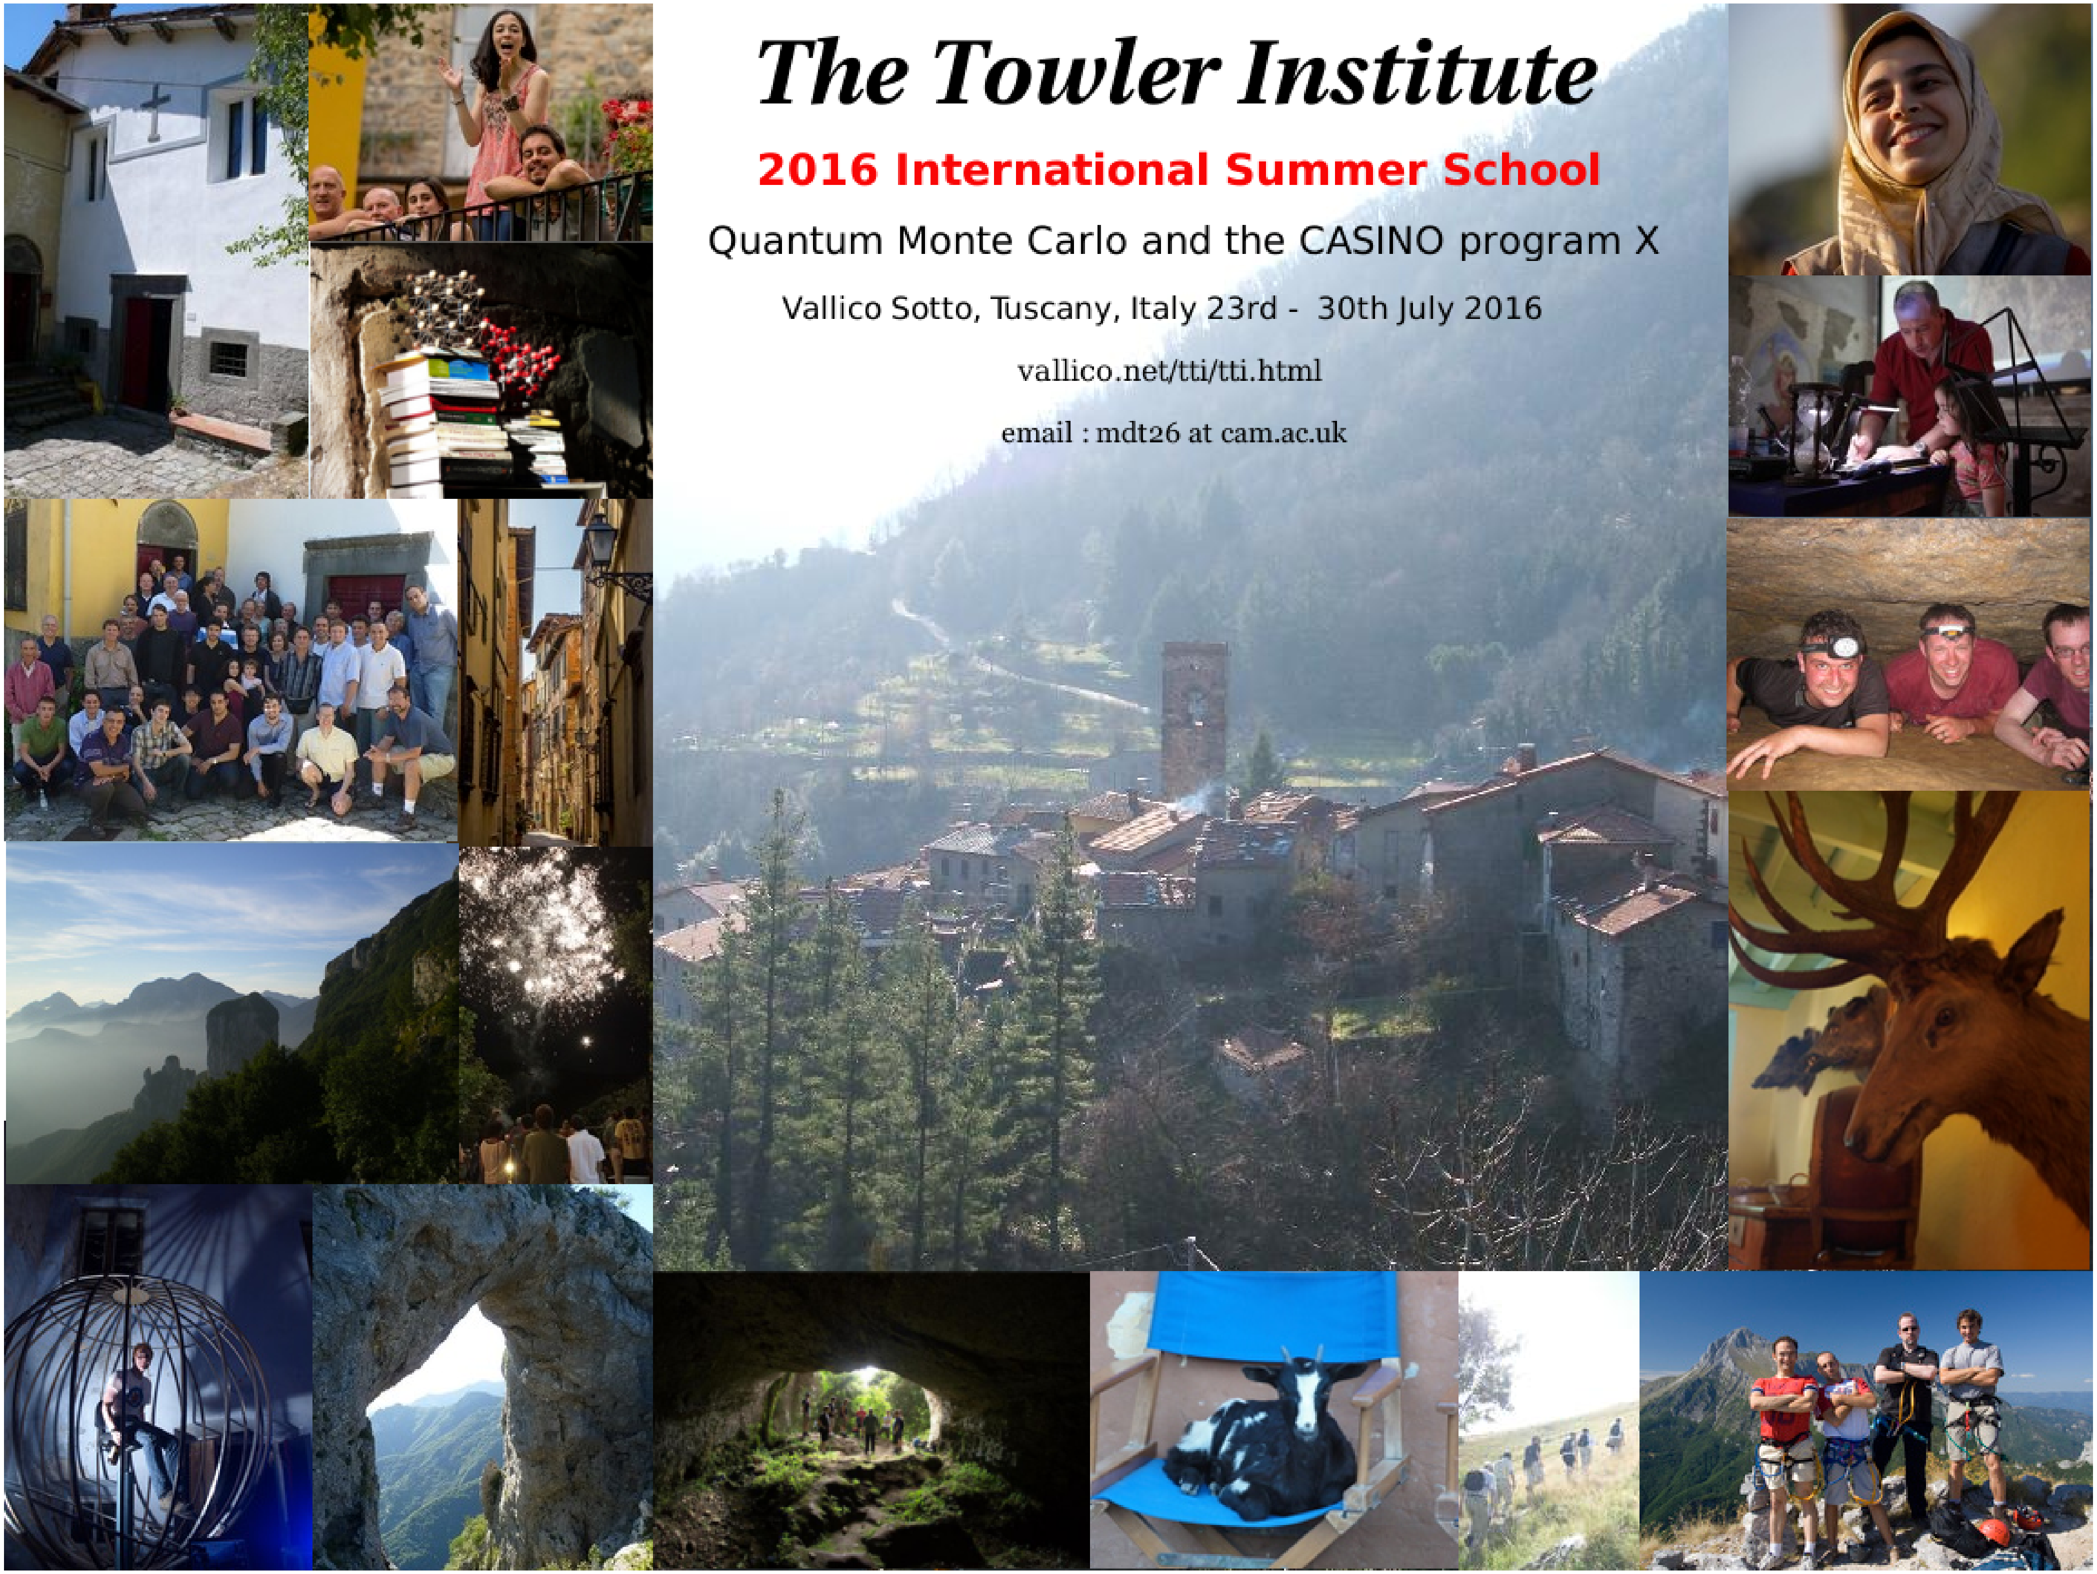

The Towler Institute

2016 International Summer School

Quantum Monte Carlo and the CASINO program X

Vallico Sotto, Tuscany, Italy 23rd - 30th July 2016

vallico.net/tti/tti.html

email : mdt26 at cam.ac.uk

Pseudopotentials for quantum Monte Carlo

A necessary evil

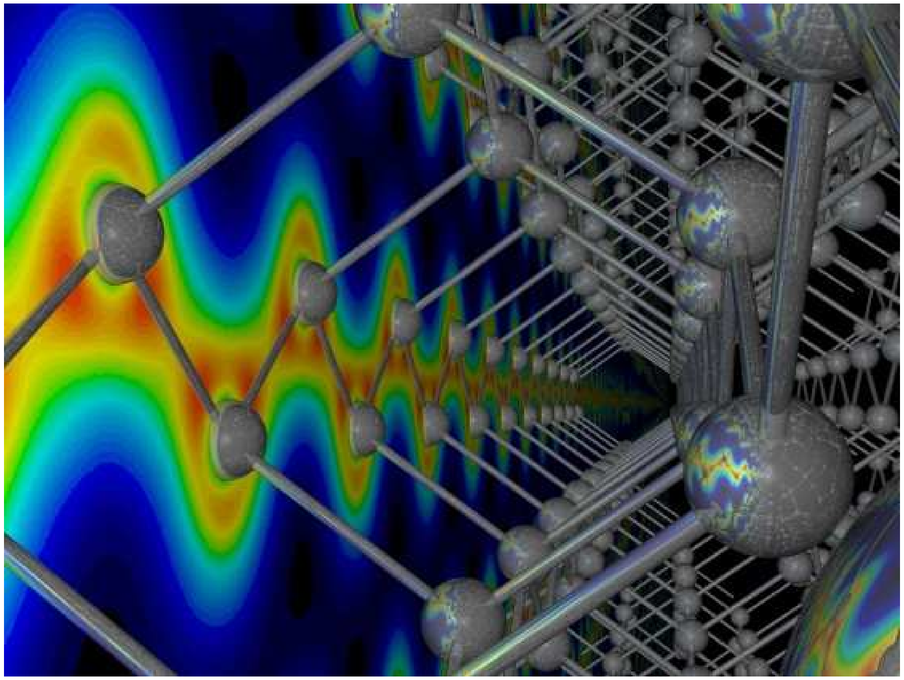

Mike Towler

TCM Group, Cavendish Laboratory, University of Cambridge

QMC web page: `vallico.net/casinoqmc`

Email: `mdt26@cam.ac.uk`

The need for pseudopotentials

"*A pseudopotential is an attempt to replace the complicated effects of the motion of the core (i.e. non-valence) electrons of an atom and its nucleus with an effective potential, or pseudopotential, so that the Schrödinger equation contains a modified effective potential term instead of the Coulombic potential term for core electrons normally found in the Schrödinger equation. The pseudopotential approximation was first introduced by Hans Hellmann in 1934.*"

Used in our context, e.g. in plane-wave DFT calculations, the pseudo-wavefunctions can be described with far fewer Fourier modes and the use of plane-wave basis sets is thus significantly more practical.

