

NWChem User Documentation

Release 4.6

**High Performance Computational Chemistry Group
W.R. Wiley Environmental Molecular Sciences Laboratory
Pacific Northwest National Laboratory
P.O. Box 999, Richland, WA 99352**

May 2004

DISCLAIMER

This material was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the United States Department of Energy, nor Battelle, nor any of their employees, MAKES ANY WARRANTY, EXPRESS OR IMPLIED, OR ASSUMES ANY LEGAL LIABILITY OR RESPONSIBILITY FOR THE ACCURACY, COMPLETENESS, OR USEFULNESS OF ANY INFORMATION, APPARATUS, PRODUCT, SOFTWARE, OR PROCESS DISCLOSED, OR REPRESENTS THAT ITS USE WOULD NOT INFRINGE PRIVATELY OWNED RIGHTS.

LIMITED USE

This software (including any documentation) is being made available to you for your internal use only, solely for use in performance of work directly for the U.S. Federal Government or work under contracts with the U.S. Department of Energy or other U.S. Federal Government agencies. This software is a version which has not yet been evaluated and cleared for commercialization. Adherence to this notice may be necessary for the author, Battelle Memorial Institute, to successfully assert copyright in and commercialize this software. This software is not intended for duplication or distribution to third parties without the permission of the Manager of Software Products at Pacific Northwest National Laboratory, Richland, Washington, 99352.

ACKNOWLEDGMENT

This software and its documentation were produced with Government support under Contract Number DE-AC06-76RLO-1830 awarded by the United States Department of Energy. The Government retains a paid-up non-exclusive, irrevocable worldwide license to reproduce, prepare derivative works, perform publicly and display publicly by or for the Government, including the right to distribute to other Government contractors.

AUTHOR DISCLAIMER

This software contains proprietary information of the authors, Pacific Northwest National Laboratory (PNNL), and the US Department of Energy (USDOE). The information herein shall not be disclosed to others, and shall not be reproduced whole or in part, without written permission from PNNL or USDOE. The information contained in this document is provided "AS IS" without guarantee of accuracy. Use of this software is prohibited without **written** permission from PNNL or USDOE. The authors, PNNL, and USDOE make no representations or warranties whatsoever with respect to this software, including the implied warranty of merchant-ability or fitness for a particular purpose. The user assumes all risks, including consequential loss or damage, in respect to the use of the software. In addition, PNNL and the authors shall not be obligated to correct or maintain the program, or notify the user community of modifications or updates that will be made over the course of time.

Contents

1	Introduction	19
1.1	Citation	19
1.2	User Feedback	20
2	Getting Started	23
2.1	Input File Structure	23
2.2	Simple Input File — SCF geometry optimization	24
2.3	Water Molecule Sample Input File	25
2.4	Input Format and Syntax for Directives	27
2.4.1	Input Format	27
2.4.2	Format and syntax of directives	28
3	NWChem Architecture	31
3.1	Database Structure	31
3.2	Persistence of data and restart	33
4	Functionality	35
4.1	Molecular electronic structure	35
4.2	Relativistic effects	37
4.3	Pseudopotential plane-wave electronic structure	37
4.4	Molecular dynamics	38
4.5	Python	38
4.6	Parallel tools and libraries (ParSoft)	39
5	Top-level directives	41
5.1	START and RESTART — Start-up mode	41
5.2	SCRATCH_DIR and PERMANENT_DIR — File directories	43

5.3	MEMORY — Control of memory limits	44
5.4	ECHO — Print input file	45
5.5	TITLE — Specify job title	46
5.6	PRINT and NOPRINT — Print control	46
5.7	SET — Enter data in the RTDB	47
5.8	UNSET — Delete data in the RTDB	49
5.9	STOP — Terminate processing	49
5.10	TASK — Perform a task	49
5.10.1	TASK Directive for Electronic Structure Calculations	50
5.10.2	TASK Directive for Special Operations	51
5.10.3	TASK Directive for the Bourne Shell	52
5.10.4	TASK Directive for QM/MM simulations	53
5.11	CHARGE — Total system charge	53
5.12	ECCE_PRINT — Print information for Ecce	54
6	Geometries	57
6.1	Keywords on the GEOMETRY directive	58
6.2	SYMMETRY — Symmetry Group Input	60
6.3	Cartesian coordinate input	61
6.4	ZMATRIX — Z-matrix input	62
6.5	ZCOORD — Forcing internal coordinates	67
6.6	Applying constraints in geometry optimizations	68
6.7	SYSTEM — Lattice parameters for periodic systems	69
7	Basis sets	71
7.1	Basis set library	73
7.2	Explicit basis set definition	74
7.3	Combinations of library and explicit basis set input	75
8	Effective Core Potentials	77
8.1	Scalar ECPs	78
8.2	Spin-orbit ECPs	80
9	Relativistic All-electron Approximations	83
9.1	Douglas-Kroll approximation	84
9.2	Dyall's Modified Dirac Hamiltonian approximation	85

10 Hartree-Fock or Self-consistent Field	89
10.1 Wavefunction type	89
10.2 SYM — use of symmetry	90
10.3 ADAPT — symmetry adaptation of MOs	90
10.4 TOL2E — integral screening threshold	91
10.5 VECTORS — input/output of MO vectors	91
10.5.1 Superposition of fragment molecular orbitals	94
10.5.2 Atomic guess orbitals with charged atoms	99
10.6 Accuracy of initial guess	100
10.7 THRESH — convergence threshold	100
10.8 MAXITER — iteration limit	100
10.9 PROFILE — performance profile	101
10.10 DIIS — DIIS convergence	101
10.11 DIRECT and SEMIDIRECT — recomputation of integrals	101
10.11.1 Integral File Size and Format for the SCF Module	103
10.12 SCF Convergence Control Options	104
10.13 NR — controlling the Newton-Raphson	105
10.14 LEVEL — level-shifting the orbital Hessian	105
10.15 Orbital Localization	106
10.16 Printing Information from the SCF Module	108
10.17 Hartree-Fock or SCF, MCSCF and MP2 Gradients	109
11 DFT for Molecules (DFT)	111
11.1 Specification of Basis Sets for the DFT Module	113
11.2 VECTORS and MAX_OVL — KS-MO Vectors	114
11.3 XC and DECOMP — Exchange-Correlation Potentials	114
11.3.1 Exchange-Correlation Functionals	115
11.3.2 Combined Exchange and Correlation Functionals	115
11.3.3 Meta-GGA Functionals	116
11.4 LB94 and CS00 — Asymptotic correction	116
11.5 Sample input file	117
11.6 ITERATIONS — Number of SCF iterations	120
11.7 CONVERGENCE — SCF Convergence Control	120
11.8 SMEAR — Fractional Occupation of the Molecular Orbitals	122
11.9 GRID — Numerical Integration of the XC Potential	122

11.9.1	Angular grids	125
11.9.2	Partitioning functions	127
11.9.3	Radial grids	127
11.9.4	Disk usage for Grid	127
11.10	TOLERANCES — Screening tolerances	127
11.11	DIRECT and NOIO — Hardware Resource Control	128
11.12	ODFT and MULT — Open shell systems	129
11.13	SIC — Self-Interaction Correction	129
11.14	MULLIKEN — Mulliken analysis	129
11.15	BSSE — Basis Set Superposition Error	130
11.16	Print Control	130
12	Spin-Orbit DFT (SODFT)	133
13	COSMO	137
14	CIS, TDHF, and TDDFT	141
14.1	Overview	141
14.2	Performance of CIS, TDHF, and TDDFT methods	142
14.3	Input syntax	143
14.4	Keywords of TDDFT input block	144
14.4.1	CIS and RPA — the Tamm–Dancoff approximation	144
14.4.2	NROOTS — the number of excited states	144
14.4.3	MAXVECS — the subspace size	144
14.4.4	SINGLET and NOSINGLET — singlet excited states	144
14.4.5	TRIPLET and NOTRIPLET — triplet excited states	145
14.4.6	THRESH — the convergence threshold of Davidson iteration	145
14.4.7	MAXITER — the maximum number of Davidson iteration	145
14.4.8	TARGET and TARGETSYM— the target root and its symmetry	145
14.4.9	SYMMETRY — restricting the excited state symmetry	145
14.4.10	ALGORITHM — algorithms for tensor contractions	145
14.4.11	FREEZE — the frozen core/virtual approximation	146
14.4.12	PRINT — the verbosity	146
14.5	Sample input	146
15	Tensor Contraction Engine Module: CI, MBPT, and CC	151

15.1	Overview	151
15.2	Performance of CI, MBPT, and CC methods	152
15.3	Algorithms of CI, MBPT, and CC methods	153
15.3.1	Spin, spatial, and index permutation symmetry	153
15.3.2	Runtime orbital range tiling	153
15.3.3	Dynamic load balancing parallelism	153
15.3.4	Parallel I/O schemes	153
15.4	Input syntax	153
15.5	Keywords of TCE input block	156
15.5.1	HF, SCF, or DFT — the reference wave function	156
15.5.2	CCSD,CCSDT,CCSDTQ,CISD,CISDT,CISDTQ, MBPT2,MBPT3,MBPT4, etc. — the correlation model	158
15.5.3	THRESH — the convergence threshold of iterative solutions of amplitude equations	159
15.5.4	MAXITER — the maximum number of iterations	159
15.5.5	IO — parallel I/O scheme	159
15.5.6	DIIS — the convergence acceleration	160
15.5.7	FREEZE — the frozen core/virtual approximation	160
15.5.8	NROOTS — the number of excited states	160
15.5.9	TARGET and TARGETSYM — the target root and its symmetry	161
15.5.10	SYMMETRY — restricting the excited state symmetry	161
15.5.11	DIPOLE — the ground- and excited-state dipole moments	161
15.5.12	(NO)FOCK — (not) recompute Fock matrix	161
15.5.13	PRINT — the verbosity	161
15.6	Sample input	162
16	MP2	165
16.1	FREEZE — Freezing orbitals	166
16.2	TIGHT — Increased precision	167
16.3	SCRATCHDISK — Limiting I/O usage	167
16.4	PRINT and NOPRINT	167
16.5	VECTORS — MO vectors	167
16.6	RI-MP2 fitting basis	168
16.7	FILE3C — RI-MP2 3-center integral filename	169
16.8	RIAPPROX — RI-MP2 Approximation	169
16.9	Advanced options for RI-MP2	169

16.9.1	Control of linear dependence	169
16.9.2	Reference Spin Mapping for RI-MP2 Calculations	169
16.9.3	Batch Sizes for the RI-MP2 Calculation	170
16.9.4	Energy Memory Allocation Mode: RI-MP2 Calculation	171
16.9.5	Local Memory Usage in Three-Center Transformation	171
16.10	One-electron properties and natural orbitals	171
17	Multiconfiguration SCF	173
17.1	ACTIVE — Number of active orbitals	173
17.2	ACTELEC — Number of active electrons	174
17.3	MULTIPLICITY	174
17.4	SYMMETRY — Spatial symmetry of the wavefunction	174
17.5	STATE — Symmetry and multiplicity	174
17.6	VECTORS — Input/output of MO vectors	175
17.7	HESSIAN — Select preconditioner	175
17.8	LEVEL — Level shift for convergence	175
17.9	PRINT and NOPRINT	175
18	Selected CI	177
18.1	Background	177
18.2	Files	178
18.3	Configuration Generation	179
18.3.1	Specifying the reference occupation	179
18.3.2	Applying creation-annihilation operators	180
18.3.3	Uniform excitation level	181
18.4	Number of roots	181
18.5	Accuracy of diagonalization	181
18.6	Selection thresholds	182
18.7	Mode	182
18.8	Memory requirements	182
18.9	Forcing regeneration of the MO integrals	182
18.10	Disabling update of the configuration list	182
18.11	Orbital locking in CI geometry optimization	183
19	Coupled Cluster Calculations	185

<i>CONTENTS</i>	11
19.1 MAXITER — Maximum number of iterations	185
19.2 THRESH — Convergence threshold	186
19.3 TOL2E — integral screening threshold	186
19.4 DIISBAS — DIIS subspace dimension	186
19.5 FREEZE — Freezing orbitals	186
19.6 IPRT — Debug printing	186
19.7 PRINT and NOPRINT	187
19.8 Methods (Tasks) Recognized	187
19.9 Debugging and Development Aids	187
19.9.1 Switching On and Off Terms	187
20 Geometry Optimization with DRIVER	189
20.1 Convergence criteria	190
20.2 Available precision	190
20.3 Controlling the step length	191
20.4 Maximum number of steps	191
20.5 Discard restart information	191
20.6 Regenerate internal coordinates	191
20.7 Initial Hessian	191
20.8 Mode or variable to follow to saddle point	192
20.9 Optimization history as XYZ files	192
20.10 Print options	193
21 Geometry Optimization with STEPPER	195
21.1 MIN and TS — Minimum or transition state search	195
21.2 TRACK — Mode selection	196
21.3 MAXITER — Maximum number of steps	196
21.4 TRUST — Trust radius	196
21.5 CONVGGM, CONVGG and CONVGE — Convergence criteria	196
21.6 Backstepping in STEPPER	197
21.7 Initial Nuclear Hessian Options	197
22 Constraints for Geometry Optimization	199
23 Hybrid Calculations with ONIOM	201
23.1 Real, model and intermediate geometries	202

23.1.1	Link atoms	203
23.1.2	Numbering of the link atoms	204
23.2	High, medium and low theories	204
23.2.1	Basis specification	204
23.2.2	Effective core potentials	204
23.2.3	General input strings	204
23.3	Use of symmetry	205
23.4	Molecular orbital files	205
23.5	Restarting	206
23.6	Examples	206
23.6.1	Hydrocarbon bond energy	206
23.6.2	Optimization and frequencies	207
23.6.3	A three-layer example	209
23.6.4	DFT with and without charge fitting	210
24	Hessians	213
24.1	Hessian Module Input	213
24.1.1	Defining the wavefunction threshold	213
24.1.2	Profile	214
24.1.3	Print Control	214
25	Vibrational frequencies	215
25.1	Vibrational Module Input	215
25.1.1	Hessian File Reuse	216
25.1.2	Redefining Masses of Elements	216
25.1.3	Animation	217
25.1.4	An Example Input Deck	218
26	DPLOT	219
26.1	GAUSSIAN — Gaussian Cube format	219
26.2	TITLE — Title directive	219
26.3	LIMITXYZ — Plot limits	219
26.4	SPIN — Density to be plotted	220
26.5	OUTPUT — Filename	220
26.6	VECTORS — MO vector file name	220

<i>CONTENTS</i>	13
26.7 WHERE — Density evaluation	220
26.8 ORBITAL — Orbital sub-space	220
26.9 Examples	221
27 Electron Transfer Calculations with ET	223
27.1 VECTORS — input of MO vectors for ET reactant and product states	224
27.2 FOCK/NOFOCK — method for calculating the two-electron contribution to V_{RP}	224
27.3 TOL2E — integral screening threshold	224
27.4 Example	224
28 Properties	227
28.1 Subdirectives	228
28.1.1 Nbofile	228
29 Electrostatic potentials	229
29.1 Grid specification	229
29.2 Constraints	230
29.3 Restraints	231
30 Prepare	233
30.1 Default database directories	235
30.2 System name and coordinate source	235
30.3 Sequence file generation	236
30.4 Topology file generation	237
30.5 Appending to an existing topology file	239
30.6 Generating a restart file	240
31 Molecular dynamics	245
31.1 Introduction	245
31.1.1 Spacial decomposition	245
31.1.2 Topology	246
31.1.3 Files	246
31.1.4 Databases	247
31.1.5 Force fields	248
31.2 Format of fragment files	248
31.3 Creating segment files	248

31.4	Creating sequence files	248
31.5	Creating topology files	248
31.6	Creating restart files	249
31.7	Molecular simulations	249
31.8	System specification	249
31.9	Restarting and continuing simulations	250
31.10	Parameter set	250
31.11	Energy minimization algorithms	250
31.12	Multi-configuration thermodynamic integration	251
31.13	Time and integration algorithm directives	252
31.14	Ensemble selection	252
31.15	Velocity reassignments	253
31.16	Cutoff radii	253
31.17	Polarization	253
31.18	External electrostatic field	254
31.19	Constraints	254
31.20	Long range interaction corrections	254
31.21	Fixing coordinates	255
31.22	Special options	255
31.23	Autocorrelation function	256
31.24	Print options	256
31.25	Periodic updates	257
31.26	Recording	258
31.27	Program control options	259
32	Analysis	263
32.1	System specification	263
32.2	Reference coordinates	263
32.3	File specification	264
32.4	Selection	265
32.5	Coordinate analysis	266
32.6	Essential dynamics analysis	267
32.7	Trajectory format conversion	268
32.8	Electrostatic potentials	269

<i>CONTENTS</i>	15
33 Combined quantum and molecular mechanics	271
34 File formats	273
35 Pseudopotential plane-wave density functional theory (NWPW)	283
35.1 PSPW Tasks	284
35.1.1 Simulation Cell	287
35.1.2 Unit Cell Optimization	288
35.1.3 DPLOT	289
35.1.4 Wannier	290
35.1.5 Self-Interaction Corrections	291
35.1.6 Point Charge Analysis	291
35.1.7 Car-Parrinello	292
35.1.8 Adding Geometry Constraints To A Car-Parrinello Simulation	294
35.1.9 PSP_GENERATOR	295
35.1.10 WAVEFUNCTION_INITIALIZER	298
35.1.11 V_WAVEFUNCTION_INITIALIZER	299
35.1.12 WAVEFUNCTION_EXPANDER	300
35.1.13 STEEPEST_DESCENT	301
35.2 Band Tasks	303
35.2.1 Brillouin Zone	305
35.3 PAW Tasks	305
35.4 Pseudopotential and PAW basis Libraries	307
35.5 NWPW RTDB Entries and DataFiles	308
35.5.1 Ion Positions	308
35.5.2 Ion Velocities	308
35.5.3 Wavefunction Datafile	308
35.5.4 Velocity Wavefunction Datafile	308
35.5.5 Formatted Pseudopotential Datafile	308
35.5.6 One-Dimensional Pseudopotential Datafile	308
35.5.7 PSPW Car-Parrinello Output Datafiles	309
35.6 Car-Parrinello Scheme for Ab Initio Molecular Dynamics	311
35.6.1 Verlet Algorithm for Integration	312
35.6.2 Constant Temperature Simulations: Nose-Hoover Thermostats	312
35.7 PSPW Tutorial 1: Minimizing the geometry for a C ₂ molecule	313

35.8	PSPW Tutorial 2: Running a Car-Parrinello Simulation	314
35.9	PSPW Tutorial 3: optimizing a unit cell and geometry for Silicon-Carbide	315
35.10	Band Tutorial 1: Minimizing the energy of a silicon-carbide crystal by running a PSPW and Band simulation in tandem	316
35.11	PAW Tutorial	317
35.12	NWPW Capabilities and Limitations	317
35.13	Questions and Difficulties	318
36	Controlling NWChem with Python	319
36.1	How to input and run a Python program inside NWChem	319
36.2	NWChem extensions	320
36.3	Examples	321
36.3.1	Hello world	321
36.3.2	Scanning a basis exponent	321
36.3.3	Scanning a basis exponent revisited.	322
36.3.4	Scanning a geometric variable	323
36.3.5	Scan using the BSSE counterpoise corrected energy	323
36.3.6	Scan the geometry and compute the energy and gradient	324
36.3.7	Reaction energies varying the basis set	325
36.3.8	Using the database	326
36.3.9	Handling exceptions from NWChem	326
36.3.10	Accessing geometry information — a temporary hack	327
36.3.11	Scanning a basis exponent yet again — plotting and handling child processes	328
36.4	Troubleshooting	329
37	Interfaces to Other Programs	331
37.1	NBO — Natural Bond Orbital Analysis	331
37.2	DIRDYVTST — DIRect Dynamics for Variational Transition State Theory	332
37.2.1	Introduction	333
37.2.2	Files	333
37.2.3	Detailed description of the input	333
38	Acknowledgments	341
A	Standard Basis Sets	343
B	Sample input files	359

B.1	Water SCF calculation and geometry optimization in a 6-31g basis	359
B.1.1	Job 1. Single point SCF energy	359
B.1.2	Job 2. Restarting and perform a geometry optimization	359
B.2	Compute the polarizability of Ne using finite field	360
B.2.1	Job 1. Compute the atomic energy	360
B.2.2	Job 2. Compute the energy with applied field	360
B.3	SCF energy of H ₂ CO using ECPs for C and O	360
B.4	MP2 optimization and CCSD(T) on nitrogen	362
C	Examples of geometries using symmetry	365
C.1	C _s methanol	365
C.2	C _{2v} water	365
C.3	D _{2h} acetylene	366
C.4	D _{2h} ethylene	366
C.5	T _d methane	366
C.6	I _h buckminsterfullerene	366
C.7	S ₄ porphyrin	367
C.8	D _{3h} iron penta-carbonyl	367
C.9	D _{3d} sodium crown ether	368
C.10	C _{3v} ammonia	368
C.11	D _{6h} benzene	368
C.12	C _{3h} BO ₃ H ₃	368
C.13	D _{5d} ferrocene	369
C.14	C _{4v} SF ₅ Cl	369
C.15	C _{2h} trans-dichloroethylene	369
C.16	D _{2d} CH ₂ CCH ₂	370
C.17	D _{5h} cyclopentadiene anion	370
C.18	D _{4h} gold tetrachloride	370
D	Running NWChem	371
D.1	Sequential execution	371
D.2	Parallel execution on UNIX-based parallel machines including workstation clusters using TCGMSG	371
D.3	Parallel execution on UNIX-based parallel machines including workstation clusters using MPI	372
D.4	Parallel execution on MPPs	373
D.5	IBM SP	373

D.6 Cray T3E	374
D.7 Alpha systems with Quadrics switch	374
D.8 Windows 98 and NT	374
D.9 Tested Platforms and O/S versions	375

Chapter 1

Introduction

NWChem is a computational chemistry package designed to run on high-performance parallel supercomputers. Code capabilities include the calculation of molecular electronic energies and analytic gradients using Hartree-Fock self-consistent field (SCF) theory, Gaussian density function theory (DFT), and second-order perturbation theory. For all methods, geometry optimization is available to determine energy minima and transition states. Classical molecular dynamics capabilities provide for the simulation of macromolecules and solutions, including the computation of free energies using a variety of force fields.

NWChem is scalable, both in its ability to treat large problems efficiently, and in its utilization of available parallel computing resources. The code uses the parallel programming tools TCGMSG and the Global Array (GA) library developed at PNNL for the High Performance Computing and Communication (HPCC) grand-challenge software program and the Environmental Molecular Sciences Laboratory (EMSL) Project. NWChem has been optimized to perform calculations on large molecules using large parallel computers, and it is unique in this regard.

This document is intended as an aid to chemists using the code for their own applications. Users are not expected to have a detailed understanding of the code internals, but some familiarity with the overall structure of the code, how it handles information, and the nature of the algorithms it contains will generally be helpful. The following sections describe the structure of the input file, and give a brief overview of the code architecture. All input directives recognized by the code are described in detail, with options, defaults, and recommended usages, where applicable. The appendices present additional information on the molecular geometry and basis function libraries included in the code.

1.1 Citation

The EMSL Software Agreement stipulates that the use of NWChem will be acknowledged in any publications which use results obtained with NWChem. The acknowledgment should be of the form:

NWChem Version 4.6, as developed and distributed by Pacific Northwest National Laboratory, P. O. Box 999, Richland, Washington 99352 USA, and funded by the U. S. Department of Energy, was used to obtain some of these results.

The words “A modified version of” should be added at the beginning, if appropriate. *Note: Your EMSL Software Agreement contains the complete specification of the required acknowledgment.*

Please use the following citation when publishing results obtained with NWChem:

Straatsma, T.P.; Aprà, E.; Windus, T.L.; Bylaska, E.J.; de Jong, W.; Hirata, S.; Valiev, M.; Hackler, M.; Pollack, L.; Harrison, R.; Dupuis, M.; Smith, D.M.A.; Nieplocha, J.; Tipparaju V.; Krishnan, M.; Auer, A.A.; Brown, E.; Cisneros, G.; Fann, G.; Früchtl, H.; Garza, J.; Hirao, K.; Kendall, R.; Nichols, J.; Tsemekhman, K.; Wolinski, K.; Anchell, J.; Bernholdt, D.; Borowski, P.; Clark, T.; Clerc, D.; Dachsel, H.; Deegan, M.; Dyllal, K.; Elwood, D.; Glendening, E.; Gutowski, M.; Hess, A.; Jaffe, J.; Johnson, B.; Ju, J.; Kobayashi, R.; Kutteh, R.; Lin, Z.; Littlefield, R.; Long, X.; Meng, B.; Nakajima, T.; Niu, S.; Rosing, M.; Sandrone, G.; Stave, M.; Taylor, H.; Thomas, G.; van Lenthe, J.; Wong, A.; Zhang, Z.; *NWChem, A Computational Chemistry Package for Parallel Computers, Version 4.6* (2004), Pacific Northwest National Laboratory, Richland, Washington 99352-0999, USA.

High Performance Computational Chemistry: an Overview of NWChem a Distributed Parallel Application, Kendall, R.A.; Aprà, E.; Bernholdt, D.E.; Bylaska, E.J.; Dupuis, M.; Fann, G.I.; Harrison, R.J.; Ju, J.; Nichols, J.A.; Nieplocha, J.; Straatsma, T.P.; Windus, T.L.; Wong, A.T. *Computer Phys. Comm.*, 2000, **128**, 260–283 .

If you use the DIRDYVTST portion of NWChem, please also use the additional citation:

DIRDYVTST, Yao-Yuan Chuang and Donald G. Truhlar, Department of Chemistry and Super Computer Institute, University of Minnesota; Ricky A. Kendall, Scalable Computing Laboratory, Ames Laboratory and Iowa State University; Bruce C. Garrett and Theresa L. Windus, Environmental Molecular Sciences Laboratory, Pacific Northwest Laboratory.

1.2 User Feedback

This software comes without warranty or guarantee of support, but we do try to meet the needs of our user community. Please send bug reports, requests for enhancement, or other comments to

`nwchem-users@emsl.pnl.gov`

When reporting problems, please provide as much information as possible, including:

- detailed description of problem
- site name
- platform you are running on, including
 - vendor name
 - computer model
 - operating system
 - compiler
- input file
- output file

Users can also subscribe to the `nwchem-users@emsl.pnl.gov` electronic mailing list itself. This is intended as a general forum through which code users can contact one another and the developers, to share experience with the code and discuss problems. Announcements of new releases and bug fixes will also be made to this list.

To subscribe to the user list, send a message to

```
majordomo@emsl.pnl.gov
```

The body of the message must contain the line

```
subscribe nwchem-users
```

The automated list manager is capable of recognizing a number of commands, including ; “subscribe”, “unsubscribe”, “get”, “index”, “which”, “who”, “info” and “lists”. The command “end” halts processing of commands. It will provide some help if the message includes the line `help` in the body.

Chapter 2

Getting Started

This section provides an overview of NWChem input and program architecture, and the syntax used to describe the input. See Sections 2.2 and 2.3 for examples of NWChem input files with detailed explanation.

NWChem consists of independent modules that perform the various functions of the code. Examples of modules include the input parser, SCF energy, SCF analytic gradient, DFT energy, etc.. Data is passed between modules and saved for restart using a disk-resident database or dumpfile (see Section 3).

The input to NWChem is composed of commands, called directives, which define data (such as basis sets, geometries, and filenames) and the actions to be performed on that data. Directives are processed in the order presented in the input file, with the exception of certain start-up directives (see Section 2.1) which provide critical job control information, and are processed before all other input. Most directives are specific to a particular module and define data that is used by that module only. A few directives (see Section 5) potentially affect all modules, for instance by specifying the total electric charge on the system.

