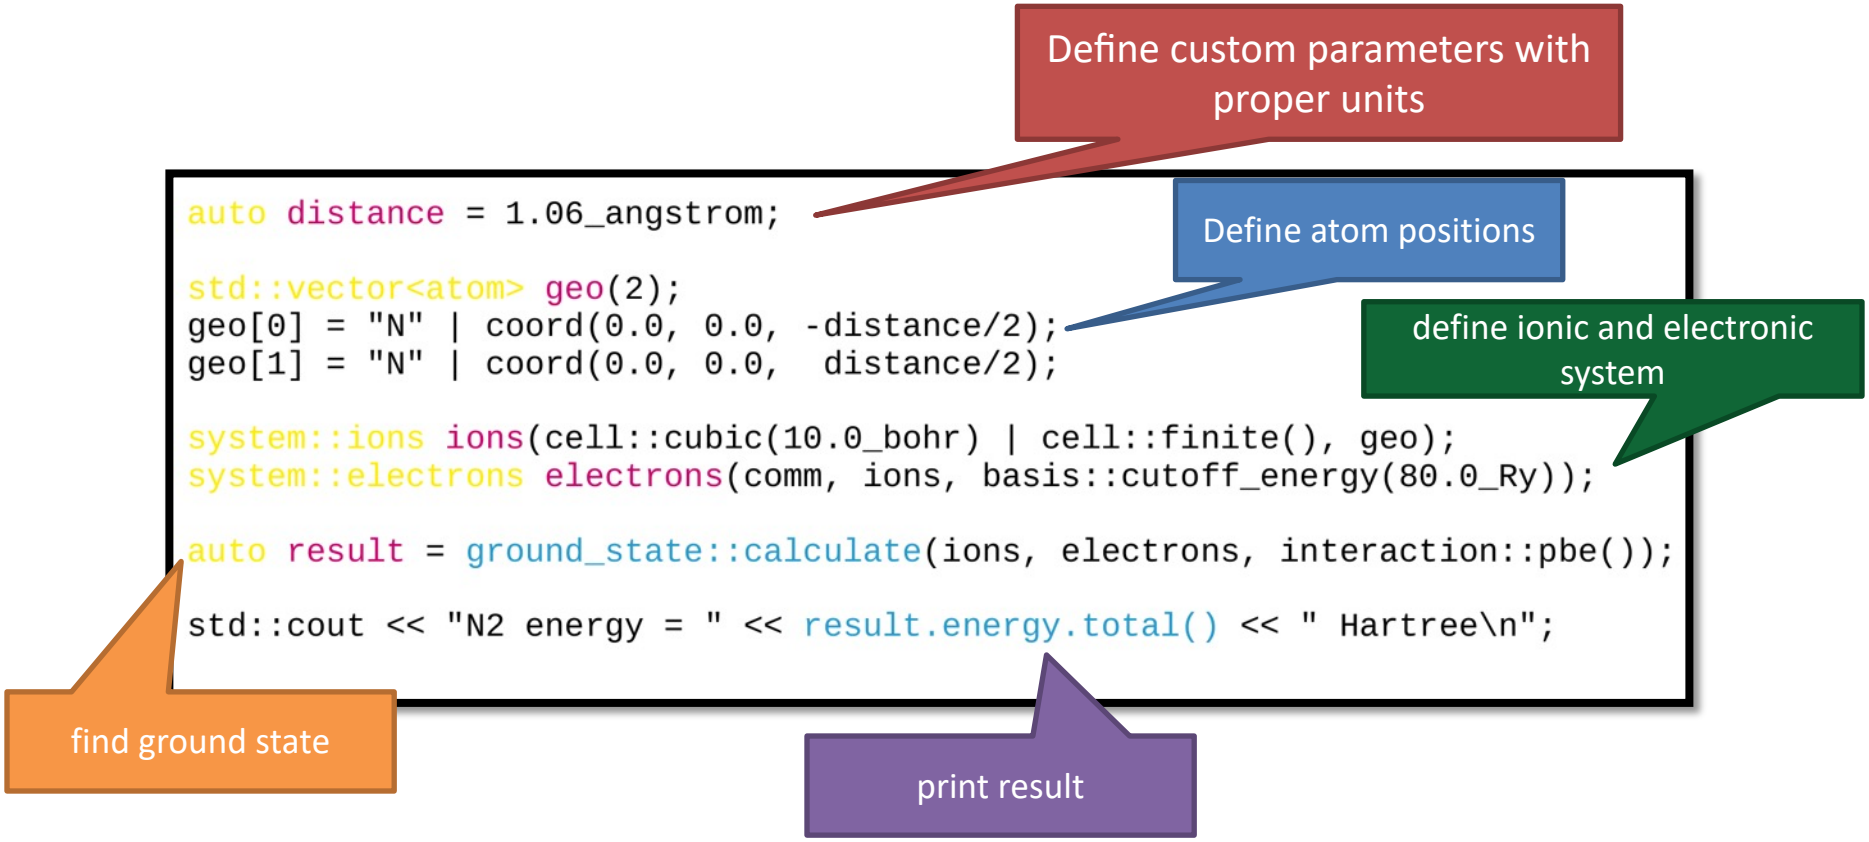


```
auto distance = 1.06_angstrom;

std::vector<atom> geo(2);
geo[0] = "N" | coord(0.0, 0.0, -distance/2);
geo[1] = "N" | coord(0.0, 0.0,  distance/2);

system::ions ions(cell::cubic(10.0_bohr) | cell::finite(), geo);
system::electrons electrons(comm, ions, basis::cutoff_energy(80.0_Ry));

auto result = ground_state::calculate(ions, electrons, interaction::pbe());

std::cout << "N2 energy = " << result.energy.total() << " Hartree\n";
```