What about QMC?

QMC scales badly with atomic number - at best roughly $\propto Z^{4.5}$ (up to $Z^{6.5}$ if you want to remove any residual timestep error from individual total energies). In practice *all-electron* calculations only really feasible for systems containing first-row atoms (and even then only using Gaussian basis sets).

Scaling due to:

- More electrons = more time!
- Step size must be smaller than minimum length scale of wave function ('Bohr radius' $\sim 1/Z$) for good DMC acceptance ratios; otherwise get significant serial correlation. More steps required to target error bar = more time.
- Larger fluctuations in $e-n$ potential energy lead to larger variance for higher Z.

Scaling improved by:

- Removing core electrons and replacing $e-n$ Coulomb interaction with the effective core-valence electron interaction of a decent pseudopotential.

Constructing pseudopotentials: one-electron theory

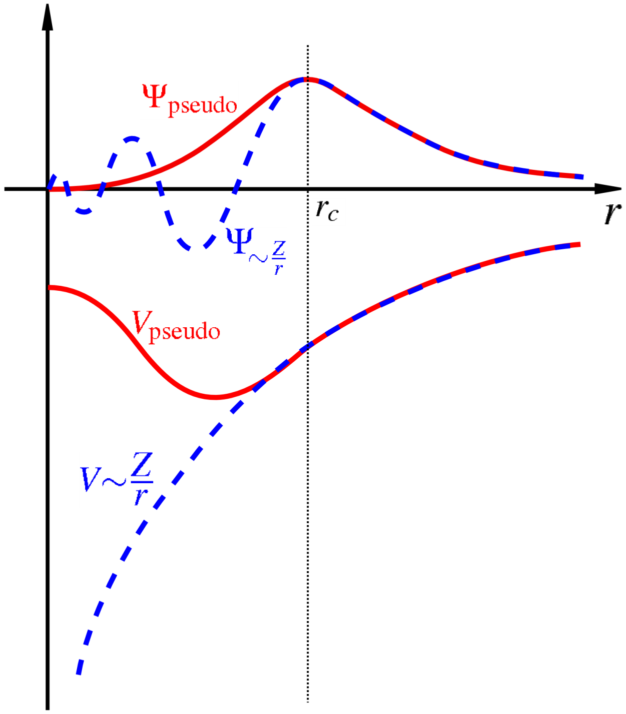

- Electrons deep in energy classified as core electrons; those shallow in energy classified as valence electrons.
- Core electrons considered inert i.e. their (large) contribution to total binding energy doesn't change when atoms brought together to form molecule or crystal. They are removed and the valence electrons feel a compensating pseudopotential.
- Pseudopotential different for each valence orbital (non-local pseudopotential depends on the angular momentum).
- Good pseudopotentials should:
  → be reasonably smooth.
  → behave as $-(Z - N_c)/r$ far from nucleus, where N_c is the number of removed core electrons.
  → have same eigenvalues as the all-electron orbitals.
  → have same orbitals as all-electron orbitals for large r.

Notes:

- '*local*' means that the potential depends only on one space position (like the Coulomb interaction).
- Space non-locality means that it depends on $\mathbf{r}$ and $\mathbf{r}'$, but this can be given a simplified picture.
- '*Semi-local*' means the pseudopotential is expressed in terms of angular momentum dependence and not on $\mathbf{r}$ and $\mathbf{r}'$ separately. A truly non-local or '*separable*' pseudopotential is defined to be non-local both in angular momentum and space coordinates.

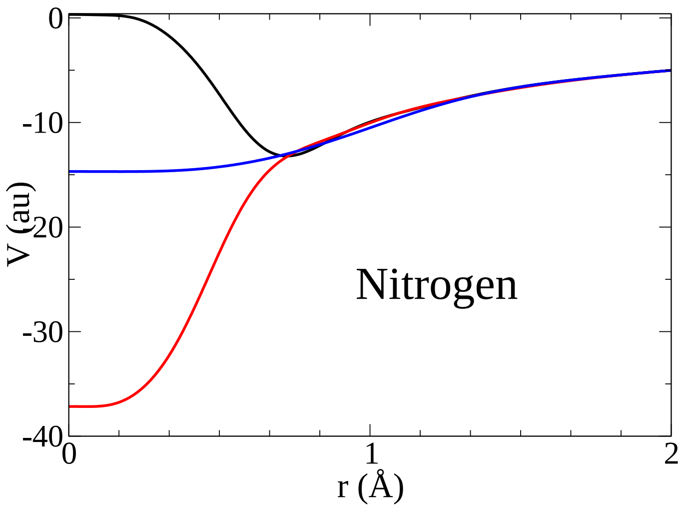

Constructing pseudopotentials: one-electron theory

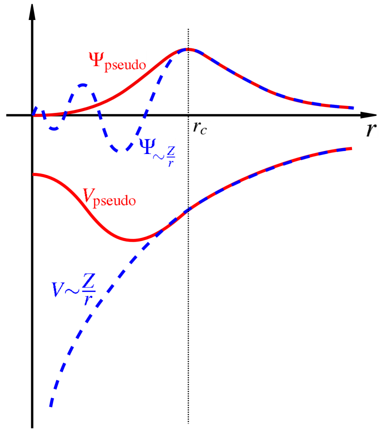

Example recipe:

- Do an all-electron atom Hartree-Fock calculation.
- Ignore 'core' orbitals.
- Construct a pseudo-orbital that is different inside r_c.
- Invert the Hartree-Fock equations.
- Inversion gives pseudopotential whose ground state is the pseudo-orbital.

Details:

- Desirable to reproduce scattering properties of atom (requires *norm conservation*, i.e. we enforce condition that, inside the cut-off radius, norm of each pseudo-wavefunction identical to its corresponding all-electron wavefunction so they produce the same core charge).
- More than one potential: $V_s \neq V_p \neq V_d$
- Project out s, p, and d parts with a projector operator $\hat{P}_l$:

$$\begin{aligned}\hat{V}^{pp}\phi &= \sum_l V_l(r)\left[\sum_m \int Y^*_{lm}(\Omega')\phi(r,\Omega')d\Omega'\right]Y_{lm}(\Omega)\\ &= \sum_l V_l(r)\hat{P}_l\phi\end{aligned}$$

- Redo with V_d as local: $\hat{V}^{pp}\phi = V_d(r) + \sum_l \left[V_l(r) - V_d(r)\right]\hat{P}_l\phi$
- $\hat{V}^{pp}$ is a one-body potential and is a non-local potential.

Constructing pseudopotentials: one-electron theory

Example: silicon atom one-electron(LDA-DFT) energy levels are

nl	Occ.	Eigenvalue (Ha)	
3p	2	-0.153526025	Valence
3s	2	-0.398313865	
2p	6	-3.514381690	Core
2s	2	-5.074463805	
1s	2	-65.184556915	

Pseudopotentials in QMC

Choice of one-electron theory

- Can try DFT or Hartree-Fock pseudopotentials.
- Hartree-Fock seems to give systematically better results in QMC.
  Why? DFT orbitals only represent density and are not components of a many-body wave function.
- Hartree-Fock neglects correlations, but in such a way that:
  1) Valence-valence correlation is done by QMC
  2) Core-core correlation is small and indirect
  3) Core-valence correlation is small
  4) Core-polarization potentials are available
- Hartree-Fock pseudopotentials are the best we have (for now).

Relativistic effects

- We can in fact go beyond this and approximately incorporate relativistic effects using Dirac-Fock theory → 'Dirac-Fock Average Relativistic Effective Potential (AREP)'.
- Solving the Schrödinger equation with Dirac-Fock AREP pseudopotentials will result in the inclusion of scalar relativistic effects (i.e. mass polarization term, mass velocity term, Darwin terms, retardation terms, but not spin-orbit potentials).
- It is good to have HF pseudopotentials available as well since sometimes one wants to compare QMC results with experimental data where relativistic effects have been explicitly subtracted.

Pseudopotentials in VMC: non-local integration

- Must evaluate the expectation value of $\hat{V}_{nl} = V_{loc} + \sum_l \Delta V_l(r_{iI})\hat{P}_l$ where $\hat{P}_l$ projects out the l-th spherical harmonic component of the wave function
- At each $\mathbf{R}$ the quantity $\Psi^{-1}(\mathbf{R}) \sum_i \sum_l \Delta V_l(r_{iI})\hat{P}_l \Psi(\mathbf{R})$ is computed
- Integration of Ψ over the surface of sphere $r_{iI}\equiv$constant for **each** electron

Non-local integration:

- Integrate over sphere surface using *quadrature grids* defined using standard published 'rules'. Number of points in grid set by 'rule number' parameter **non_local_grid** in CASINO input file (generally) or in pseudopotential file (for individual atom types). Integration error decreases with increasing **non_local_grid** value.

```
+------------------------------------------------------+
| NON_LOCAL_GRID    Exactly integrates l=...   No. points |
+------------------------------------------------------+
|       1                    0                    1    |
|       2                    2                    4    |
|       3                    3                    6    |
|       4                    5                   12    |
|       5                    5                   18    |
|       6                    7                   26    |
|       7                   11                   50    |
+------------------------------------------------------+
```

Number of grid points increases roughly as square of rule number, and each rule integrates an expansion in spherical harmonics Y_{lm} up to $l = l_{max}$ exactly, where l_{max} increases approximately linearly with rule number.

- In a QMC calculation the spherical grid is typically rotated randomly before each numerical integration, so that the numerical integrals are unbiased random estimates of the angular-momentum components of the trial wave function.
- Using integration scheme with *small* number of grid points reduces the cost of each local-energy evaluation in a QMC calculation; on the other hand using a smaller number of grid points increases the random error in the integration, and hence the standard error in the mean energy. It is therefore expected that there is an optimal numerical integration rule.

Optimal choice of non-local integration grid

Two examples of the effects of using different non-local integration rules:

TABLE III. VMC total energy, variance per electron, and efficiency with different nonlocal pseudopotential integration rules for an isolated O_2 molecule. The efficiency is defined as the reciprocal of the product of the average walltime per iteration, the variance, and the mean decorrelation period. The calculations used a plane-wave cutoff energy of 120 Ha and a Slater-Jastrow trial wave function with electron-electron (u), electron-nucleus (χ), and electron-electron-nucleus (f) terms, which was optimized by unreweighted variance minimization.

