

Computing Atomic Nuclei on the Cray XT5



Hai Ah Nam


Scientific Computing Group
National Center for Computational Sciences

In collaboration with


David J. Dean, Oak Ridge National Laboratory
James P. Vary and Pieter Maris, Iowa State University
Petr Navratil and W. Erich Ormand, LLNL

Presentation for

Cray User Group Meeting, May 2009

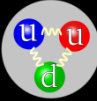


Nuclear Physics 101




- Atom = Protons, neutrons, and electrons
- Protons & neutrons



 QCD is a theory of the fundamental strong force describing the interactions of quarks & gluons



Strong nuclear force




Hadron
(bound state of quarks held together by the strong force)

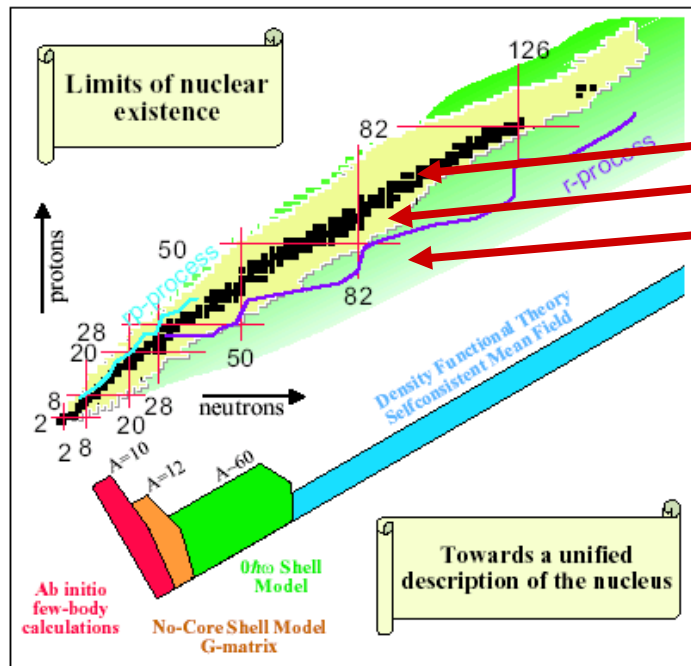
Baryon
(3 quarks)
protons & neutrons

Meson
(quark – anti quark pair)
pions, kaons, etc.

What do we want?



- To understand nuclear properties in terms of the interactions between nucleons.
- Consistent microscopic theory of nuclei and their reactions.



stable nuclei
known nuclei
terra incognita

TRICKS OF THE TRADE... Methods


No-Core Shell Model (NCSM)
Green's Function Monte Carlo (GFMC)
Coupled Cluster Methods (CC)

Density Functional Theory (DFT)

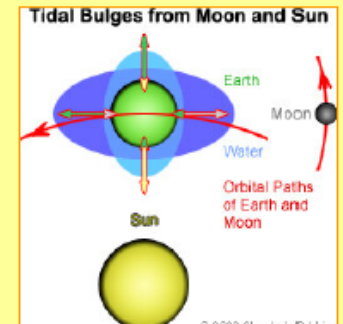
Weapon of Choice

- **Ab initio No-Core Shell Model with 3-nucleon forces**
 - **Why *ab initio* (first principles)?**
 - Satisfaction at the end of the day
 - **Why no-core shell model (NCSM)?**
 - Proven successful ab initio approach to nuclear structure
 - Only method capable of employing ab initio Chiral EFT interactions for $A > 4$
 - **Why 3-nucleon forces?**
 - Nucleons are not point particles (i.e. not elementary)
 - We neglect some internal degrees of freedom (e.g Δ -resonance, polarization effects, ...)