Input consists in a normal C++ program using inq functionality

Validation with other codes

molecules

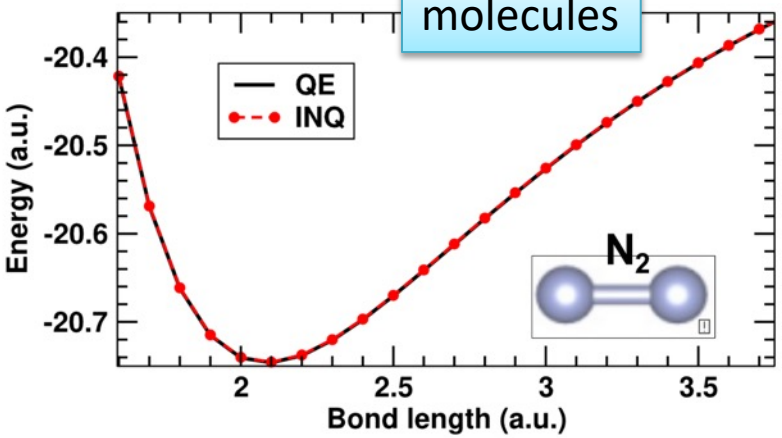

solids

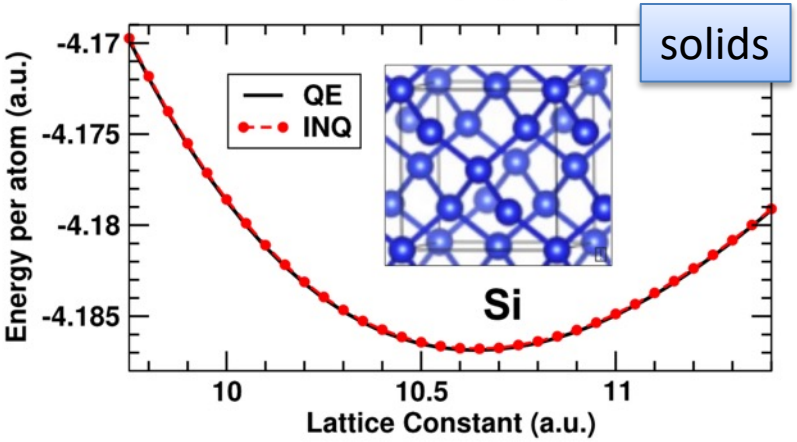

electron dynamics (TDDFT)