There are two types of directives. Simple directives consist of one line of input, which may contain multiple fields. Compound directives group together multiple simple directives that are in some way related and are terminated with an END directive. See the sample inputs (Sections 2.2, 2.3) and the input syntax specification (Section 2.4).

All input is free format and case is ignored except for actual data (e.g., names/tags of centers, titles). Directives or blocks of module-specific directives (i.e., compound directives) can appear in any order, with the exception of the TASK directive (see sections 2.1 and 5.10) which is used to invoke an NWChem module. All input for a given task must precede the TASK directive. This input specification rule allows the concatenation of multiple tasks in a single NWChem input file.

To make the input as short and simple as possible, most options have default values. The user needs to supply input only for those items that have no defaults, or for items that must be different from the defaults for the particular application. In the discussion of each directive, the defaults are noted, where applicable.

The input file structure is described in the following sections, and illustrated with two examples. The input format and syntax for directives is also described in detail.

2.1 Input File Structure

The structure of an input file reflects the internal structure of NWChem. At the beginning of a calculation, NWChem needs to determine how much memory to use, the name of the database, whether it is a new or restarted job, where to put scratch/permanent files, etc.. It is not necessary to put this information at the top of the input file, however.

NWChem will read through the *entire* input file looking for the start-up directives. In this first pass, all other directives are ignored.

The start-up directives are

- START
- RESTART
- SCRATCH_DIR
- PERMANENT_DIR
- MEMORY
- ECHO

After the input file has been scanned for the start-up directives, it is rewound and read sequentially. Input is processed either by the top-level parser (for the directives listed in Section 5, such as `TITLE`, `SET`, ...) or by the parsers for specific computational modules (e.g., `SCF`, `DFT`, ...). Any directives that have already been processed (e.g., `MEMORY`) are ignored. Input is read until a `TASK` directive (see Section 5.10) is encountered. A `TASK` directive requests that a calculation be performed and specifies the level of theory and the operation to be performed. Input processing then stops and the specified task is executed. The position of the `TASK` directive in effect marks the end of the input for that task. Processing of the input resumes upon the successful completion of the task, and the results of that task are available to subsequent tasks in the same input file.

The name of the input file is usually provided as an argument to the execute command for NWChem. That is, the execute command looks something like the following;

```
nwchem input_file
```

The default name for the input file is `nwchem.nw`. If an input file name `input_file` is specified without an extension, the code assumes `.nw` as a default extension, and the input filename becomes `input_file.nw`. If the code cannot locate a file named either `input_file` or `input_file.nw` (or `nwchem.nw` if no file name is provided), an error is reported and execution terminates. The following section presents two input files to illustrate the directive syntax and input file format for NWChem applications.

2.2 Simple Input File — SCF geometry optimization

A simple example of an NWChem input file is an SCF geometry optimization of the nitrogen molecule, using a Dunning cc-pvdz basis set. This input file contains the bare minimum of information the user must specify to run this type of problem — fewer than ten lines of input, as follows:

```
title "Nitrogen cc-pvdz SCF geometry optimization"
geometry
  n 0 0 0
  n 0 0 1.08
end
basis
  n library cc-pvdz
end
task scf optimize
```

Examining the input line by line, it can be seen that it contains only four directives; TITLE, GEOMETRY, BASIS, and TASK. The TITLE directive is optional, and is provided as a means for the user to more easily identify outputs from different jobs. An initial geometry is specified in Cartesian coordinates and Angströms by means of the GEOMETRY directive. The Dunning cc-pvdz basis is obtained from the NWChem basis library, as specified by the BASIS directive input. The TASK directive requests an SCF geometry optimization.

The GEOMETRY directive (Section 6) defaults to Cartesian coordinates and Angströms (options include atomic units and Z-matrix format; see Section 6.4). The input blocks for the BASIS and GEOMETRY directives are structured in similar fashion, i.e., name, keyword, ..., end (In this simple example, there are no keywords). The BASIS input block *must* contain basis set information for every atom type in the geometry with which it will be used. Refer to Sections 7 and 8, and Appendix A for a description of available basis sets and a discussion of how to define new ones.

The last line of this sample input file (`task scf optimize`) tells the program to optimize the molecular geometry by minimizing the SCF energy. (For a description of possible tasks and the format of the TASK directive, refer to Section 5.10.)

If the input is stored in the file `n2.nw`, the command to run this job on a typical UNIX workstation is as follows:

```
nwchem n2
```

NWChem output is to UNIX standard output, and error messages are sent to both standard output and standard error.

2.3 Water Molecule Sample Input File

A more complex sample problem is the optimization of a positively charged water molecule using second-order Møller-Plesset perturbation theory (MP2), followed by a computation of frequencies at the optimized geometry. A preliminary SCF geometry optimization is performed using a computationally inexpensive basis set (STO-3G). This yields a good starting guess for the optimal geometry, and any Hessian information generated will be used in the next optimization step. Then the optimization is finished using MP2 and a basis set with polarization functions. The final task is to calculate the MP2 vibrational frequencies. The input file to accomplish these three tasks is as follows:

```
start h2o_freq

charge 1

geometry units angstroms
  O      0.0  0.0  0.0
  H      0.0  0.0  1.0
  H      0.0  1.0  0.0
end

basis
  H library sto-3g
  O library sto-3g
end

scf
  uhf; doublet
  print low
```

```
end

title "H2O+ : STO-3G UHF geometry optimization"

task scf optimize

basis
  H library 6-31g**
  O library 6-31g**
end

title "H2O+ : 6-31g** UMP2 geometry optimization"

task mp2 optimize

mp2; print none; end
scf; print none; end

title "H2O+ : 6-31g** UMP2 frequencies"

task mp2 freq
```

The `START` directive (Section 5.1) tells NWChem that this run is to be started from the beginning. This directive need not be at the beginning of the input file, but it is commonly placed there. Existing database or vector files are to be ignored or overwritten. The entry `h2o_freq` on the `START` line is the prefix to be used for all files created by the calculation. This convention allows different jobs to run in the same directory or to share the same scratch directory (see Section 5.2), as long as they use different prefix names in this field.

As in the first sample problem, the geometry is given in Cartesian coordinates. In this case, the units are specified as Angströms. (Since this is the default, explicit specification of the units is not actually necessary, however.) The `CHARGE` directive defines the total charge of the system. This calculation is to be done on an ion with charge +1.

A small basis set (STO-3G) is specified for the initial geometry optimization. Next, the multiple lines of the first SCF directive in the `scf . . . end` block specify details about the SCF calculation to be performed. Unrestricted Hartree-Fock is chosen here (by specifying the keyword `uhf`), rather than the default, restricted open-shell high-spin Hartree-Fock (ROHF). This is necessary for the subsequent MP2 calculation, because only UMP2 is currently available for open-shell systems (see Section 4). For open-shell systems, the spin multiplicity has to be specified (using `doublet` in this case), or it defaults to `singlet`. The print level is set to `low` to avoid verbose output for the starting basis calculations.

All input up to this point affects only the settings in the runtime database. The program takes its information from this database, so the sequence of directives up to the first `TASK` directive is irrelevant. An exchange of order of the different blocks or directives would not affect the result. The `TASK` directive, however, must be specified after all relevant input for a given problem. The `TASK` directive causes the code to perform the specified calculation using the parameters set in the preceding directives. In this case, the first task is an SCF calculation with geometry optimization, specified with the input `scf` and `optimize`. (See Section 5.10 for a list of available tasks and operations.)

After the completion of any task, settings in the database are used in subsequent tasks without change, unless they are overridden by new input directives. In this example, before the second task (`task mp2 optimize`), a better basis set (6-31G**) is defined and the title is changed. The second `TASK` directive invokes an MP2 geometry optimization.

Once the MP2 optimization is completed, the geometry obtained in the calculation is used to perform a frequency calculation. This task is invoked by the keyword `freq` in the final `TASK` directive, `task mp2 freq`. The second derivatives of the energy are calculated as numerical derivatives of analytical gradients. The intermediate energies and gradients are not of interest in this case, so output from the SCF and MP2 modules is disabled with the `PRINT` directives.

2.4 Input Format and Syntax for Directives

This section describes the input format and the syntax used in the rest of this documentation to describe the format of directives. The input format for the directives used in NWChem is similar to that of UNIX shells, which is also used in other chemistry packages, most notably GAMESS-UK. An input line is parsed into whitespace (blanks or tabs) separating tokens or fields. Any token that contains whitespace must be enclosed in double quotes in order to be processed correctly. For example, the basis set with the descriptive name `modified Dunning DZ` must appear in a directive as `"modified Dunning DZ"`, since the name consists of three separate words.

2.4.1 Input Format

A (physical) line in the input file is terminated with a newline character (also known as a ‘return’ or ‘enter’ character). A semicolon (;) can be also used to indicate the end of an input line, allowing a single physical line of input to contain multiple logical lines of input. For example, five lines of input for the `GEOMETRY` directive can be entered as follows;

```
geometry
O 0 0 0
H 0 1.430 1.107
H 0 -1.430 1.107
end
```

These same five lines could be entered on a single line, as

```
geometry; O 0 0 0; H 0 1.430 1.107; H 0 -1.430 1.107; end
```

This one physical input line comprises five logical input lines. Each logical or physical input line must be no longer than 1023 characters.

In the input file:

- a string, token, or field is a sequence of ASCII characters (NOTE: if the string includes blanks or tabs (i.e., white space), the entire string must be enclosed in double quotes).
- \ (backslash) at the end of a line concatenates it with the next line. Note that a space character is automatically inserted at this point so that it is *not* possible to split tokens across lines. A backslash is also used to quote special characters such as whitespace, semi-colons, and hash symbols so as to avoid their special meaning (NOTE: these special symbols must be quoted with the backslash even when enclosed within double quotes).
- ; (semicolon) is used to mark the end of a logical input line within a physical line of input.
- # (the hash or pound symbol) is the comment character. All characters following # (up to the end of the physical line) are ignored.

- If *any* input line (excluding Python programs, Section 36) begins with the string `INCLUDE` (ignoring case) and is followed by a valid file name, then the data in that file are read as if they were included into the current input file at the current line. Up to three levels of nested include files are supported. The user should note that inputting a basis set from the standard basis library (Section 7) uses one level of include.
- Data is read from the input file until an end-of-file is detected, or until the string `EOF` (ignoring case) is encountered at the beginning of an input line.

2.4.2 Format and syntax of directives

Directives consist of a directive name, keywords, and optional input, and may contain one line or many. Simple directives consist of a single line of input with one or more fields. Compound directives can have multiple input lines, and can also include other optional simple and compound directives. A compound directive is terminated with an `END` directive. The directives `START` (see Section 5.1) and `ECHO` (see Section 5.4) are examples of simple directives. The directive `GEOMETRY` (see Section 6) is an example of a compound directive.

Some limited checking of the input for self-consistency is performed by the input module, but most defaults are imposed by the application modules at runtime. It is therefore usually impossible to determine beforehand whether or not all selected options are consistent with each other.

In the rest of this document, the following notation and syntax conventions are used in the generic descriptions of the NWChem input.

- a directive name always appears in all-capitals, and in computer typeface (e.g., `GEOMETRY`, `BASIS`, `SCF`). Note that the case of directives and keywords is ignored in the actual input.
- a keyword always appears in lower case, in computer typeface (e.g., `swap`, `print`, `units`, `bqbg`).
- variable names always appear in lower case, in computer typeface, and enclosed in angle brackets to distinguish them from keywords (e.g., `<input_filename>`, `<basishname>`, `<tag>`).
- `$variable$` is used to indicate the substitution of the value of a variable.
- `()` is used to group items (the parentheses and other special symbols should not appear in the input).
- `| |` separate exclusive options, parameters, or formats.
- `[]` enclose optional entries that have a default value.
- `< >` enclose a type, a name of a value to be specified, or a default value, if any.
- `\` is used to concatenate lines in a description.
- `. . .` is used to indicate indefinite continuation of a list.

An input parameter is identified in the description of the directive by prefacing the name of the item with the type of data expected, i.e.,

- `string` – an ASCII character string
- `integer` – integer value(s) for a variable or an array
- `logical` – true/false logical variable
- `real` – real floating point value(s) for a variable or an array

- `double` – synonymous with `real`

If an input item is not prefaced by one of these type names, it is assumed to be of type “string”.

In addition, integer lists may be specified using Fortran triplet notation, which interprets `lo:hi:inc` as `lo`, `lo+inc`, `lo+2*inc`, ..., `hi`. For example, where a list of integers is expected in the input, the following two lines are equivalent

```
7 10 21:27:2 1:3 99
7 10 21 23 25 27 1 2 3 99
```

(In Fortran triplet notation, the increment, if unstated, is 1; e.g., `1:3 = 1:3:1`.)

The directive `VECTORS` (Section 10.5) is presented here as an example of an NWChem input directive. The general form of the directive is as follows:

```
VECTORS [input (<string input_movecs default atomic>) || \
        (project <string basisname> <string filename>)] \
[swap [(alpha||beta)] <integer vec1 vec2> ...] \
[output <string output_movecs default $file_prefix$.movecs>]
```

This directive contains three optional keywords, as indicated by the three main sets of square brackets enclosing the keywords `input`, `swap`, and `output`. The keyword `input` allows the user to specify the source of the molecular orbital vectors. There are two mutually exclusive options for specifying the vectors, as indicated by the `||` symbol separating the option descriptions;

```
(<string input_movecs default atomic>) || \
        (project <string basisname> <string filename>) \
```

The first option, (`<string input_movecs default atomic>`), allows the user to specify an ASCII character string for the parameter `input_movecs`. If no entry is specified, the code uses the default `atomic` (i.e., `atomic guess`). The second option, (`project <string basisname> <string filename>`), contains the keyword `project`, which takes two string arguments. When this keyword is used, the vectors in file `<filename>` will be projected from the (smaller) basis `<basisname>` into the current atomic orbital (AO) basis.

The second keyword, `swap`, allows the user to re-order the starting vectors, specifying the pairs of vectors to be swapped. As many pairs as the user wishes to have swapped can be listed for `<integer vec1 vec2 ... >`. The optional keywords `alpha` and `beta` allow the user to swap the alpha or beta spin orbitals.

The third keyword, `output`, allows the user to tell the code where to store the vectors, by specifying an ASCII string for the parameter `output_movecs`. If no entry is specified for this parameter, the default is to write the vectors back into either the user-specified MO vectors input file or, if this is not available, the file `$file_prefix$.movecs`.

A particular example of the `VECTORS` directive is shown below. It specifies both the `input` and `output` keywords, but does not use the `swap` option.

```
vectors input project "small basis" small_basis.movecs \
        output large_basis.movecs
```

This directive tells the code to generate input vectors by projecting from vectors in a smaller basis named "small basis", which is stored in the file `small_basis.movecs`. The output vectors will be stored in the file `large_basis.movecs`.

The order of keyed optional entries within a directive should not matter, unless noted otherwise in the specific instructions for a particular directive.

Chapter 3

NWChem Architecture

As noted above, NWChem consists of independent modules that perform the various functions of the code. Examples include the input parser, self-consistent field (SCF) energy, SCF analytic gradient, and density functional theory (DFT) energy modules. The independent NWChem modules can share data only through a disk-resident database, which is similar to the GAMESS-UK dumpfile or the Gaussian checkpoint file. This allows the modules to share data, or to share access to files containing data.

It is not necessary for the user to be intimately familiar with the contents of the database in order to run NWChem. However, a nodding acquaintance with the design of the code will help in clarifying the logic behind the input requirements, especially when restarting jobs or performing multiple tasks within one job. Section 3.1 gives a general description of the database.

As described above (Section 2.1), all start-up directives are processed at the beginning of the job by the main program, and then the input module is invoked. Each input directive usually results in one or more entries being made in the database. When a `TASK` directive is encountered, control is passed to the appropriate module, which extracts relevant data from the database and any associated files. Upon completion of the task, the module will store significant results in the database, and may also modify other database entries in order to affect the behavior of subsequent computations.

3.1 Database Structure

Data is shared between modules of NWChem by means of the database. Three main types of information are stored in the data base: (1) arrays of data, (2) names of files that contain data, and (3) objects. Arrays are stored directly in the database, and contain the following information:

1. the name of the array, which is a string of ASCII characters (e.g., "reference energies")
2. the type of the data in the array (i.e., real, integer, logical, or character)
3. the number (N) of data items in the array (Note: A scalar is stored as an array of unit length.)
4. the N items of data of the specified type

It is possible to enter data directly into the database using the `SET` directive (see Section 5.7). For example, to store a (64-bit precision) three-element real array with the name "reference energies" in the database, the directive is as follows:

```
set "reference energies" 0.0 1.0 -76.2
```

NWChem determines the data to be real (based on the type of the first element, 0.0), counts the number of elements in the array, and enters the array into the database.

Much of the data stored in the database is internally managed by NWChem and should not be modified by the user. However, other data, including some NWChem input options, can be freely modified.

Objects are built in the database by storing associated data as multiple entries, using an internally consistent naming convention. This data is managed exclusively by the subroutines (or methods) that are associated with the object. Currently, the code has two main objects: basis sets and geometries. Sections 6 and 7 present a complete discussion of the input to describe these objects.

As an illustration of what comprises a geometry object, the following table contains a partial listing of the database contents for a water molecule geometry named "test geom". Each entry contains the field test geom, which is the unique name of the object.

Contents of RTDB h2o.db

Entry	Type[nelem]
-----	-----
geometry:test geom:efield	double[3]
geometry:test geom:coords	double[9]
geometry:test geom:ncenter	int[1]
geometry:test geom:charges	double[3]
geometry:test geom:tags	char[6]
...	

Using this convention, multiple instances of objects may be stored with different names in the same database. For example, if a user needed to do calculations considering alternative geometries for the water molecule, an input file could be constructed containing all the geometries of interest by storing them in the database under different names.

The runtime database contents for the file h2o.db listed above were generated from the user-specified input directive,

```
geometry "test geom"
  O    0.00000000    0.00000000    0.00000000
  H    0.00000000    1.43042809   -1.10715266
  H    0.00000000   -1.43042809   -1.10715266
end
```

The GEOMETRY directive allows the user to specify the coordinates of the atoms (or centers), and identify the geometry with a unique name. (Refer to Section 6 for a complete description of the GEOMETRY directive.)

Unless a specific name is defined for the geometry, (such as the name "test geom" shown in the example), the default name of geometry is assigned. This is the geometry name that computational modules will look for when executing a calculation. The SET directive can be used in the input to force NWChem to look for a geometry with a name other than geometry. For example, to specify use of the geometry with the name "test geom" in the example above, the SET directive is as follows:

```
set geometry "test geom"
```

NWChem will automatically check for such indirections when loading geometries. Storage of data associated with basis sets, the other database resident object, functions in a similar fashion, using the default name "ao basis".

3.2 Persistence of data and restart

The database is persistent, meaning that all input data and output data (calculation results) that are not destroyed in the course of execution are permanently stored. These data are therefore available to subsequent tasks or jobs. This makes the input for restart jobs very simple, since only new or changed data must be provided. It also makes the behavior of successive restart jobs *identical* to that of multiple tasks within one job.

Sometimes, however, this persistence is undesirable, and it is necessary to return an NWChem module to its default behavior by restoring the database to its input-free state. In such a case, the `UNSET` directive (see Section 5.8) can be used to delete all database entries associated with a given module (including both inputs and outputs).

Chapter 4

Functionality

NWChem provides many methods to compute the properties of molecular and periodic systems using standard quantum mechanical descriptions of the electronic wavefunction or density. In addition, NWChem has the capability to perform classical molecular dynamics and free energy simulations. These approaches may be combined to perform mixed quantum-mechanics and molecular-mechanics simulations.

NWChem is available on almost all high performance computing platforms, workstations, PCs running LINUX, as well as clusters of desktop platforms or workgroup servers. NWChem development has been devoted to providing maximum efficiency on massively parallel processors. It achieves this performance on the 1960 processors HP Itanium2 system in the EMSL's MSCF. It has not been optimized for high performance on single processor desktop systems.

4.1 Molecular electronic structure

The following quantum mechanical methods are available to calculate energies, analytic first derivatives and second derivatives with respect to atomic coordinates.

- Self Consistent Field (SCF) or Hartree Fock (RHF, UHF).
- Gaussian Density Functional Theory (DFT), using many local, non-local (gradient-corrected), and hybrid (local, non-local, and HF) exchange-correlation potentials (spin-restricted) with formal N^3 and N^4 scaling.

The following methods are available to calculate energies and analytic first derivatives with respect to atomic coordinates. Second derivatives are computed by finite difference of the first derivatives.

- Self Consistent Field (SCF) or Hartree Fock (ROHF).
- Gaussian Density Functional Theory (DFT), using many local, non-local (gradient-corrected), and hybrid (local, non-local, and HF) exchange-correlation potentials (spin-unrestricted) with formal N^3 and N^4 scaling.
- Spin-orbit DFT (SODFT), using many local and non-local (gradient-corrected) exchange-correlation potentials (spin-unrestricted).
- MP2 including semi-direct using frozen core and RHF and UHF reference.
- Complete active space SCF (CASSCF).

The following methods are available to compute energies only. First and second derivatives are computed by finite difference of the energies.

- CCSD, CCSD(T), CCSD+T(CCSD), with RHF reference.
- Selected-CI with second-order perturbation correction.
- MP2 fully-direct with RHF reference.
- Resolution of the identity integral approximation MP2 (RI-MP2), with RHF and UHF reference.
- CIS, TDHF, TDDFT, and Tamm–Dancoff TDDFT for excited states with RHF, UHF, RDFT, or UDFT reference.
- CCSD(T) and CCSD[T] for closed- and open-shell systems (TCE module)
- UCCD, ULCCD, UCCSD, ULCCSD, UQCISD, UCCSDT, and UCCSDTQ with RHF, UHF, or ROHF reference.
- UCISD, UCISDT, and UCISDTQ with RHF, UHF, or ROHF reference.
- Non-canonical UMP2, UMP3, and UMP4 with RHF or UHF reference.
- EOM-CCSD, EOM-CCSDT, EOM-CCSDTQ for excitation energies, transition moments, and excited-state dipole moments of closed- and open-shell systems
- CCSD, CCSDT, CCSDTQ for dipole moments of closed- and open-shell systems

For all methods, the following operations may be performed:

- Single point energy
- Geometry optimization (minimization and transition state)
- Molecular dynamics on the fully *ab initio* potential energy surface
- Numerical first and second derivatives automatically computed if analytic derivatives are not available
- Normal mode vibrational analysis in cartesian coordinates
- ONIOM hybrid method of Morokuma and co-workers
- Generation of the electron density file for graphical display
- Evaluation of static, one-electron properties.
- Electrostatic potential fit of atomic partial charges (CHELPG method with optional RESP restraints or charge constraints)

For closed and open shell SCF and DFT:

- COSMO energies - the continuum solvation ‘COnductor-like Screening MOdel’ of A. Klamt and G. Schüürmann to describe dielectric screening effects in solvents.

In addition, automatic interfaces are provided to

- The natural bond orbital (NBO) package
- Python

4.2 Relativistic effects

The following methods for including relativity in quantum chemistry calculations are available:

- The spin-free one-electron Douglas-Kroll approximation is available for all quantum mechanical methods and their gradients.
- Dyall's spin-free Modified Dirac Hamiltonian approximation is available for the Hartree-Fock method and its gradients.
- One-electron spin-orbit effects can be included via spin-orbit potentials. This option is available for DFT and its gradients, but has to be run without symmetry.

4.3 Pseudopotential plane-wave electronic structure

Two modules are available to compute the energy, optimize the geometry, numerical second derivatives, and perform ab initio molecular dynamics using pseudopotential plane-wave DFT.

- PSPW - (Pseudopotential plane-wave) A gamma point code for calculating molecules, liquids, crystals, and surfaces.
- Band - A prototype band structure code for calculating crystals and surfaces with small band gaps (e.g. semiconductors and metals)

With

- Conjugate gradient and limited memory BFGS minimization
- Car-Parrinello (extended Lagrangian dynamics)
- Constant energy and constant temperature Car-Parrinello simulations
- Fixed atoms in cartesian and SHAKE constraints in Car-Parrinello
- Pseudopotential libraries
- Hamann and Troullier-Martins norm-conserving pseudopotentials with optional semicore corrections
- Automated wavefunction initial guess, now with LCAO
- Vosko and PBE96 exchange-correlation potentials (spin-restricted and unrestricted)
- Orthorhombic simulation cells with periodic and free space boundary conditions.
- Modules to convert between small and large plane-wave expansions
- Interface to DRIVER, STEPPER, and VIB modules
- Polarization through the use of point charges
- Mulliken, point charge, DPLOT (wavefunction, density and electrostatic potential plotting) analysis

4.4 Molecular dynamics

The following functionality is available for classical molecular simulations:

- Single configuration energy evaluation
- Energy minimization
- Molecular dynamics simulation
- Free energy simulation (multistep thermodynamic perturbation (MSTP) or multiconfiguration thermodynamic integration (MCTI) methods with options of single and/or dual topologies, double wide sampling, and separation-shifted scaling)

The classical force field includes:

- Effective pair potentials (functional form used in AMBER, GROMOS, CHARMM, etc.)
- First order polarization
- Self consistent polarization
- Smooth particle mesh Ewald (SPME)
- Twin range energy and force evaluation
- Periodic boundary conditions
- SHAKE constraints
- Consistent temperature and/or pressure ensembles

NWChem also has the capability to combine classical and quantum descriptions in order to perform:

- Mixed quantum-mechanics and molecular-mechanics (QM/MM) minimizations and molecular dynamics simulation, and
- Quantum molecular dynamics simulation by using any of the quantum mechanical methods capable of returning gradients.

By using the DIRDYVTST module of NWChem, the user can write an input file to the POLYRATE program, which can be used to calculate rate constants including quantum mechanical vibrational energies and tunneling contributions.

4.5 Python

The Python programming language has been embedded within NWChem and many of the high level capabilities of NWChem can be easily combined and controlled by the user to perform complex operations.

4.6 Parallel tools and libraries (ParSoft)

- Global arrays (GA)
- Agregate Remote Memory Copy Interface (ARMCI)
- Linear Algebra (PeIGS) and FFT
- ParIO
- Memory allocation (MA)

Chapter 5

Top-level directives

Top-level directives are directives that can affect all modules in the code. Some specify molecular properties (e.g., total charge) or other data that should apply to all subsequent calculations with the current database. However, most top-level directives provide the user with the means to manage the resources for a calculation and to start computations. As the first step in the execution of a job, NWChem scans the entire input file looking for start-up directives, which NWChem must process before all other input. The input file is then rewound and processed sequentially, and each directive is processed in the order in which it is encountered. In this second pass, start-up directives are ignored.

The following sections describe each of the top-level directives in detail, noting all keywords, options, required input, and defaults.

5.1 START and RESTART — Start-up mode

A `START` or `RESTART` directive is optional. If one of these two directives is not specified explicitly, the code will infer one, based upon the name of the input file and the availability of the database. When allowing NWChem to infer the start-up directive, the user must be quite certain that the contents of the database will result in the desired action. It is usually more prudent to specify the directive explicitly, using the following format:

```
(RESTART || START) \
    [<string file_prefix default $input_file_prefix$>] \
    [rtdb <string rtdb_file_name default $file_prefix$.db>]
```

The `START` directive indicates that the calculation is one in which a new database is to be created. Any relevant information that already exists in a previous database of the same name is destroyed. The string variable `<file_prefix>` will be used as the prefix to name any files created in the course of the calculation.

E.g., to start a new calculation on water, one might specify

```
start water
```

which will make all files begin with "water."

If the user does not specify an entry for `<file_prefix>` on the `START` directive (or omits the `START` directive altogether), the code uses the base-name of the input file as the file prefix. That is, the variable `<file_prefix>` is assigned the name of the input file (not its full pathname), but without the last "dot-suffix". For example, the input

file name `/home/dave/job.2.nw` yields `job.2` as the file prefix, if a name is not assigned explicitly using the `START` directive.

