UNEDF project: some perceived software issues and collaboration opportunities

K. Roche
Future Technologies Group
Computer Science and Mathematics Division
Oak Ridge National Laboratory

*comments for: CScADS Workshop on Libraries and Algorithms for Petascale Applications; Snowbird, Utah ; July 30 – August 2

*I am not able to represent all the efforts happening in this project.

UNEDF := Universal Nuclear Energy Density Functional

- SciDAC2 Application Project
  - http://www.scidac.gov/physics/unedf.html
  - http://www.unedf.org
  - PI, George Bertsch (University of Washington)
  - ASCR, NNSA, NP :: $3M per year , 5 years

Ames National Laboratory - Masha Sosonkina
Argonne National Laboratory - Steven Pieper, Robert Wiringa, E. (Rusty) Lusk, Jorge Moré, Boyana Norris
Lawrence Berkeley National Laboratory - Esmond Ng, Chao Yang
Lawrence Livermore National Laboratory - Juliet Escher, Petr Navratil, Erich Ormand, Ian Thompson
Los Alamos National Laboratory - Joseph Carlson, Toshihiko Kawano
Oak Ridge National Laboratory - Goran Arbanas, ***David Dean***, Witek Nazarewicz, George Fann, Kenneth Roche, William Shelton
Central Michigan University - Mihai Horoi
Iowa State University - James Vary
Michigan State University - B. Alex Brown, Thomas Duguet
University of North Carolina at Chapel Hill - Jonathan Engel
Ohio State University - Richard Furnstahl
San Diego State University - Calvin Johnson
University of Tennessee - Carlos Bertulani, Thomas Papenbrock
University of Washington - *George Bertsch* (PI), ***Aurel Bulgac***

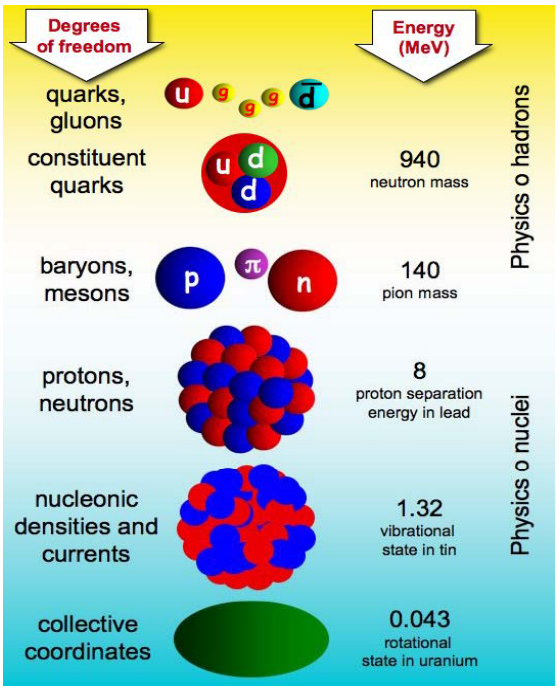
Degrees of freedom
Energy (MeV)
quarks, gluons
constituent quarks
940
neutron mass
baryons, mesons
140
pion mass
protons, neutrons
8
proton separation energy in lead
nucleonic densities and currents
1.32
vibrational state in tin
collective coordinates
0.043
rotational state in uranium
Physics o hadrons
Physics o nuclei

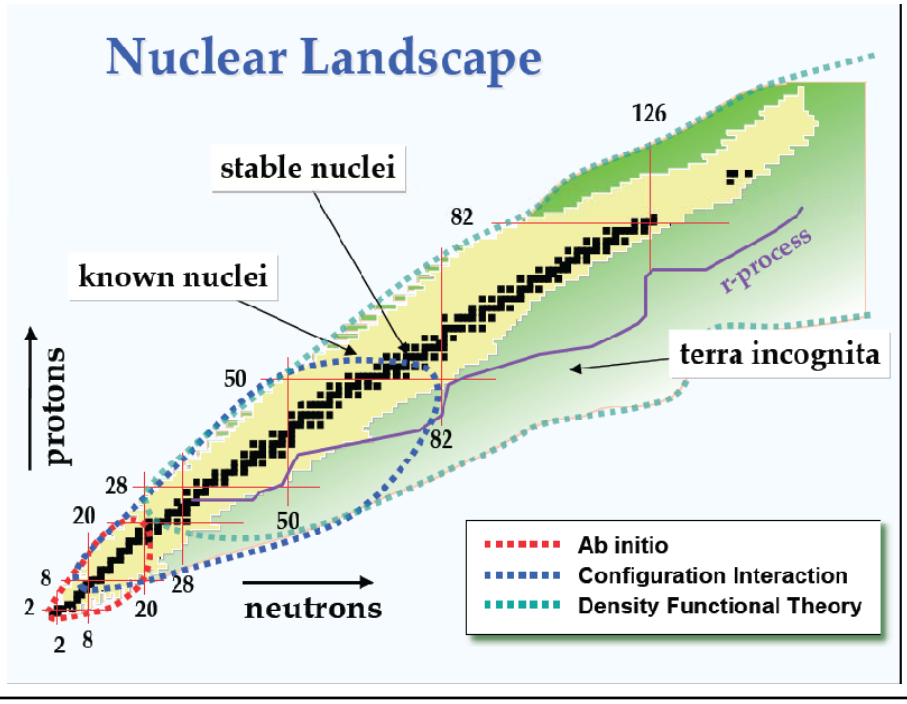
Nuclear Landscape
stable nuclei
known nuclei
terra incognita
r-process
protons
neutrons
126
82
50
28
20
8
2
Ab initio
Configuration Interaction
Density Functional Theory