Rule	VMC en. (Ha)	Var. per e. (Ha^2)	Effic. ($Ha^{-2}s^{-1}$)
1	−31.74(1)	86.4	7
2	−31.765(5)	12.3	49
3	−31.7646(5)	0.077	5050
4	−31.7649(4)	0.059	4510
5	−31.7649(4)	0.059	3660
6	−31.7649(4)	0.059	2520
7	−31.7649(4)	0.059	1770

TABLE IV. As Table III but for a 3 × 3 simulation supercell of monolayer hexagonal indium selenide subject to periodic boundary conditions.

Rule	VMC en. (Ha)	Var. per e. (Ha^2)	Effic. ($Ha^{-2}s^{-1}$)
1	−22.852(9)	1.03	0.09
2	−22.870(1)	0.0152	3.91
3	−22.869(1)	0.0120	2.80
4	−22.870(1)	0.0115	1.95
5	−22.870(1)	0.0115	2.12
6	−22.870(1)	0.0115	1.62
7	−22.870(1)	0.0115	0.94

For the oxygen molecule the most efficient integration rule is Rule 3, which uses six points and would be exact for an expansion in spherical harmonics up to $l_{max} = 3$. For indium selenide the optimal rule is Rule 2. However, since the efficiency falls off very much more steeply when the integration rule is too small than when it is too large, we recommend that Rule 4 be chosen to be the default. Rule 4 uses 12 points and would be exact for an expansion in spherical harmonics up to $l_{max} = 5$. Nevertheless, QMC practitioners should be aware of the possibility of substantially increasing the efficiency of their calculations by choosing Rule 3 instead of Rule 4.

[Note: don't use Rule 5 - gives same results as Rule 4 but is more expensive]

Pseudopotentials in QMC: implementation in DMC

- Non-local operators problematic in DMC: must evaluate $\hat{V}_{nl}\Phi$, but Φ is unknown, and the standard DMC algorithm assumes the potential energy operator to be local.
- Pseudopotential localization approximation (PLA): $\Phi^{-1}\hat{V}_{nl}\Phi \approx \Psi_T^{-1}\hat{V}_{nl}\Psi_T$
- PLA effective potential is many-body and local.
- Error incurred is proportional to square of error in trial wave function Ψ_T. However, sign of error is arbitrary, undermining the variational principle for the fixed-node DMC ground-state energy.
- Divergences in the localized pseudopotential due to nodes in the trial wave function can result in instabilities in the DMC algorithm.
- The latter two problems can largely be removed by means of a partial locality approximation known as the “T-move” scheme, although this may increase the absolute magnitude of the pseudopotential errors.

```
24x64:brutus% casinohelp use_tmove
 CASINO HELP SYSTEM
 ==================

 DESCRIPTION
 -----------
 If USE_TMOVE is T then the Casula nonlocal pseudopotential scheme will be used in DMC. So-called 'T-moves' will be
 performed in order to give a DMC energy that is greater than or equal to the ground-state energy. This violates the
 detailed-balance principle at finite time steps, but greatly improves the stability of the DMC algorithm when nonlocal
 pseudopotentials are used. The advantages of T-moves are that they restore the variational principle and help to
 prevent population explosions; the disadvantages of T-moves are that the magnitude of the error due to the locality
 approximation is generally larger, although always positive, and the time-step bias is generally worse. [This latter
 problem is alleviated, to some extent, by using a symmetric branching factor (Casula 2010) as opposed to the
 asymmetric one suggested in his 2006 paper. This advice was implemented in CASINO in June 2014.]. A further
 disadvantage is that this option requires a truly enormous amount of memory in systems with large numbers of particles
 (seeing if this can be reduced remains a project).
```

Trail-Needs pseudopotentials

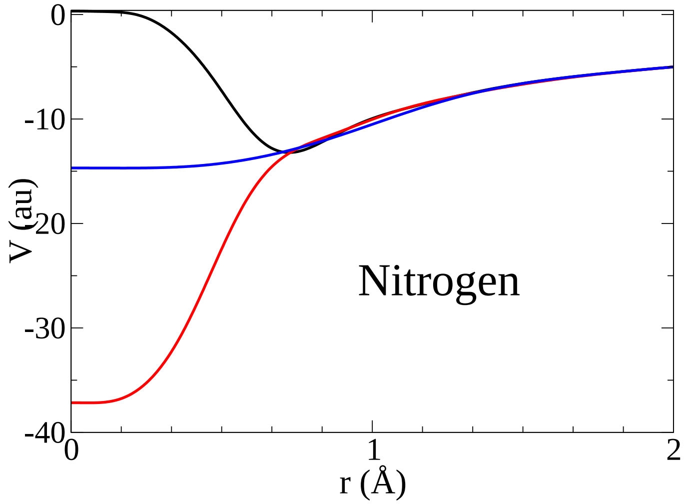

Developed by John Trail and Richard Needs in our group in Cambridge.

- The TN pseudopotentials are Dirac-Fock average relativistic effective pseudopotentials optimized for QMC calculations.
- The core is as large as possible in each case, and pseudopotential data are provided for the s, p, and d angular-momentum channels only.
- The issue of the need for higher-angular momentum channels (f,g,..) in some cases was discussed by Tipton, Drummond, and Hennig in *Phys. Rev. B* **90**, 125110 (2014).
- Either s, p, or d must be chosen to be the local channel, the potential for which is then applied to all higher angular-momentum components of the wave function. By default the d channel is chosen to be local in the TN pseudopotentials (but see further discussion later).

Where to get TN pseudopotentials for CASINO

The CASINO pseudopotential library

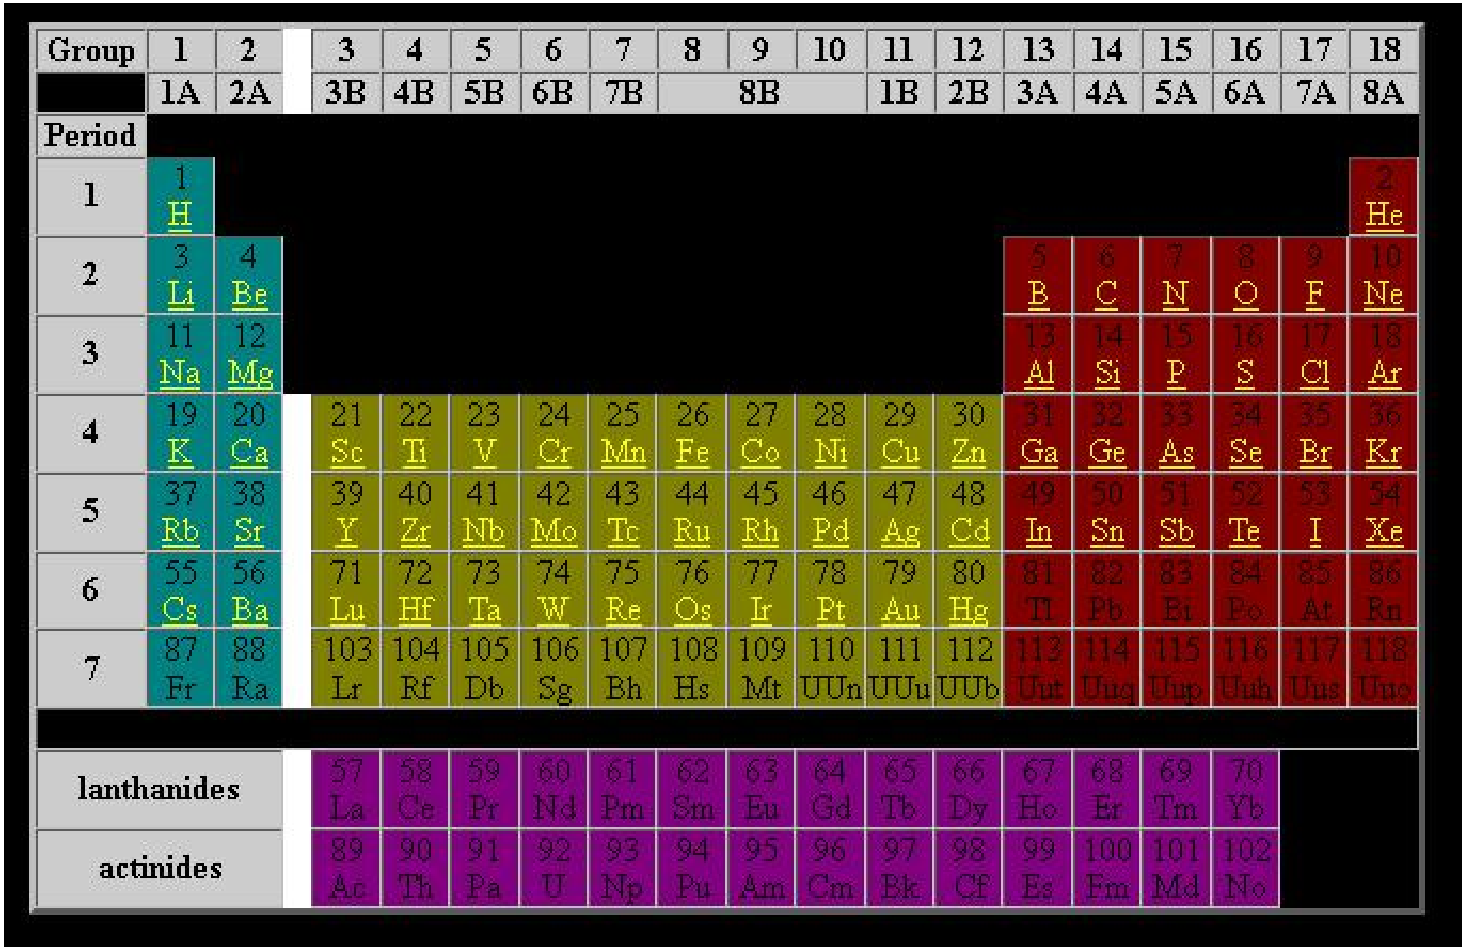

Group	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
	1A	2A	3B	4B	5B	6B	7B	8B			1B	2B	3A	4A	5A	6A	7A	8A
Period																		
1	1 H																	2 He
2	3 Li	4 Be											5 B	6 C	7 N	8 O	9 F	10 Ne
3	11 Na	12 Mg											13 Al	14 Si	15 P	16 S	17 Cl	18 Ar
4	19 K	20 Ca	21 Sc	22 Ti	23 V	24 Cr	25 Mn	26 Fe	27 Co	28 Ni	29 Cu	30 Zn	31 Ga	32 Ge	33 As	34 Se	35 Br	36 Kr
5	37 Rb	38 Sr	39 Y	40 Zr	41 Nb	42 Mo	43 Tc	44 Ru	45 Rh	46 Pd	47 Ag	48 Cd	49 In	50 Sn	51 Sb	52 Te	53 I	54 Xe
6	55 Cs	56 Ba	71 Lu	72 Hf	73 Ta	74 W	75 Re	76 Os	77 Ir	78 Pt	79 Au	80 Hg	81 Tl	82 Pb	83 Bi	84 Po	85 At	86 Rn
7	87 Fr	88 Ra	103 Lr	104 Rf	105 Db	106 Sg	107 Bh	108 Hs	109 Mt	110 UUn	111 UUu	112 UUb	113 Uut	114 Uuq	115 Uup	116 Uuh	117 Uus	118 Uuo
lanthanides			57 La	58 Ce	59 Pr	60 Nd	61 Pm	62 Sm	63 Eu	64 Gd	65 Tb	66 Dy	67 Ho	68 Er	69 Tm	70 Yb		
actinides			89 Ac	90 Th	91 Pa	92 U	93 Np	94 Pu	95 Am	96 Cm	97 Bk	98 Cf	99 Es	100 Fm	101 Md	102 No		

`https://vallico.net/casinoqmc/pplib/`

(*Note colours and design have changed - probably for the worse - on the latest website - but you get the idea..*)

The CASINO pseudopotential library

Click on an element to see..