The user also has the option of specifying a unique name for the database, using the keyword `rtdb`. When this keyword is entered, the string entered for `rtdb_file_name` is used as the database name. If the keyword `rtdb` is omitted, the name of the database defaults to `<file_prefix>$.db` in the directory for permanent files.

If a calculation is to start from a previous calculation and go on using the existing database, the `RESTART` directive must be used. In such a case, the previous database must already exist. The name specified for `<file_prefix>` usually should not be changed when restarting a calculation. If it is changed, NWChem will not be able to find needed files when going on with the calculation.

In the most common situation, the previous calculation was completed (with or without an error condition), and it is desired to perform a new task or restart the previous one, perhaps with some input changes. In these instances, the `RESTART` directive should be used. This reuses the previous database and associated files, and reads the input file for new input and task information.

The `RESTART` directive looks immediately for new input and task information, deleting information about previous incomplete tasks. For example, when doing a `RESTART` there is no need to specify geometry or basis set declaration because the program will detect this information since it is stored in the run-time database.

If a calculation runs out of time, for example because it is on a queuing system, this is another instance where doing a `RESTART` is advisable. Simply include nothing after the `RESTART` directive except those tasks that are unfinished.

NOTE: Due to changes in the runtime database structure, `RESTART` will not work on database files generated by NWChem versions 4.0.1 and older.

To summarize the default options for this start-up directive, if the input file does *not* contain a `START` or a `RESTART` directive, then

- the variable `<file_prefix>` is assigned the name of the input file for the job, without the suffix (which is usually `.nw`)
- the variable `<rtdb_file_name>` is assigned the default name, `$file_prefix$.db`

If the database with name `$file_prefix$.db` does *not* already exist, the calculation is carried out as if a `START` directive had been encountered. If the database with name `$file_prefix$.db` *does* exist, then the calculation is performed as if a `RESTART` directive had been encountered.

For example, NWChem can be run using an input file with the name `water.nw` by typing the UNIX command line,

```
nwchem water.nw
```

If the NWChem input file `water.nw` does not contain a `START` or `RESTART` directive, the code sets the variable `<file_prefix>` to `water`. Files created by the job will have this prefix, and the database will be named `water.db`. If the database `water.db` does *not* exist already, the code behaves as if the input file contains the directive,

```
start water
```

If the database `water.db` *does* exist, the code behaves as if the input file contained the directive,

```
restart water
```

5.2 SCRATCH_DIR and PERMANENT_DIR — File directories

These are start-up directives that allow the user to specify the directory location of scratch and permanent files created by NWChem. NWChem distinguishes between permanent (or persistent) files and scratch (or temporary) files, and allows the user the option of putting them in different locations. In most installations, however, permanent and scratch files are all written to the current directory by default. What constitutes “local” disk space may also differ from machine to machine.

The conventions for file storage are at the discretion of the specific installation, and are quite likely to be different on different machines. When assigning locations for permanent and scratch files, the user must be cognizant of the characteristics of the installation on a particular platform. To consider just a few examples, on clusters, machine-specific or process-specific names must be supplied for both local and shared file systems, while on SMPs it is useful to specify scratch file directories with automated striping across processors with round-robin allocation. On SMP clusters (a.k.a. constellations), both of these specifications are required.

The `SCRATCH_DIR` and `PERMANENT_DIR` directives are identical in format and capability, and enable the user to specify a single directory for all processes, or different directories for different processes. The general form of the directive is as follows:

```
(PERMANENT_DIR || SCRATCH_DIR) [(<string host>||<integer process>):] \  
                                <string directory> \  
                                [...]
```

Directories are extracted from the user input by executing the following steps, in sequence:

1. Look for a directory qualified by the process ID number of the invoking process. Processes are numbered from zero. Else,
2. If there is a list of directories qualified by the name of the host machine¹, then use round-robin allocation from the list for processes executing on the given host. Else,
3. If there is a list of directories unqualified by any hostname or process ID, then use round-robin allocation from this list.

If directory allocation directive(s) are not specified in the input file, or if no match is found to the directory names specified by input using these directives, then the steps above are executed using the installation-specific defaults. If the code cannot find a valid directory name based on the input specified in either the directive(s) or the system defaults, files are automatically written to the current working directory (".").

The following is a list of examples of specific allocations of scratch directory locations:

- Put scratch files from all processes in the local scratch directory (Warning: the definition of “local scratch directory” may change from machine to machine):

```
scratch_dir /localscratch
```

- Put scratch files from Process 0 in `/piofs/rjh`, but put all other scratch files in `/scratch`:

```
scratch_dir /scratch 0:/piofs/rjh
```

- Put scratch files from Process 0 in directory `scr1`, those from Process 1 in `scr2`, and so forth, in a round-robin fashion, using the given list of directories:

¹As returned by `util_hostname()` which maps to the output of the command `hostname` on Unix workstations.

```
scratch_dir /scr1 /scr2 /scr3 /scr4 /scr5
```

- Allocate files in a round-robin fashion from host-specific lists for processes distributed across two SGI multi-processor machines (node names *coho* and *bohr*):

```
scratch_dir coho:/xfs1/rjh coho:/xfs2/rjh coho:/xfs3/rjh \
bohr:/disk01/rjh bohr:/disk02/rjh bohr:/disk13/rjh
```

5.3 MEMORY — Control of memory limits

This is a start-up directive that allows the user to specify the amount of memory that NWChem can use for the job. If this directive is not specified, memory is allocated according to installation-dependent defaults. *The defaults should generally suffice for most calculations, since the defaults usually correspond to the total amount of memory available on the machine. It should usually be unnecessary to provide a memory directive!!!*

The general form of the directive is as follows:

```
MEMORY [[total] <integer total_size>] \
[stack <integer stack_size>] \
[heap <integer heap_size>] \
[global <integer global_size>] \
[units <string units default real>] \
[(verify|noverify)] \
[(nohardfail|hardfail)] \
```

NWChem recognizes the following memory units:

- real and double (synonyms)
- integer
- byte
- kb (kilobytes)
- mb (megabytes)
- mw (megawords, 64-bit word)

In most cases, the user need specify only the total memory limit to adjust the amount of memory used by NWChem. The following specifications all provide for eight megabytes of total memory (assuming 64-bit floating point numbers), which will be distributed according to the default partitioning:

```
memory 1048576
memory 1048576 real
memory 1 mw
memory 8 mb
memory total 8 mb
memory total 1048576
```

In NWChem there are three distinct regions of memory: stack, heap, and global. Stack and heap are node-private, while the union of the global region on all processors is used to provide globally-shared memory. The allowed limits on each category are determined from a default partitioning (currently 25% heap, 25% stack, and 50% global). Alternatively, the keywords `stack`, `heap`, and `global` can be used to define specific allocations for each of these categories. If the user sets only one of the stack, heap, or global limits by input, the limits for the other two categories are obtained by partitioning the remainder of the total memory available in proportion to the weight of those two categories in the default memory partitioning. If two of the category limits are given, the third is obtained by subtracting the two given limits from the total limit (which may have been specified or may be a default value). If all three category limits are specified, they determine the total memory allocated. However, if the total memory is also specified, it must be larger than the sum of all three categories. The code will abort if it detects an inconsistent memory specification.

The following memory directives also allocate 8 megabytes, but specify a complete partitioning as well:

```
memory total 8 stack 2 heap 2 global 4 mb
memory stack 2 heap 2 global 4 mb
```

The optional keywords `verify` and `noverify` in the directive give the user the option of enabling or disabling automatic detection of corruption of allocated memory. The default is `verify`, which enables the feature. This incurs some overhead (which can be around 10% increase in walltime on some platforms), which can be eliminated by specifying `noverify`.

The keywords `hardfail` and `nohardfail` give the user the option of forcing (or not forcing) the local memory management routines to generate an internal fatal error if any memory operation fails. The default is `nohardfail`, which allows the code to continue past any memory operation failure, and perhaps generate a more meaningful error message before terminating the calculation. Forcing a hard-fail can be useful when poorly coded applications do not check the return status of memory management routines.

When assigning the specific memory allocations using the keywords `stack`, `heap`, and `global` in the `MEMORY` directive, the user should be aware that some of the distinctions among these categories of memory have been blurred in their actual implementation in the code. The memory allocator (MA) allocates both the heap and the stack from a single memory region of size `heap+stack`, without enforcing the partition. The heap vs. stack partition is meaningful only to applications developers, and can be ignored by most users. Further complicating matters, the global array (GA) toolkit is allocated from within the MA space on distributed memory machines, while on shared-memory machines it is separate².

On distributed memory platforms, the MA region is actually the total size of

```
stack+heap+global
```

All three types of memory allocation compete for the same pool of memory, with no limits except on the total available memory. This relaxation of the memory category definitions usually benefits the user, since it can allow allocation requests to succeed where a stricter memory model would cause the directive to fail. These implementation characteristics must be kept in mind when reading program output that relates to memory usage.

Standard default for memory is currently 400 MB.

5.4 ECHO — Print input file

This start-up directive is provided as a convenient way to include a listing of the input file in the output of a calculation. It causes the entire input file to be printed to Fortran unit six (standard output). It has no keywords, arguments, or

²This is because on true shared-memory machines there is no choice but to allocate GAs from within a shared-memory segment, which is managed differently by the operating system.

options, and consists of the single line:

```
ECHO
```

The ECHO directive is processed only once, by Process 0 when the input file is read.

5.5 TITLE — Specify job title

This top-level directive allows the user to identify a job or series of jobs that use a particular database. It is an optional directive, and if omitted, the character string containing the input title will be empty. Multiple TITLE directives may appear in the input file (e.g., the example file in Section 2.3) in which case a task will use the one most recently specified. The format for the directive is as follows:

```
TITLE <string title>
```

The character string <title> is assigned to the contents of the string following the TITLE directive. If the string contains white space, it must be surrounded by double quotes. For example,

```
title "This is the title of my NWChem job"
```

The title is stored in the database and will be used in all subsequent tasks/jobs until redefined in the input.

5.6 PRINT and NOPRINT — Print control

The PRINT and NOPRINT directives allow the user to control how much output NWChem generates. These two directives are special in that the compound directives for *all* modules are supposed to recognize them. Each module can control both the overall print level (general verbosity) and the printing of individual items which are identified by name (see below). The standard form of the PRINT directive is as follows:

```
PRINT [(none || low || medium || high || debug) default medium] \  
      [<string list_of_names ... >]
```

```
NOPRINT <string list_of_names ... >
```

The default print level is medium.

Every output that is printed by NWChem has a print threshold associated with it. If this threshold is equal to or lower than the print level requested by the user, then the output is generated. For example, the threshold for printing the SCF energy at convergence is low (Section 10.16). This means that if the user-specified print level on the PRINT directive is low, medium, high, or debug, then the SCF energy will be printed at convergence.

The overall print level specified using the PRINT directive is a convenient tool for controlling the verbosity of NWChem. Setting the print level to high might be helpful in diagnosing convergence problems. The print level of debug might also be of use in evaluating problem cases, but the user should be aware that this can generate a huge amount of output. Setting the print level to low might be the preferable choice for geometry optimizations that will perform many steps which are in themselves of little interest to the user.

In addition, it is possible to enable the printing of specific items by naming them in the PRINT directive in the <list_of_names>. Items identified in this way will be printed, regardless of the overall print level specified. Similarly, the NOPRINT directive can be used to suppress the printing of specific items by naming them in its <list_of_names>. These items will *not* be printed, regardless of the overall print level, or the specific print level of the individual items.

The list of items that can be printed for each module is documented as part of the input instructions for that module. The items recognized by the top level of the code, and their thresholds, are:

Name	Print Level	Description
"total time"	medium	Print cpu and wall time at job end
"task time"	high	Print cpu and wall time for each task
"rtdb"	high	Print names of RTDB entries
"rtdbvalues"	high	Print name and values of RTDB entries
"ga summary"	medium	Summarize GA allocations at job end
"ga stats"	high	Print GA usage statistics at job end
"ma summary"	medium	Summarize MA allocations at job end
"ma stats"	high	Print MA usage statistics at job end
"version"	debug	Print version number of all compiled routines
"tcgmsg"	never	Print TCGMSG debug information

Table 5.1: Top Level Print Control Specifications

The following example shows how a PRINT directive for the top level process can be used to limit printout to only essential information. The directive is

```
print none "ma stats" rtdb
```

This directive instructs the NWChem main program to print nothing, except for the memory usage statistics (ma stats) and the names of all items stored in the database at the end of the job.

The print level within a module is inherited from the calling layer. For instance, by specifying the print to be low within the MP2 module will cause the SCF, CPHF and gradient modules when invoked from the MP2 to default to low print. Explicit user input of print thresholds overrides the inherited value.

5.7 SET — Enter data in the RTDB

This top-level directive allows the user to enter data directly into the run-time database (see Section 3.1 for a description of the database). The format of the directive is as follows:

```
SET <string name> [<string type default automatic>] <$type$ data>
```

The entry for variable <name> is the name of data to be entered into the database. This must be specified; there is no default. The variable <type>, which is optional, allows the user to define a string specifying the type of data in the array <name>. The data type can be explicitly specified as integer, real, double, logical, or string. If no entry for <type> is specified on the directive, its value is inferred from the data type of the *first* datum. In such a case, floating-point data entered using this directive must include either an exponent or a decimal point, to ensure that the correct default type will be inferred. The correct default type will be inferred for logical values if logical-true

values are specified as `.true.`, `true`, or `t`, and logical-false values are specified as `.false.`, `false`, or `f`. One exception to the automatic detection of the data type is that the data type **must** be explicitly stated to input integer ranges, unless the first element in the list is an integer that is not a range (c.f., 2.4). For example,

```
set atomid 1 3:7 21
```

will be interpreted as a list of integers. However,

```
set atomid 3:7 21
```

will not work since the first element will be interpreted as a string and not an integer. To work around this feature, use instead

```
set atomid integer 3:7 21
```

which says to write three through seven, as well as twenty-one.

The SET directive is useful for providing indirection by associating the name of a basis set or geometry with the standard object names (such as `"ao basis"` or `geometry`) used by NWChem. The following input file shows an example using the SET directive to direct different tasks to different geometries. The required input lines are as follows:

```
title "Ar dimer BSSE corrected MP2 interaction energy"
geometry "Ar+Ar"
  Ar1 0 0 0
  Ar2 0 0 2
end

geometry "Ar+ghost"
  Ar1 0 0 0
  Bq2 0 0 2
end

basis
  Ar1 library aug-cc-pvdz
  Ar2 library aug-cc-pvdz
  Bq2 library Ar aug-cc-pvdz
end

set geometry "Ar+Ar"
task mp2

scf; vectors atomic; end

set geometry "Ar+ghost"
task mp2
```

This input tells the code to perform MP2 energy calculations on an argon dimer in the first task, and then on the argon atom in the presence of the “ghost” basis of the other atom.

The SET directive can also be used as an indirect means of supplying input to a part of the code that does not have a separate input module (e.g., the atomic SCF, Section 10.5.2). Additional examples of applications of this directive can

be found in the sample input files (see Section 2.3), and its usage with basis sets (Section 7) and geometries (Section 6). Also see Section 3.1 for an example of how to store an array in the database.

5.8 UNSET — Delete data in the RTDB

This directive gives the user a way to delete simple entries from the database. The general form of the directive is as follows:

```
UNSET <string name>[*]
```

This directive cannot be used with complex objects such as geometries and basis sets³. A wild-card (*) specified at the end of the string <name> will cause *all* entries whose name begins with that string to be deleted. This is very useful as a way to reset modules to their default behavior, since modules typically store information in the database with names that begin with `module:`. For example, the SCF program can be restored to its default behavior by deleting all database entries beginning with `scf:`, using the directive

```
unset scf:*
```

Section 10.5.1 has an example using `unset` on a water dimer calculation.

The following example makes an entry in the database using the `SET` directive, and then immediately deletes it using the `UNSET` directive:

```
set mylist 1 2 3 4
unset mylist
```

5.9 STOP — Terminate processing

This top-level directive provides a convenient way of verifying an input file without actually running the calculation. It consists of the single line,

```
STOP
```

As soon as this directive is encountered, all processing ceases and the calculation terminates with an error condition.

5.10 TASK — Perform a task

The `TASK` directive is used to tell the code what to do. The input directives are parsed sequentially until a `TASK` directive is encountered, as described in Section 2.1. At that point, the calculation or operation specified in the `TASK` directive is performed. When that task is completed, the code looks for additional input to process until the next `TASK` directive is encountered, which is then executed. This process continues to the end of the input file. NWChem expects the last directive before the end-of-file to be a `TASK` directive. If it is not, a warning message is printed. Since the

³Complex objects are stored using a structured naming convention that is not matched by a simple wild card.

database is persistent, multiple tasks within one job behave *exactly* the same as multiple restart jobs with the same sequence of input.

There are four main forms of the the TASK directive. The most common form is used to tell the code at what level of theory to perform an electronic structure calculation, and which specific calculations to perform. The second form is used to specify tasks that do not involve electronic structure calculations or tasks that have not been fully implemented at all theory levels in NWChem, such as simple property evaluations. The third form is used to execute UNIX commands on machines having a Bourne shell. The fourth form is specific to combined quantum-mechanics and molecular-mechanics (QM/MM) calculations.

By default, the program terminates when a task does not complete successfully. The keyword `ignore` can be used to prevent this termination, and is recognized by all forms of the TASK directive. When a TASK directive includes the keyword `ignore`, a warning message is printed if the task fails, and code execution continues with the next task. An example of this feature is given in the sample input file in Section 11.5.

The input options, keywords, and defaults for each of these four forms for the TASK directive are discussed in the following sections.

5.10.1 TASK Directive for Electronic Structure Calculations

This is the most commonly used version of the TASK directive, and it has the following form:

```
TASK <string theory> [<string operation default energy>] [ignore]
```

The string `<theory>` specifies the level of theory to be used in the calculations for this task. NWChem currently supports ten different options. These are listed below, with the corresponding entry for the variable `<theory>`:

- `scf` — Hartree-Fock
- `dft` — Density functional theory for molecules
- `sodft` — Spin-Orbit Density functional theory
- `mp2` — MP2 using a semi-direct algorithm
- `direct_mp2` — MP2 using a full-direct algorithm
- `rimp2` — MP2 using the RI approximation
- `ccsd` — Coupled-cluster single and double excitations
- `ccsd(t)` — Coupled-cluster linearized triples approximation
- `ccsd+t(ccsd)` — Fourth order triples contribution
- `mcscf` — Multiconfiguration SCF
- `selci` — Selected configuration interaction with perturbation correction
- `md` — Classical molecular dynamics simulation
- `pspw` — Pseudopotential plane-wave density functional theory for molecules and insulating solids using NWPW
- `band` — Pseudopotential plane-wave density functional theory for solids using NWPW
- `tce` — Tensor Contraction Engine (please see Section 15.4 for a complete description of this task directive)

The string `<operation>` specifies the calculation that will be performed in the task. The default operation is a single point energy evaluation. The following list gives the selection of operations currently available in NWChem:

- `energy` — Evaluate the single point energy.
- `gradient` — Evaluate the derivative of the energy with respect to nuclear coordinates.
- `optimize` — Minimize the energy by varying the molecular structure. By default, this geometry optimization is presently driven by the Driver module (see Section 20), but the Stepper module (see Section 21) may also be used.
- `saddle` — Conduct a search for a transition state (or saddle point) using either Driver (Section 20, the default) or Stepper (Section 21).
- `hessian` — Compute second derivatives. See Section 24 for analytic Hessians.
- `frequencies` or `freq` — Compute second derivatives and print out an analysis of molecular vibrations. See Section 25 for controls for vibration calculations.
- `dynamics` — Perform classical molecular dynamics.
- `thermodynamics` — Perform multi-configuration thermodynamic integration using classical MD

NOTE: See Section 35.1 for the complete list of operations that accompany the NWPW module.

The user should be aware that some of these operations (`gradient`, `optimize`, `dynamics`, `thermodynamics`) require computation of derivatives of the energy with respect to the molecular coordinates. If analytical derivatives are not available (Section 4), they must be computed numerically, which can be very computationally intensive.

Here are some examples of the TASK directive, to illustrate the input needed to specify particular calculations with the code. To perform a single point energy evaluation using any level of theory, the directive is very simple, since the energy evaluation is the default for the string `operation`. For an SCF energy calculation, the input line is simply

```
task scf
```

Equivalently, the operation can be specified explicitly, using the directive

```
task scf energy
```

Similarly, to perform a geometry optimization using density functional theory, the TASK directive is

```
task dft optimize
```

The optional keyword `ignore` can be used to allow execution to continue even if the task fails, as discussed above. An example with the keyword `ignore` can be found in Section 11.5.

5.10.2 TASK Directive for Special Operations

This form of the TASK directive is used in instances where the task to be performed does not fit the model of the previous version (such as execution of a Python program, Section 36), or if the operation has not yet been implemented in a fashion that applies to a wide range of theories (e.g., property evaluation). Instead of requiring `theory` and `operation` as input, the directive needs only a string identifying the task. The form of the directive in such cases is as follows:

```
TASK <string task> [ignore]
```

The supported tasks that can be accessed with this form of the `TASK` directive are listed below, with the corresponding entries for string variable `<task>`.

- `python` — Execute a Python program (Section 36).
- `rtdbprint` — Print the contents of the database.
- `cphf` — Invoke the CPHF module.
- `property` — Perform miscellaneous property calculations.
- `dplot` — Execute a DPLOT run (Section 26)
- `nbo` — Execute a NBO run (Section 37.1)

This directive also recognizes the keyword `ignore`, which allows execution to continue after a task has failed.

5.10.3 TASK Directive for the Bourne Shell

This form of the `TASK` directive is supported only on machines with a fully UNIX-style operating system. This directive causes specified processes to be executed using the Bourne shell. This form of the task directive is:

```
TASK shell [( <integer-range process = 0> | all ) ] \
          <string command>
```

The keyword `shell` is required for this directive. It specifies that the given command will be executed in the Bourne shell. The user can also specify which process(es) will execute this command by entering values for `process` on the directive. The default is for only process zero to execute the command. A range of processes may be specified, using Fortran triplet notation⁴. Alternatively, all processes can be specified simply by entering the keyword `all`. The input entered for `command` must form a single string, and must consist of valid UNIX command(s). If the string includes white space, it must be enclosed in double quotes.

For example, the `TASK` directive to tell process zero to copy the molecular orbitals file to a backup location `/piofs/save` can be input as follows:

```
task shell "cp *.movecs /piofs/save"
```

The `TASK` directive to tell all processes to list the contents of their `/scratch` directories is as follows:

```
task shell all "ls -l /scratch"
```

The `TASK` directive to tell processes 0 to 10 to remove the contents of the current directory is as follows:

```
task shell 0:10:1 "/bin/rm -f *"
```

Note that NWChem's ability to quote special input characters is *very* limited when compared with that of the Bourne shell. To execute all but the simplest UNIX commands, it is usually much easier to put the shell script in a file and execute the file from within NWChem.

⁴The notation `lo:hi:inc` denotes the integers `lo`, `lo+inc`, `lo+2*inc`, ..., `hi`

5.10.4 TASK Directive for QM/MM simulations

This is very similar to the most commonly used version of the TASK directive described in Section 5.10.1, and it has the following form;

```
TASK QMMM <string theory> [<string operation default energy>] [ignore]
```

The string `<theory>` specifies the QM theory to be used in the QM/MM simulation⁵. The level of theory may be any QM method that can compute gradients but those algorithms in NWChem that do not support analytic gradients should be avoided (c.f., Section 4).

The string `<operation>` is used to specify the calculation that will be performed in the QM/MM task. The default operation is a single point energy evaluation. The following list gives the selection of operations currently available in the NWChem QM/MM module;

- `energy` — single point energy evaluation
- `optimize` — minimize the energy by variation of the molecular structure.
- `dynamics` — molecular dynamics using nwARGOS

Here are some examples of the TASK directive for QM/MM simulations. To perform a single point energy of a QM/MM system using any QM level of theory, the directive is very simple. As with the general task directive, the QM/MM energy evaluation is the default. For a DFT energy calculation the task directive input is,

```
task qmmm dft
```

or completely as

```
task qmmm dft energy
```

To do a molecular dynamics simulation of a QM/MM system using the SCF level of theory the task directive input would be

```
task qmmm scf dynamics
```

The optional keyword `ignore` can be used to allow execution to continue even if the task fails, as discussed above.

5.11 CHARGE — Total system charge

This is an optional top-level directive that allows the user to specify the total charge of the system. The form of the directive is as follows:

```
CHARGE <real charge default 0>
```

⁵If theory is “md” this is not a QM/MM simulation and will result in an appropriate error

The default charge⁶ is zero if this directive is omitted. An example of a case where the directive would be needed is for a calculation on a doubly charged cation. In such a case, the directive is simply,

```
charge 2
```

If centers with fractional charge have been specified (Section 6) the net charge of the system should be adjusted to ensure that there are an integral number of electrons.

The charge may be changed between tasks, and is used by all wavefunction types. For instance, in order to compute the first two vertical ionization energies of *LiH*, one might optimize the geometry of *LiH* using a UHF SCF wavefunction, and then perform energy calculations at the optimized geometry on *LiH*⁺ and *LiH*²⁺ in turn. This is accomplished with the following input:

```
geometry; Li 0 0 0; H 0 0 1.64; end
basis; Li library 3-21g; H library 3-21g; end

scf; uhf; singlet; end
task scf optimize

charge 1
scf; uhf; doublet; end
task scf

charge 2
scf; uhf; singlet; end
task scf
```

The GEOMETRY, BASIS, and SCF directives are described below (Sections 6, 7 and 10 respectively) but their intent should be clear. The TASK directive is described above (Section 5.10).

5.12 ECCE_PRINT — Print information for Ecce

The ECCE_PRINT directive allows the user to print out a file, usually called *ecce.out*, that will allow the calculation and its results to be imported into Ecce.

```
ECCE_PRINT <string name>
```

The entry for variable <name> is the name of the file that will contain the Ecce import information and should include the full path to the directory where you want that file. For example

```
ecce_print /home/user/job/ecce.out
```

If the full path is not given and only the file name is given, the file will be located in whatever directory the job is started in. For example, if the line

```
ecce_print ecce.out
```

⁶The charge directive, in conjunction with the charges of atomic nuclei (which can be changed via the geometry input, cf. Section 6.3), determines the total number of electrons in the chemical system. Therefore, a `charge n` specification removes "n" electrons from the chemical system. Similarly, `charge -n` adds "n" electrons.

is in the input file, the file could end up in the scratch directory if the user is using a batch script that copies the input file to a local scratch directory and then launches NWChem from there. If the system then automatically removes files in the scratch space at the end of the job, the ecce.out file will be lost. So, the best practice is to include the full path name for the file.