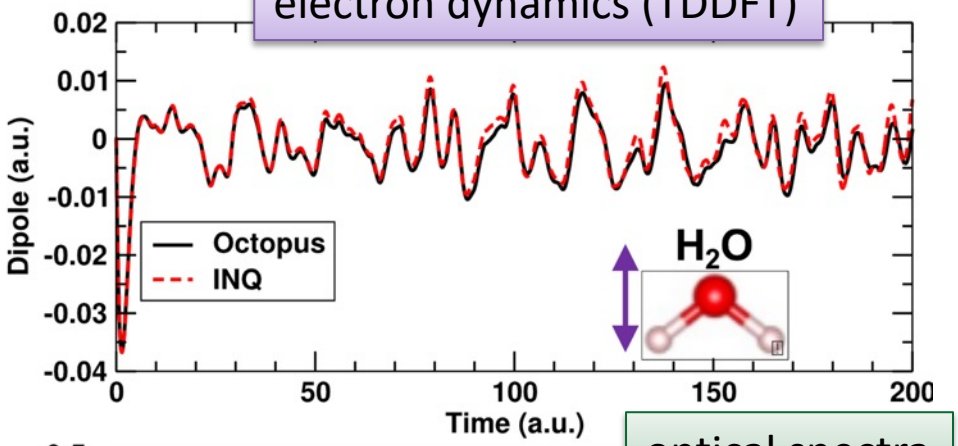

optical spectra

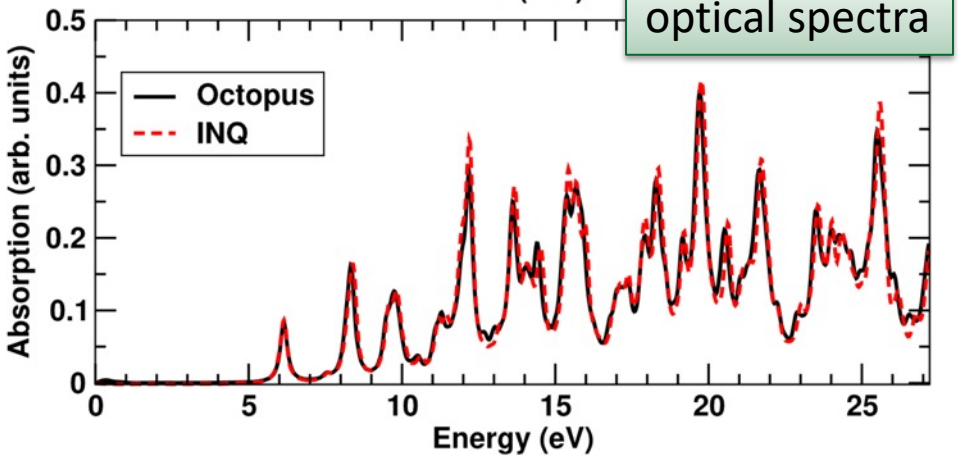

inq reproduces accurately the physical results of established codes

Large systems validation: Stopping Power

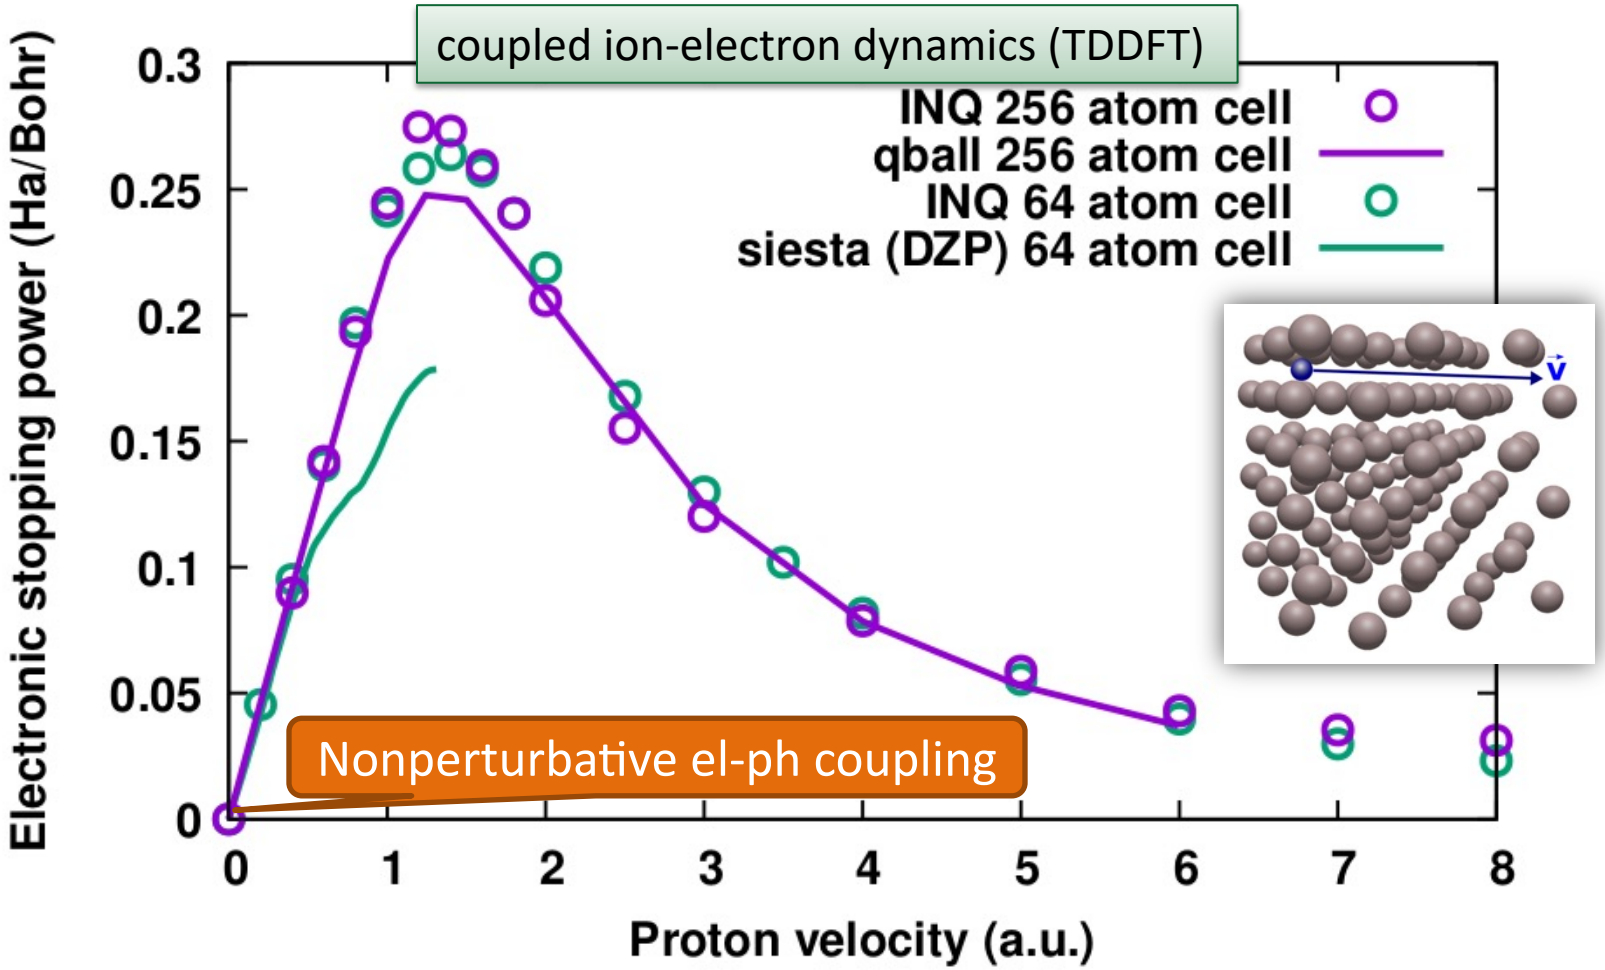

Electronic stopping power of a proton in an FCC aluminum supercell

TDDFT Computational scaling

(LLNL Lassen GPU supercomputer)

Strong scaling:
reduce time to solution for fixed problem size
Aluminum supercell with 1024 atoms