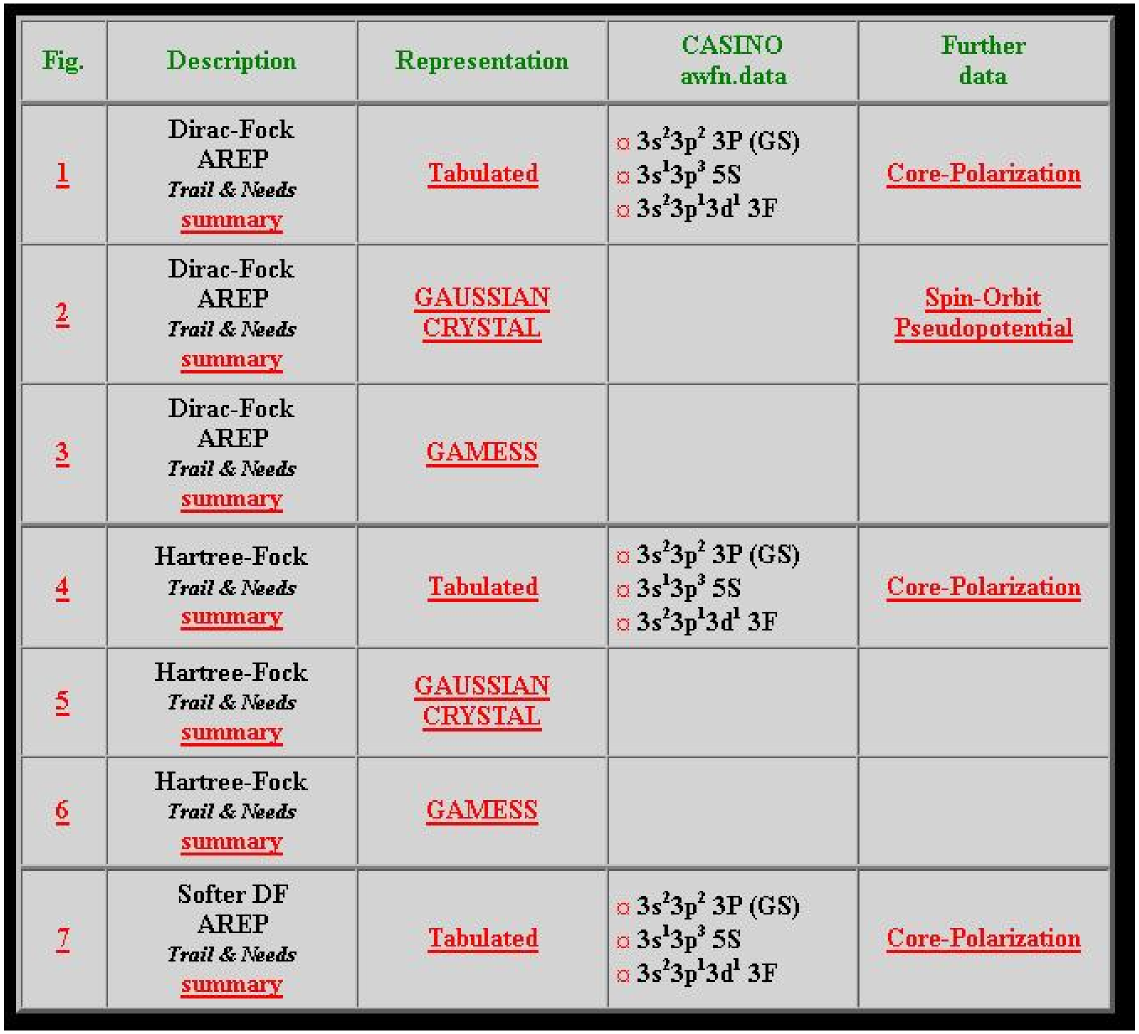

Fig.	Description	Representation	CASINO awfn.data	Further data
1	Dirac-Fock AREP *Trail & Needs* summary	Tabulated	○ $3s^2 3p^2$ 3P (GS) ○ $3s^1 3p^3$ 5S ○ $3s^2 3p^1 3d^1$ 3F	Core-Polarization
2	Dirac-Fock AREP *Trail & Needs* summary	GAUSSIAN CRYSTAL		Spin-Orbit Pseudopotential
3	Dirac-Fock AREP *Trail & Needs* summary	GAMESS		
4	Hartree-Fock *Trail & Needs* summary	Tabulated	○ $3s^2 3p^2$ 3P (GS) ○ $3s^1 3p^3$ 5S ○ $3s^2 3p^1 3d^1$ 3F	Core-Polarization
5	Hartree-Fock *Trail & Needs* summary	GAUSSIAN CRYSTAL		
6	Hartree-Fock *Trail & Needs* summary	GAMESS		
7	Softer DF AREP *Trail & Needs* summary	Tabulated	○ $3s^2 3p^2$ 3P (GS) ○ $3s^1 3p^3$ 5S ○ $3s^2 3p^1 3d^1$ 3F	Core-Polarization

All words underlined in red are links to further data.

- Pseudopotential input files for other codes
- Pseudopotential plots
- Pseudopotential properties (total energies, etc)
- Atomic wave functions in CASINO format
- CPPs and spin-orbit potentials

The CASINO pseudopotential

Choosing a pseudopotential

Flow chart of the seven pseudopotentials

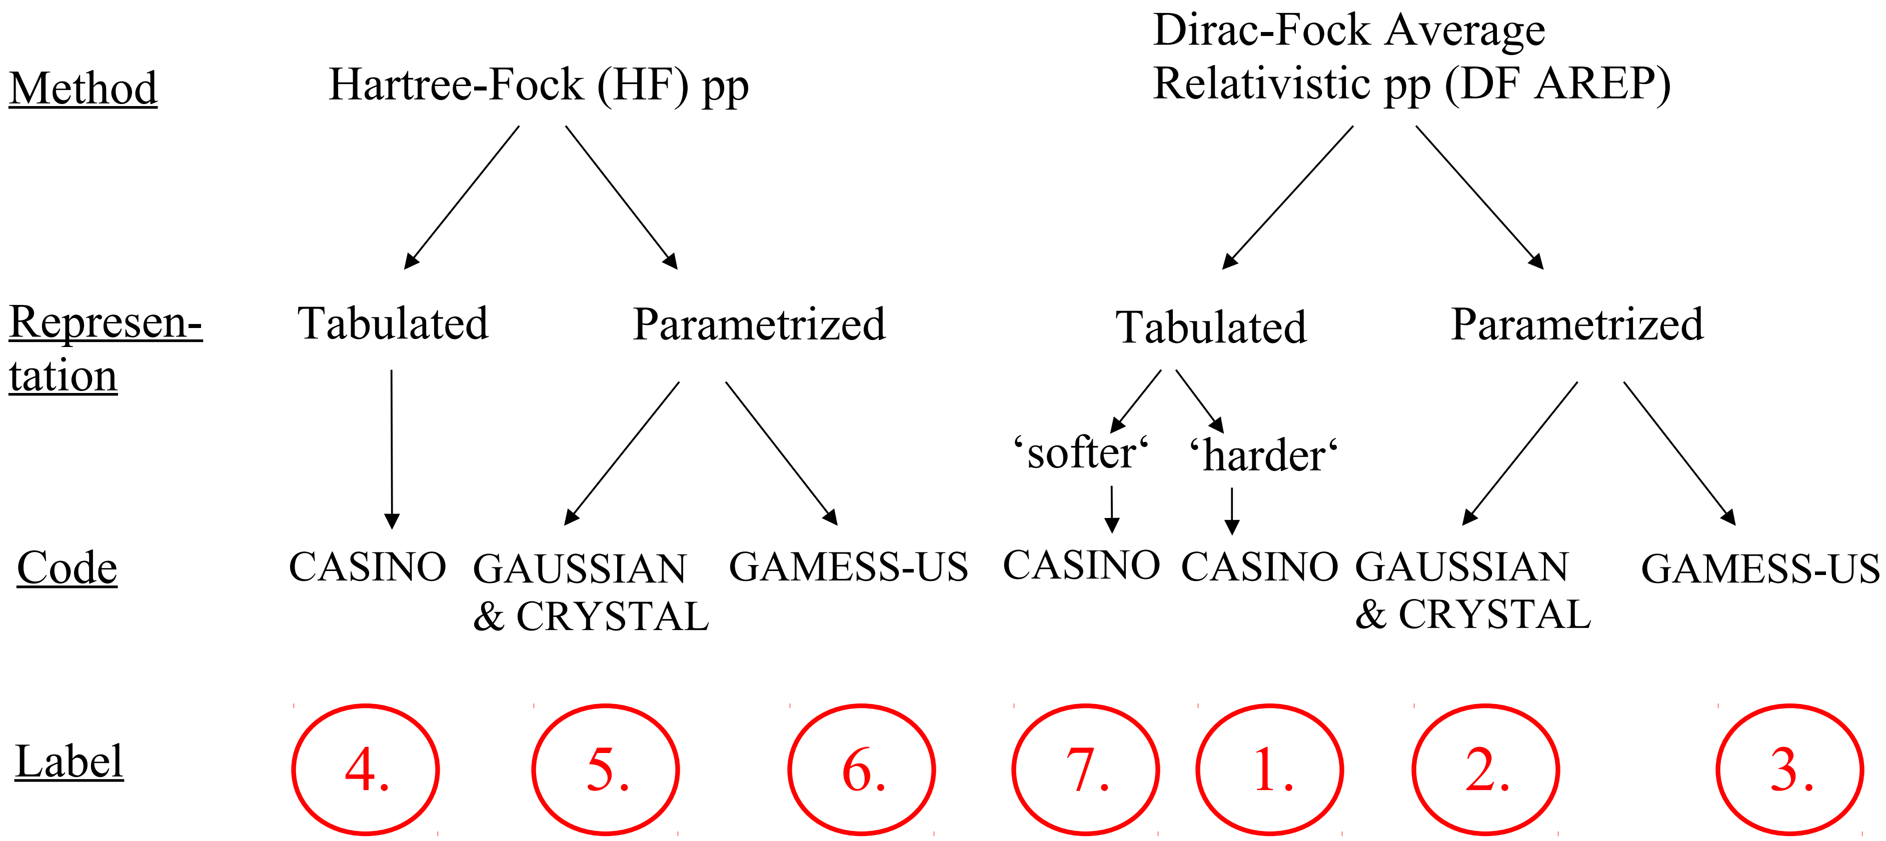

Method

Hartree-Fock (HF) pp

Dirac-Fock Average
Relativistic pp (DF AREP)

Represen-
tation

Tabulated

Parametrized

Tabulated

Parametrized

‘softer‘

‘harder‘

Code

CASINO

GAUSSIAN
& CRYSTAL

GAMESS-US

CASINO

CASINO

GAUSSIAN
& CRYSTAL

GAMESS-US

Label

4.

5.

6.

7.

1.

2.

3.

Pseudopotentials in CASINO

Which is best?

- Pseudopotentials are not unique and there is no “best” pseudopotential.

- User must choose the most appropriate pseudopotential, and this depends on what is being calculated.

Hartree-Fock or Dirac-Fock?

- Hartree-Fock includes no relativistic effects.

- Dirac-Fock includes some relativistic effects.

What results do you compare with?

- Compare with experiment → use DF pseudopotential.

- Compare with non-relativistic DFT → use HF pseudopotential.

Pseudopotentials in CASINO

Tabulated or parameterized

What code do you use to generate the trial wave function?

- Parameterisation is neccessary for many packages.
- GAUSSIAN, CRYSTAL, GAMESS for these packages and more.
- Format conversion routines in `CASINO/utils/pseudo_converters`.
- Consistency - if we use Dirac-Fock for GAMESS, we should use Dirac-Fock pseudopotential for CASINO.