Chapter 6

Geometries

The GEOMETRY directive is a compound directive that allows the user to define the geometry to be used for a given calculation. The directive allows the user to specify the geometry with a relatively small amount of input, but there are a large number of optional keywords and additional subordinate directives that the user can specify, if needed. The directive therefore appears to be rather long and complicated when presented in its general form, as follows:

```
GEOMETRY [<string name default geometry>] \  
  [units <string units default angstroms>] \  
  [(angstrom_to_au || ang2au) \  
    <real angstrom_to_au default 1.8897265>] \  
  [print [xyz] || noprint] \  
  [center || nocenter] \  
  [bq|bq] \  
  [autosym [real tol default 1d-2]] \  
  [autoz || noautoz] \  
  [adjust] \  
  [(nuc || nucl || nucleus) <string nucmodel>]  
  
[SYMMETRY [group] <string group_name> [print] \  
  [tol <real tol default 1d-2>]]  
  
<string tag> <real x y z> [vx vy vz] [charge <real charge>] \  
  [mass <real mass>] \  
  [(nuc || nucl || nucleus) <string nucmodel>]  
... ]  
  
[ZMATRIX || ZMT || ZMAT  
  <string tagn> <list_of_zmatrix_variables>  
  ...  
  
[VARIABLES  
  <string symbol> <real value>  
  ... ]
```

```

[CONSTANTS
  <string symbol> <real value>
  ... ]

(END || ZEND)]

[ZCOORD
  CVR_SCALING <real value>
  BOND <integer i> <integer j> \
    [<real value>] [<string name>] [constant]
  ANGLE <integer i> <integer j> \
    [<real value>] [<string name>] [constant]
  TORSION <integer i> <integer j> <integer k> <integer l> \
    [<real value>] [<string name>] [constant]
END]

[SYSTEM surface <molecule polymer surface crystal default molecule>
  lat_a <real lat_a> lat_b <real lat_b> lat_c <real lat_c>
  alpha <real alpha> beta <real beta> gamma <real gamma>
END]

END

```

The three main parts of the GEOMETRY directive are:

- keywords on the first line of the directive (to specify such optional input as the geometry name, input units, and print level for the output)
- symmetry information
- Cartesian coordinates or Z-matrix input to specify the locations of the atoms and centers
- lattice parameters (needed only for periodic systems)

The following sections present the input for this compound directive in detail, describing the options available and the usages of the various keywords in each of the three main parts.

6.1 Keywords on the GEOMETRY directive

This section presents the options that can be specified using the keywords and optional input on the main line of the GEOMETRY directive. As described above, the first line of the directive has the general form,

```

GEOMETRY [<string name default geometry>] \
  [units <string units default angstroms>] \
  [bqbq] \
  [print [xyz] || noprint] \
  [center || nocenter] \

```

```
[autosym [real tol default 1d-2]] \
[autoz || noautoz] \
[adjust] \
[(nuc || nucl || nucleus) <string nucmodel>]
```

All of the keywords and input on this line are optional. The following list describes all options and their defaults.

- `<name>` – user-supplied name for the geometry; the default name is `geometry`, and all NWChem modules look for a geometry with this name. However, multiple geometries may be specified by using a different name for each. Subsequently, the user can direct a module to a named geometry by using the `SET` directive (see the example in Section 5.7) to associate the default name of `geometry` with the alternate name.
- `units` – keyword specifying that a value will be entered by the user for the string variable `<units>`. The default units for the geometry input are Angströms (Note: atomic units or Bohr are used within the code, regardless of the option specified for the input units. The default conversion factor used in the code to convert from Angströms to Bohr is 1.8897265 which may be overridden with the `angstrom_to_au` keyword described below.). The code recognizes the following possible values for the string variable `<units>`:
 - `angstroms` or `an` — Angstroms (Å), the default (converts to A.U. using the Å to A.U. conversion factor)
 - `au` or `atomic` or `bohr` — Atomic units (A.U.)
 - `nm` or `nanometers` — nanometers (converts to A.U. using a conversion factor computed as 10.0 times the Å to A.U. conversion factor)
 - `pm` or `picometers` — picometers (converts to A.U. using a conversion factor computed as 0.01 times the Å to A.U. conversion factor)
- `angstrom_to_au` – may also be specified as `ang2au`. This enables the user to modify the conversion factors used to convert between Å and A.U.. The default value is 1.8897265.
- `bq bq` – keyword to specify the treatment of interactions between dummy centers. The default in NWChem is to ignore such interactions when computing energies or energy derivatives. These interactions will be included if the keyword `bq bq` is specified.
- `print` and `noprint` – complementary keyword pair to enable or disable printing of the geometry. The default is to print the output associated with the geometry. In addition, the keyword `print` may be qualified by the additional keyword `xyz`, which specifies that the coordinates should be printed in the XYZ format of molecular graphics program XMol.
- `center` and `nocenter` – complementary keyword pair to enable or disable translation of the center of nuclear charge to the origin. With the origin at this position, all three components of the nuclear dipole are zero. The default is to move the center of nuclear charge to the origin.
- `autosym` – keyword to specify that the symmetry of the geometric system should be automatically determined. This option is on by default. Only groups up to and including O_h are recognized. Occasionally NWChem will be unable to determine the full symmetry of a molecular system, but will find a proper subgroup of the full symmetry. The default tolerance is set to work for most cases, but may need to be decreased to find the full symmetry of a geometry. Note that `autosym` will be turned off if the `SYMMETRY` group input is given (See section 6.2).
- `noautoz` – by default NWChem (release 3.3 and later) will generate redundant internal coordinates from user input Cartesian coordinates. The internal coordinates will be used in geometry optimizations. The `noautoz` keyword disables use of internal coordinates. The `autoz` keyword is provided only for backward compatibility. See Section 6.5 for a more detailed description of redundant internal coordinates, including how to force the definition of specific internal variables in combination with automatically generated variables.

- `adjust` – This indicates that an existing geometry is to be adjusted. Only new input for the redundant internal coordinates may be provided (Section 6.5). It is not possible to define new centers or to modify the point group using this keyword. See Section 6.5 for an example of its usage.
- `nucleus` – keyword to specify the default model for the nuclear charge distribution. The following values are recognized:
 - `point` or `pt` — point nuclear charge distribution. This is the default.
 - `finite` or `fi` — finite nuclear charge distribution with a Gaussian shape. The RMS radius of the Gaussian is determined from the nuclear mass number A by the expression $r_{\text{RMS}} = 0.836 * A^{1/3} + 0.57$ fm.

NOTE: If you specify a finite nuclear size, you should ensure that the basis set you use is contracted for a finite nuclear size. See the Section 7 for more information.

The following examples illustrate some of the various options that the user can specify on the first input line of the GEOMETRY directive, using the keywords and input options described above.

The following directives all specify the same geometry for H_2 (a bond length of 0.732556 Å):

```

geometry                               geometry units nm
  h 0 0 0                               h 0 0 0
  h 0 0 0.732556                       h 0 0 0.0732556
end                                     end

geometry units pm                       geometry units atomic
  h 0 0 0                               h 0 0 0
  h 0 0 73.2556                         h 0 0 1.3843305
end                                     end

```

6.2 SYMMETRY — Symmetry Group Input

The SYMMETRY directive is used (optionally) within the compound GEOMETRY directive to specify the point group for the molecular geometry. The general form of the directive, as described above within the general form of the GEOMETRY directive, is as follows:

```
[SYMMETRY [group] <string group_name> [print] \
  [tol <real tol default 1d-2>]]
```

The keyword `group` is optional, and can be omitted without affecting how the input for this directive is processed¹. However, if the SYMMETRY directive is used, a group name must be specified by supplying an entry for the string variable `<group_name>`. The group name should be specified as the standard Schöflies symbol. Examples of expected input for the variable `group_name` include such entries as:

- `c2v` – for molecular symmetry C_{2v}
- `d2h` – for molecular symmetry D_{2h}
- `Td` – for molecular symmetry T_d

¹For periodic systems, there are additional keywords within this directive (not yet documented), so having a keyword for the group name is useful.

- d6h – for molecular symmetry D_{6h}

The SYMMETRY directive is optional. The default is no symmetry (i.e., C_1 point group). Automatic detection of point group symmetry is available through the use of `autosym` in the GEOMETRY directive main line (discussed in Section 6.1). Note: if the SYMMETRY directive is present the `autosym` keyword is ignored.

If only symmetry-unique atoms are specified, the others will be generated through the action of the point group operators, but the user is free to specify all atoms. The user must know the symmetry of the molecule being modeled, and be able to specify the coordinates of the atoms in a suitable orientation relative to the rotation axes and planes of symmetry. Appendix C lists a number of examples of the GEOMETRY directive input for specific molecules having symmetry patterns recognized by NWChem. The exact point group symmetry will be forced upon the molecule, and atoms within 10^{-3} A.U. of a symmetry element (e.g., a mirror plane or rotation axis) will be forced onto that element. Thus, it is not necessary to specify to a high precision those coordinates that are determined solely by symmetry.

The keyword `print` gives information concerning the point group generation, including the group generators, a character table, the mapping of centers, and the group operations.

The keyword `tol` relates to the accuracy with which the symmetry-unique atoms should be specified. When the atoms are generated, those that are within the tolerance, `tol`, are considered the same.

6.3 Cartesian coordinate input

The default in NWChem is to specify the geometry information entirely in Cartesian coordinates, and examples of this format have appeared above (e.g. Section 2.3). Each center (usually an atom) is identified on a line of the following form:

```
<string tag> <real x y z> [vx vy vz] \
  [charge <real charge>] [mass <real mass>] \
  [(nuc || nucl || nucleus) <string nucmodel>]
```

The string `<tag>` is the name of the atom or center, and its case (upper or lower) is important. The tag is limited to 16 characters and is interpreted as follows:

- If the entry for `<tag>` begins with either the symbol or name of an element (regardless of case), then the center is treated as an atom of that type. The default charge is the atomic number (adjusted for the presence of ECPs by the ECP `NELEC` directive ; see Section 8). Additional characters can be added to the string, to distinguish between atoms of the same element (For example, the tags `oxygen`, `O`, `o34`, `olonepair`, and `Oxygen-ether`, will all be interpreted as oxygen atoms).
- If the entry for `<tag>` begins with the characters `bq` or `x` (regardless of case), then the center is treated as a dummy center with a default zero charge (Note: a tag beginning with the characters `xe` will be interpreted as a xenon atom rather than as a dummy center.). Dummy centers may optionally have basis functions or non-zero charge. See Section B.2 for a sample input using dummy centers with charges.

It is *important* to be aware of the following points regarding the definitions and usage of the values specified for the variable `<tag>` to describe the centers in a system:

- If the tag begins with characters that cannot be matched against an atom, and those characters are not `BQ` or `X`, then a fatal error is generated.

- The tag of a center is used in the BASIS (Section 7) and ECP (Section 8) directives to associate functions with centers.
- All centers with the same tag will have the same basis functions.
- When using automatic symmetry detection, only centers with the same tag will be candidates for testing for symmetry equivalence.
- The user-specified charges (of all centers, atomic and dummy) and any net total charge of the system (Section 5.11) are used to determine the number of electrons in the system.

The Cartesian coordinates of the atom in the molecule are specified as real numbers supplied for the variables x , y , and z following the characters entered for the tag. The values supplied for the coordinates must be in the units specified by the value of the variable `<units>` on the first line of the GEOMETRY directive input.

After the Cartesian coordinate input, optional velocities may be entered as real numbers for the variables v_x , v_y , and v_z . The velocities should be given in atomic units and are used in QMD and PSPW calculations.

The Cartesian coordinate input line also contains the optional keywords `charge`, `mass` and `nucleus`, which allow the user to specify the charge of the atom (or center) and its mass (in atomic mass units), and the nuclear model. The default charge for an atom is its atomic number, adjusted for the presence of ECPs (see Section 8). In order to specify a different value for the charge on a particular atom, the user must enter the keyword `charge`, followed by the desired value for the variable `<charge>`.

The default mass for an atom is taken to be the mass of its most abundant naturally occurring isotope or of the isotope with the longest half-life. To model some other isotope of the element, its mass must be defined explicitly by specifying the keyword `mass`, followed by the value (in atomic mass units) for the variable `<mass>`.

The default nuclear model is a point nucleus. The keyword `nucleus` (or `nucl` or `nuc`) followed by the model name `<nucmodel>` overrides this default. Allowed values of `<nucmodel>` are `point` or `pt` and `finite` or `fi`. The `finite` option is a nuclear model with a Gaussian shape. The RMS radius of the Gaussian is determined by the atomic mass number via the formula $r_{\text{RMS}} = 0.836 * A^{1/3} + 0.57$ fm. The mass number A is derived from the variable `<mass>`.

The geometry of the system can be specified entirely in Cartesian coordinates by supplying a `<tag>` line of the type described above for each atom or center. The user has the option, however, of supplying the geometry of some or all of the atoms or centers using a Z-matrix description. In such a case, the user supplies the input tag line described above for any centers to be described by Cartesian coordinates, and then specifies the remainder of the system using the optional ZMATRIX directive described below in Section 6.4.

6.4 ZMATRIX — Z-matrix input

The ZMATRIX directive is an optional directive that can be used within the compound GEOMETRY directive to specify the structure of the system with a Z-matrix, which can include both internal and Cartesian coordinates. The ZMATRIX directive is itself a compound directive that can include the VARIABLES and CONSTANTS directives, depending on the options selected. The general form of the compound ZMATRIX directive is as follows:

```
[ZMATRIX || ZMT || ZMAT
  <string tagn> <list_of_zmatrix_variables>
  ...
  [VARIABLES
    <string symbol> <real value>
```

```

    ... ]

[CONSTANTS
  <string symbol> <real value>
  ... ]

(END || ZEND)]

```

The input module recognizes three possible spellings of this directive name. It can be invoked with ZMATRIX, ZMT, or ZMAT. The user can specify the molecular structure using either Cartesian coordinates or internal coordinates (bond lengths, bond angles and dihedral angles. The Z-matrix input for a center defines connectivity, bond length, and bond or torsion angles. Cartesian coordinate input for a center consists of three real numbers defining the x,y,z coordinates of the atom.

Within the Z-matrix input, bond lengths and Cartesian coordinates must be input in the user-specified units, as defined by the value specified for the variable <units> on the first line of the GEOMETRY directive. All angles are specified in degrees.

The individual centers (denoted as i, j, and k below) used to specify Z-matrix connectivity may be designated either as integers (identifying each center by number) or as tags (*If tags are used, the tag must be unique for each center.*) The use of “dummy” atoms is possible, by using X or BQ at the start of the tag.

Bond lengths, bond angles and dihedral angles (denoted below as R, alpha, and beta, respectively) may be specified either as numerical values or as symbolic strings that must be subsequently defined using the VARIABLES or CONSTANTS directives. The numerical values of the symbolic strings labeled VARIABLES may be subject to changes during a geometry optimization say, while the numerical values of the symbolic strings labeled CONSTANTS will stay frozen to the value given in the input. The same symbolic string can be used more than once, and any mixture of numeric data and symbols is acceptable. Bond angles (α) must be in the range $0 < \alpha < 180$.

The Z-matrix input is specified sequentially as follows:

```

tag1
tag2 i R
tag3 i R j alpha
tag4 i R j alpha k beta [orient]
...

```

The structure of this input is described in more detail below. In the following discussion, the tag or number of the center being currently defined is labeled as C (“C” for current). The values entered for these tags for centers defined in the Z-matrix input are interpreted in the same way as the <tag> entries for Cartesian coordinates described above (see Section 6.3). Figures 6.1, 6.2 and 6.3 display the relationships between the input data and the definitions of centers and angles.

The Z-matrix input shown above is interpreted as follows:


1. tag1

Only a tag is required for the first center.

2. tag2 i R


The second center requires specification of its tag and the bond length (R_{Ci}) distance to a previous atom, which is identified by i.

3. tag3 i R j alpha




Input line: C i R j alpha k beta

Figure 6.1: Relationships between the centers, bond angle and dihedral angle in Z-matrix input.



Input line: C i R j alpha k beta +1

Figure 6.2: Relationships between the centers and two bond angles in Z-matrix input with optional parameter specified as +1.



Input line: C i R j α k beta -1

Figure 6.3: Relationships between the centers and two bond angles in Z-matrix input with optional parameter specified as -1 .

The third center requires specification of its tag, its bond length distance (R_{Ci}) to one of the two previous centers (identified by the value of i), and the bond angle $\alpha = \widehat{Cij}$.

4. tag i R j alpha k beta [<integer orient default 0>]

The fourth, and all subsequent centers, require the tag, a bond length (R_{Ci}) relative to center i , the bond angle with centers i and j ($\alpha = \widehat{Cij}$), and *either*

- (a) the dihedral angle (β) between the current center and centers i , j , and k (Figure 6.1), or
- (b) a second bond angle $\beta = \widehat{Cik}$ and an orientation to the plane containing the other three centers (Figure 6.2 and 6.3).

By default, β is interpreted as a dihedral angle (see Figure 6.1), but if the optional final parameter (<orient>) is specified with the value ± 1 , then β is interpreted as the angle \widehat{Cik} . The sign of <orient> specifies the direction of the bond angle relative to the plane containing the three reference atoms. If <orient> is $+1$, then the new center (C) is above the plane (Figure 6.2); and if <orient> is -1 , then C is below the plane (Figure 6.3).

Following the Z-matrix center definitions described above, the user can specify initial values for any symbolic variables used to define the Z-matrix tags. This is done using the optional VARIABLES directive, which has the general form:

```
VARIABLES
  <string symbol> <real value>
  ...
```

Each line contains the name of a variable followed by its value. Optionally, an equals sign (=) can be included between the symbol and its value, for clarity in reading the input file.

Following the VARIABLES directive, the CONSTANTS directive may be used to define any Z-matrix symbolic variables that remain unchanged during geometry optimizations. To freeze the Cartesian coordinates of an atom, refer to Section 6.6. The general form of this directive is as follows:

```
CONSTANTS
  <string symbol> <real value>
  ...
```

Each line contains the name of a variable followed by its value. As with the VARIABLES directive, an equals sign (=) can be included between the symbol and its value.

The end of the Z-matrix input using the compound ZMATRIX directive is signaled by a line containing either END or ZEND, following all input for the directive itself and its associated optional directives.

A simple example is presented for water. All Z-matrix parameters are specified numerically, and symbolic tags are used to specify connectivity information. This requires that all tags be unique, and therefore different tags are used for the two hydrogen atoms, which may or may not be identical.

```
geometry
  zmatrix
    O
    H1 O 0.95
    H2 O 0.95 H1 108.0
  end
end
```

The following example illustrates the Z-matrix input for the molecule CH_3CF_3 . This input uses the numbers of centers to specify the connectivity information (i, j, and k), and uses symbolic variables for the Z-matrix parameters R, alpha, and beta, which are defined in the inputs for the VARIABLES and CONSTANTS directives.

```
geometry
  zmatrix
    C
    C 1 CC
    H 1 CH1 2 HCH1
    H 1 CH2 2 HCH2 3 TOR1 0
    H 1 CH3 2 HCH3 3 -TOR2 0
    F 2 CF1 1 CCF1 3 TOR3 0
    F 2 CF2 1 CCF2 6 FCH1 1
    F 2 CF3 1 CCF3 6 FCH2 -1
  variables
    CC 1.4888
    CH1 1.0790
    CH2 1.0789
    CH3 1.0789
    CF1 1.3667
    CF2 1.3669
    CF3 1.3669
  constants
```

```

HCH1  104.28
HCH2  104.74
HCH3  104.7
CCF1  112.0713
CCF2  112.0341
CCF3  112.0340
TOR1  109.3996
TOR2  109.3997
TOR3  180.0000
FCH1  106.7846
FCH2  106.7842
end
end

```

The input for any centers specified with Cartesian coordinates must be specified using the format of the <tag> lines described in Section 6.3 above. However, in order to correctly specify these Cartesian coordinates within the Z-matrix, the user must understand the orientation of centers specified using internal coordinates. These are arranged as follows:

- The first center is placed at the origin.
- The second center is placed along the positive z-axis.
- The third center is placed in the z-x plane.

6.5 ZCOORD — Forcing internal coordinates

By default redundant internal coordinates are generated for use in geometry optimizations. Connectivity is inferred by comparing inter-atomic distances with the sum of the van der Waals radii of the two atoms involved in a possible bond, times a scaling factor. The scaling factor is an input parameter of ZCOORD which maybe changed from its default value of 1.3. Under some circumstances (unusual bonding, bond dissociation, ...) it will be necessary to augment the automatically generated list of internal coordinates to force some specific internal coordinates to be included in among the internal coordinates. This is accomplished by including the optional directive ZCOORD within the geometry directive. The general form of the ZCOORD directive is as follows:

```

ZCOORD
  CVR_SCALING <real value>
  BOND      <integer i> <integer j> \
            [<real value>] [<string name>] [constant]
  ANGLE     <integer i> <integer j> <integer k> \
            [<real value>] [<string name>] [constant]
  TORSION   <integer i> <integer j> <integer k> <integer l> \
            [<real value>] [<string name>] [constant]
END

```

The centers *i*, *j*, *k* and *l* *must* be specified using the numbers of the centers, as supplied in the input for the Cartesian coordinates. The ZCOORD input parameters are defined as follows:

- `cvr_scaling` — scaling factor applied to van der Waals radii.

- `bond` — a bond between the two centers.
- `angle` — a bond angle \widehat{ijk} .
- `torsion` — a torsion (or dihedral) angle. The angle between the planes $i-j-k$ and $j-k-l$.

A value may be specified for a user-defined internal coordinate, in which case it is forced upon the input Cartesian coordinates while attempting to make only small changes in the other internal coordinates. If no value is provided the value implicit in the input coordinates is kept. If the keyword `constant` is specified, then that internal variable is not modified during a geometry optimization with `DRIVER` (Section 20). Each internal coordinate may also be named either for easy identification in the output, or for the application of constraints (Section 6.6).

If the keyword `adjust` is specified on the main `GEOMETRY` directive, only `ZCOORD` data may be specified and it can be used to change the user-defined internal coordinates, including adding/removing constraints and changing their values.

6.6 Applying constraints in geometry optimizations

Internal coordinates specified as `constant` in a `ZCOORD` directive or in the constants section of a `ZMATRIX` directive, will be frozen at their initial values if a geometry optimization is performed with `DRIVER` (Section 20).

If internal coordinates have the same name (give or take an optional sign for torsions) then they are forced to have the same value. This may be used to force bonds or angles to be equal even if they are not related by symmetry.

When atoms have been specified by their Cartesian coordinates, *and* internal coordinates are not being used, it is possible to freeze the cartesian position of selected atoms. This is useful for such purposes as optimizing a molecule absorbed on the surface of a cluster with fixed geometry. Only the gradients associated with the active atoms are computed. This can result in a big computational saving, since gradients associated with frozen atoms are forced to zero (Note, however, that this destroys the translational and rotational invariance of the gradient. This is not yet fully accommodated by the `STEPPER` geometry optimization software, and can sometimes result in slower convergence of the optimization. The `DRIVER` optimization package does not suffer from this problem).

The `SET` directive (Section 5.7) is used to freeze atoms, by specifying a directive of the form:

```
set geometry:actlist <integer list_of_center_numbers>
```

This defines only the centers in the list as active. All other centers will have zero force assigned to them, and will remain frozen at their starting coordinates during a geometry optimization.

For example, the following directive specifies that atoms numbered 1, 5, 6, 7, 8, and 15 are active and all other atoms are frozen:

```
set geometry:actlist 1 5:8 15
```

or equivalently,

```
set geometry:actlist 1 5 6 7 8 15
```

If this option is not specified by entering a `SET` directive, the default behavior in the code is to treat all atoms as active. To revert to this default behavior after the option to define frozen atoms has been invoked, the `UNSET` directive must be used (since the database is persistent, see Section 3.2). The form of the `UNSET` directive is as follows:

```
unset geometry:actlist
```

6.7 SYSTEM — Lattice parameters for periodic systems

This keyword is needed only for 1-, 2-, and 3-dimensional periodic systems.

The `system` keyword can assume the following values

- `polymer` — system with 1-d translational symmetry.
- `surface` — system with 2-d translational symmetry.
- `crystal` — system with 3-d translational symmetry.
- `molecule` — no translational symmetry (this is the default)

When the system possess translational symmetry, **fractional** coordinates are used in the directions where translational symmetry exists. This means that for crystals x , y and z are fractional, for surfaces x and y are fractional, whereas for polymers only z is fractional. For example, in the following H₂O layer input (a 2-d periodic system), x and y coordinates are fractional, whereas z is expressed in Å.

```
geometry units angstrom
O      0.353553    0.353553        2.100000000
H      0.263094    0.353553        2.663590000
H      0.444007    0.353553        2.663590000
```

Since no space group symmetry is available yet other than $P1$, input of cell parameters is relative to the primitive cell. For example, this is the input required for the cubic face-centered type structure of bulk MgO.

```
system crystal
lat_a 2.97692 lat_b 2.97692 lat_c 2.97692
alpha 60.00 beta 60.00 gamma 60.00
end
```


Chapter 7

Basis sets

NWChem currently supports basis sets consisting of generally contracted¹ Cartesian Gaussian functions up to a maximum angular momentum of six (*h* functions), and also *sp* (or *L*) functions². The `BASIS` directive is used to define these, and also to specify use of an effective core potential (ECP) that is associated with a basis set; see Section 8.

The basis functions to be used for a given calculation can be drawn from a standard set in the EMSL basis set library that is included in the release of NWChem (See Appendix A for a list of the standard basis sets currently supplied with the release of the code). Alternatively, the user can specify particular functions explicitly in the input, to define a particular basis set.

The general form of the `BASIS` directive is as follows:

```
BASIS [<string name default "ao basis">] \  
      [(spherical || cartesian) default cartesian] \  
      [(segment || nosegment) default segment] \  
      [(print || noprint) default print] \  
      [rel] \  
  
      <string tag> library [<string tag_in_lib>] \  
                    <string standard_set> [file <filename>] \  
                    [except <string tag list>] [rel] \  
  
      ... \  
  
      <string tag> <string shell_type> [rel] \  
                <real exponent> <real list_of_coefficients> \  
                ... \  
  
END
```

Examining the keywords on the first line of the `BASIS` directive:

- name

¹Generally contracted meaning that the same primitive, Gaussian functions are contracted into multiple contracted functions using different contraction coefficients. Reuse of the radial functions increases the efficiency of integral generation.

²An *sp* shell is two-component general contraction. However, the first component specifies an *s* shell and the second a *p* shell. Again, reuse of the radial functions increases the efficiency of integral generation.

By default, the basis set is stored in the database with the name "ao basis". Another name may be specified in the `BASIS` directive, thus, multiple basis sets may be stored simultaneously in the database. Also, the DFT (Section 11) and RI-MP2 (Section 16) modules and the Dyll-modified-Dirac relativistic method (Section 9.2) require multiple basis sets with specific names.

The user can associate the "ao basis" with another named basis using the `SET` directive (see Section 5.7).

- `SPHERICAL` or `CARTESIAN`

The keywords `spherical` and `cartesian` offer the option of using either spherical-harmonic (5 d, 7 f, 9 g, ...) or Cartesian (6 d, 10 f, 15 g, ...) angular functions. The default is Cartesian.

Note that the correlation-consistent basis sets were designed using spherical harmonics and to use these, the `spherical` keyword should be present in the `BASIS` directive. The use of spherical functions also helps eliminate problems with linear dependence.

- `SEGMENT` or `NOSEGMENT`

By default, NWChem forces all basis sets to be segmented, even if they are input with general contractions or *L* or *sp* shells. This is because the current derivative integral program cannot handle general contractions. If a calculation is computing energies only, a performance gain can result from exploiting generally contracted basis sets, in which case `NOSEGMENT` should be specified.

- `PRINT` or `NOPRINT`

The default is for the input module to print all basis sets encountered. Specifying the keyword `noprnt` allows the user to suppress this output.

- `REL`

This keyword marks the entire basis as a relativistic basis for the purposes of the Dyll-modified-Dirac relativistic integral code. The marking of the basis set is necessary for the code to make the proper association between the relativistic shells in the ao basis and the shells in the large and/or small component basis. This is only necessary for basis sets which are to be used as the ao basis. The user is referred to Section 9.2 for more details.

Basis sets are associated with centers by using the tag of a center in a geometry that has either been input by the user (Section 6) or is available elsewhere. Each atom or center with the same `tag` will have the same basis set. All atoms must have basis functions assigned to them — only dummy centers (X or Bq) may have no basis functions. To facilitate the specification of the geometry and the basis set for any chemical system, the matching process of a basis set tag to a geometry tag first looks for an exact match. If no match is found, NWChem will attempt to match, ignoring case, the name or symbol of the element. E.g., all hydrogen atoms in a system could be labeled "H1", "H2", ..., in the geometry but only one basis set specification for "H" or "hydrogen" is necessary. If desired, a special basis may be added to one or more centers (e.g., "H1") by providing a basis for that tag. If the matching mechanism fails then NWChem stops with an appropriate error message.