Lassen supercomputer

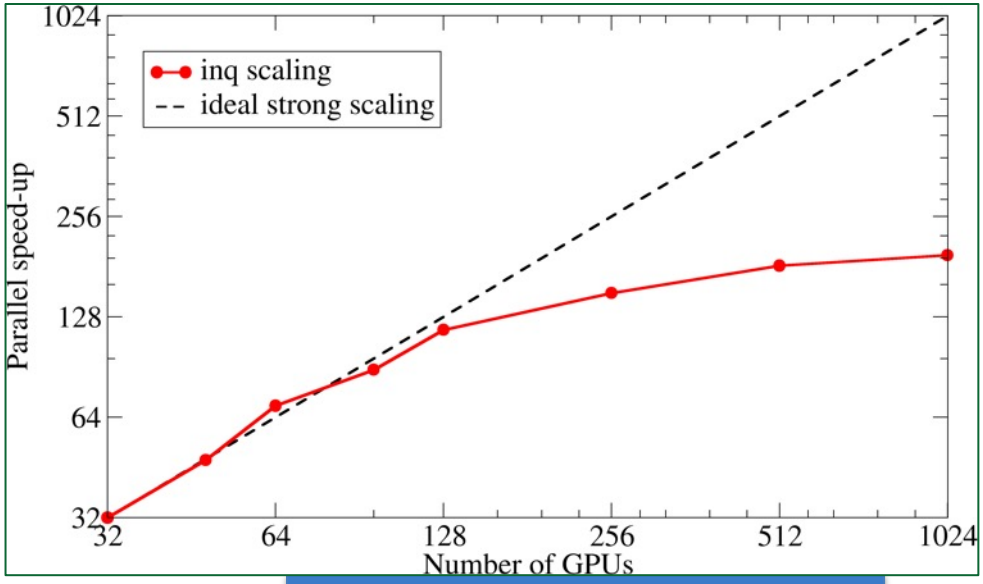

Efficient scaling up to 1000 GPUs

Under 1 second per step is practical for production

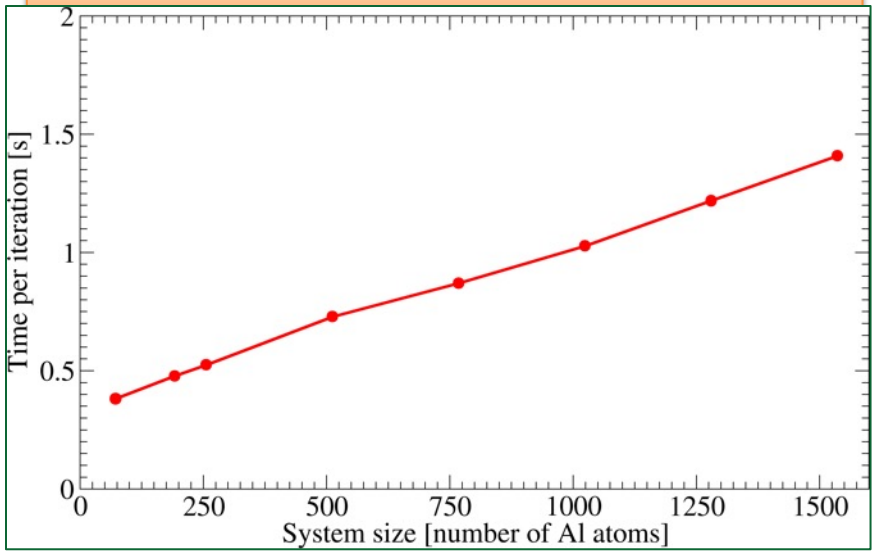

Weak scaling:
variable system size, optimal resources
Aluminum supercell

INQ development plan

	FY21	FY22	FY23
INQ	CPU+GPU+MPI opt.	Adopting to AMD/Intel GPUs	
	Fast hybrid exc-corr	Spiral Boundary Condition	Electric/magnetic field
	Spin-Orbit Coupling	Beyond Ehrenfest	Nuclear spin
Outreach	CCMS/ TMF Users Meeting	Benasque TDDFT (TBD)	CCMS, CECAM (TBD), CCMS

- CCMS: LLNL's summer school
- Online tutorial could be arranged upon request
- Your participation to our regular online meetings will be welcomed
  - Software use, programming, scientific collaboration

Thank you for your attention!

Project URL: https://sc-programs.llnl.gov/npneq
Download INQ: https://gitlab.com/npneq/inq
Contact: ogitsu1@llnl.gov