Use tabulated pseudopotentials if possible.

Core polarization potentials (CPP)

CPPs account for the polarization of the pseudo-ion cores by the fields of the other charged particles in the system. The polarization of the pseudo-ion cores by the fields of the valence electrons is a many-body effect which includes some of the core-valence correlation energy.

Derivation

- From electrostatic theory and an approximation
- A core J feels $\mathbf{E}$ due to cores, I, and electrons, i.

$$\mathbf{E} = -\sum_{I \neq J} Z_I \frac{\mathbf{R_I} - \mathbf{R_J}}{|\mathbf{R_I} - \mathbf{R_J}|^3} + \sum_i \frac{\mathbf{r}_i - \mathbf{R_J}}{|\mathbf{r}_i - \mathbf{R_J}|^3}$$

- $\mathbf{E}$ polarizes core J by $\mathbf{P} = \alpha_J \mathbf{E}$
- Polarization energy is $-1/2\alpha_J \mathbf{E}.\mathbf{E}$
- Add up energy of **all** cores, and add to Hamiltonian

$$H_{CPP} = \sum_{iJ} V_e(\mathbf{r}_{iJ}) + \sum_{ijJ} V_{e-e}(\mathbf{r}_{iJ}, \mathbf{r}_{jJ}) + \sum_{iIJ} V_{e-n}(\mathbf{R}_{IJ}, \mathbf{r}_{iJ}) + \sum_{IJ} V_n(\mathbf{R}_{IJ})$$

- Potentials are many-body, local and $\propto 1/\mathrm{distance}^4$

Core polarization potentials (CPP)

Results

Energies (eV) for excitations of a Si atom

	$3s^2 3p^2 \rightarrow 3s^1 3p^3$	$3s^2 3p^1 \rightarrow 3s^1 3p^2$	$3s^2 \rightarrow 3s^1 3p^1$
LDA	3.827(10)	4.994(10)	6.232(16)
HF	3.909(13)	5.096(9)	6.363(15)
HF+CPP	4.052(10)	5.264(9)	6.571(6)
DF	3.955(13)	5.146(9)	6.434(7)
DF+CPP	4.069(9)	5.297(9)	6.578(6)
Exp.	4.11	5.30	6.56

Pseudopotentials in CASINO

Core polarization potentials (CPPs)

- Used in addition to pseudopotentials.
- CPPs can only be included in CASINO and MOLPRO.
- Should make results more accurate.
- BUT largely untested so it's up to you to test them.

Spin-orbit (SO) potentials

- Used in addition to pseudopotentials to describe fine structure.
- Include fine structure effects.
- Should make results more accurate.
- BUT almost no packages use them.

Ignore CPP and SO unless your research is actually about this.

Plane-wave basis and ‘ghost states’

Most QMC calculations for real materials have featured either first- or second-row atoms. An issue that quickly emerges when one attempts to use TN pseudopotentials for transition metals ‘off the shelf’ in plane-wave DFT orbital-generation calculations is the presence of *ghost states* due to the Kleinman-Bylander representation of the pseudopotentials in the DFT code.

The presence of ghost states gives rise to some or all of the following symptoms:

- the failure of the DFT self-consistent-field (SCF) process to converge
- a large difference between the DFT energies obtained with plane-wave and Gaussian basis sets
- the existence of an absurdly low Kohn-Sham eigenvalue
- an absurdly high (unbound) energy when the orbitals are used in VMC calculations
- a very large energy variance
- enormous difficulty optimizing a trial wave function in VMC
- enormous difficulty controlling the configuration population in a DMC simulation.

Furthermore, these difficulties may change or disappear when the local channel is changed.

TN pseudopotentials and ghost states

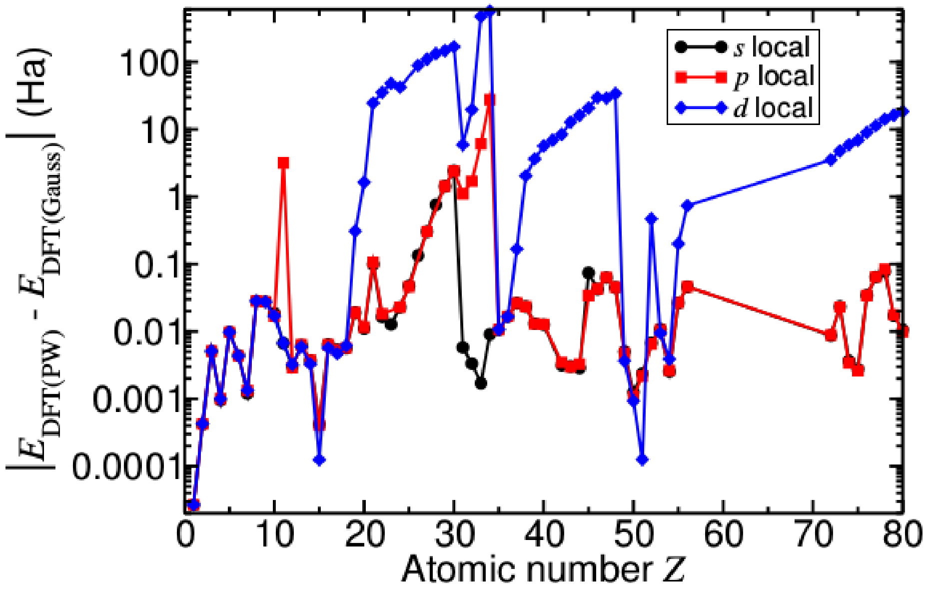

Here the difference between the DFT energy obtained using a plane-wave basis with a Kleinman-Bylander representation of the pseudopotential and the DFT energy obtained using a Gaussian basis for each TN pseudopotential.

- Choosing the d channel to be local often leads to a relatively enormous difference between the plane-wave and Gaussian DFT results, strongly suggesting a problem caused by a ghost state.
- Choosing the s channel to be local avoids this problem in *every* case apart from niobium.
- Choosing the p channel to be local avoids the problem with ghost states in most cases.

TN pseudopotentials and ghost states

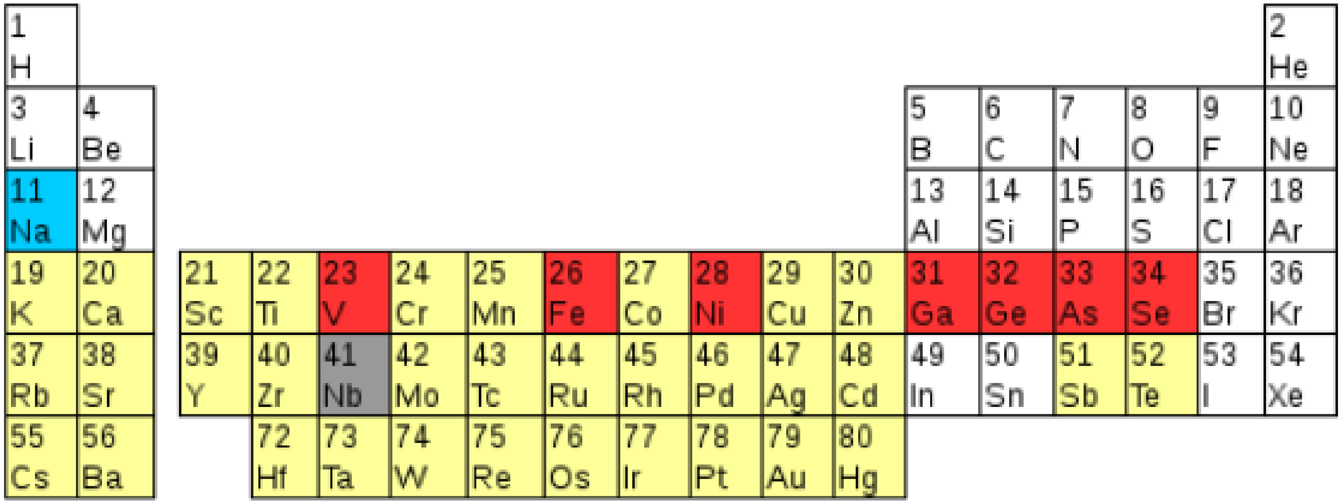

FIG. 2. (Color online) TN pseudopotentials that are affected by ghost states for different choices of local channel in plane-wave DFT calculations. For elements in white cells, there are no ghost states; for elements in blue cells, choosing p (but not s or d) to be local results in ghost states; for elements in yellow cells, choosing d (but not s or p) to be local results in ghost states; for elements in red cells, choosing p or d (but not s) to be local leads to ghost states; and elements in grey cells are haunted by ghost states for any choice of local channel. Gaussian DFT results are not available for lutetium ($Z = 71$), which is therefore omitted.

Shows which TN pseudopotentials are likely to be affected by ghost states for different choices of local channel, based on the analysis of the DFT energies on the previous slide.

The presence of ghost states makes QMC work meaningless or impossible; however, inexperienced users may wrongly ascribe the problems encountered to the general difficulty of optimizing QMC trial wave functions. Eliminating ghost states from DFT orbital-generation calculations is a necessary but not sufficient condition for accurate QMC work. Even if the orbitals generated in the DFT calculation are unaffected by ghost states, the choice of local channel may still affect the behavior of the subsequent QMC calculations.

TN pseudopotentials and ghost states

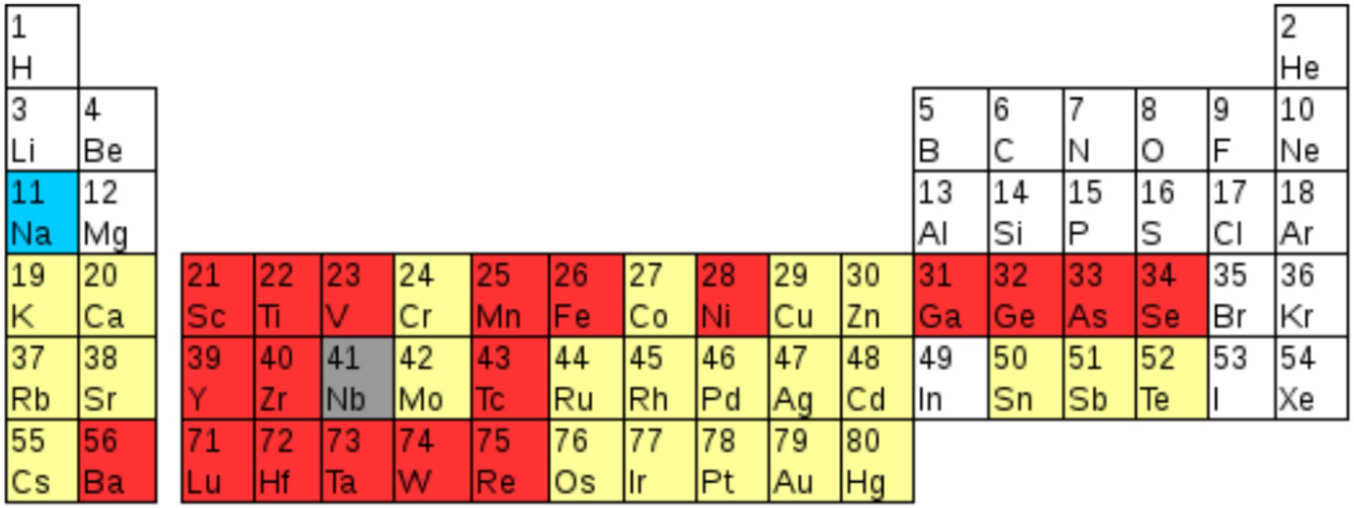

1 H																	2 He
3 Li	4 Be											5 B	6 C	7 N	8 O	9 F	10 Ne
🟦 11 Na	12 Mg											13 Al	14 Si	15 P	16 S	17 Cl	18 Ar
🟨 19 K	🟨 20 Ca	🟥 21 Sc	🟥 22 Ti	🟥 23 V	🟨 24 Cr	🟥 25 Mn	🟥 26 Fe	🟨 27 Co	🟥 28 Ni	🟨 29 Cu	🟨 30 Zn	🟥 31 Ga	🟥 32 Ge	🟥 33 As	🟥 34 Se	35 Br	36 Kr
🟨 37 Rb	🟨 38 Sr	🟥 39 Y	🟥 40 Zr	▨ 41 Nb	🟨 42 Mo	🟥 43 Tc	🟨 44 Ru	🟨 45 Rh	🟨 46 Pd	🟨 47 Ag	🟨 48 Cd	49 In	🟨 50 Sn	🟨 51 Sb	🟨 52 Te	53 I	54 Xe
🟨 55 Cs	🟥 56 Ba	🟥 71 Lu	🟥 72 Hf	🟥 73 Ta	🟥 74 W	🟥 75 Re	🟨 76 Os	🟨 77 Ir	🟨 78 Pt	🟨 79 Au	🟨 80 Hg						

FIG. 3. (Color online) TN pseudopotentials that are affected by ghost states or similar difficulties for different choices of local channel in QMC calculations using plane-wave DFT orbitals. For elements in white cells, there are no ghost-state-like symptoms; for elements in blue cells, choosing p (but not s or d) to be local results in ghost-state-like symptoms; for elements in yellow cells, choosing d (but not s or p) to be local results in ghost-state-like symptoms; for elements in red cells, choosing p or d (but not s) to be local leads to ghost-state-like symptoms; and elements in grey cells are affected by ghost-state-like symptoms for any choice of local channel.