A special set of tags, "*" and tags ending with a "*" (E.g. "H*") can be used in combination with the keyword `library` (see section below). These tags facilitate the definition of a certain type of basis set of all atoms, or a group of atoms, in a geometry using only a single or very few basis set entries. The "*" tag will not place basis sets on dummy atoms, Bq* can be used for that if necessary.

Examined next is how to reference standard basis sets in the basis set library, and finally, how to define a basis set using exponents and coefficients.

7.1 Basis set library

The keyword `library` associated with each specific `tag` entry specifies that the calculation will use the standard basis set in NWChem for that center. The variable `<standard_set>` is the name that identifies the functions in the library. The names of standard basis sets are not case sensitive. See Appendix A for a complete list of the basis sets in the NWChem library and their specifications.

The general form of the input line requesting basis sets from the NWChem basis set library is:

```
<string tag> library [<string tag_in_lib>] \  
                <string standard set> [file < filename> \  
                [except <string tag list>] [rel]  
...  

```

For example, the NWChem basis set library contains the Dunning cc-pvdz basis set. These may be used as follows

```
basis  
  oxygen library cc-pvdz  
  hydrogen library cc-pvdz  
end
```

A default path of the NWChem basis set libraries is provided on installation of the code, but a different path can be defined by specifying the keyword `file`, and one can explicitly name the file to be accessed for the basis functions. For example,

```
basis  
  o library 3-21g file /usr/d3g681/nwchem/library  
  si library 6-31g file /usr/d3g681/nwchem/libraries/  
end
```

This directive tells the code to use the basis set 3-21g in the file `/usr/d3g681/nwchem/library` for atom `o` and to use the basis set 6-31g in the directory `/usr/d3g681/nwchem/libraries/` for atom `si`, rather than look for them in the default libraries. When a directory is defined the code will search for the basis set in a file with the name 6-31g.

The “*” tag can be used to efficiently define basis set input directives for large numbers of atoms. An example is:

```
basis  
  * library 3-21g  
end
```

This directive tells the code to assign the basis sets 3-21g to all the atom tags defined in the geometry. If one wants to place a different basis set on one of the atoms defined in the geometry, the following directive can be used:

```
basis  
  * library 3-21g except H  
end
```

This directive tells the code to assign the basis sets 3-21g to all the atoms in the geometry, except the hydrogen atoms. Remember that the user will have to explicitly define the hydrogen basis set in this directive! One may also define tags that end with a “*”:

```
basis
  oxy* library 3-21g
end
```

This directive tells the code to assign the basis sets 3-21g to all atom tags in the geometry that start with “oxy”.

If standard basis sets are to be placed upon a dummy center, the variable `<tag_in_lib>` must also be entered on this line, to identify the correct atom type to use from the basis function library (see the ghost atom example in Section 5.7 and below). For example: To specify the cc-pvdz basis for a calculation on the water monomer in the dimer basis, where the dummy oxygen and dummy hydrogen centers have been identified as `bqo` and `bqh` respectively, the BASIS directive is as follows:

```
basis
  o library cc-pvdz
  h library cc-pvdz
  bqo library o cc-pvdz
  bqh library h cc-pvdz
end
```

A special dummy center tag is `bq*`, which will assign the same basis set to all `bq` centers in the geometry. Just as with the “*” tag, the `except` list can be used to assign basis sets to unique dummy centers.

The library basis sets can also be marked as relativistic by adding the `rel` keyword to the tag line. See Section 9.2 for more details. The correlation consistent basis sets have been contracted for relativistic effects and are included in the standard library.

There are also contractions in the standard library for both a point nucleus and a finite nucleus of Gaussian shape. These are usually distinguished by the suffixes `_pt` and `_fi`. It is the user’s responsibility to ensure that the contraction matches the nuclear type specified in the geometry object. The specification of a finite nucleus basis set does NOT automatically set the nuclear type for that atom to be finite. See Section 6 for information.

7.2 Explicit basis set definition

If the basis sets in the library or available in other external files are not suitable for a given calculation, the basis set may be explicitly defined. A generally contracted Gaussian basis function is associated with a center using an input line of the following form:

```
<string tag> <string shell_type> [rel]
  <real exponent> <real list_of_coefficients>
  ...
```

The variable `<shell_type>` identifies the angular momentum of the shell, s, p, d, \dots . NWChem is configured to handle up to h shells. The keyword `rel` marks the shell as relativistic — see Section 9.2 for more details. Subsequent lines define the primitive function exponents and contraction coefficients. General contractions are specified by including multiple columns of coefficients.

The following example defines basis sets for the water molecule:

```
basis spherical nosegment
  oxygen s
```

```

11720.0000    0.000710   -0.000160
 1759.0000    0.005470   -0.001263
  400.8000    0.027837   -0.006267
 113.7000     0.104800   -0.025716
  37.0300     0.283062   -0.070924
 13.2700      0.448719   -0.165411
  5.0250      0.270952   -0.116955
  1.0130      0.015458    0.557368
  0.3023     -0.002585    0.572759
oxygen s
  0.3023      1.000000
oxygen p
 17.7000      0.043018
  3.8540      0.228913
  1.0460      0.508728
  0.2753      0.460531
oxygen p
  0.2753      1.000000
oxygen d
  1.1850      1.000000
hydrogen s
 13.0100      0.019685
  1.9620      0.137977
  0.4446      0.478148
  0.1220      0.501240
hydrogen s
  0.1220      1.000000
hydrogen p
  0.7270      1.000000
oxygen s
  0.01         1.0
hydrogen s
  0.02974     1.0
hydrogen p
  0.141       1.0
end

```

7.3 Combinations of library and explicit basis set input

The user can use a mixture of library basis and explicit basis set input to define the basis sets used on the various atoms.

For example, the following BASIS directive augments the Dunning cc-pvdz basis set for the water molecule with a diffuse s-shell on oxygen and adds the aug-cc-pVDZ diffuse functions onto the hydrogen.

```

basis spherical nosegment
oxygen library cc-pvdz
hydrogen library cc-pvdz
oxygen s
  0.01 1.0

```

```
hydrogen library "aug-cc-pVDZ Diffuse"  
end
```

The resulting basis set defined is identical to the one defined above in the explicit basis set input.

Chapter 8

Effective Core Potentials

Effective core potentials (ECPs) are a useful means of replacing the core electrons in a calculation with an effective potential, thereby eliminating the need for the core basis functions, which usually require a large set of Gaussians to describe them. In addition to replacing the core, they may be used to represent relativistic effects, which are largely confined to the core. In this context, both the scalar (spin-free) relativistic effects and spin-orbit (spin-dependent) relativistic effects may be included in effective potentials. NWChem has the facility to use both, and these are described in the next two sections.

A brief recapitulation of the development of RECPs is given here, following Pacios and Christiansen¹. The process can be viewed as starting from an atomic Dirac-Hartree-Fock calculation, done in *jj* coupling, and producing relativistic effective potentials (REPs) for each *l* and *j* value, U_{lj}^{REP} . From these, a local potential is extracted, which for example contains the Coulomb potential of the core electrons balanced by the part of the nuclear attraction which cancels the core electron charge. The residue is expressed in a semi-local form,

$$U^{\text{REP}} = U_{Lj}^{\text{REP}}(r) + \sum_{l=0}^{L-1} \sum_{j=|l-1/2}^{l+1/2} [U_{lj}^{\text{REP}}(r) - U_{Lj}^{\text{REP}}(r)] \sum_m |ljm\rangle \langle ljm| \quad (8.1)$$

where *L* is one larger than the maximum angular momentum in the atom. The scalar potential is obtained by averaging the REPs for each *j* for a given *l* to give an averaged relativistic effective potential, or AREP,

$$U_l^{\text{AREP}}(r) = \frac{1}{2l+1} [lU_{l-1/2}^{\text{REP}}(r) + (l+1)U_{l+1/2}^{\text{REP}}(r)]. \quad (8.2)$$

These are summed into the full potential.

The spin-orbit potential is obtained from the difference between the REPs for the two *j* values for a given *l*, and may be represented in terms of an effective spin-orbit operator,

$$H^{\text{SO}} = \mathbf{s} \cdot \sum_{l=1}^{L-1} \frac{2}{2l+1} \Delta U_l^{\text{REP}} \sum_{mm'} |lm\rangle \langle lm|\hat{l}|lm'\rangle \langle lm'|. \quad (8.3)$$

where

$$\Delta U_l^{\text{REP}} = U_{l+1/2}^{\text{REP}}(r) - U_{l-1/2}^{\text{REP}}(r). \quad (8.4)$$

¹ F. Pacios and P. A. Christiansen, J. Chem. Phys. **82**, 2664 (1985)

The spin-orbit integrals generated by NWChem are the integrals over the sum, including the factor of $2/(2l+1)$, so that they may be treated as an effective spin-orbit operator without further factors introduced.

The effective potentials, both scalar and spin-orbit, are fitted to Gaussians with the form

$$r^2 U_l(r) = \sum_k A_{lk} r^{n_{lk}} e^{-B_{lk} r^2}$$

where A_{lk} is the contraction coefficient, n_{lk} is the exponent of the “r” term (r-exponent), and B_{lk} is the Gaussian exponent. The n_{lk} is shifted by 2, in accordance with most of the ECP literature and implementations, i.e., an $n_{lk} = 0$ implies r^{-2} . The current implementation allows n_{lk} values of only 0, 1, or 2.

8.1 Scalar ECPs

The optional directive ECP allows the user to describe an effective core potential (ECP) in terms of contracted Gaussian functions as given above. Potentials using these functions must be specified explicitly by user input in the ECP directive. This directive has essentially the same form and properties as the standard BASIS directive, except for essential differences required for ECPs. Because of this, the ECP is treated internally as a basis set. The form of the input for the ECP directive is as follows:

```
ECP [<string name default "ecp basis">] \
    [print || noprint default print]

    <string tag> library [<string tag_in_lib>] \
        <string standard_set> [file <filename>] \
        [except <string tag list>]

    <string tag> [nelec] <integer number_of_electrons_replaced>

    ...

    <string tag> <string shell_type>
    <real r-exponent> <real Gaussian-exponent> <real list_of_coefficients>
    ...

END
```

ECPs are automatically segmented, even if general contractions are input. The projection operators defined in an ECP are spherical by default, so there is no need to include the CARTESIAN or SPHERICAL keyword as there is for a standard basis set. ECPs are associated with centers in geometries through tags or names of centers. These tags must match in the same manner as for basis sets the tags in a GEOMETRY and ECP directives, and are limited to sixteen (16) characters. Each center with the same tag will have the same ECP. By default, the input module prints each ECP that it encounters. The NOPRINT option can be used to disable printing. There can be only one active ECP, even though several may exist in the input deck. The ECP modules load “ecp basis” inputs along with any “ao basis” inputs present. ECPs may be used in both energy and gradient calculations.

ECPs are named in the same fashion as geometries or regular basis sets, with the default name being “ecp basis”. It should be clear from the above discussion on geometries and database entries how indirection is supported. All directives that are in common with the standard Gaussian basis set input have the same function and syntax.

As for regular basis sets, ECPs may be obtained from the standard library. The names of the sets of ECPs available in the standard library (their coverage is described in Appendix A) are

- "Hay-Wadt MB (n+1) ECP"
- "Hay-Wadt VDZ (n+1) ECP"
- "LANL2DZ ECP"
- "SBKJC VDZ ECP"
- "Stuttgart RLC ECP"
- "Stuttgart RSC ECP"
- "CRENBL ECP"
- "CRENBS ECP"

The keyword `nelec` allows the user to specify the number of core electrons replaced by the ECP. Additional input lines define the specific coefficients and exponents. The variable `<shell_type>` is used to specify the components of the ECP. The keyword `ul` entered for `<shell_type>` denotes the local part of the ECP. This is equivalent to the highest angular momentum functions specified in the literature for most ECPs. The standard entries (`s`, `p`, `d`, etc.) for `shell_type` specify the angular momentum projector onto the local function. The shell type label of `s` indicates the `ul-s` projector input, `p` indicates the `ul-p`, etc.

For example, the Christiansen, Ross and Ermler ARECPs are available in the standard basis set library named `{crenbl_ecp}`. To perform a calculation on uranyl (UO_2^{2+}) with all-electron oxygen (aug-cc-pvdz basis), and uranium with an ARECP and using the corresponding basis the following input can be used

```
geometry
  U 0 0 0
  O 0 0 1.65
  O 0 0 -1.65
end
basis
  U library crenbl_ecp
  O library aug-cc-pvdz
end
ecp
  U library crenbl_ecp
end
```

The following is an example of explicit input of an ECP for H_2CO . It defines an ECP for the carbon and oxygen atoms in the molecule.

```
ecp
  C nelec 2 # ecp replaces 2 electrons on C
  C ul      # d
      1      80.0000000      -1.60000000
      1      30.0000000      -0.40000000
      2      0.5498205       -0.03990210
  C s      # s - d
```

```

      0      0.7374760      0.63810832
      0      135.2354832      11.00916230
      2      8.5605569      20.13797020
C p      # p - d
      2      10.6863587      -3.24684280
      2      23.4979897      0.78505765
O nelec 2 # ecp replaces 2 electrons on O
O ul      # d
      1      80.0000000      -1.60000000
      1      30.0000000      -0.40000000
      2      1.0953760      -0.06623814
O s      # s - d
      0      0.9212952      0.39552179
      0      28.6481971      2.51654843
      2      9.3033500      17.04478500
O p      # p - s
      2      52.3427019      27.97790770
      2      30.7220233      -16.49630500
end

```

8.2 Spin-orbit ECPs

The Spin-orbit ECPs can be used with the Density Functional Approach, but one has to run the calculations without symmetry. Note: when a Hartree-Fock method is specified the spin-orbit input will be ignored.

Spin-orbit ECPs are fitted in precisely the same functional form as the scalar RECPs and have the same properties, with the exception that there is no local potential ul, no *s* potential and no effective charge has to be defined. Spin-orbit potentials are specified in the same way as ECPs except that the directive SO is used instead of ECP. Note that there currently are no spin-orbit ECPs defined in the standard NWChem library. The SO directive is as follows:

```

SO [<string name default "so basis">] \
  [print || noprint default print]

  <string tag> library [<string tag_in_lib>] \
    <string standard_set> [file <filename>]
    [except <string tag list>]
  ...

  <string tag> <string shell_type>
  <real r-exponent> <real Gaussian-exponent> <real list_of_coefficients>
  ...

END

```

Note: in the literature the coefficients of the spin-orbit potentials are NOT always defined in the same manner. The NWChem code assumes that the spin-orbit potential defined in the input is of the form:

$$\Delta U_l^{\text{NWChem}} = \frac{2}{2l+1} \Delta U_l \quad (8.5)$$

For example, in the literature the Stuttgart potentials are defined as ΔU_l and, hence, have to be multiplied by $2/(2l+1)$. On the other hand, the CRENBL potentials in the published papers are defined as $\frac{l}{2l+1}\Delta U_l$ and, hence, have to be multiplied by $2/l$ (Warning: on the CRENBL website the spin-orbit potentials already have been corrected with the $2/l$ factor).

Chapter 9

Relativistic All-electron Approximations

All methods which include treatment of relativistic effects are ultimately based on the Dirac equation, which has a four component wave function. The solutions to the Dirac equation describe both positrons (the “negative energy” states) and electrons (the “positive energy” states), as well as both spin orientations, hence the four components. The wave function may be broken down into two-component functions traditionally known as the large and small components; these may further be broken down into the spin components.

The implementation of approximate all-electron relativistic methods in quantum chemical codes requires the removal of the negative energy states and the factoring out of the spin-free terms. Both of these may be achieved using a transformation of the Dirac Hamiltonian known in general as a Foldy-Wouthuysen transformation. Unfortunately this transformation cannot be represented in closed form for a general potential, and must be approximated. One popular approach is that originally formulated by Douglas and Kroll¹ and developed by Hess². This approach decouples the positive and negative energy parts to second order in the external potential (and also fourth order in the fine structure constant, α). Another approach is based on a modification of the Dirac equation by Dyll³, and involves an exact FW transformation on the atomic basis set level⁴.

Since these approximations only modify the integrals, they can in principle be used at all levels of theory. At present the Douglas-Kroll implementation can be used at all levels of theory whereas Dyll’s approach is currently available at the Hartree-Fock level. The derivatives have been implemented, allowing both methods to be used in geometry optimizations and frequency calculations.

The RELATIVISTIC directive provides input for the implemented relativistic approximations and is a compound directive that encloses additional directives specific to the approximations:

```
RELATIVISTIC
[DOUGLAS-KROLL [<string (ON|OFF) default ON> \
  <string (FPP|DKH|DKFULL|DK3|DK3FULL) default DKH>] ||
  DYALL-MOD-DIRAC [ (ON || OFF) default ON ]
  [ (NESC1E || NESC2E) default NESC1E ] ]
[CLIGHT <real clight default 137.0359895>]
END
```

Only one of the methods may be chosen at a time. If both methods are found to be on in the input block, NWChem

¹M. Douglas and N. M. Kroll, *Ann. Phys. (N.Y.)* **82**, 89 (1974)

²B.A. Hess, *Phys. Rev. A* **32**, 756 (1985); **33**, 3742 (1986)

³K. G. Dyll, *J. Chem. Phys.* **100**, 2118 (1994)

⁴K. G. Dyll, *J. Chem. Phys.* **106**, 9618 (1997); K. G. Dyll and T. Enevoldsen, *J. Chem. Phys.* **111**, 10000 (1999).

will stop and print an error message. There is one general option for both methods, the definition of the speed of light in atomic units:

```
CLIGHT <real clight default 137.0359895>
```

The following sections describe the optional sub-directives that can be specified within the RELATIVISTIC block.

9.1 Douglas-Kroll approximation

The (spin-free) one-electron Douglas-Kroll approximation has been implemented. The use of relativistic effects from this Douglas-Kroll approximation can be invoked by specifying:

```
DOUGLAS-KROLL [<string (ON|OFF) default ON> \
               <string (FPP|DKH|DKFULL|DK3|DK3FULL) default DKH>]
```

The ON|OFF string is used to turn on or off the Douglas-Kroll approximation. By default, if the DOUGLAS-KROLL keyword is found, the approximation will be used in the calculation. If the user wishes to calculate a non-relativistic quantity after turning on Douglas-Kroll, the user will need to define a new RELATIVISTIC block and turn the approximation OFF. The user could also simply put a blank RELATIVISTIC block in the input file and all options will be turned off.

The FPP is the approximation based on free-particle projection operators⁵ whereas the DKH and DKFULL approximations are based on external-field projection operators⁶. The latter two are considerably better approximations than the former. DKH is the Douglas-Kroll-Hess approach and is the approach that is generally implemented in quantum chemistry codes. DKFULL includes certain cross-product integral terms ignored in the DKH approach (see for example Häberlen and Rösch⁷). The third-order Douglas-Kroll approximation has been implemented by T. Nakajima and K. Hirao⁸. This approximation can be called using DK3 (DK3 without cross-product integral terms) or DK3FULL (DK3 with cross-product integral terms).

The contracted basis sets used in the calculations should reflect the relativistic effects, i.e. one should use contracted basis sets which were generated using the Douglas-Kroll Hamiltonian. Basis sets that were contracted using the non-relativistic (Schödinger) Hamiltonian WILL PRODUCE ERRONEOUS RESULTS for elements beyond the first row. See appendix A for available basis sets and their naming convention.

NOTE: we suggest that spherical basis sets are used in the calculation. The use of high quality cartesian basis sets can lead to numerical inaccuracies.

In order to compute the integrals needed for the Douglas-Kroll approximation the implementation makes use of a fitting basis set (see literature given above for details). The current code will create this fitting basis set based on the given "ao basis" by simply uncontracting that basis. This again is what is commonly implemented in quantum chemistry codes that include the Douglas-Kroll method. Additional flexibility is available to the user by explicitly specifying a Douglas-Kroll fitting basis set. This basis set must be named "D-K basis" (see Chapter 7).

⁵B.A. Hess, Phys. Rev. A **32**, 756 (1985)

⁶B.A. Hess, Phys. Rev. A **33**, 3742 (1986)

⁷O.D. Häberlen, N. Rösch, Chem. Phys. Lett. **199**, 491 (1992)

⁸T. Nakajima and K. Hirao, Chem. Phys. Lett. **329**, 5111 (2000); T. Nakajima and K. Hirao, J. Chem. Phys. **113**, 7786 (2000)

9.2 Dyall's Modified Dirac Hamiltonian approximation

The approximate methods described in this section are all based on Dyall's modified Dirac Hamiltonian. This Hamiltonian is entirely equivalent to the original Dirac Hamiltonian, and its solutions have the same properties. The modification is achieved by a transformation on the small component, extracting out $\sigma \cdot \mathbf{p}/2mc$. This gives the modified small component the same symmetry as the large component, and in fact it differs from the large component only at order α^2 . The advantage of the modification is that the operators now resemble the operators of the Breit-Pauli Hamiltonian, and can be classified in a similar fashion into spin-free, spin-orbit and spin-spin terms. It is the spin-free terms which have been implemented in NWChem, with a number of further approximations.

The first is that the negative energy states are removed by a normalized elimination of the small component (NESC), which is equivalent to an exact Foldy-Wouthuysen (EFW) transformation. The number of components in the wave function is thereby effectively reduced from 4 to 2. NESC on its own does not provide any advantages, and in fact complicates things because the transformation is energy-dependent. The second approximation therefore performs the elimination on an atom-by-atom basis, which is equivalent to neglecting blocks which couple different atoms in the EFW transformation. The advantage of this approximation is that all the energy dependence can be included in the contraction coefficients of the basis set. The tests which have been done show that this approximation gives results well within chemical accuracy. The third approximation neglects the commutator of the EFW transformation with the two-electron Coulomb interaction, so that the only corrections that need to be made are in the one-electron integrals. This is the equivalent of the Douglas-Kroll(-Hess) approximation as it is usually applied.

The use of these approximations can be invoked with the use of the `DYALL-MOD-DIRAC` directive in the `RELATIVISTIC` directive block. The syntax is as follows.

```
DYALL-MOD-DIRAC [ (ON || OFF) default ON ]
                 [ (NESC1E || NESC2E) default NESC1E ]
```

The `ON|OFF` string is used to turn on or off the Dyall's modified Dirac approximation. By default, if the `DYALL-MOD-DIRAC` keyword is found, the approximation will be used in the calculation. If the user wishes to calculate a non-relativistic quantity after turning on Dyall's modified Dirac, the user will need to define a new `RELATIVISTIC` block and turn the approximation `OFF`. The user could also simply put a blank `RELATIVISTIC` block in the input file and all options will be turned off.

Both one- and two-electron approximations are available `NESC1E || NESC2E`, and both have analytic gradients. The one-electron approximation is the default. The two-electron approximation specified by `NESC2E` has some sub options which are placed on the same logical line as the `DYALL-MOD-DIRAC` directive, with the following syntax:

```
NESC2E [ (SS1CENT [ (ON || OFF) default ON ] || SSALL) default SSALL ]
         [ (SSSS [ (ON || OFF) default ON ] || NOSSSS) default SSSS ]
```

The first sub-option gives the capability to limit the two-electron corrections to those in which the small components in any density must be on the same center. This reduces the $(LL|SS)$ contributions to at most three-center integrals and the $(SS|SS)$ contributions to two centers. For a case with only one relativistic atom this option is redundant. The second controls the inclusion of the $(SS|SS)$ integrals which are of order α^4 . For light atoms they may safely be neglected, but for heavy atoms they should be included.

In addition to the selection of this keyword in the `RELATIVISTIC` directive block, it is necessary to supply basis sets in addition to the `ao basis`. For the one-electron approximation, three basis sets are needed: the atomic FW basis set, the large component basis set and the small component basis set. The atomic FW basis set should be included in the `ao basis`. The large and small components should similarly be incorporated in basis sets named `large component` and `small component`, respectively. For the two-electron approximation, only two basis sets are needed. These are the large component and the small component. The large component should be included

in the `ao basis` and the small component is specified separately as `small component`, as for the one-electron approximation. This means that the two approximations can *not* be run correctly without changing the `ao basis`, and it is up to the user to ensure that the basis sets are correctly specified.

There is one further requirement in the specification of the basis sets. In the `ao basis`, it is necessary to add the `rel` keyword either to the `basis` directive or the library tag line (See below for examples). The former marks the basis functions specified by the tag as relativistic, the latter marks the whole basis as relativistic. The marking is actually done at the unique shell level, so that it is possible not only to have relativistic and nonrelativistic atoms, it is also possible to have relativistic and nonrelativistic shells on a given atom. This would be useful, for example, for diffuse functions or for high angular momentum correlating functions, where the influence of relativity was small. The marking of shells as relativistic is necessary to set up a mapping between the `ao basis` and the large and/or small component basis sets. For the one-electron approximation the large and small component basis sets **MUST** be of the same size and construction, i.e. differing only in the contraction coefficients.

It should also be noted that the relativistic code will **NOT** work with basis sets that contain `sp` shells, nor will it work with ECPs. Both of these are tested and flagged as an error.

Some examples follow. The first example sets up the data for relativistic calculations on water with the one-electron approximation and the two-electron approximation, using the library basis sets.

```

start h2o-dmd

geometry units bohr
symmetry c2v
  O      0.000000000    0.000000000   -0.009000000
  H      1.515260000    0.000000000   -1.058900000
  H     -1.515260000    0.000000000   -1.058900000
end

basis "fw" rel
  oxygen library cc-pvdz_pt_sf_fw
  hydrogen library cc-pvdz_pt_sf_fw
end

basis "large"
  oxygen library cc-pvdz_pt_sf_lc
  hydrogen library cc-pvdz_pt_sf_lc
end

basis "large2" rel
  oxygen library cc-pvdz_pt_sf_lc
  hydrogen library cc-pvdz_pt_sf_lc
end

basis "small"
  oxygen library cc-pvdz_pt_sf_sc
  hydrogen library cc-pvdz_pt_sf_sc
end

set "ao basis" fw
set "large component" large
set "small component" small

```

```
relativistic
  dyall-mod-dirac
end

task scf

set "ao basis" large2
unset "large component"
set "small component" small
```

```
relativistic
  dyall-mod-dirac nesc2e
end
```

```
task scf
```

The second example has oxygen as a relativistic atom and hydrogen nonrelativistic.

```
start h2o-dmd2

geometry units bohr
symmetry c2v
  O      0.000000000    0.000000000   -0.009000000
  H      1.515260000    0.000000000   -1.058900000
  H     -1.515260000    0.000000000   -1.058900000
end

basis "ao basis"
  oxygen library cc-pvdz_pt_sf_fw rel
  hydrogen library cc-pvdz
end

basis "large component"
  oxygen library cc-pvdz_pt_sf_lc
end

basis "small component"
  oxygen library cc-pvdz_pt_sf_sc
end

relativistic
  dyall-mod-dirac
end

task scf
```


Chapter 10

Hartree-Fock or Self-consistent Field

The NWChem self-consistent field (SCF) module computes closed-shell restricted Hartree-Fock (RHF) wavefunctions, restricted high-spin open-shell Hartree-Fock (ROHF) wavefunctions, and spin-unrestricted Hartree-Fock (UHF) wavefunctions.

The SCF directive provides input to the SCF module and is a compound directive that encloses additional directives specific to the SCF module:

```
SCF
...
END
```

10.1 Wavefunction type

A spin-restricted, closed shell RHF calculation is performed by default. An error results if the number of electrons is inconsistent with this assumption. The number of electrons is inferred from the total charge on the system and the sum of the effective nuclear charges of all centers (atoms and dummy atoms, Section 6). The total charge on the system is zero by default, unless specified at some value by input on the CHARGE directive (Section 5).