Example from celestial mechanics:
Earth-Moon system: point masses and modified two-body interaction



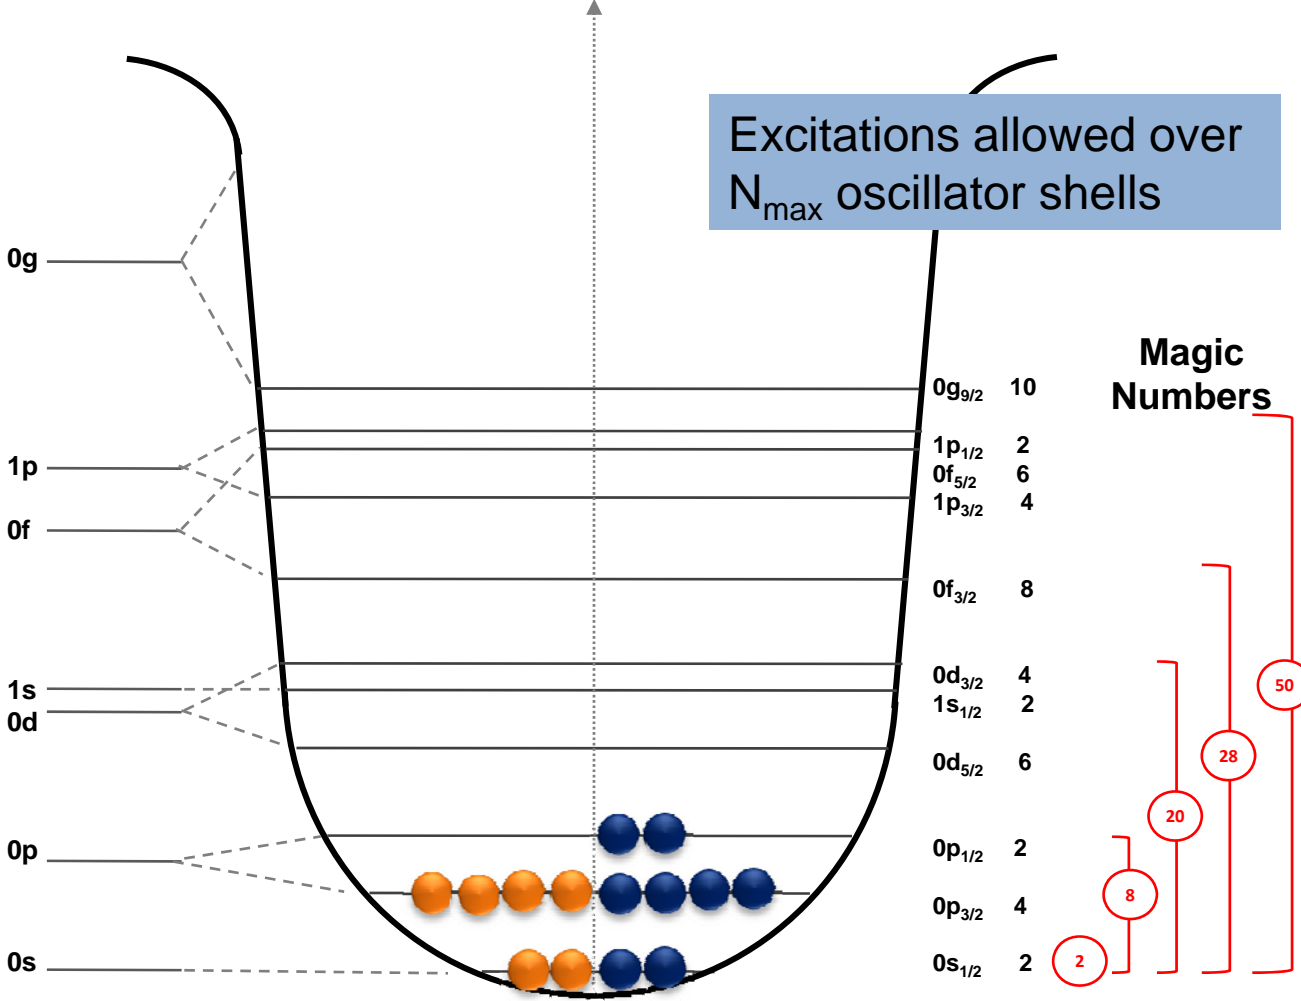
Other tidal effects cannot be included in the two-body interaction! Three-body force unavoidable for point masses.