This shows which TN pseudopotentials are affected by ghost-state-like symptoms in QMC calculations using a plane-wave/blip basis for different choices of local channel (used in both the DFT and the subsequent QMC calculations). The criteria used for judging that a particular pseudopotential with a particular choice of local channel is problematic are the symptoms listed three slides back.

- Only one TN pseudopotential, niobium, is apparently completely unusable in plane-wave calculations. The problem with niobium is not currently understood.
- In every other case, the problem of ghost-state-like symptoms can be avoided by choosing the s channel to be local in the plane-wave DFT calculation.

What should I use as a local channel?

- In QMC calculations that use orbitals generated with the s channel chosen to be local, the local channel can either be left as s or (preferably) changed to d; no symptoms of ghost states occur in either case.
- By contrast, there are several elements for which choosing the p channel to be local does not cause any problems in the DFT calculation but does adversely affect the subsequent QMC calculations (see previous figures).
- Applying the potential for the d channel to higher angular-momentum components in the QMC calculation is expected to be more accurate in principle than applying the potential for the s channel; furthermore, it can be shown that choosing d to be local in QMC reduces the achievable variance in many cases.
- Experience with plane-wave DFT calculations for $3d$ transition metals suggests that choosing p to be local provides reasonable energies, but is very vulnerable to unstable convergence and/or apparent convergence to different final energies for different initial conditions. A likely explanation for this behavior is a ghost state of similar energy to the actual state. This may well be the cause of the difference between the elements highlighted as being problematic when p is local in the previous figures.
- For the transition metal pseudopotentials the s and d characters of the wave functions are expected to be dominant near the nuclei, with the p character being either minimal or zero. This suggests that taking s rather than p to be local in DFT calculations will be the most accurate choice, as it avoids an approximate projector representation for this channel.

Plane wave basis convergence

Is it important to use highly converged plane-wave orbitals in QMC calculations, i.e., is it crucial to use a tight tolerance for self-consistency in the orbital-generation DFT calculations?

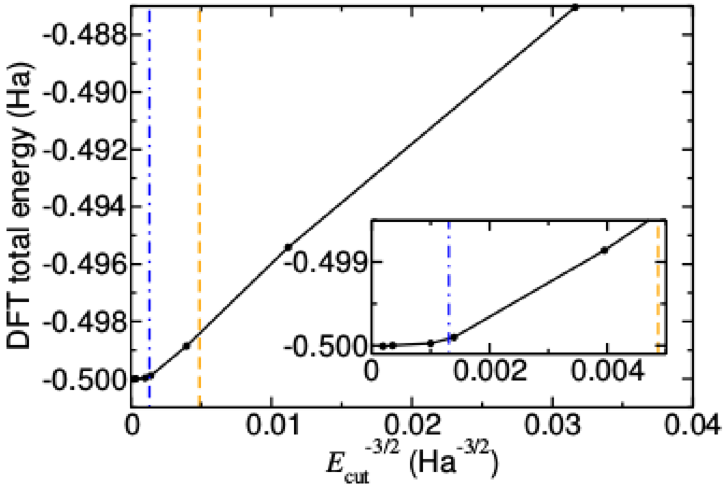

FIG. 4. (Color online) DFT total energy against plane-wave cutoff energy E_{cut} for an isolated hydrogen pseudoatom. The inset shows the higher-cutoff region in finer detail. The dashed and dash-dotted vertical lines show the cutoff energies at which the DFT total energy is converged to within chemical accuracy and to within 0.1 mHa, respectively.

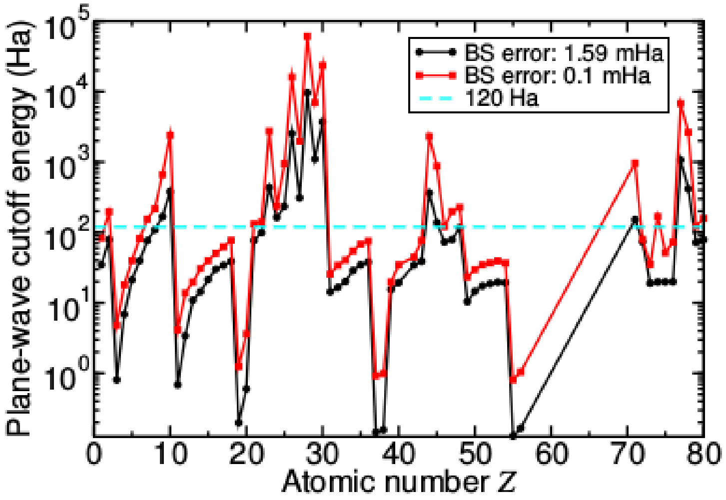

FIG. 5. (Color online) Plane-wave cutoff energy required to achieve a given level of convergence in the DFT energy per atom for each TN pseudopotential. The s channel is chosen to be local in each case. 120 Ha is the cutoff energy used for the tests reported in Fig. 8.

- This is plot of the plane-wave cutoff energy required to achieve a given level of convergence in the DFT total energy against atomic number.
- The cutoff energy required to converge the total energy of each atom to within so-called chemical accuracy (1 kcal/mol or 1.59 mHa) is shown, as is the cutoff energy required to converge the total energy to 0.1 mHa, an order of magnitude tighter than chemical accuracy.
- Interesting observation: the required cutoff energies for the *3d transition metals* are in many cases impractically large; hence any attempt to perform plane-wave-DFT-QMC calculations using the TN pseudopotentials for those atoms will inevitably encounter problems associated with large variances, and the outcome will at best rely on a cancellation of errors.

Plane-wave basis convergence

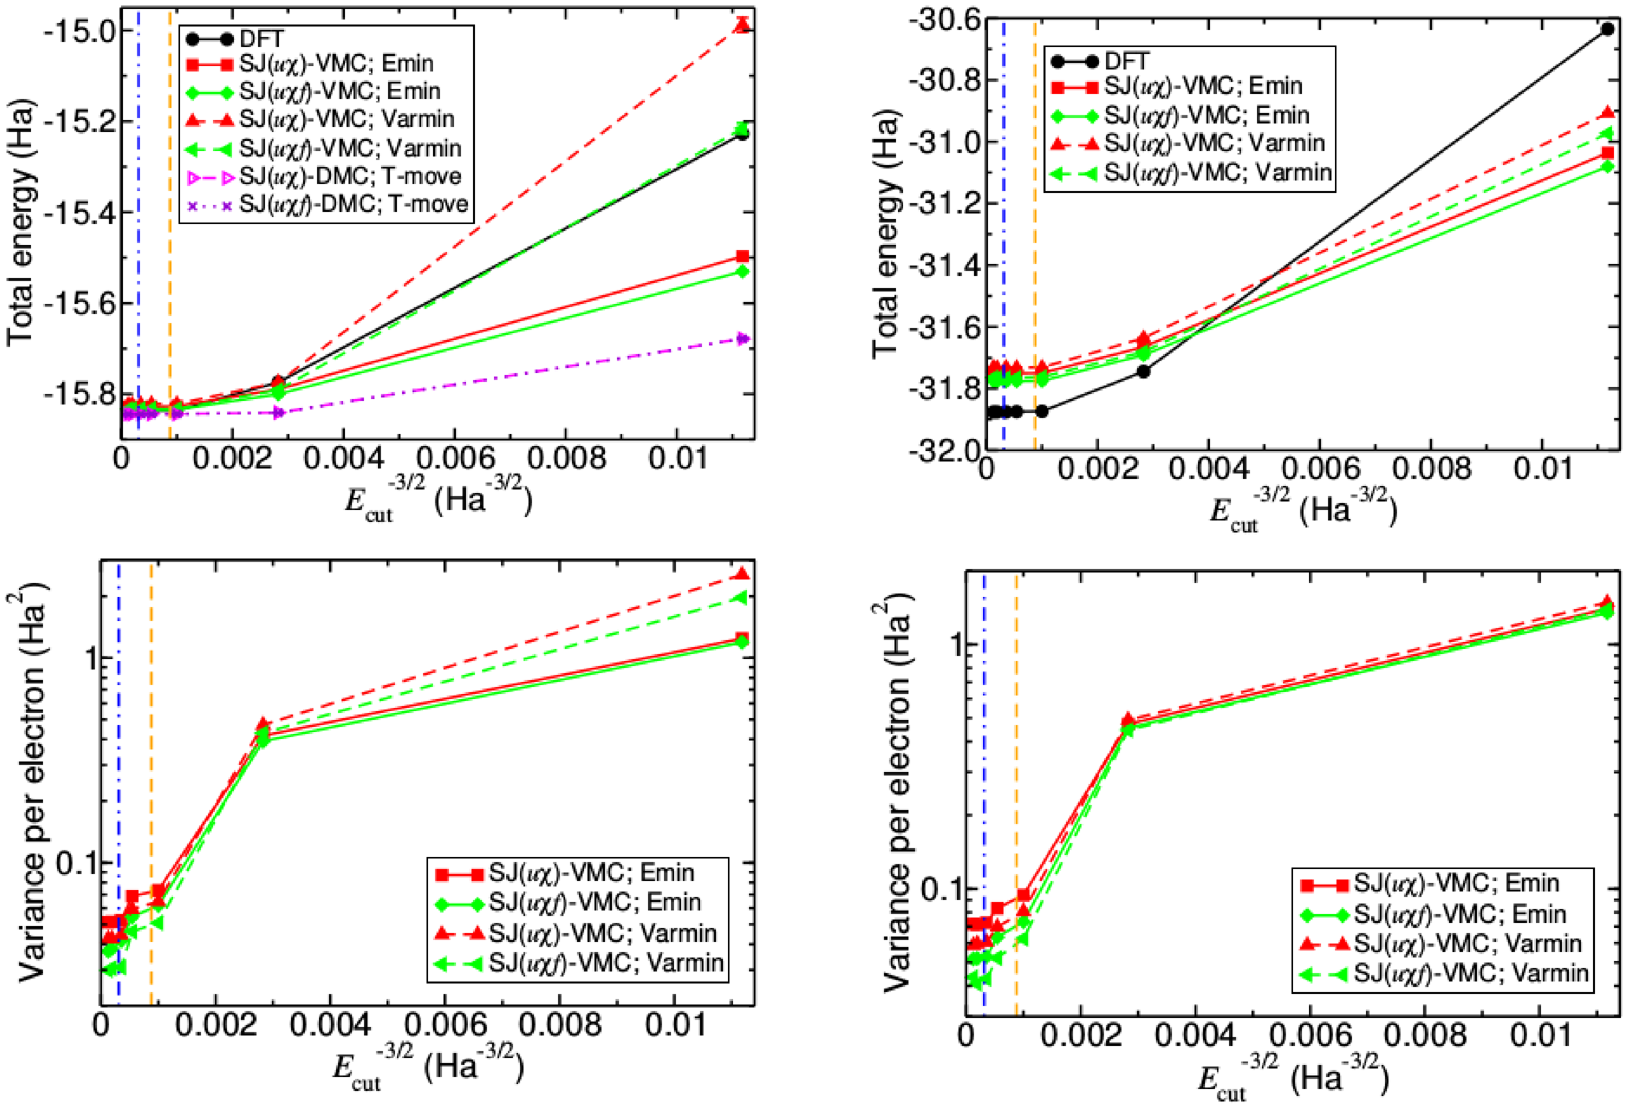

Isolated oxygen atom (left) and an oxygen molecule (right). Top panels: DFT, VMC, and DMC energies as a function of plane-wave cutoff energy E_{cut} Bottom panel: VMC energy variance against plane-wave cutoff energy. Two different levels of correlated wave function were used in the VMC calculations, and linear-least-squares energy minimization ('Emin') and unreweighted variance minimization ('Varmin') were used to optimize the wave functions. The dashed and dash-dotted vertical lines show the cutoff energies at which the DFT total energy is converged to within chemical accuracy and to within 0.1 mHa, respectively. The DMC calculations used the T-move scheme.