The options available to define the SCF wavefunction and multiplicity are as follows:

```
SINGLET
DOUBLET
TRIPLET
QUARTET
QUINTET
SEXTET
SEPTET
OCTET
NOPEN <integer nopen default 0>
RHF
ROHF
UHF
```

The optional keywords SINGLET, DOUBLET, ..., OCTET and NOPEN allow the user to specify the number of

singly occupied orbitals for a particular calculation. SINGLET is the default, and specifies a closed shell; DOUBLET specifies one singly occupied orbital; TRIPLET specifies two singly occupied orbitals; and so forth. If there are more than seven singly occupied orbitals, the keyword NOOPEN must be used, with the integer nopen defining the number of singly occupied orbitals (sometimes referred to as open shells).

If the multiplicity is any value other than SINGLET, the default calculation will be a spin-restricted, high-spin, open-shell SCF calculation (keyword ROHF). The open-shell orbitals must be the highest occupied orbitals. If necessary, any starting vectors may be rearranged through the use of the SWAP keyword on the VECTORS directive (see Section 10.5) to accomplish this.

A spin-unrestricted solution can also be performed by specifying the keyword UHF. In UHF calculations, it is assumed that the number of singly occupied orbitals corresponds to the difference between the number of alpha-spin and beta-spin orbitals. For example, a UHF calculation with 2 more alpha-spin orbitals than beta-spin orbitals can be obtained by specifying

```
scf
  triplet ; uhf      # (Note: two logical lines of input)
  ...
end
```

The user should be aware that, by default, molecular orbitals are symmetry adapted in NWChem. This may not be desirable for fully unrestricted wavefunctions. In such cases, the user has the option of defeating the defaults by specifying the keywords ADAPT OFF (see Section 10.3) and SYM OFF (see Section 10.2).

The keywords RHF and ROHF are provided in the code for completeness. It may be necessary to specify these in order to modify the behavior of a previous calculation (see Section 3.2 for restart behavior).

10.2 SYM — use of symmetry

```
SYM <string (ON|OFF) default ON>
```

This directive enables/disables the use of symmetry to speed up Fock matrix construction (via the petite-list or skeleton algorithm) in the SCF, if symmetry was used in the specification of the geometry. Symmetry adaptation of the molecular orbitals is not affected by this option. The default is to use symmetry if it is specified in the geometry directive (Section 6).

For example, to disable use of symmetry in Fock matrix construction:

```
sym off
```

10.3 ADAPT – symmetry adaptation of MOs

```
ADAPT <string (ON|OFF) default ON>
```

The default in the SCF module calculation is to force symmetry adaption of the molecular orbitals. This does not affect the speed of the calculation, but without explicit adaption the resulting orbitals may be symmetry contaminated for some problems. This is especially likely if the calculation is started using orbitals from a distorted geometry.

The underlying assumption in the use of symmetry in Fock matrix construction is that the density is totally symmetric. If the orbitals are symmetry contaminated, this assumption may not be valid — which could result in incorrect energies and poor convergence of the calculation. It is thus advisable when specifying `ADAPT OFF` to also specify `SYM OFF` (Section 10.2).

10.4 TOL2E — integral screening threshold

```
TOL2E <real tol2e default min(10e-7 , 0.01*$thresh$)>
```

The variable `tol2e` is used in determining the integral screening threshold for the evaluation of the energy and related Fock-like matrices. The Schwarz inequality is used to screen the product of integrals and density matrices in a manner that results in an accuracy in the energy and Fock matrices that approximates the value specified for `tol2e`.

It is generally not necessary to set this parameter directly. Specify instead the required precision in the wavefunction, using the `THRESH` directive (Section 10.7). The default threshold is the minimum of 10^{-7} and 0.01 times the requested convergence threshold for the SCF calculation (Section 10.7).

The input to specify the threshold explicitly within the SCF directive is, for example:

```
tol2e 1e-9
```

For very diffuse basis sets, or for high-accuracy calculations it might be necessary to set this parameter. A value of 10^{-12} is sufficient for nearly all such purposes.

10.5 VECTORS — input/output of MO vectors

```
VECTORS [[input] (<string input_movecs default atomic>) || \
              (project <string basisname> <string filename>) || \
              (fragment <string file1> [<string file2> ...])] \
[swap [alpha||beta] <integer vec1 vec2> ...] \
[reorder <integer atom1 atom2> ...] \
[output <string output_filename default input_movecs>] \
[lock]
[rotate <string input_geometry> <string input_movecs>]
```

The `VECTORS` directive allows the user to specify the source and destination of the molecular orbital vectors. In a startup calculation (see Section 5.1), the default source for guess vectors is a diagonalized Fock matrix constructed from a superposition of the atomic density matrices for the particular problem. This is usually a very good guess. For a restarted calculation, the default is to use the previous MO vectors.

The optional keyword `INPUT` allows the user to specify the source of the input molecular orbital vectors as any of the following:

- `ATOMIC` — eigenvectors of a Fock-like matrix formed from a superposition of the atomic densities (the default guess). See Sections 10.5.2 and 10.6.
- `HCORE` — eigenvectors of the bare-nucleus Hamiltonian or the one-electron Hamiltonian.

- `filename` — the name of a file containing the MO vectors from a previous calculation. Note that unless the path is fully qualified, or begins with a dot (“.”), then it is assumed to reside in the directory for permanent files (see Section 5.2).
- `PROJECT basisname filename` — projects the existing MO vectors in the file `filename` from the smaller basis with name `basisname` into the current basis. The definition of the basis `basisname` must be available in the current database, and the basis must be smaller than the current basis. In addition, the geometry used for the previous calculations must have the atoms in the same order and in the same orientation as the current geometry.
- `FRAGMENT file1 ...` — assembles starting MO vectors from previously performed calculations on fragments of the system and is described in more detail in Section 10.5.1. Even though there are some significant restrictions in the use of the initial implementation of this method (see Section 10.5.1), this is the most powerful initial guess option within the code. It is particularly indispensable for open shell metallic systems.
- `ROTATE input_geometry input_movecs` — rotates MO vectors generated at a previous geometry to the current active geometry.

The molecular orbitals are saved every iteration if more than 600 seconds have elapsed, and also at the end of the calculation. At completion (converged or not), the SCF module always canonically transforms the molecular orbitals by *separately* diagonalizing the closed–closed, open–open, and virtual–virtual blocks of the Fock matrix.

The name of the file used to store the MO vectors is determined as follows:

- if the `OUTPUT` keyword was specified on the `VECTORS` directive, then the filename that follows this keyword is used, or
- if the input vectors were read from a file, this file is reused for the output vectors (overwriting the input vectors); else,
- a default file name is generated in the directory for permanent files (Section 5.2) by prepending “`.movecs`” with the file prefix, i.e., “`<file_prefix>.movecs`”.

The name of this file is stored in the database so that a subsequent SCF calculation will automatically restart from these MO vectors.

Applications of this directive are illustrated in the following examples.

Example 1:

```
vectors output h2o.movecs
```

Assuming a start-up calculation, this directive will result in use of the default atomic density guess, and will output the vectors to the file `h2o.movecs`.

Example 2:

```
vectors input initial.movecs output final.movecs
```

This directive will result in the initial vectors being read from the file “`initial.movecs`”. The results will be written to the file `final.movecs`. The contents of “`initial.movecs`” will not be changed.

Example 3:

```
vectors input project "small basis" small.movecs
```

This directive will cause the calculation to start from vectors in the file "small.movecs" which are in a basis named "small basis". The output vectors will be written to the default file "<file_prefix.movecs>".

Once starting vectors have been obtained using any of the possible options, they may be reordered through use of the SWAP keyword. This optional keyword requires a list of orbital pairs that will be swapped. For UHF calculations, separate SWAP keywords may be provided for the alpha and beta orbitals, as necessary.

An example of use of the SWAP directive:

```
vectors input try1.movecs swap 173 175 174 176 output try2.movecs
```

This directive will cause the initial orbitals to be read from the file "try1.movecs". The vectors for the orbitals within the pairs 173–175 will be swapped with those within 174–176, so the resulting order is 175, 176, 173, 174. The final orbitals obtained in the calculation will be written to the file "try2.movecs".

The swapping of orbitals occurs as a sequential process in the order (left to right) input by the user. Thus, regarding each pair as an elementary transposition it is possible to construct arbitrary permutations of the orbitals. For instance, to apply the permutation (6789)¹ we note that this permutation is equal to (67)(78)(89), and thus may be specified as

```
vectors swap 8 9 7 8 6 7
```

Another example, now illustrating this feature for a UHF calculation, is the directive

```
vectors swap beta 4 5 swap alpha 5 6
```

This input will result in the swapping of the 5–6 alpha orbital pair and the 4–5 beta orbital pair. (All other items in the input use the default values.)

The LOCK keyword allows the user to specify that the ordering of orbitals will be locked to that of the initial vectors, insofar as possible. The default is to order by ascending orbital energies within each orbital space. One application where locking might be desirable is a calculation where it is necessary to preserve the ordering of a previous geometry, despite flipping of the orbital energies. For such a case, the LOCK directive can be used to prevent the SCF calculation from changing the ordering, even if the orbital energies change.

The mapping of the MO's to the nuclei can be changed using the REORDER keyword. Once starting vectors have been obtained using any of the possible options, the REORDER keyword moves the MO coefficients between atoms listed in the integer list. This keyword is particularly useful for calculating localized electron and hole states.

This optional keyword requires a list containing the new atom ordering. It is not necessary to provide separate lists for alpha and beta orbitals.

An example of use of the REORDER keyword:

```
vectors input try1.movecs reorder 2 1 output try2.movecs
```

This directive will cause the initial orbitals to be read from the file "try1.movecs". The MO coefficients for the basis functions on atom 2 will be swapped with those on atom 1. The final orbitals obtained in the calculation will be written to the file "try2.movecs".

The following example shows how the ROTATE keyword can be used to rotate MO vectors calculated at geometry geom1 to geometry geom2, which has a different rotational orientation:

```
set geometry geom1
```

¹The cyclic permutation (6789) maps the ordered list 6 7 8 9 into 9 6 7 8.

```

dft
  vectors input atomic output geom1.mo
end
task dft

set geometry geom2
dft
  vectors input rotate geom1 geom1.mo output geom2.mo
end
task dft

```

10.5.1 Superposition of fragment molecular orbitals

The fragment initial guess is particularly useful in the following instances:

- The system naturally decomposes into molecules that can be treated individually, e.g., a cluster.
- One or more fragments are particularly hard to converge and therefore much time can be saved by converging them independently.
- A fragment (e.g., a metal atom) must be prepared with a specific occupation. This can often be readily accomplished with a calculation on the fragment using dummy charges to model a ligand field.
- The molecular occupation predicted by the atomic initial guess is often wrong for systems with heavy metals which may have partially occupied orbitals with lower energy than some doubly occupied orbitals. The fragment initial guess avoids this problem.

```
VECTORS [input] fragment <string file1> [<string file2> ...]
```

The molecular orbitals are formed by superimposing the previously generated orbitals of fragments of the molecule being studied. These fragment molecular orbitals must be in the same basis as the current calculation. The input specifies the files containing the fragment molecular orbitals. For instance, in a calculation on the water dimer, one might specify

```
vectors fragment h2o1.movecs h2o2.movecs
```

where `h2o1.movecs` contains the orbitals for the first fragment, and `h2o2.movecs` contains the orbitals for the second fragment.

A complete example of the input for a calculation on the water dimer using the fragment guess is as follows:

```

start dimer

title "Water dimer SCF using fragment initial guess"

geometry dimer
  O   -0.595   1.165  -0.048
  H    0.110   1.812  -0.170
  H   -1.452   1.598  -0.154
  O    0.724  -1.284   0.034

```

```

    H    0.175  -2.013   0.348
    H    0.177  -0.480   0.010
end

geometry h2o1
  O   -0.595   1.165  -0.048
  H    0.110   1.812  -0.170
  H   -1.452   1.598  -0.154
end

geometry h2o2
  O    0.724  -1.284   0.034
  H    0.175  -2.013   0.348
  H    0.177  -0.480   0.010
end

basis
  o library 3-21g
  h library 3-21g
end

set geometry h2o1
scf; vectors input atomic output h2o1.movecs; end
task scf

set geometry h2o2
scf; vectors input atomic output h2o2.movecs; end
task scf

set geometry dimer
scf
vectors input fragment h2o1.movecs h2o2.movecs \
              output dimer.movecs
end
task scf

```

First, the geometry of the dimer and the two monomers are specified and given names. Then, after the basis specification, calculations are performed on the fragments by setting the geometry to the appropriate fragment (Section 5.7) and redirecting the output molecular orbitals to an appropriately named file. Note also that use of the atomic initial guess is forced, since the default initial guess is to use any existing MOs which would not be appropriate for the second fragment calculation. Finally, the dimer calculation is performed by specifying the dimer geometry, indicating use of the fragment guess, and redirecting the output MOs.

The following points are important in using the fragment initial guess:

1. The fragment calculations must be in the same basis set as the full calculation.
2. The order of atoms in the fragments and the order in which the fragment files are specified must be such that when the fragment basis sets are concatenated all the basis functions are in the same order as in the full system. This is readily accomplished by first generating the full geometry with atoms for each fragment contiguous, splitting this into numbered fragments and specifying the fragment MO files in the correct order on the VECTORS directive.

3. The occupation of orbitals is preserved when they are merged from the fragments to the full molecule and the resulting occupation must match the requested occupation for the full molecule. E.g., a triplet ROHF calculation must be comprised of fragments that have a total of exactly two open-shell orbitals.
4. Because of these restrictions, it is not possible to introduce additional atoms (or basis functions) into fragments for the purpose of cleanly breaking real bonds. However, it is possible, and highly recommended, to introduce additional point charges to simulate the presence of other fragments.
5. MO vectors of partially occupied or strongly polarized systems are very sensitive to orientation. While it is possible to specify the same fragment MO vector file multiple times in the `VECTORS` directive, it is usually much better to do a separate calculation for each fragment.
6. Linear dependencies which were present in a fragment calculation may be magnified in the full calculation. When this occurs, some of the fragment's highest virtual orbitals will not be copied to the full system, and a warning will be printed.

A more involved example is now presented. We wish to model the sextet state of Fe(III) complexed with water, imidazole and a heme with a net unit positive charge. The default atomic guess does not give the correct d^5 occupation for the metal and also gives an incorrect state for the double anion of the heme. The following performs calculations on all of the fragments. Things to note are:

1. The use of a dummy +2 charge in the initial guess on the heme which in part simulates the presence of the metal ion, and also automatically forces an additional two electrons to be added to the system (the default net charge being zero).
2. The iron fragment calculation (charge +3, d^5 , sextet) will yield the correct open-shell occupation for the full system. If, instead, the d -orbitals were partially occupied (e.g., the doublet state) it would be useful to introduce dummy charges around the iron to model the ligand field and thereby lift the degeneracy to obtain the correct occupation.
3. C_s symmetry is used for all of the calculations. It is not necessary that the same symmetry be used in all of the calculations, provided that the order and orientation of the atoms is preserved.
4. The `unset scf:*` directive is used immediately before the calculation on the full system so that the default name for the output MO vector file can be used, rather than having to specify it explicitly.

```
start heme6a1
title "heme-H2O (6A1) from M.Dupuis"

#####
# Define the geometry of the full system and the fragments #
#####

geometry full-system
  symmetry cs

  H      0.438   -0.002   4.549
  C      0.443   -0.001   3.457
  C      0.451   -1.251   2.828
  C      0.452    1.250   2.828
  H      0.455    2.652   4.586
  H      0.461   -2.649   4.586
  N1     0.455   -1.461   1.441
```

N1	0.458	1.458	1.443
C	0.460	2.530	3.505
C	0.462	-2.530	3.506
C	0.478	2.844	1.249
C	0.478	3.510	2.534
C	0.478	-2.848	1.248
C	0.480	-3.513	2.536
C	0.484	3.480	0.000
C	0.485	-3.484	0.000
H	0.489	4.590	2.664
H	0.496	-4.592	2.669
H	0.498	4.573	0.000
H	0.503	-4.577	0.000
H	-4.925	1.235	0.000
H	-4.729	-1.338	0.000
C	-3.987	0.685	0.000
N	-3.930	-0.703	0.000
C	-2.678	1.111	0.000
C	-2.622	-1.076	0.000
H	-2.284	2.126	0.000
H	-2.277	-2.108	0.000
N	-1.838	0.007	0.000
Fe	0.307	0.000	0.000
O	2.673	-0.009	0.000
H	3.238	-0.804	0.000
H	3.254	0.777	0.000

end

geometry ring-only

symmetry cs

H	0.438	-0.002	4.549
C	0.443	-0.001	3.457
C	0.451	-1.251	2.828
C	0.452	1.250	2.828
H	0.455	2.652	4.586
H	0.461	-2.649	4.586
N1	0.455	-1.461	1.441
N1	0.458	1.458	1.443
C	0.460	2.530	3.505
C	0.462	-2.530	3.506
C	0.478	2.844	1.249
C	0.478	3.510	2.534
C	0.478	-2.848	1.248
C	0.480	-3.513	2.536
C	0.484	3.480	0.000
C	0.485	-3.484	0.000
H	0.489	4.590	2.664
H	0.496	-4.592	2.669

```
Bq    0.307    0.0    0.0    charge 2 # simulate the iron
end
```

```
geometry imid-only
  symmetry cs
  H    0.498    4.573    0.000
  H    0.503   -4.577    0.000
  H   -4.925    1.235    0.000
  H   -4.729   -1.338    0.000
  C   -3.987    0.685    0.000
  N   -3.930   -0.703    0.000
  C   -2.678    1.111    0.000
  C   -2.622   -1.076    0.000
  H   -2.284    2.126    0.000
  H   -2.277   -2.108    0.000
  N   -1.838    0.007    0.000
end
```

```
geometry fe-only
  symmetry cs
  Fe   .307    0.000    0.000
end
```

```
geometry water-only
  symmetry cs
  O    2.673   -0.009    0.000
  H    3.238   -0.804    0.000
  H    3.254    0.777    0.000
end
```

```
#####
# Basis set for everything #
#####
```

```
basis nosegment
  O library 6-31g*
  N library 6-31g*
  C library 6-31g*
  H library 6-31g*
  Fe library "Ahlichs pVDZ"
end
```

```
#####
# SCF on the fragments for initial guess for full system #
#####
```

```
scf; thresh 1e-2; end
```

```
set geometry ring-only
scf; vectors atomic swap 80 81 output ring.mo; end
```

```

task scf

set geometry water-only
scf; vectors atomic output water.mo; end
task scf

set geometry imid-only
scf; vectors atomic output imid.mo; end
task scf

charge 3
set geometry fe-only
scf; sextet; vectors atomic output fe.mo; end
task scf

#####
# SCF on the full system #
#####

unset scf:*      # This restores the defaults

charge 1

set geometry full-system

scf
  sextet
  vectors fragment ring.mo imid.mo fe.mo water.mo
  maxiter 50
end

task scf

```

10.5.2 Atomic guess orbitals with charged atoms

As noted above, the default guess vectors are based on superimposing the density matrices of the neutral atoms. If some atoms are significantly charged, this default guess may be improved upon by modifying the atomic densities. This is done by setting parameters that add fractional charges to the occupation of the valence atomic orbitals. Since the atomic SCF program does not have its own input block, the `SET` directive (Section 5.7) must be used to set these parameters.

The input specifies a list of tags (i.e., names of atoms in a geometry, see Section 6) and the charges to be added to those centers. Two parameters must be set as follows:

```

set atomscf:tags_z <string list_of_tags>
set atomscf:z      <real list_of_charges>

```

The array of strings `atomscf:tags_z` should be set to the list of tags, and the array `atomscf:z` should be set to the list of charges which must be real numbers (not integers). All atoms that have a tag specified in the list of tags will be assigned the corresponding charge from the list of charges.

For example, the following specifies that all oxygen atoms with tag `O` be assigned a charge of `-1` and all iron atoms with tag `Fe` be assigned a charge of `+2`

```
set atomscf:z          -1  2.0
set atomscf:tags_z    O  Fe
```

There are some limitations to this feature. It is not possible to add electrons to closed shell atoms, nor is it possible to remove all electrons from a given atom. Attempts to do so will cause the code to report an error, and it will not report further errors in the input for modifying the charge even when they are detected.

Finally, recall that the database is persistent (Section 3.2) and that the modified settings will be used in subsequent atomic guess calculations unless the data is deleted from the database with the `UNSET` directive (Section 5.8).

10.6 Accuracy of initial guess

For SCF, the initial Fock-matrix construction from the atomic guess is now (starting from version 3.3) performed to a default precision of $1e-7$. However, other wavefunctions, notably DFT, use a lower precision. In charged, or diffuse basis sets, this precision may not be sufficient and could result in incorrect ordering of the initial orbitals. The accuracy may be increased with the following directive which should be inserted in the top-level of input (i.e., outside of the SCF input block) and before the `TASK` directive.

```
set tolguess 1e-7
```

10.7 THRESH — convergence threshold

```
THRESH <real thresh default 1.0e-4>
```

This directive specifies the convergence threshold for the calculation. The convergence threshold is the norm of the orbital gradient, and has a default value in the code of 10^{-4} .

The norm of the orbital gradient corresponds roughly to the precision available in the wavefunction, and the energy should be converged to approximately the square of this number. It should be noted, however, that the precision in the energy will not exceed that of the integral screening tolerance. This tolerance (Section 10.4) is automatically set from the convergence threshold, so that sufficient precision is usually available by default.

The default convergence threshold suffices for most SCF energy and geometry optimization calculations, providing about 6–8 decimal places in the energy, and about four significant figures in the density and energy derivative with respect to nuclear coordinates. However, greater precision may be required for calculations involving weakly interacting systems, floppy molecules, finite-difference of gradients to compute the Hessian, and for post-Hartree-Fock calculations. A threshold of 10^{-6} is adequate for most such purposes, and a threshold of 10^{-8} might be necessary for very high accuracy or very weak interactions. A threshold of 10^{-10} should be regarded as the best that can be attained in most circumstances.

10.8 MAXITER — iteration limit

```
MAXITER <integer maxiter default 8>
```

The maximum number of iterations for the SCF calculation defaults to 20 for both ROHF/RHF and UHF calculations. For most molecules, this number of iterations is more than sufficient for the quadratically convergent SCF algorithm to obtain a solution converged to the default threshold (see Section 10.7 above). If the SCF program detects that the quadratically convergent algorithm is not efficient, then it will resort to a linearly convergent algorithm and increase the maximum number of iterations by 10.

Convergence may not be reached in the maximum number of iterations for many reasons, including input error (e.g., an incorrect geometry or a linearly dependent basis), a very low convergence threshold, a poor initial guess, or the fact that the system is intrinsically hard to converge due to the presence of many states with similar energies.

The following sets the maximum number of SCF iterations to 50:

```
maxiter 50
```

10.9 PROFILE — performance profile

This directive allows the user to obtain timing and parallel execution information about the SCF module. It is specified by the simple keyword

```
PROFILE
```

This option can be helpful in understanding the computational performance of an SCF calculation. However, it can introduce a significant overhead on machines that have expensive timing routines, such as the SUN.

10.10 DIIS — DIIS convergence

This directive allows the user to specify DIIS convergence rather than second-order convergence for the SCF calculation. The form of the directive is as follows:

```
DIIS
```

The implementation of this option is currently fairly rudimentary. It does not have level-shifting and damping, and does not support open shells or UHF. It is provided on an “as is” basis, and should be used with caution.

When the DIIS directive is specified in the input, the user has the additional option of specifying the size of the subspace for the DIIS extrapolation. This is accomplished with the DIISBAS directive, which is of the form:

```
DIISBAS <integer diisbas default 5>
```

The default of 5 should be adequate for most applications, but may be increased if convergence is poor. On large systems, it may be necessary to specify a lower value for `diisbas`, to conserve memory.

10.11 DIRECT and SEMIDIRECT — recomputation of integrals

In the context of SCF calculations `direct` means that all integrals are recomputed as required and none are stored. The other extreme are disk- or memory-resident (sometimes termed conventional) calculations in which all integrals

are computed once and stored. Semi-direct calculations are between these two extremes with some integrals being precomputed and stored, and all other integrals being recomputed as necessary.

The default behavior of the SCF module is

- If enough memory is available, the integrals are computed once and are cached in memory.
- If there is not enough memory to store all the integrals at once, then 95% of the available disk space in the scratch directory (see Section 5.2) is assumed to be available for this purpose, and as many integrals as possible are cached on disk (with no memory being used for caching). Some attempt is made to store the most expensive integrals in the cache.
- If there is not enough room in memory or on disk for all the integrals, then the ones that are not cached are recomputed in a semidirect fashion.

The integral file is deleted at the end of a calculation, so it is not possible to restart a semidirect calculation when the integrals are cached in memory or on disk. Many computer systems (e.g., the EMSL IBM SP) clear the fast scratch space at the end of each job, adding a further complication to the problem of restarting a *parallel* semidirect calculation.

On the IBM SP or any other computer with fast disks local to each processor, semidirect calculation offers the best behavior. It can result in *quadratic speedup* as more processors are added.

A fully direct calculation (with recomputation of the integrals at each iteration) is forced by specifying the directive

```
DIRECT
```

Alternatively, the SEMIDIRECT directive can be used to control the default semidirect calculation by defining the amount of disk space and the cache memory size. The form of this directive is as follows:

```
SEMIDIRECT [filesize <integer filesize default disksize>]
           [memsize <integer memsize default available>]
           [filename <string filename default $file_prefix.aoints$>]
```

The keyword FILESIZE allows the user to specify the amount of disk space to be used per process for storing the integrals in 64-bit words. Similarly, the keyword MEMSIZE allows the user to specify the number of 64-bit words to be used per process for caching integrals in memory. (Note: If the amount of storage space specified by the entry for memsize is not available, the code cuts the value in half and checks again for available space. This process is repeated until the request is satisfied.)

By default, the integral files are placed into the scratch directory (see Section 5.2). Specifying the keyword FILENAME overrides this default. The user-specified name entered in the string filename has the process number appended to it, so that each process has a distinct file but with a common base-name and directory. Therefore, it is not possible to use this keyword to specify different disks for different processes. The SCRATCH_DIR directive (see Section 5.2) can be used for this purpose.

For example, to force full recomputation of all integrals:

```
direct
```

Exactly the same result could be obtained by entering the directive:

```
semidirect filesize 0 memsize 0
```

To disable the use of memory for caching integrals and limit disk usage by each process to 100 megawords (MW):

```
semidirect memsize 0 filesize 100000000
```

The integral records are typically 32769 words long and any non-zero value for `filesize` or `memsize` should be enough to hold at least one record.

10.11.1 Integral File Size and Format for the SCF Module

The file format is rather complex, since it accommodates a variety of packing and compression options and the distribution of data. This section presents some information that may help the user understand the output, and illustrates how to use the output information to estimate file sizes.

If integrals are stored with a threshold of greater than 10^{-10} , then the integrals are stored in a 32-bit fixed-point format (with appropriate treatment for large values to retain precision). If integrals are stored with a threshold less than 10^{-10} , however, the values are stored in 64-bit floating-point format. If a replicated-data calculation is being run, then 8 bits are used for each basis function label, unless there are more than 256 functions, in which case 16 bits are used. If distributed data is being used, then the labels are always packed to 8 bits (the distributed blocks always being less than 256; labels are relative to the start of the block).

Thus, the number (W) of 64-bit words required to store N integrals, may be computed as

$$W = \begin{cases} N, & \text{8-bit labels and 32-bit values} \\ \frac{3}{2}N, & \text{16-bit labels and 32-bit values} \\ \frac{3}{2}N, & \text{8-bit labels and 64-bit values} \\ 2N, & \text{16-bit labels and 64-bit values} \end{cases}$$

The actual number of words required can exceed this computed value by up to one percent, due to bookkeeping overhead, and because the file itself is organized into fixed-size records.