Nuclear Shell Model


 protons

 neutrons



What's your problem man?

- **Physics Problem** $H\Psi_i = E_i\Psi_i$
 - Given a 2- or 3-body interaction, # of protons & neutrons, calculate the energy spectrum (E_i) and wavefunctions (Ψ_i) for different states of the system
 - Use the wavefunctions to calculate observables i.e. rms radii, moments, transition rates between ground state/excited states, nuclear reactions, ...
- **Computational Problem**
 - Construct large ($10^9 \times 10^9$) sparse symmetric real matrix H
 - Obtain the lowest eigenvalues & eigenfunctions (Lanczos)



Pick Your Poison


- **Store matrix elements in memory**
 - I feel the need for speed
 - Limited by available memory
- **Store matrix elements on disk**
 - It just doesn't all fit
 - It's soooooooooooooo slow
- **Re-compute on-the-fly**
 - Efficient determination of non-zero matrix elements
 - Also slow... i.e. ${}^9\text{Be}$: 4064 CPU-hrs: 8128 cores @ 30 min or 48 cores @ 3.5 days
- **All of the above**




- **Many Fermion Dynamics – nuclear**

- Platform-independent F90 code with MPI
- Scalable (has run successfully on 30k+ processors)
- Load-balanced

Scaled to
30,000 cores
on Jaguar XT4



on single processor

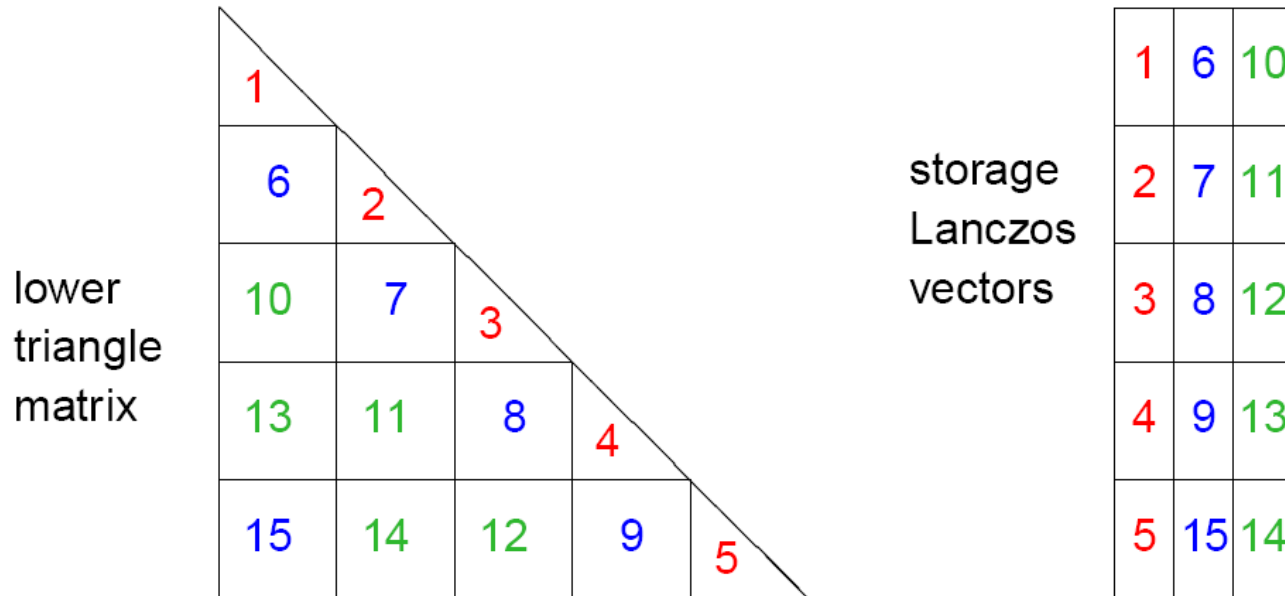


on 10 processors

- round-robin distribution of many-body states over d procs
- however, no (apparent) structure in sparse matrix

MFDn Distributes the Fun

- Store lower half of symmetric matrix, distributed over $n = d(d + 1)/2$ processors with d “diagonal” processors
- Store lanczos vectors on one of $(d+1)/2$ groups of d procs



We need more memory...