Efficiency considerations

- The dependence of the DMC energy on the orbitals in the Slater wave function is in general very weak, because the DMC energy only depends on the trial wave function via the fixed-node approximation and the pseudopotential locality approximation.
- However, the VMC energy, the VMC energy variance, and hence the efficiency of QMC calculations can depend significantly on the orbitals.
- It has been shown that, for a given wave-function form, the efficiency of the importance-sampled DMC algorithm is maximized when the trial wave function is optimized by energy minimization.
- Reducing the finite-basis error in the DFT total energy per atom provides a better starting point for optimization of the correlated part of the trial wave function, and the reduction in the DFT energy with increasing basis-set size translates directly into a reduction in the VMC energy, as shown for an oxygen atom and oxygen molecule on next slide.
- Since QMC calculations are generally intended to achieve chemical accuracy or higher, it is desirable for the finite-basis error in the DFT energy to be substantially less than chemical accuracy.
- Considerable reductions in the VMC variance can be achieved by increasing the plane-wave cutoff energy up to the value suggested by the convergence of the DFT total energy to chemical accuracy. The performance of unreweighted variance minimization improves significantly once the basis set becomes adequate.
- The results on the previous slide show that, apart from the lowest plane wave cutoff energy studied, the DMC energy is almost independent of the cutoff energy.

Further observations

- Comparing results for the oxygen atom and the oxygen molecule it is clear that there is a significant cancellation of finite-basis errors in the VMC binding energy of an oxygen molecule, suggesting that one could ‘get away with’ relatively low plane-wave cutoff energies.
- However, it is more efficient to use a high plane-wave cutoff energy (such that the DFT total energy is converged to 0.1 mHa), because it leads to an enormous reduction in the variance of the energies of the atom and the molecule, and hence gives smaller statistical error bars in the binding energy.
- When the orbitals are represented in a blip basis in QMC calculations, the cost of the calculation is only weakly dependent on the plane-wave cutoff energy. The cost of the DFT orbital-generation calculation clearly depends on the cutoff energy, but orbital generation is usually a negligibly small fraction of the computational expense of a QMC project.

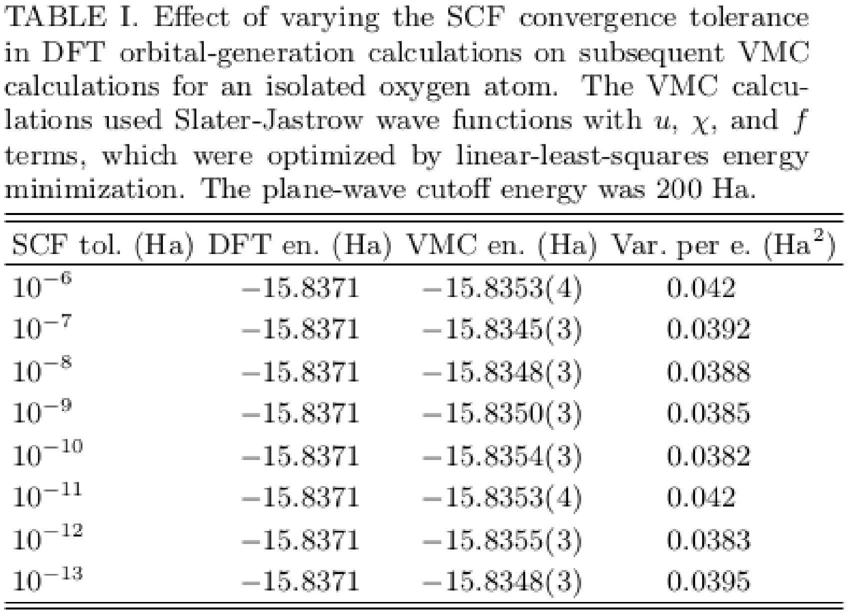

TABLE I. Effect of varying the SCF convergence tolerance in DFT orbital-generation calculations on subsequent VMC calculations for an isolated oxygen atom. The VMC calculations used Slater-Jastrow wave functions with u, χ, and f terms, which were optimized by linear-least-squares energy minimization. The plane-wave cutoff energy was 200 Ha.

SCF tol. (Ha)	DFT en. (Ha)	VMC en. (Ha)	Var. per e. (Ha^2)
10^{-6}	-15.8371	$-15.8353(4)$	0.042
10^{-7}	-15.8371	$-15.8345(3)$	0.0392
10^{-8}	-15.8371	$-15.8348(3)$	0.0388
10^{-9}	-15.8371	$-15.8350(3)$	0.0385
10^{-10}	-15.8371	$-15.8354(3)$	0.0382
10^{-11}	-15.8371	$-15.8353(4)$	0.042
10^{-12}	-15.8371	$-15.8355(3)$	0.0383
10^{-13}	-15.8371	$-15.8348(3)$	0.0395

No evidence of any need for tight SCF tolerances. That said, since DFT calculations are generally negligibly cheap compared with DMC calculations, there is no good reason for not using a tight convergence criterion.

Achievable energy variance with TN pseudopotentials

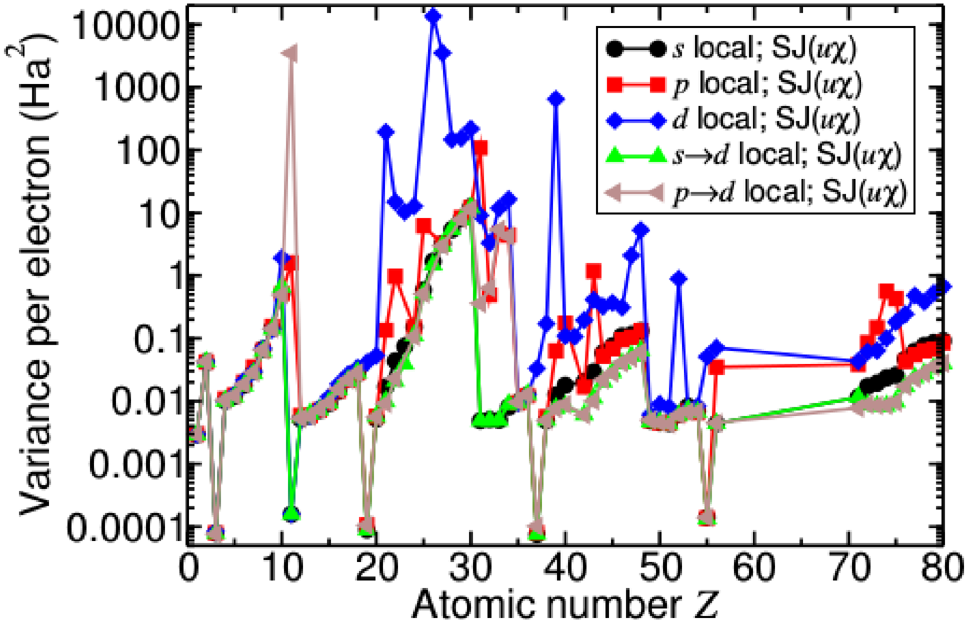

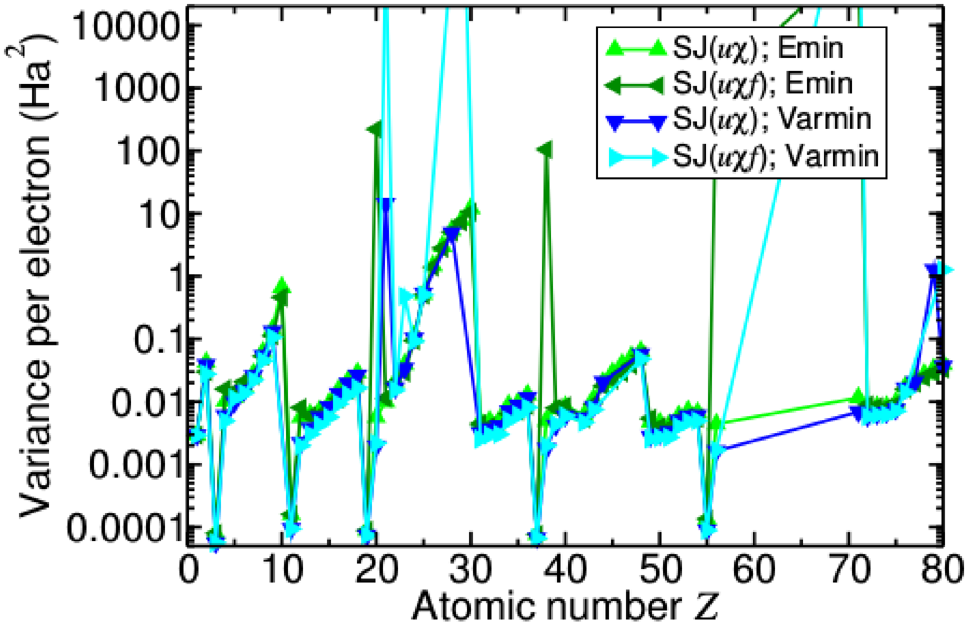

- Here is plotted the VMC energy variance achieved for each element using a fixed plane-wave cutoff energy of 120 Ha. Lower variances can nearly always be achieved by setting s as local channel during DFT orbital-generation calculation.
- Variance either unchanged or lowered further if d channel subsequently chosen to be local in the QMC calculation.
- Again clear that the most difficult cases by far are the 3d transition metals, where the variance is orders of magnitude larger than for other elements.
- As results obtained with fixed 120 Ha cutoff energy this is partly due to cutoff being too small. However, for real calculations, it becomes impractical to use the cutoffs required for these elements.
- The lower panel shows that, although unreweighted variance minimization generally gives a lower variance than linear-least-squares energy minimization, there are more cases in which variance minimization catastrophically fails.

Worst case scenario case study: the copper atom..