With at least the default print level, all semidirect (not direct) calculations will print out information about the integral file and the number of integrals computed. The form of this output is as follows:

```
Integral file           = ./c6h6.aoints.0
Record size in doubles = 32769           No. of integs per rec = 32768
Max. records in memory = 3               Max. records in file  = 5
No. of bits per label  = 8               No. of bits per value = 32
```

```
#quartets = 2.0D+04 #integrals = 7.9D+05 direct = 63.6% cached = 36.4%
```

The file information above relates only to process 0. The line of information about the number of quartets, integrals, etc., is a sum over all processes.

When the integral file is closed, additional information of the following form is printed:

```
-----
EAF file 0: "./c6h6.aoints.0" size=262152 bytes
-----
           write      read      awrite      aread      wait
           -----      ----      -
calls:           6          12           0           0           0
```

```

data(b): 1.57e+06  3.15e+06  0.00e+00  0.00e+00
time(s): 1.09e-01  3.12e-02  0.00e+00
rate(mb/s): 1.44e+01  1.01e+02
-----

```

```

Parallel integral file used          4 records with          0 large values

```

Again, the detailed file information relates just to process 0, but the final line indicates the total number of integral records stored by all processes.

This information may be used to optimize subsequent calculations, for instance by assigning more memory or disk space.

10.12 SCF Convergence Control Options

Note to users: It is desired that the SCF program converge reliably with the default options for a wide variety of molecules. In addition, it should be guaranteed to converge for any system, with sufficient iterations. Please report significant convergence problems to nwchem-support@emsl.pnl.gov, and include the input file.

The SCF program uses a preconditioned conjugate gradient (PCG) method that is unconditionally convergent. Basically, a search direction is generated by multiplying the orbital gradient (the derivative of the energy with respect to the orbital rotations) by an approximation to the inverse of the level-shifted orbital Hessian. In the initial iterations (see Section 10.13), an inexpensive one-electron approximation to the inverse orbital Hessian is used. Closer to convergence, the full orbital Hessian is used, which should provide quadratic convergence. For both the full or one-electron orbital Hessians, the inverse-Hessian matrix-vector product is formed iteratively. Subsequently, an approximate line search is performed along the new search direction. If the exact Hessian is being employed, then the line search should require a single step (of unity). Preconditioning with approximate Hessians may require additional steps, especially in the initial iterations. It is the (approximate) line search that provides the convergence guarantee. The iterations required to solve the linear equations are referred to as micro-iterations. A macro-iteration comprises both the iterative solution and a line search.

Level-shifting plays the same role in this algorithm as it does in the conventional iterative solution of the SCF equations. The approximate Hessian used for preconditioning should be positive definite. If this is not the case, then level-shifting by a positive constant (Δ) serves to make the preconditioning matrix positive definite, by adding Δ to all of its eigenvalues. The level-shifts employed for the RHF orbital Hessian should be approximately four times (only twice for UHF) the value that one would employ in a conventional SCF². Level-shifting is automatically enabled in the early iterations, and the default options suffice for most test cases.

So why do things go wrong and what can be done to fix convergence problems? Most problems encountered so far arise either poor initial guesses or from small or negative eigenvalues of the orbital Hessian. The atomic orbital guess is usually very good. However, in calculations on charged systems, especially with open shells, incorrect initial occupations may result. The SCF might then converge very slowly since very large orbital rotations might be required to achieve the correct occupation or move charge large distances in the molecule. Possible actions are

- Modify the atomic guess by assigning charges to the atoms known to carry substantial charges (Section 10.5.2)
- Examining an analysis of the initial orbitals (Section 10.16) and then swapping them to attain the desired occupation (Section 10.5).

²This can be seen by considering a one-electron approximation to the closed-shell RHF Hessian in canonical orbitals, $A_{ia,jb} = 4\delta_{ij}\delta_{ab}(\epsilon_a - \epsilon_i)$. Similarly, the level shift should be twice as large for UHF.

- Converging the calculation in a minimal basis set, which is usually easier, and then projecting into a larger basis set (Section 10.5).
- Using the fragment orbital initial guess (Section 10.5.1).

Small or negative Hessian eigenvalues can occur even though the calculation seem to be close to convergence (as measured by the gradient norm, or the off-diagonal Fock matrix elements). Small eigenvalues will cause the iterative linear equation solver to converge slowly, resulting in an excessive number of micro-iterations. This makes the SCF expensive in terms of computation time, and it is possible to exceed the maximum number of iterations without achieving the accuracy required for quadratic convergence — which causes more macro-iterations to be performed.

Two main options are available when a problem will not converge: Newton-Raphson can be disabled temporarily or permanently (see Section 10.13), and level-shifting can be applied to the matrix (see Section 10.14). In some cases, both options may be necessary to achieve final convergence.

If there is reason to suspect a negative eigenvalue, the first course is to disable the Newton-Raphson iteration until the solution is closer to convergence. It may be necessary to disable it completely. At some point close to convergence, the Hessian will be positive definite, so disabling Newton-Raphson should yield a solution with approximately the same convergence rate as DIIS.

If temporarily disabling Newton-Raphson is not sufficient to achieve convergence, it may be necessary to disable it entirely and apply a small level-shift to the approximate Hessian. This should improve the convergence rate of the micro-iterations and stabilize the macro-iterations. The level-shifting will destroy exact quadratic convergence, but the optimization process is automatically adjusted to reflect this by enforcing conjugacy and reducing the accuracy to which the linear equations are solved. The net result of this is that the solution will do more macro-iterations, but each one should take less time than it would with the unshifted Hessian.

The following sections describe the directives needed to disable the Newton-Raphson iteration and specify level-shifting.

10.13 NR — controlling the Newton-Raphson

```
NR <real nr_switch default 0.1>
```

The exact orbital Hessian is adopted as the preconditioner when the maximum element of the orbital gradient is below the value specified for `nr_switch`. The default value is 0.1, which means that Newton-Raphson will be disabled until the maximum value of the orbital gradient (twice the largest off-diagonal Fock matrix element) is less than 0.1. To disable the Newton-Raphson entirely, the value of `nr_switch` must be set to zero. The directive to accomplish this is as follows:

```
nr 0
```

10.14 LEVEL — level-shifting the orbital Hessian

This directive allows the user to specify level-shifting to obtain a positive-definite preconditioning matrix for the SCF solution procedure. Separate level shifts can be set for the first-order convergent one-electron approximation to the Hessian used with the preconditioned conjugate gradient (PCG) method, and for the full Hessian used with the Newton-Raphson (NR) approach. It is also possible to change the level-shift automatically as the solution attains some specified accuracy. The form of the directive is as follows:

```

LEVEL [pcg <real initial default 20.0> \
      [<real tol default 0.5> <real final default 0.0>]] \
      [nr <real initial default 0.0> \
      [<real tol default 0.0> <real final default 0.0>]]

```

This directive contains only two keywords: one for the PCG method and the other for the exact Hessian (Newton Raphson, or NR). Use of PCG or NR is determined by the input specified for `nr_switch` on the NR directive, Section 10.13 above.

Specifying the keyword `pcg` on the LEVEL directive allows the user to define the level shifting for the approximate (i.e., PCG) method. Specifying the keyword `nr` allows the user to define the level shifting for the exact Hessians. In both options, the initial level shift is defined by the value specified for the variable `initial`. Optionally, `tol` can be specified independently with each keyword to define the level of accuracy that must be attained in the solution before the level shifting is changed to the value specified by input in the real variable `final`. Level shifts and gradient thresholds are specified in atomic units.

For the PCG method (as specified using the keyword `pcg`), the defaults for this input are 20.0 for `initial`, 0.5 for `tol`, and 0.0 for `final`. This means that the approximate Hessian will be shifted by 20.0 until the maximum element of the gradient falls below 0.5, at which point the shift will be set to zero.

For the exact Hessian (as specified using the keyword `nr`), the defaults are all zero. The exact Hessian is usually not shifted since this destroys quadratic convergence. An example of an input directive that applies a shift of 0.2 to the exact Hessian is as follows:

```
level nr 0.2
```

To apply this shift to the exact Hessian only until the maximum element of the gradient falls below 0.005, the required input directive is as follows:

```
level nr 0.2 0.005 0
```

Note that in both of these examples, the parameters for the PCG method are at the default values. To obtain values different from the defaults, the keyword `pcg` must also be specified. For example, to specify the level shifting in the above example for the exact Hessian *and* non-default shifting for the PCG method, the directive would be something like the following:

```
level pcg 20 0.3 0.0 nr 0.2 0.005 0.0
```

This input will cause the PCG method to be level-shifted by 20.0 until the maximum element of the gradient falls below 0.3, then the shift will be zero. For the exact Hessian, the level shifting is initially 0.2, until the maximum element falls below 0.005, after which the shift is zero.

The default options correspond to

```
level pcg 20 0.5 0 nr 0 0 0
```

10.15 Orbital Localization

The SCF module includes an *experimental* implementation of orbital localization, including Foster-Boys and Pipek-Mezey which only works for closed-shell (RHF) wavefunctions. There is currently no input in the SCF block to control this so the SET directive (Section 5.7) must be used.

The directive

```
set scf:localize t
```

will separately localize the core, valence, and virtual orbital spaces using the Pipek-Mezey algorithm. If the additional directive

```
set scf:loctype FB
```

is included, then the Foster-boys algorithm is used. The partitioning of core-orbitals is performed using the atomic information described in Section 16.1.

In the next release, this functionality will be extended to included all wavefunctions using molecular orbitals.

10.16 Printing Information from the SCF Module

All output from the SCF module is controlled using the `PRINT` directive described in Section 5.6. The following list describes the items from SCF that are currently under direct print control, along with the print level for each one.

Name	Print Level	Description
“atomic guess density”	debug	guess density matrix
“atomic scf”	debug	details of atomic SCF
“mo guess”	default	brief info from mo guess
“information”	low	results
“initial vectors”	debug	
“intermediate vectors”	debug	
“final vectors”	debug	
“final vectors analysis”	default	
“initial vectors analysis”	never	
“intermediate evals”	debug	
“final evals”	default	
“schwarz”	high	integral screening info & stats at completion
“screening statistics”	debug	display stats after every Fock build
“geometry”	high	
“symmetry”	debug	detailed symmetry info
“basis”	high	
“geombas”	debug	detailed basis map info
“vectors i/o”	default	report vectors I/O
“parameters”	default	convergence parameters
“convergence”	default	info each iteration
“mulliken ao”	never	Mulliken population of basis functions

Table 10.1: SCF Print Control Specifications

10.17 Hartree-Fock or SCF, MCSCF and MP2 Gradients

The input for this directive allows the user to adjust the print control for the SCF, UHF, ROHF, MCSCF and MP2 gradients. The form of the directive is as follows:

```
GRADIENTS
  [print || noprint] ...
END
```

The complementary keyword pair `print` and `noprint` allows the user some additional control on the information that can be included in the print output from the SCF calculation. Currently, only a few items can be explicitly invoked via print control. These are as follows:

Name	Print Level	Description
“information”	low	calculation info
“geometry”	high	geometry information
“basis”	high	basis set(s) used
“forces”	low	details of force components
“timing”	default	timing for each phase

Table 10.2: Gradient Print Control Specifications

Chapter 11

DFT for Molecules (DFT)

The NWChem density functional theory (DFT) module uses the Gaussian basis set approach to compute closed shell and open shell densities and Kohn-Sham orbitals in the:

- local density approximation (LDA),
- non-local density approximation (NLDA),
- local spin-density approximation (LSD),
- non-local spin-density approximation (NLSD),
- any empirical mixture of local and non-local approximations (including exact exchange), and
- asymptotically corrected exchange-correlation potentials.

The formal scaling of the DFT computation can be reduced by choosing to use auxiliary Gaussian basis sets to fit the charge density (CD) and/or fit the exchange-correlation (XC) potential.

DFT input is provided using the compound `DFT` directive

```
DFT
...
END
```

The actual DFT calculation will be performed when the input module encounters the `TASK` directive (Section 5.10).

```
TASK DFT
```

Once a user has specified a geometry and a Kohn-Sham orbital basis set the DFT module can be invoked with no input directives (defaults invoked throughout). There are sub-directives which allow for customized application; those currently provided as options for the DFT module are:

```
VECTORS [[input] (<string input_movecs default atomic>) || \
          (project <string basisname> <string filename>)] \
        [swap [alpha||beta] <integer vec1 vec2> ...] \
        [output <string output_filename default input_movecs>] \
```

```

XC [[acm] [b3lyp] [beckehandh] [pbe0]\
  [becke97] [becke97-1] [becke97-2] [becke98] [hcth] [hcth120] [hcth147]\
  [hcth407] [becke97gga1] [hcth407p]\
  [mpw91] [mpw1k] [xft97] [cft97] [ft97] [xpkzb99] [cpkzb99]\
  [HFexch <real prefactor default 1.0>] \
  [becke88 [nonlocal] <real prefactor default 1.0>] \
  [xperdew91 [nonlocal] <real prefactor default 1.0>] \
  [xpb96 [nonlocal] <real prefactor default 1.0>] \
  [gill96 [nonlocal] <real prefactor default 1.0>] \
  [lyp <real prefactor default 1.0>] \
  [perdew81 <real prefactor default 1.0>] \
  [perdew86 [nonlocal] <real prefactor default 1.0>] \
  [perdew91 [nonlocal] <real prefactor default 1.0>] \
  [cpbe96 [nonlocal] <real prefactor default 1.0>] \
  [pw91lda <real prefactor default 1.0>] \
  [slater <real prefactor default 1.0>] \
  [vwn_1 <real prefactor default 1.0>] \
  [vwn_2 <real prefactor default 1.0>] \
  [vwn_3 <real prefactor default 1.0>] \
  [vwn_4 <real prefactor default 1.0>] \
  [vwn_5 <real prefactor default 1.0>] \
  [vwn_1_rpa <real prefactor default 1.0>]]

```

```

CONVERGENCE [[energy <real energy default 1e-7>] \
  [density <real density default 1e-5>] \
  [gradient <real gradient default 5e-4>] \
  [dampon <real dampon default 0.0>] \
  [dampoff <real dampoff default 0.0>] \
  [diison <real diison default 0.0>] \
  [diisoff <real diisoff default 0.0>] \
  [levlon <real levlon default 0.0>] \
  [levloff <real levloff default 0.0>] \
  [ncydp <integer ncydp default 2>] \
  [ncyds <integer ncyds default 30>] \
  [ncysh <integer ncysh default 30>] \
  [damp <integer ndamp default 0>] [nodamping] \
  [diis [nfock <integer nfock default 10>]] \
  [nodiis] [lshift <real lshift default 0.5>] \
  [nolevelshifting] \
  [hl_tol <real hl_tol default 0.1>] \
  [rabuck [n_rabuck <integer n_rabuck default 25>]]

```

```

GRID [(xcoarse||coarse|medium|fine|xfine) default medium] \
  [(gausleg||lebedev) default lebedev] \
  [(becke||erf1||erf2||ssf) default erf1] \
  [(euler||mura||treutler) default mura] \
  [rm <real rm default 2.0>] \

```

```

[nodisk]

TOLERANCES [[tight] [tol_rho <real tol_rho default 1e-10>] \
            [accCoul <integer accCoul default 8>] \
            [radius <real radius default 25.0>]]

[(LB94| |CS00 <real shift default none>)]

DECOMP
ODFT
DIRECT
INCORE
ITERATIONS <integer iterations default 30>
MAX_OVL
MULLIKEN
MULT <integer mult default 1>
NOIO
PRINT| |NOPRINT

```

The following sections describe these keywords and optional sub-directives that can be specified for a DFT calculation in NWChem.

11.1 Specification of Basis Sets for the DFT Module

The DFT module requires at a minimum the basis set for the Kohn-Sham molecular orbitals. This basis set must be in the default basis set named "ao basis", or it must be assigned to this default name using the SET directive (see Section 5.7).

In addition to the basis set for the Kohn-Sham orbitals, the charge density fitting basis set can also be specified in the input directives for the DFT module. This basis set is used for the evaluation of the Coulomb potential in the Dunlap scheme¹. The charge density fitting basis set must have the name "cd basis". This can be the actual name of a basis set, or a basis set can be assigned this name using the SET directive, as described in Section 5.7. If this basis set is not defined by input, the $O(N^4)$ exact Coulomb contribution is computed.

The user also has the option of specifying a third basis set for the evaluation of the exchange-correlation potential. This basis set must have the name "xc basis". If this basis set is not specified by input, the exchange contribution (XC) is evaluated by numerical quadrature. In most applications, this approach is efficient enough, so the "xc basis" basis set is not generally required.

For the DFT module, the input options for defining the basis sets in a given calculation can be summarized as follows;

- "ao basis" – Kohn-Sham molecular orbitals; required for all calculations
- "cd basis" – charge density fitting basis set; optional, but recommended for evaluation of the Coulomb potential
- "xc basis" – exchange-correlation (XC) fitting basis set; optional, and usually not recommended

¹B.I. Dunlap, J.W.D. Connolly and J.R. Sabin, J. Chem. Phys. **71**, 4993 (1979)

11.2 VECTORS and MAX_OVL — KS-MO Vectors

The VECTORS directive is the same as that in the SCF module (Section 10.5). Currently, the LOCK keyword is not supported by the DFT module, however the directive

```
MAX_OVL
```

has the same effect.

11.3 XC and DECOMP — Exchange-Correlation Potentials

```
XC [[acm] [b3lyp] [beckehandh] [pbe0]\
  [becke97] [becke97-1] [becke97-2] [becke98] [hcth] [hcth120] [hcth147] \
  [hcth407] [becke97gga1] [hcth407p] \
  [optx] [hcthp14] [mpw91] [mpw1k] [xft97] [cft97] [ft97]\
  [HFexch <real prefactor default 1.0>] \
  [becke88 [nonlocal] <real prefactor default 1.0>] \
  [xperdew91 [nonlocal] <real prefactor default 1.0>] \
  [xpbe96 [nonlocal] <real prefactor default 1.0>] \
  [gill96 [nonlocal] <real prefactor default 1.0>] \
  [lyp <real prefactor default 1.0>] \
  [perdew81 <real prefactor default 1.0>] \
  [perdew86 [nonlocal] <real prefactor default 1.0>] \
  [perdew91 [nonlocal] <real prefactor default 1.0>] \
  [cpbe96 [nonlocal] <real prefactor default 1.0>] \
  [pw91lda <real prefactor default 1.0>] \
  [slater <real prefactor default 1.0>] \
  [vwn_1 <real prefactor default 1.0>] \
  [vwn_2 <real prefactor default 1.0>] \
  [vwn_3 <real prefactor default 1.0>] \
  [vwn_4 <real prefactor default 1.0>] \
  [vwn_5 <real prefactor default 1.0>] \
  [vwn_1_rpa <real prefactor default 1.0>]]
```

The user has the option of specifying the exchange-correlation treatment in the DFT Module (see table 11.1). The default exchange-correlation functional is defined as the local density approximation (LDA) for closed shell systems and its counterpart the local spin-density (LSD) approximation for open shell systems. Within this approximation the exchange functional is the Slater $\rho^{1/3}$ functional (from J.C. Slater, *Quantum Theory of Molecules and Solids, Vol. 4: The Self-Consistent Field for Molecules and Solids* (McGraw-Hill, New York, 1974)), and the correlation functional is the Vosko-Wilk-Nusair (VWN) functional (functional V) (S.J. Vosko, L. Wilk and M. Nusair, *Can. J. Phys.* **58**, 1200 (1980)). The parameters used in this formula are obtained by fitting to the Ceperley and Alder² Quantum Monte-Carlo solution of the *homogeneous electron gas*.

These defaults can be invoked explicitly by specifying the following keywords within the DFT module input directive, XC slater vwn_5.

That is, this statement in the input file

²D.M. Ceperley and B.J. Alder, *Phys. Rev. Lett.* **45**, 566 (1980).

```
dft
  XC slater vwn_5
end
task dft
```

is equivalent to the simple line

```
task dft
```

The DECOMP directive causes the components of the energy corresponding to each functional to be printed, rather than just the total exchange-correlation energy which is the default. You can see an example of this directive in the sample input in Section 11.5.

Many alternative exchange and correlation functionals are available to the user as listed in table 11.1. The following sections describe how to use these options.

11.3.1 Exchange-Correlation Functionals

There are several Exchange and Correlation functionals in addition to the default `slater` and `vwn_5` functionals. These are either local or gradient-corrected functionals (GCA); a full list can be found in table 11.1.

The Hartree-Fock exact exchange functional, (which has $O(N^4)$ computation expense), is invoked by specifying

```
XC HFexch
```

Note that the user also has the ability to include only the local or nonlocal contributions of a given functional. In addition the user can specify a multiplicative prefactor (the variable `<prefactor>` in the input) for the local/nonlocal component or total. An example of this might be,

```
XC becke88 nonlocal 0.72
```

The user should be aware that the Becke88 local component is simply the Slater exchange and should be input as such.

Any combination of the supported exchange functional options can be used. For example the popular Gaussian B3 exchange could be specified as:

```
XC slater 0.8 becke88 nonlocal 0.72 HFexch 0.2
```

Any combination of the supported correlation functional options can be used. For example B3LYP could be specified as:

```
XC vwn_1_rpa 0.19 lyp 0.81 HFexch 0.20 slater 0.80 becke88 nonlocal 0.72
```

11.3.2 Combined Exchange and Correlation Functionals

In addition to the options listed above for the exchange and correlation functionals, the user has the alternative of specifying combined exchange and correlation functionals. A complete list of the available functionals appears in table 11.1.

The available hybrid functionals (where a Hartree-Fock Exchange component is present) consist of the Becke “*half and half*” (see A.D. Becke, J. Chem. Phys. 98, 1372 (1992)), the adiabatic connection method (see A.D. Becke,

J. Chem. Phys. 98, 5648 (1993)), B3LYP (popularized by Gaussian9X), Becke 1997 (“Becke V” paper: A.D.Becke, J. Chem. Phys., **107**, 8554 (1997)).

The keyword `beckehandh` specifies that the exchange-correlation energy will be computed as

$$E_{XC} \approx \frac{1}{2}E_X^{\text{HF}} + \frac{1}{2}E_X^{\text{Slater}} + \frac{1}{2}E_C^{\text{PW91LDA}}$$

We know this is NOT the correct Becke prescribed implementation which requires the XC potential in the energy expression. But this is what is currently implemented as an approximation to it.

The keyword `acm` specifies that the exchange-correlation energy is computed as

$$\begin{aligned} E_{XC} &= a_0 E_X^{\text{HF}} + (1 - a_0) E_X^{\text{Slater}} + a_X \Delta E_X^{\text{Becke88}} + E_C^{\text{VWN}} + a_C \Delta E_C^{\text{Perdew91}} \\ &\text{where} \\ a_0 &= 0.20, a_X = 0.72, a_C = 0.81 \end{aligned}$$

and Δ stands for a non-local component.

The keyword `b3lyp` specifies that the exchange-correlation energy is computed as

$$\begin{aligned} E_{XC} &= a_0 E_X^{\text{HF}} + (1 - a_0) E_X^{\text{Slater}} + a_X \Delta E_X^{\text{Becke88}} + (1 - a_C) E_C^{\text{VWN}_1\text{-RPA}} + a_C E_C^{\text{LYP}} \\ &\text{where} \\ a_0 &= 0.20, a_X = 0.72, a_C = 0.81 \end{aligned}$$

11.3.3 Meta-GGA Functionals

One way to calculate meta-GGA energies is to use orbitals and densities from fully self-consistent GGA or LDA calculations and run them in one iteration in the meta-GGA functional. It is expected that meta-GGA energies obtained this way will be close to fully self consistent meta-GGA calculations. This can be easily accomplished in NWChem, and is illustrated in the example below. This kind of calculation will obviously not converge the energy. To avoid an error in the standard Unix output of NWChem, you must tell NWChem to ignore the returned result of the task, which may be accomplished in the input file with `task dft ignore`. You may still get a warning in the output to the effect, `!! warning: dft energy failed`. This simply means the energy failed to converge since you ran only one iteration of the functional.

(For more information, see S. Kurth, J. Perdew, P. Blaha, Int. J. Quant. Chem 75, 889 (1999) for a brief description of meta-GGAs, and citations 14-27 therein for thorough background)

Note: both TPSS and PKZB correlation require the PBE GGA CORRELATION (which is itself dependent on an LDA). The decision has been made to use these functionals with the accompanying local PW91LDA. The user does not have the ability to set the local part of these metaGGA functionals.

11.4 LB94 and CS00 — Asymptotic correction

The keyword `LB94` will correct the asymptotic region of the XC definition of exchange-correlation *potential* by the van-Leeuwen–Baerends exchange-correlation *potential* that has the correct $-1/r$ asymptotic behavior. The total energy will be computed by the XC definition of exchange-correlation functional. This scheme is known to tend to overcorrect the deficiency of most uncorrected exchange-correlation potentials.

The keyword CS00, when supplied with a real value of shift (in atomic units), will perform Casida–Salahub '00 asymptotic correction. This is primarily intended for use in conjunction with TDDFT and the background of this method is given in more detail in Chapter 14. The shift is normally positive (which means that the original uncorrected exchange-correlation potential must be shifted down).

When the keyword CS00 is specified without the value of shift, the program will automatically supply it according to the semi-empirical formula of Zhan, Nichols, and Dixon (again, see Chapter 14 for more details and references). As the Zhan's formula is calibrated against B3LYP results, it is most meaningful to use this in conjunction with the B3LYP functional, although the program does not prohibit (or even warn) the use of any other functional.

Sample input files of asymptotically corrected TDDFT calculations can be found in Chapter 14.