Estimates of aggregate memory needed for storage of sparse symmetric Hamiltonian matrix in compressed column format

nucleus	N_{\max}	dimension	2-body	3-body	4-body
${}^6\text{Li}$	12	$4.9 \cdot 10^6$	0.6 GB	33 TB	590 TB
${}^{12}\text{C}$	8	$6.0 \cdot 10^8$	4 TB	180 TB	4 PB
${}^{12}\text{C}$	10	$7.8 \cdot 10^9$	80 TB	5 PB	140 PB
${}^{16}\text{O}$	8	$9.9 \cdot 10^8$	5 TB	300 TB	5 PB
${}^{16}\text{O}$	10	$2.4 \cdot 10^{10}$	230 TB	12 PB	350 PB
${}^8\text{He}$	12	$4.3 \cdot 10^8$	7 TB	300 TB	7 PB
${}^{11}\text{Li}$	10	$9.3 \cdot 10^8$	11 TB	390 TB	10 PB
${}^{14}\text{Be}$	8	$2.8 \cdot 10^9$	32 TB	1100 TB	28 PB
${}^{20}\text{C}$	8	$2 \cdot 10^{11}$	2 PB	150 PB	6 EB
${}^{28}\text{O}$	8	$1 \cdot 10^{11}$	1 PB	56 PB	2 EB

(presented at *Extreme Scale Computing Workshop – nuclear physics* Washington DC Jan 2009)

Petascale Early Science

COMPUTE THE FUTURE

Cosmic Radiation

Energetic Neutron

Neutron capture
by ^{14}N



Reacts with oxygen
to form CO_2

Biosphere absorbs ^{14}C



The carbon
in buried matter
decays and is not
replaced with new ^{14}C

Due to its long half-life,
 ^{14}C has been used in
dating organic materials,
up to 60,000 years old,
since the 1950's.



Puzzling to Scientists...

What is the nuclear structure of ^{14}C that leads to its anomalously long half-life?

$\tau_{1/2} = 5730 \text{ years}$

^{10}Be and ^{14}C have extremely long half-lives compared to other light nuclei ($1.6 \times 10^6 \text{ years} / 5,730 \text{ years}$). Their long half-lives make both isotopes useful for radioactive dating.

	Experimental	Calculated
$B(\text{GT}) : ^{10}\text{Be} \rightarrow ^{10}\text{B}$	0.08	0.06 (3-body: 0.066)
$M_{\text{GT}} \ ^{14}\text{C} \rightarrow ^{14}\text{N}$	0.002	0.07