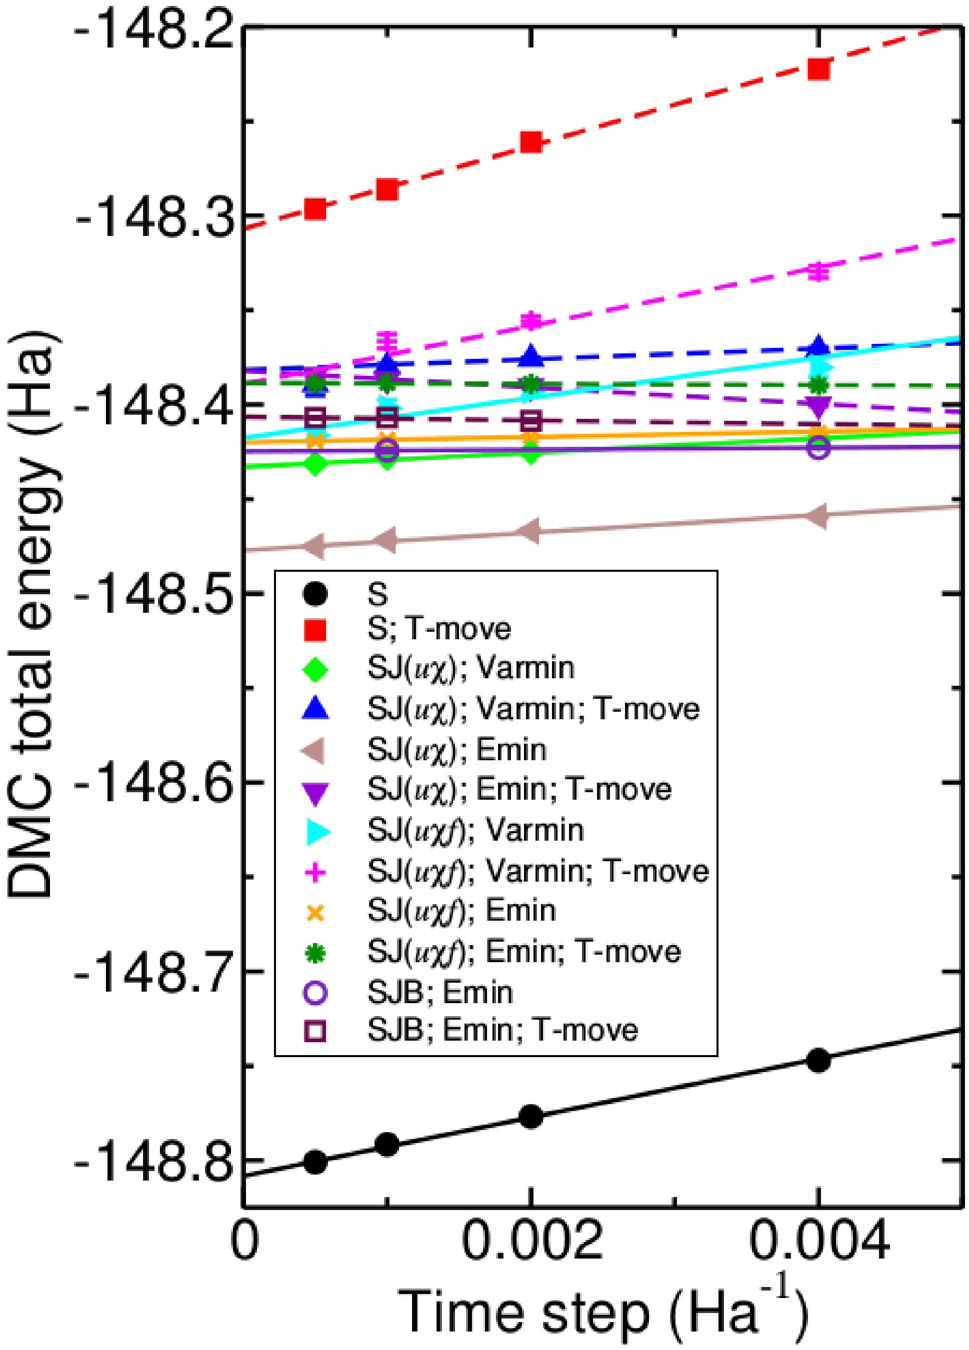

- Clear that the choice of trial wave function and optimization method affects not only the behavior at finite time step, but also the final DMC energies extrapolated to zero time step.
- In all-electron fixed-node DMC calculations, the DMC energies at zero time step obtained using Slater and Slater-Jastrow wave functions are identical, because the nodal surface is not affected by the Jastrow factor.
- However, for copper, pseudopotential locality errors lead to differences on the scale of several eV between the DMC energies obtained with different Jastrow factors and without a Jastrow factor.
- This problem is significantly ameliorated by the use of the T-move scheme.

The Cu atom continued..

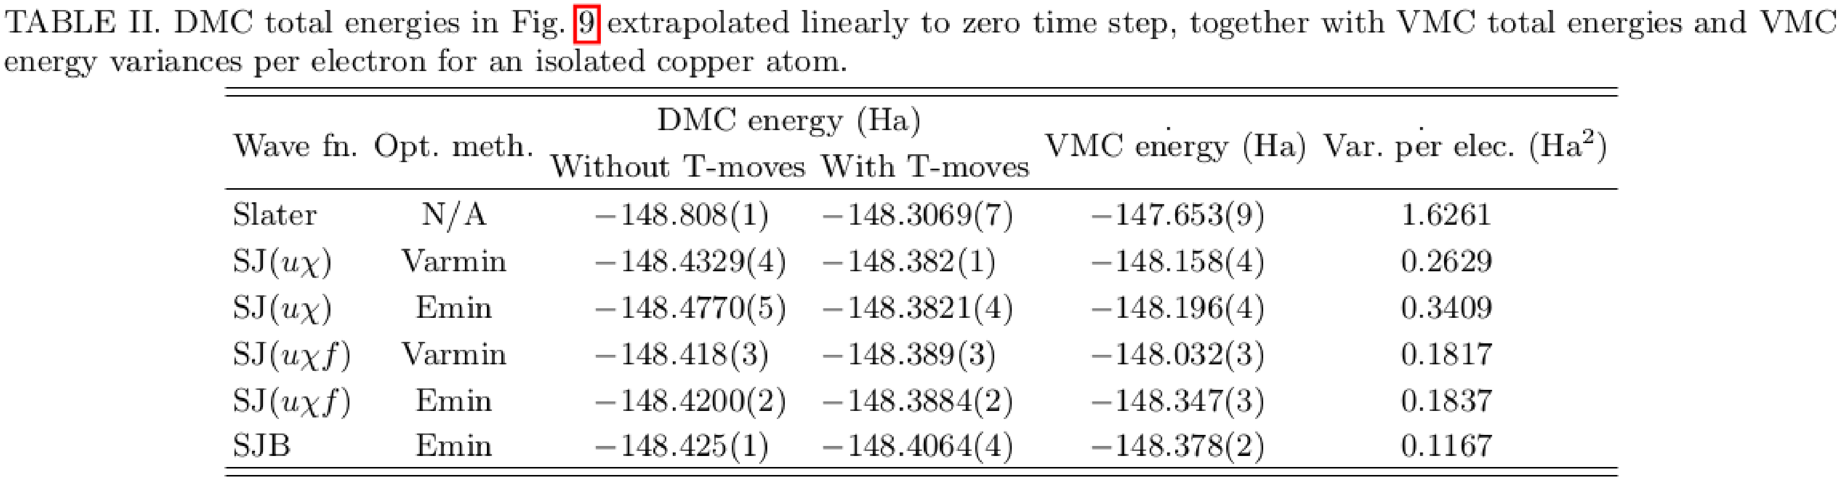

TABLE II. DMC total energies in Fig. 9 extrapolated linearly to zero time step, together with VMC total energies and VMC energy variances per electron for an isolated copper atom.

Wave fn.	Opt. meth.	DMC energy (Ha) Without T-moves	DMC energy (Ha) With T-moves	VMC energy (Ha)	Var. per elec. (Ha^2)
Slater	N/A	$-148.808(1)$	$-148.3069(7)$	$-147.653(9)$	1.6261
SJ($u\chi$)	Varmin	$-148.4329(4)$	$-148.382(1)$	$-148.158(4)$	0.2629
SJ($u\chi$)	Emin	$-148.4770(5)$	$-148.3821(4)$	$-148.196(4)$	0.3409
SJ($u\chi f$)	Varmin	$-148.418(3)$	$-148.389(3)$	$-148.032(3)$	0.1817
SJ($u\chi f$)	Emin	$-148.4200(2)$	$-148.3884(2)$	$-148.347(3)$	0.1837
SJB	Emin	$-148.425(1)$	$-148.4064(4)$	$-148.378(2)$	0.1167

- Table shows the DMC energies extrapolated to zero time step and infinite population, together with the corresponding VMC energies and variances.
- Without T-moves the DMC energies depend significantly on the trial wave function and can be nonvariational.
- The spread of DMC energies is significantly reduced by the T- move scheme, and in particular the spuriously low energies obtained with poorer wave functions are eliminated.
- The change in the T-move DMC energy resulting from the inclusion of f terms in the Jastrow factor is small compared with the difference in energy resulting from the inclusion of backflow.
- By contrast, including f terms has a much larger effect on VMC energies than including backflow.
- This suggests that, if one has a Slater-Jastrow wave function with u, χ, and f terms, locality errors are small compared with fixed-node errors. When the wave function is optimized by energy minimization, the standard error in the DMC energy is, as expected, significantly lower in general (this is always the case when T-moves are used).

Conclusions for this case..

The copper atom is an extreme case that highlights a number of important issues:

- Pseudopotential locality errors in the DMC total energy can be as large as several eV per atom, and they manifest themselves by significant dependence of the DMC energy on the Jastrow factor;
- pseudopotential locality errors can be greatly ameliorated by the use of the T-move scheme
- if the plane-wave cutoff energy is smaller than ideal then energy minimization appears to be more reliable than unreweighted variance minimization and also results in greater efficiency in subsequent DMC calculations;
- including extra correlation in the trial wave function using, e.g., electron-electron-nucleus (f) terms reduces both locality errors and time-step bias.

Strategies that prevent numerical problems with copper can be expected to work for all the other transition metals, because copper has the deepest d-channel.

For some elements..

..we need better pseudopotentials!

Constructing pseudopotentials: many-body theory

- Core and valence are ill-defined concepts
- Different attempts to derive many-body pseudopotentials:
  - Approximate $\Psi = \hat{A}\{\Psi_{\mathrm{core}}\Psi_{\mathrm{valence}}\}$ → no core-valence correlation, not much faster.
  - Generalise norm-conservation condition to many-body case → current research.
  - Pseudo-hamiltonians → inaccurate when transferred.
  - Use VMC for 'core' and DMC for 'valence' → does not change the scaling with atomic number.

Thus up to now usually considered best to construct pseudopotentials from one-particle theories.

But! John Trail and Richard Needs in Cambridge are working on '*correlated electron pseudopotentials* which are now sufficiently well developed to be published:

J.R. Trail and R.J. Needs, *J. Chem. Phys.* **139**, 014101 (2013)
J.R. Trail and R.J. Needs, *J. Chem. Phys.* **142**, 064110 (2013)

The latter was published with 6 correlated-electron pseudopotentials for Sc-Fe.

John email to me 15/7/2016 says: "*A new type (yes, yet another one) of pseudopotential should be available soon (test data + paper are in preparation). These improve the accuracy of the correlated-electron pseudopotentials further by combining correlated-electron 'norm-conservation' and energy consistency. These will be available for 1st row, and the 3d-transitions Sc-Fe and Cu.*"

The CASINO website will be reengineered soon to incorporate these new pseudopotentials. For now, no-one other than John has ever had access to them - hopefully soon!

Conclusions

Positive:

- Pseudopotentials reduce number of electrons.
- Pseudopotential make everything smoother
  $\longrightarrow$ more samples so higher accuracy.

Negative:

- Uncontrolled but small error.
- One must choose carefully and check.

References:

J.R. Trail and R.J. Needs, *J. Chem. Phys.* **122**, 174109 (2005)
J.R. Trail and R.J. Needs, *J. Chem. Phys.* **122**, 014112 (2005)
E.L. Shirley and R.M. Martin, *Phys. Rev. B* **47**, 15413 (1993)
N.D. Drummond, J.R. Trail, and R.J. Needs, `http://arxiv.org/abs/1607.01544`