Universal Nuclear Energy Density Functional

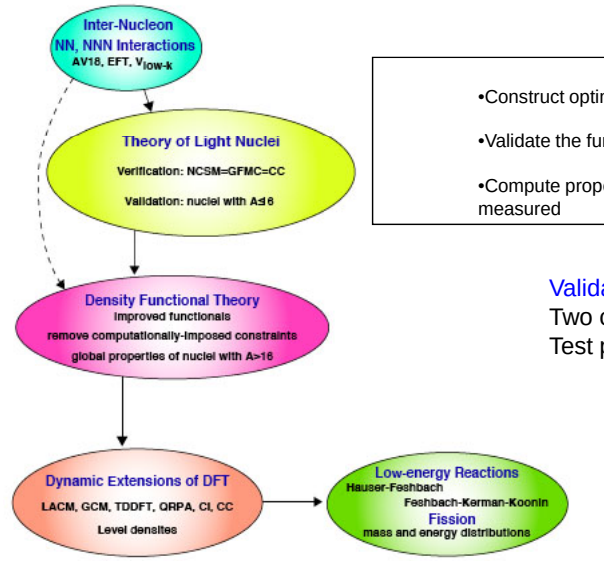

- Construct optimal energy density functional
- Validate the functional against structure data
- Compute properties that have not been (cannot be) measured

Validation:
Two center potential
Test problem @unedf.org

Simple example: HF for fermionic system

H = $\sum_{i=1,N}(t_i + u_i) + .5 * \sum_{i,j=1;i.ne.j} v(x_i,x_j)$

- approximate 2-particle operator by effective 1-particle potential

H_{hf} = (N)^(-.5) det$\{\psi_i(x_j)\}$

- do not know HF potential → can't compute 1particle orbitals
- apply Ritz' variational theorem and minimize total energy of system wrt 1particle orbitals

$\delta_{\psi_i(x)}$ [<H> - $\sum_{i=1,N}\lambda_i$ Integral[dy $\psi_i(y)$ $\psi_i(y)$]] = 0

- Lagrange multipliers enforce normalized 1particle orbitals

$v_{hf}\psi_i$ = u(x)$\psi_i(x)$ + integral[dy $\rho(y)$ v(x,y) $\psi_i(x)$] –integral[dy $\rho(x,y)$ v(x,y) $\psi_i(y)$

$\rho(y)=\sum_i$ ψ_i (y) ψ_i (y) and $\rho(x,y)=\sum_i$ ψ_i (y) ψ_i (x

Simple example: HF for fermionic system (continued)

- Self-consistent iteration
  - prescribe the 1particle orbitals ψ_i i=1,N (in chosen basis)
  - (loop)compute HF potential v_{hf} from density matrix $\rho(x,y)=\sum_i \psi_i(x)\psi_i^*(y)$
  - solve eigenequation $(t + v_{hf})\psi_\alpha = \varepsilon_\alpha \psi_\alpha$
  - choose N orbitals ψ_α with N lowest eigenvalues ε_α
  - loop until self-consistent convergence
- Approximate HF total energy of the system from the occupied 1particle orbitals
  - $E_{hf} = .5 \sum_{i=1,N} \varepsilon_i + \langle i | t + u | i \rangle$ (2nd quantized notation)

Kohn-Sham Theory:

- take the variational functions to be the HF orbitals
- calculate kinetic energy and density w/ orbitals
- introduce the exchange correlation energy, E_{xc}
- Minimize the resulting functional $E_{KS}[\psi_1 \psi_2 \psi_3 \psi_4 \psi_N]$
- depends on nonlocal term, $\int \rho(x) | \text{grad}\, \psi_i(x) |^{**}2$

Some computational issues:

- Solve the poisson equation
  - periodic lattice, 3D Fourier transforms ~ O(N log N)
    - $\psi(x)$ = F^-1 * 4pi/k**2 * F $\rho(x)$
  - can use multigrid on laplacian in coordinate space
- Solve Schroedinger equation → multiply H_{KS} on each wavefunction
  - Split operator
    - $H_{KS}\psi_x$ = F^-1 * k**2 / 2m * F ψ (x) $V_{KS}(x)\psi$ (x)
  - Iterative refrinement of wavefunctions
    - Imaginary time
    - Conjugate gradient (large KE → preconditioning)

DFT Challenges for nuclear physics: LDA, GGA, SLDA , …

- Two kinds of fermions (isospin)
- Need to include pairing
- Broken symmetry of self bound systems

- Hartree Fock Bogoliubov –a generalized single particle model
- HFB equation has the following symplectic structure (no time to dissect this)

```
| h11-mu  h12          0          Delta  | U+      U+
| h21       h22-mu    -Delta       0     | U-  = E  U-
| 0         -Delta^*   -h11+mu   -h12    | V+      V+
| Delta^*    0          -h21     -h22+mu | V-      V-
```

Occupied
Not occupied
Spin up
Spin down

$$\begin{vmatrix} h-\mu & \text{pair\_field} \\ & \\ -\text{pair\_field}^* & -h+\mu \end{vmatrix} \begin{vmatrix} U \\ \\ V \end{vmatrix} = E \begin{vmatrix} U \\ \\ V \end{vmatrix}$$

Current software dependencies for the HFB based (SLDA) TDDFT:

- Work between coordinate and k spaces in a plane wave basis
- Each MPI process has a local copy of both lattices
- Quasiparticle wavefunctions are distributed over each MPI process
- 3D FFTs → FFTW3.1.2
  - Used extensively each time step on the single particle orbitals
- Formation of density based terms depends only on 1particle qpwf
  - Requires reduction over all processes each time step for the normal density, kinetic energy density, density gradients
- Compute gradients of both kinetic and normal densities as corrections to the edf
- Time evolution is now with ABM PMC technique
  - Could use Crank-Nicholson
  - Loss of orthonormality is an issue here → Gram-Schmidt

Discussion of memory concerns:

```
ALLOCATE (wavf(nwfip,-3:1,4,Nx,Nx,Nx),STAT=ierr)
IF (ierr .NE. 0) THEN
  WRITE(*,*) 'ERROR: cannot ALLOCATE wavf()'
ENDIF

ALLOCATE (wavf_t_der(nwfip,-3:1,4,Nx,Nx,Nx),STAT=ierr)
IF (ierr .NE. 0) THEN
  WRITE(*,*) 'ERROR: cannot ALLOCATE wavf_t_der()'
ENDIF

ALLOCATE (wavf_modifier(nwfip,-3:1,4,Nx,Nx,Nx),STAT=ierr)
IF (ierr .NE. 0) THEN
  WRITE(*,*) 'ERROR: cannot ALLOCATE wavf_modifier()'
ENDIF

ALLOCATE (wavf_predictor(nwfip,-3:1,4,Nx,Nx,Nx),STAT=ierr)
IF (ierr .NE. 0) THEN
  WRITE(*,*) 'ERROR: cannot ALLOCATE wavf_predictor()'
ENDIF

ALLOCATE (wavf_corrector(nwfip,-3:1,4,Nx,Nx,Nx),STAT=ierr)
IF (ierr .NE. 0) THEN
  WRITE(*,*) 'ERROR: cannot ALLOCATE wavf_corrector()'
ENDIF

ALLOCATE (rho(Nx,Nx,Nx),STAT=ierr)
IF (ierr .NE. 0) THEN
  WRITE(*,*) 'ERROR: cannot ALLOCATE rho()'
ENDIF

IF (np>1) THEN
  IF (ip==0) THEN
    ALLOCATE (trho(Nx,Nx,Nx),STAT=ierr)
    IF (ierr .NE. 0) THEN
      WRITE(*,*) 'ERROR: cannot ALLOCATE trho()'
    ENDIF
  ENDIF
ENDIF
```

Pursuits:

- Check-point restart
- rma
- Pthread based io overlap with computation , data reorganization

128**3 ~ 2,097,152

Coupled Cluster example:

- ab initio technique
- calculates ground-state properties of closed-shell (or sub-shell) nuclei
- solves coupled nonlinear sets of equations (largest ~10M unknowns)
- CCSD ~ O(nu^2 * No^4) (nu:=unoccupied orbitals; No:=occupied)
- CCSDT ~ O(nu^3 * No^5)

SAMPLE TERM: (fully reduced in 2nd quantization notation)

For i=1,n ~ 100
 For j=1,n
  For a=n+1,N ~ 1000
   For b=n+1,N

$$F(a,b,i,j) = \sum_{c=n+1,\,N} \sum_{d=n+1,\,N} \sum_{k=1,n} \sum_{k=1,n} V(k,l,c,d) * T(c,d,i,j) * T(a,b,k,l)$$