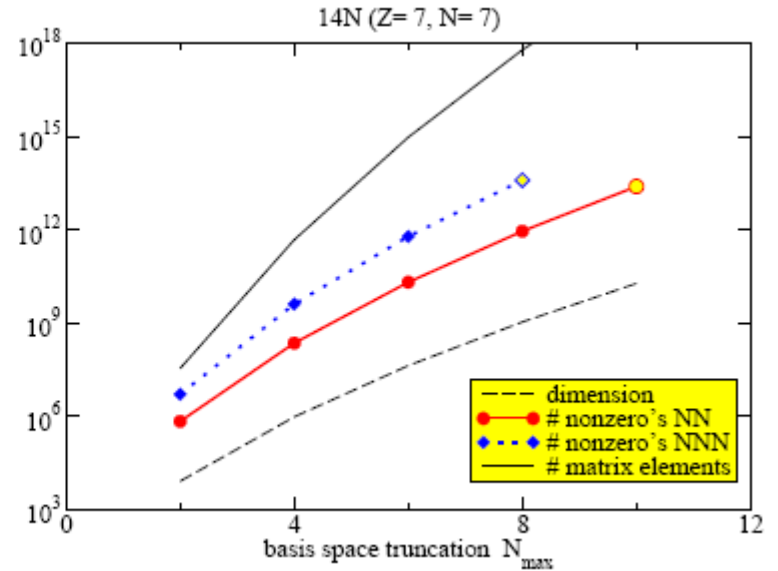

Chart of light nuclei that decay via beta emissions



$$B(\text{GT}) \sim |M_{\text{GT}}|^2 \sim \frac{1}{\tau_{1/2}}$$


$$M_{\text{GT}} \sim \langle \psi_f | \hat{O} | \psi_i \rangle$$

Growing Pains



14C		Est. Non-Zero M.E.	
Nmax	Dim	2B	3B
0			
2	5.80E+03	4.00E+05	2.90E+06
4	7.32E+05	1.62E+08	2.80E+09
6	3.37E+07	1.55E+10	4.42E+11
8	8.73E+08	6.97E+11	2.90E+13
10	1.54E+10	1.94E+13	

14N		Est. Non-Zero M.E.	
Nmax	Dim	2B	3B
0			
2	8.40E+03	7.00E+05	5.20E+06
4	9.75E+05	2.29E+08	4.10E+09
6	4.32E+07	2.07E+10	6.08E+11
8	1.09E+09	9.01E+11	3.90E+13
10	1.89E+10	2.45E+13	



Back of the envelope...

$$\frac{\text{memory}}{\text{core}} = 2(4) \frac{m.e.}{\text{cores}} + 5(4) \frac{\text{dim}}{\text{diag}}$$

Matrix elements input/output
vectors

$$\text{cores} = \frac{d(d+1)}{2}, d = \text{diagonal}$$

148,785 cores / d = 545

PAST DUE

3-body, Nmax=8

$$^{14}\text{C}: \frac{\text{memory}}{\text{core}} = 2(4) \frac{2.9e13}{\text{cores}} + 5(4) \frac{8.73e8}{\text{diag}} \approx 1.59\text{GB}$$

$$^{14}\text{N}: \frac{\text{memory}}{\text{core}} = 2(4) \frac{3.9e13}{\text{cores}} + 5(4) \frac{1.09e9}{\text{diag}} \approx 2.14\text{GB}$$

+ ~.15GB overhead

Needed: 260 – 340 TB

Fitting in

- **Sucking in our breathe**
 - Integer compression (integer*4 → integer3?)
- **Exercise & Diet**
 - Out-of-core
- **Exorcism**
 - New algorithms
 - Return to the physics
- **Have we jumped the shark?**
 - Wait for the next upgrade & cross our fingers

Out-of-core... Need More Envelopes!

3-body, Nmax=8

$${}^{14}\text{N}: \frac{\text{memory}}{\text{core}} = 2(4) \frac{3.9e13}{\text{cores}} + 5(4) \frac{1.09e9}{\text{diag}} \approx 2.14\text{GB}$$

2.10 GB 0.04 GB

Move 0.5GB/core to disk → **AGGREGATE:** 74392.5 GB read

200 GB/s maximum throughput → 6.2 min/read

ITERATIVE: 2 reads/iteration (move in/out data as needed)

12.4 min/iteration... 500 iteration for Lanczos

103 hrs for Lanczos + 4 hr for other stuff ~ 107 hrs

= 15+ Million CPU-hrs/run ... need 12 runs for full study of ${}^{14}\text{N}$!!! ... maniacal laughter!!

30 Million
CPU-HR
Allocation

What on earth am I going to do?

- **Physics don't fail me now!**
- **Rather than asking questions... your suggestions are welcome!**