Make the thing look like a sequence of matrix multiplies.

```
For i=1,n~100
 For j=1,n~100
  For a=n+1,N~1000
   For b=n+1,N
```

distributed local

$$F(a,b,i,j) = \sum_{c=n+1,\,N} \sum_{d=n+1,\,N} \sum_{k=1,n} \sum_{k=1,n} V(k,l,c,d) * T(c,d,i,j) * T(a,b,k,l)$$

Tasks

- copy data into matrix form
  - from disk (next slide)
  - in network
- impose 2D block decomposition for larger sub-blocks
- use Goto BLAS or PBLAS *gemm() subroutines
- write products

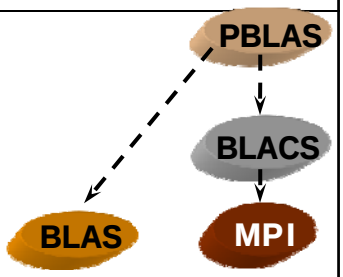

Other

- Shared and remote memory protocols, co-arrays, global arrays
- Role of multi-core is not clear

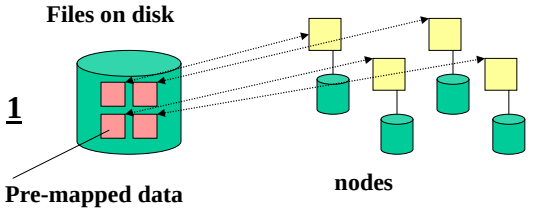

The pre-mapped data is brought in-core in parallel by each process in the logical process grid. Each process accesses its own file.

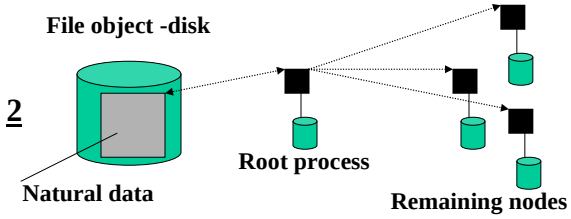

The natural data is brought in-core by a chosen process (root) from the logical process grid. The data may be distributed by root in a manner imposing the mapping during direct communication with other nodes .or. Root may broadcast the data as is and let the mapping be done by each process locally.

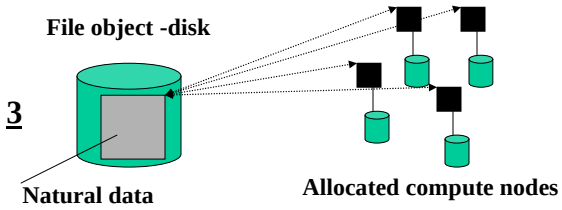

The natural data is brought in-core in parallel by each process in the logical process grid. The mapping may be imposed exactly during the load through random access to the file or bulk data can be brought over the network by each process and the mapping imposed in-core.

(talk phase transistion here) Lesson learned:

G1 –from MPI_COMM_WORLD
G3 = G1 \ G2 (the work group)
G2 := outliers

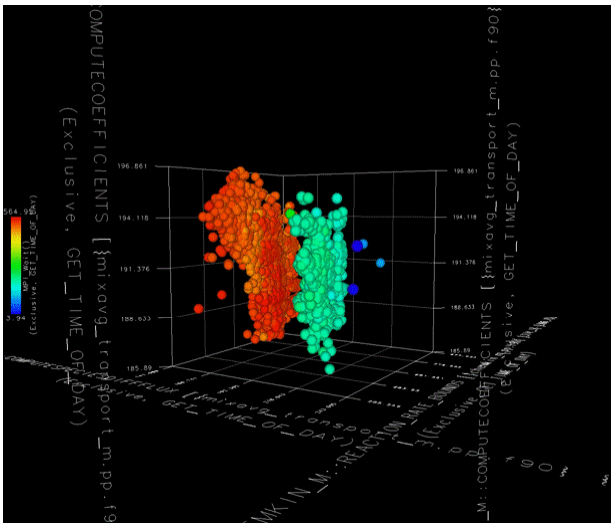

Conclusion: (Could go on but this is a good place to stop)

•Multicore, what adjustments will be needed, enhancements gained

•Collective io is coming
  •PxQ → RxS 2d block cyclic mapping research
  •Check-pt-restart

•Believe we have a need for hybrid programming model for pthreads in an MPI environment; also language interoperability

•Many of the unedf codes could not be discussed
  •Eg, multiwavelet basis based approach is interesting effort

•Gram-Schmidt is a concern

•SVD and HOSVD are being pursued in some UNEDF codes

•Dense and sparse (both iterative and direct) solvers are critical to unedf and at the petascale

•Diagonalization of the Hamiltonian → zheevd()
  •Ax=cx ; A(x+dx)=(c+dc)(x+dx)