11.5 Sample input file

A simple example calculates the meta-GGA exchange energy of water, using converged GGA orbitals and densities to evaluate the meta-GGA energy functional, and also highlights some of the print features in the DFT module:

```
title "WATER 6-311G* meta-GGA X with PBE orbitals"
echo
geometry units angstroms
  O      0.0  0.0  0.0
  H      0.0  0.0  1.0
  H      0.0  1.0  0.0
end

basis
  H library 6-311G*
  O library 6-311G*
end

dft
  print kinetic_energy
  xc xpbe96 cpbe96
  decomp
end
task dft optimize

dft
  iterations 1
  xc xpkzb99
  decomp
  print quadrature kinetic_energy
end
task dft ignore
```

Below are the results of the exchange-only meta-GGA calculation part, and as expected we are reminded the calculation was not allowed to converge:

```
Calculation failed to converge
```

Total DFT energy =	-75.948526603774
One electron energy =	-122.899234375708
Coulomb energy =	46.745269703266
Exchange energy =	-8.857824284365
Correlation energy =	0.000000000000
Nuclear repulsion energy =	9.063262353032
Numeric. integr. density =	10.000001055407

1. C. Slater, <i>Quantum Theory of Molecules and Solids, Vol. 4</i> (McGraw-Hill, New York, 1974)	Keyword	X	C	GGA	meta	Hybrid	2nds	Ref.
2. S.J. Vosko, L. Wilk and M. Nusair, <i>Can. J. Phys.</i> 58 , 1200 (1980).	slater	*					Y	[1]
3. J. P. Perdew and A. Zunger, <i>Phys. Rev. B</i> 23 , 5048 (1981).	vwn_1		*				Y	[2]
4. J.P. Perdew and Y. Wang, <i>Phys. Rev. B</i> 45 , 13244 (1992).	vwn_2		*				Y	[2]
5. A.D. Becke, <i>Phys. Rev. A</i> 88 , 3098 (1988).	vwn_3		*				Y	[2]
6. J.P. Perdew, J.A. Chevary, S.H. Vosko, K.A. Jackson, M.R. Pederson, D.J. Singh and C. Fiolhais, <i>Phys. Rev. B</i> 46 , 6671 (1992).	vwn_4		*				Y	[2]
7. J.P. Perdew, K. Burke and M. Ernzerhof, <i>Phys. Rev. Lett.</i> 77 , 3865 (1996); 78 , 1396 (1997).	vwn_5		*				Y	[2]
8. P.W.Gill, <i>Mol. Phys.</i> 89 , 433 (1996).	vwn_1_rpa		*				Y	[2]
9. J. P. Perdew, <i>Phys. Rev. B</i> 33 , 8822 (1986).	perdew81		*				Y	[3]
10. C. Lee, W. Yang and R. G. Parr, <i>Phys. Rev. B</i> 37 , 785 (1988).	pw91lda		*				Y	[4]
11. F.A.Hamprecht, A.J.Cohen, D.J.Tozer and N.C. Handy, <i>J. Chem. Phys.</i> 109 , 6264 (1998).	becke88	*		*			Y	[5]
12. A.D.Boese, N.L.Doltsinis, N.C.Handy and M.Spriker. <i>J. Chem. Phys.</i> 112 , 1670 (2000).	xperdew91	*		*			Y	[6]
13. A.D. Becke, <i>J. Chem. Phys.</i> 98 , 1372 (1992).	xpbe96	*		*			Y	[7]
14. A.D. Becke, <i>J. Chem. Phys.</i> 98 , 5648 (1993).	gill96	*		*			Y	[8]
15. A.D.Becke, <i>J. Chem. Phys.</i> 107 , 8554 (1997).	optx	*		*			N	[20]
16. H.L.Schmider and A.D. Becke, <i>J. Chem. Phys.</i> 108 , 9624 (1998).	mpw91	*		*			Y	[23]
17. C.Adamo and V.Barone, <i>J. Chem. Phys.</i> 110 , 6158 (1998).	xf97	*		*			N	[24]
18. A.J.Cohen and N.C. Handy, <i>Chem. Phys. Lett.</i> 316 , 160 (2000).	perdew86		*	*			Y	[9]
19. A.D.Boese, N.C.Handy, <i>J. Chem. Phys.</i> 114 , 5497 (2001).	lyp		*	*			Y	[10]
20. N.C.Handy, A.J. Cohen, <i>Mol. Phys.</i> 99 , 403 (2001).	perdew91		*	*			Y	[6]
21. G. Menconi, P.J. Wilson, D.J. Tozer, <i>J. Chem. Phys.</i> 114 , 3958 (2001).	cpbe96		*	*			Y	[7]
22. P.J. Wilson, T.J. Bradley, D.J. Tozer, <i>J. Chem. Phys.</i> 115 , 9233 (2001).	cft97		*	*			N	[25]
23. C. Adamo and V. Barone, <i>J. Chem. Phys.</i> 108 , 664 (1998).	hcth	*	*	*			N	[11]
24. M.Filatov and W.Thiel, <i>Mol.Phys.</i> 91 , 847 (1997).	hcth120	*	*	*			N	[12]
25. M.Filatov and W.Thiel, <i>Int.J.Quantum Chem.</i> 62 , 603 (1997).	hcth147	*	*	*			N	[12]
26. B.J.Lynch, P.L.Fast, M.Harris and D.G.Truhlar, <i>J. Phys. Chem. A</i> 104 , 4811(2000).	hcth407	*	*	*			N	[19]
27. J.P. Perdew, S. Kurth, A. Zupan and P. Blaha, <i>Phys. Rev. Lett.</i> 82 , 2544 (1999)	becke97gga1	*	*	*			N	[18]
28. A. D. Boese, A. Chandra, J. M. L. Martin and D. Marx, <i>J. Chem. Phys.</i> 119 , 5965 (2003)	hcthp14	*	*	*			N	[21]
29. J. Tao, J. Perdew, V. Staroverov and G. Scuseria, <i>Phys. Rev. Lett.</i> 91 , 146401-1 (2003)	ft97	*	*	*			N	[24,25]
	htch407p	*	*	*			N	[28]
	xpkzb99	*			*		N	[27]
	cpkzb99		*		*		N	[27]
	xtpss03	*			*		N	[29]
	ctpss03		*		*		N	[29]
	beckehandh	*	*			*	Y	[13]
	b3lyp	*	*	*		*	Y	[14]
	acm	*	*	*		*	Y	[14]
	becke97	*	*	*		*	N	[15]
	becke97-1	*	*	*		*	N	[15]
	becke97-2	*	*	*		*	N	[22]
	becke98	*	*	*		*	N	[16]
	pbe0	*	*	*		*	Y	[17]
	mpw1k	*	*	*		*	Y	[26]

Table 11.1: Table of available Exchange (X) and Correlation (C) functionals. GGA is the Generalized Gradient Approximation, and meta refers to Meta-GGAs. The column 2nds refers to second derivatives of the energy with respect to nuclear position.

11.6 ITERATIONS — Number of SCF iterations

```
ITERATIONS <integer iterations default 30>
```

The default optimization in the DFT module is to iterate on the Kohn-Sham (SCF) equations for a specified number of iterations (default 30). The keyword that controls this optimization is `ITERATIONS`, and has the following general form,

```
iterations <integer iterations default 30>
```

The optimization procedure will stop when the specified number of iterations is reached or convergence is met. See an example that uses this directive in section 11.5.

11.7 CONVERGENCE — SCF Convergence Control

```
CONVERGENCE [energy <real energy default 1e-6>] \
  [density <real density default 1e-5>] \
  [gradient <real gradient default 5e-4>] \
  [hl_tol <real hl_tol default 0.1>] \
  [dampon <real dampon default 0.0>] \
  [dampoff <real dampoff default 0.0>] \
  [ncydp <integer ncydp default 2>] \
  [ncyds <integer ncyds default 30>] \
  [ncysh <integer ncysh default 30>] \
  [damp <integer ndamp default 0>] [nodamping] \
  [diison <real diison default 0.0>] \
  [diisoff <real diisoff default 0.0>] \
  [(diis [nfock <integer nfock default 10>]) || nodiis] \
  [levlon <real levlon default 0.0>] \
  [levloff <real levloff default 0.0>] \
  [(lshift <real lshift default 0.5>) || nolevelshifting] \
  [rabuck [n_rabuck <integer n_rabuck default 25>]]
```

Convergence is satisfied by meeting any or all of three criteria;

- convergence of the total energy; this is defined to be when the total DFT energy at iteration N and at iteration $N-1$ differ by a value less than some value (the default is $1e-6$). This value can be modified using the key word,

```
CONVERGENCE energy <real energy default 1e-6>
```

- convergence of the total density; this is defined to be when the total DFT density matrix at iteration N and at iteration $N-1$ have a RMS difference less than some value (the default is $1e-5$). This value can be modified using the key word,

```
CONVERGENCE density <real density default 1e-5>
```

- convergence of the orbital gradient; this is defined to be when the DIIS error vector becomes less than some value (the default is $5e-4$). This value can be modified using the key word,

```
CONVERGENCE gradient <real gradient default 5e-4>
```

The default optimization strategy is to immediately begin direct inversion of the iterative subspace³. Damping is also initiated (using 70% of the previous density) for the first 2 iteration. In addition, if the HOMO - LUMO gap is small and the Fock matrix somewhat diagonally dominant, then level-shifting is automatically initiated. There are a variety of ways to customize this procedure to whatever is desired.

An alternative optimization strategy is to specify, by using the change in total energy (from iterations when N and N-1), when to turn damping, level-shifting, and/or DIIS on/off. Start and stop keywords for each of these is available as,

```
CONVERGENCE [dampon <real dampon default 0.0>] \
             [dampoff <real dampoff default 0.0>] \
             [diison <real diison default 0.0>] \
             [diisoff <real diisoff default 0.0>] \
             [levlon <real levlon default 0.0>] \
             [levloff <real levloff default 0.0>]
```

So, for example, damping, DIIS, and/or level-shifting can be turned on/off as desired.

Another strategy can be to simply specify how many iterations (cycles) you wish each type of procedure to be used. The necessary keywords to control the number of damping cycles (ncydp), the number of DIIS cycles (ncyds), and the number of level-shifting cycles (ncysh) are input as,

```
CONVERGENCE [ncydp <integer ncydp default 2>] \
             [ncyds <integer ncyds default 30>] \
             [ncysh <integer ncysh default 0>]
```

The amount of damping, level-shifting, time at which level-shifting is automatically imposed, and Fock matrices used in the DIIS extrapolation can be modified by the following keywords

```
CONVERGENCE [damp <integer ndamp default 0>] \
             [diis [nfock <integer nfock default 10>]] \
             [lshift <real lshift default 0.5>] \
             [hl_tol <real hl_tol default 0.1>]]
```

Damping is defined to be the percentage of the previous iterations density mixed with the current iterations density. So, for example

```
CONVERGENCE damp 70
```

would mix 30% of the current iteration density with 70% of the previous iteration density.

Level-Shifting⁴ is defined as the amount of shift applied to the diagonal elements of the unoccupied block of the Fock matrix. The shift is specified by the keyword `lshift`. For example the directive,

```
CONVERGENCE lshift 0.5
```

causes the diagonal elements of the Fock matrix corresponding to the virtual orbitals to be shifted by 0.5 a.u. By default, this level-shifting procedure is switched on whenever the HOMO-LUMO gap is small. Small is defined by

³P. Pulay, Chem. Phys. Lett. **73**, 393 (1980) and P. Pulay, J. Comp. Chem. **3**, 566 (1982)

⁴M.F. Guest and V.R. Saunders, Mol. Phys. **28**, 819 (1974)

default to be 0.05 au but can be modified by the directive `hl_tol`. An example of changing the HOMO-LUMO gap tolerance to 0.01 would be,

```
CONVERGENCE hl_tol 0.01
```

Direct inversion of the iterative subspace with extrapolation of up to 10 Fock matrices is a default optimization procedure. For large molecular systems the amount of available memory may preclude the ability to store this number of N^2 arrays in global memory. The user may then specify the number of Fock matrices to be used in the extrapolation (must be greater than three (3) to be effective). To set the number of Fock matrices stored and used in the extrapolation procedure to 3 would take the form,

```
CONVERGENCE diis 3
```

The user has the ability to simply turn off any optimization procedures deemed undesirable with the obvious keywords,

```
CONVERGENCE [nodamping] [nodiis] [nolevelshifting]
```

For systems where the initial guess is very poor, the user can try the method described in ⁵ that makes use of **fractional occupation** of the orbital levels during the initial cycles of the SCF convergence. The input has the following form

```
CONVERGENCE rabuck [n_rabuck <integer n_rabuck default 25>]
```

where the optional value `n_rabuck` determines the number of SCF cycles during which the method will be active. For example, to set equal to 30 the number of cycles where the Rabuck method is active, you need to use the following line

```
CONVERGENCE rabuck 30
```

11.8 SMEAR — Fractional Occupation of the Molecular Orbitals

The **SMEAR** keyword is useful in cases with many degenerate states near the HOMO (eg metallic clusters)

```
SMEAR <real smear default 0.001>
```

This option allows fractional occupation of the molecular orbitals. A Gaussian broadening function of exponent `smear` is used as described in the paper: R.W. Warren RW and B.I. Dunlap, Chem. Phys. Letters **262**, 384 (1996). The user must be aware that an additional energy term is added to the total energy in order to have energies and gradients consistent.

11.9 GRID — Numerical Integration of the XC Potential

```
GRID [(xcoarse|coarse|medium|fine|xfine) default medium] \
```

⁵A. D. Rabuck and G. E. Scuseria, J. Chem. Phys **110**,695 (1999)

```
[(gausleg||lebedev ) default lebedev ] \
[(becke||erf1||erf2||ssf) default erf1] \
[(euler||mura||treatler) default mura] \
[rm <real rm default 2.0>] \
[nodisk]
```

A numerical integration is necessary for the evaluation of the exchange-correlation contribution to the density functional. The default quadrature used for the numerical integration is an Euler-MacLaurin scheme for the radial components (with a modified Mura-Knowles transformation) and a Lebedev scheme for the angular components. Within this numerical integration procedure various levels of accuracy have been defined and are available to the user. The user can specify the level of accuracy with the keywords; xcoarse, coarse, medium, fine, and xfine. The default is medium.

```
GRID [xcoarse|coarse|medium|fine|xfine]
```

Our intent is to have a numerical integration scheme which would give us approximately the accuracy defined below regardless of molecular composition.

Keyword	Total Energy Target Accuracy
xcoarse	1×10^{-4}
coarse	1×10^{-5}
medium	1×10^{-6}
fine	1×10^{-7}
xfine	1×10^{-8}

In order to determine the level of radial and angular quadrature needed to give us the target accuracy we computed total DFT energies at the LDA level of theory for many homonuclear atomic, diatomic and triatomic systems in rows 1-4 of the periodic table. In each case all bond lengths were set to twice the Bragg-Slater radius. The total DFT energy of the system was computed using the converged SCF density with atoms having radial shells ranging from 35-235 (at fixed 48/96 angular quadratures) and angular quadratures of 12/24-48/96 (at fixed 235 radial shells). The error of the numerical integration was determined by comparison to a “best” or most accurate calculation in which a grid of 235 radial points 48 theta and 96 phi angular points on each atom was used. This corresponds to approximately 1 million points per atom. The following tables were empirically determined to give the desired target accuracy for DFT total energies. These tables below show the number of radial and angular shells which the DFT module will use for for a given atom depending on the row it is in (in the periodic table) and the desired accuracy. Note, differing atom types in a given molecular system will most likely have differing associated numerical grids. The intent is to generate the desired energy accuracy (with utter disregard for speed).

Table 11.2: Program default number of radial and angular shells empirically determined for Row 1 atoms (Li → F) to reach the desired accuracies.

Keyword	Radial	Angular
xcoarse	21	194
coarse	35	302
medium	49	434
fine	70	590
xfine	100	1202

Table 11.3: Program default number of radial and angular shells empirically determined for Row 2 atoms (Na \rightarrow Cl) to reach the desired accuracies.

Keyword	Radial	Angular
xcoarse	42	194
coarse	70	302
medium	88	434
fine	123	770
xfine	125	1454

Table 11.4: Program default number of radial and angular shells empirically determined for Row 3 atoms (K \rightarrow Br) to reach the desired accuracies.

Keyword	Radial	Angular
xcoarse	75	194
coarse	95	302
medium	112	590
fine	130	974
xfine	160	1454

Table 11.5: Program default number of radial and angular shells empirically determined for Row 4 atoms (Rb \rightarrow I) to reach the desired accuracies.

Keyword	Radial	Angular
xcoarse	84	194
coarse	104	302
medium	123	590
fine	141	974
xfine	205	1454

11.9.1 Angular grids

In addition to the simple keyword specifying the desired accuracy as described above, the user has the option of specifying a custom quadrature of this type in which ALL atoms have the same grid specification. This is accomplished by using the `gausleg` keyword.

Gauss-Legendre angular grid

```
GRID gausleg <integer nradpts default 50> <integer nagrid default 10>
```

In this type of grid, the number of phi points is twice the number of theta points. So, for example, a specification of,

```
GRID gausleg 80 20
```

would be interpreted as 80 radial points, 20 theta points, and 40 phi points per center (or 64000 points per center before pruning).

Lebedev angular grid A second quadrature is the Lebedev scheme for the angular components⁶. Within this numerical integration procedure various levels of accuracy have also been defined and are available to the user. The input for this type of grid takes the form,

```
GRID lebedev <integer radpts > <integer iangquad >
```

In this context the variable `iangquad` specifies a certain number of angular points as indicated by the table below.⁷ Therefore the user can specify any number of radial points along with the level of angular quadrature (1-29).

The user can also specify grid parameters specific for a given atom type: parameters that must be supplied are: atom tag and number of radial points. As an example, here is a grid input line for the water molecule

```
grid lebedev 80 11 H 70 8 O 90 11
```

⁶The subroutine for the Lebedev grid was derived from a routine supplied by M. Causà of the University of Torino and from the grid points supplied by D.N. Laikov from Moscow State University.

⁷V.I. Lebedev and D.N. Laikov, *Doklady Mathematics* **366**, 741 (1999).

<i>IANGQUAD</i>	$N_{angular}$	l
1	38	9
2	50	11
3	74	13
4	86	15
5	110	17
6	146	19
7	170	21
8	194	23
9	230	25
10	266	27
11	302	29
12	350	31
13	434	35
14	590	41
15	770	47
16	974	53
17	1202	59
18	1454	65
19	1730	71
20	2030	77
21	2354	83
22	2702	89
23	3074	95
24	3470	101
25	3890	107
26	4334	113
27	4802	119
28	5294	125
29	5810	131

Table 11.6: List of Lebedev quadratures

11.9.2 Partitioning functions

```
GRID [(becke| |erf1| |erf2| |ssf) default erf1]
```

becke A. D. Becke, J. Chem. Phys. **88**, 1053 (1988).

ssf R.E.Stratmann, G.Scuseria and M.J.Frisch, Chem. Phys. Lett. **257**, 213 (1996).

erf1 modified ssf

erf2 modified ssf

Erfn partitioning functions

$$w_A(r) = \prod_{B \neq A} \frac{1}{2} [1 - \text{erf}(\mu'_{AB})]$$

$$\mu'_{AB} = \frac{1}{\alpha} \frac{\mu_{AB}}{(1 - \mu_{AB}^2)^n}$$

$$\mu_{AB} = \frac{r_A - r_B}{|r_A - r_B|}$$

11.9.3 Radial grids

```
GRID [[euler| |mura| |treutler] default mura]
```

euler Euler-McLaurin quadrature with the transformation devised by C.W. Murray, N.C. Handy, and G.L. Laming, Mol. Phys. **78**, 997 (1993).

mura Modification of the Murray-Handy-Laming scheme by M.E.Mura and P.J.Knowles, J Chem Phys **104**, 9848 (1996) (we are not using the scaling factors proposed in this paper).

treutler Gauss-Chebyshev using the transformation suggested by O.Treutler and R.Alrhichs, J.Chem.Phys **102**, 346 (1995).

11.9.4 Disk usage for Grid

```
NODISK
```

This keyword turns off storage of grid points and weights on disk.

11.10 TOLERANCES — Screening tolerances

```
TOLERANCES [[tight] [tol_rho <real tol_rho default 1e-10>] \
             [accCoul <integer accCoul default 8>] \
             [radius <real radius default 25.0>]]
```

The user has the option of controlling screening for the tolerances in the integral evaluations for the DFT module. In most applications, the default values will be adequate for the calculation, but different values can be specified in the input for the DFT module using the keywords described below.

The input parameter `accCoul` is used to define the tolerance in Schwarz screening for the Coulomb integrals. Only integrals with estimated values greater than $10^{(-\text{accCoul})}$ are evaluated.

```
TOLERANCES accCoul <integer accCoul default 8>
```

Screening away needless computation of the XC functional (on the grid) due to negligible density is also possible with the use of,

```
TOLERANCES tol_rho <real tol_rho default 1e-10>
```

XC functional computation is bypassed if the corresponding density elements are less than `tol_rho`.

A screening parameter, `radius`, used in the screening of the Becke or Delley spatial weights is also available as,

```
TOLERANCES radius <real radius default 25.0>
```

where `radius` is the cutoff value in bohr.

The tolerances as discussed previously are insured at convergence. More sleazy tolerances are invoked early in the iterative process which can speed things up a bit. This can also be problematic at times because it introduces a discontinuity in the convergence process. To avoid use of initial sleazy tolerances the user can invoke the `tight` option:

```
TOLERANCES tight
```

This option sets all tolerances to their default/user specified values at the very first iteration.

11.11 DIRECT and NOIO — Hardware Resource Control

```
DIRECT | INCORE
NOIO
```

The inverted charge-density and exchange-correlation matrices for a DFT calculation are normally written to disk storage. The user can prevent this by specifying the keyword `noio` within the input for the DFT directive. The input to exercise this option is as follows,

```
noio
```

If this keyword is encountered, then the two matrices (inverted charge-density and exchange-correlation) are computed “on-the-fly” whenever needed.

The `INCORE` option is always assumed to be true but can be overridden with the option `DIRECT` in which case all integrals are computed “on-the-fly”.

11.12 ODFT and MULT — Open shell systems

```
ODFT
MULT <integer mult default 1>
```

Both *closed-shell* and *open-shell* systems can be studied using the DFT module. Specifying the keyword MULT within the DFT directive allows the user to define the spin multiplicity of the system. The form of the input line is as follows;

```
MULT <integer mult default 1>
```

When the keyword MULT is specified, the user can define the integer variable mult, where mult is equal to the number of alpha electrons minus beta electrons, plus 1.

The keyword ODFT is unnecessary except in the context of forcing a singlet system to be computed as an open shell system (i.e., using a spin-unrestricted wavefunction).

11.13 SIC — Self-Interaction Correction

```
sic [perturbative || oep || oep-loc <default perturbative>]
```

The Perdew and Zunger (see J. P. Perdew and A. Zunger, Phys. Rev. B 23, 5048 (1981)) method to remove the self-interaction contained in many exchange-correlation functionals has been implemented with the Optimized Effective Potential method (see R. T. Sharp and G. K. Horton, Phys. Rev. **90**, 317 (1953), J. D. Talman and W. F. Shadwick, Phys. Rev. A **14**, 36 (1976)) within the Krieger-Li-Iafrate approximation (J. B. Krieger, Y. Li, and G. J. Iafrate, Phys. Rev. A **45**, 101 (1992); **46**, 5453 (1992); 47, 165 (1993)) Three variants of these methods are included in NWChem:

- `sic perturbative` This is the default option for the sic directive. After a self-consistent calculation, the Kohn-Sham orbitals are localized with the Foster-Boys algorithm (see section 10.15) and the self-interaction energy is added to the total energy. All exchange-correlation functionals implemented in the NWChem can be used with this option.
- `sic oep` With this option the optimized effective potential is built in each step of the self-consistent process. Because the electrostatic potential generated for each orbital involves a numerical integration, this method can be expensive.
- `sic oep-loc` This option is similar to the oep option with the addition of localization of the Kohn-Sham orbitals in each step of the self-consistent process.

With oep and oep-loc options a **xfine grid** (see section 11.9) must be used in order to avoid numerical noise, furthermore the hybrid functionals can not be used with these options. More details of the implementation of this method can be found in J. Garza, J. A. Nichols and D. A. Dixon, J. Chem. Phys. 112, 7880 (2000). The components of the sic energy can be printed out using:

```
print "SIC information"
```

11.14 MULLIKEN — Mulliken analysis

Mulliken analysis of the charge distribution is invoked by the keyword:

MULLIKEN

When this keyword is encountered, Mulliken analysis of both the input density as well as the output density will occur. For example, to perform a mulliken analysis and print the explicit population analysis of the basis functions, use the following

```
dft
  mulliken
  print "mulliken ao"
end
task dft
```

11.15 BSSE — Basis Set Superposition Error

Particular care is required to compute BSSE by the counter-poise method for the DFT module. In order to include terms deriving from the numerical grid used in the XC integration, the user must label the ghost atoms not just `bq`, but `bq` followed by the given atomic symbol. For example, the first component needed to compute the BSSE for the water dimer, should be written as follows

```
geometry h2o autosym units au
O      0.00000000    0.00000000    0.22143139
H      1.43042868    0.00000000   -0.88572555
H     -1.43042868    0.00000000   -0.88572555
bqH    0.71521434    0.00000000   -0.33214708
bqH   -0.71521434    0.00000000   -0.33214708
bqO    0.00000000    0.00000000   -0.88572555
end

basis
H library aug-cc-pvdz
O library aug-cc-pvdz
bqH library H aug-cc-pvdz
bqO library O aug-cc-pvdz
end
```

Please note that the “ghost” oxygen atom has been labeled `bqO`, and not just `bq`.

11.16 Print Control

```
PRINT | |NOPRINT
```

The `PRINT | |NOPRINT` options control the level of output in the DFT. Please see some examples using this directive in section 11.5, a sample input file. Known controllable print options are:

Name	Print Level	Description
“all vector symmetries”	high	symmetries of all molecular orbitals
“alpha partner info”	high	unpaired alpha orbital analysis
“common”	debug	dump of common blocks
“convergence”	default	convergence of SCF procedure
“coulomb fit”	high	fitting electronic charge density
“dft timings”	high	
“final vectors”	high	
“final vector symmetries”	default	symmetries of final molecular orbitals
“information”	low	general information
“initial vectors”	high	
“intermediate energy info”	high	
“intermediate evals”	high	intermediate orbital energies
“intermediate fock matrix”	high	
“intermediate overlap”	high	overlaps between the alpha and beta sets
“intermediate S2”	high	values of S2
“intermediate vectors”	high	intermediate molecular orbitals
“interm vector symm”	high	symmetries of intermediate orbitals
“io info”	debug	reading from and writing to disk
“kinetic_energy”	high	kinetic energy
“mulliken ao”	high	mulliken atomic orbital population
“multipole”	default	moments of alpha, beta, and nuclear charge densities
“parameters”	default	input parameters
“quadrature”	high	numerical quadrature
“schwarz”	high	integral screening info & stats at completion
“screening parameters”	high	integral accuracies
“semi-direct info”	default	semi direct algorithm

Table 11.7: DFT Print Control Specifications

Chapter 12

Spin-Orbit DFT (SODFT)

The spin-orbit DFT module (SODFT) in the NWChem code allows for the variational treatment of the one-electron spin-orbit operator within the DFT framework. The implementation requires the definition of an effective core potential (ECP) and a matching spin-orbit potential (SO). The current implementation does NOT use symmetry.

The actual SODFT calculation will be performed when the input module encounters the `TASK` directive (Section 5.10).

```
TASK SODFT
```

Input parameters are the same as for the DFT, see section 11 for specifications. Some of the DFT options are not available in the SODFT. These are `max_ovl` and `sic`.

Besides using the standard ECP and basis sets, see Section 8 for details, one also has to specify a spin-orbit (SO) potential. The input specification for the SO potential can be found in section 8.2. At this time we have not included any spin-orbit potentials in the basis set library.

Note: One should use a combination of ECP and SO potentials that were designed for the same size core, i.e. don't use a small core ECP potential with a large core SO potential (it will produce erroneous results).

The following is an example of a calculation of UO_2 :

```
start uo2_sodft
echo

Memory 32 mw

charge 2

geometry noautoz noautosym units angstrom
U      0.00000      0.00000      0.00000
O      0.00000      0.00000      1.68000
O      0.00000      0.00000     -1.68000
end

basis "ao basis" cartesian print
U      S
```

		12.12525300	0.02192100
		7.16154500	-0.22516000
		4.77483600	0.56029900
		2.01169300	-1.07120900
U	S		
		0.58685200	1.00000000
U	S		
		0.27911500	1.00000000
U	S		
		0.06337200	1.00000000
U	S		
		0.02561100	1.00000000
U	P		
		17.25477000	0.00139800
		7.73535600	-0.03334600
		5.15587800	0.11057800
		2.24167000	-0.31726800
U	P		
		0.58185800	1.00000000
U	P		
		0.26790800	1.00000000
U	P		
		0.08344200	1.00000000
U	P		
		0.03213000	1.00000000
U	D		
		4.84107000	0.00573100
		2.16016200	-0.05723600
		0.57563000	0.23882800
U	D		
		0.27813600	1.00000000
U	D		
		0.12487900	1.00000000
U	D		
		0.05154800	1.00000000
U	F		
		2.43644100	0.35501100
		1.14468200	0.40084600
		0.52969300	0.30467900
U	F		
		0.24059600	1.00000000
U	F		
		0.10186700	1.00000000
O	S		
		47.10551800	-0.01440800
		5.91134600	0.12956800
		0.97648300	-0.56311800
O	S		
		0.29607000	1.00000000
O	P		
		16.69221900	0.04485600

		3.90070200	0.22261300
		1.07825300	0.50018800
O	P		
		0.28418900	1.00000000
O	P		
		0.07020000	1.00000000

END

ECP

U nelec 78

U	s		
	2	4.06365300	112.92010300
	2	1.88399500	15.64750000
	2	0.88656700	-3.68997100
U	p		
	2	3.98618100	118.75801600
	2	2.00016000	15.07722800
	2	0.96084100	0.55672000
U	d		
	2	4.14797200	60.85589200
	2	2.23456300	29.28004700
	2	0.91369500	4.99802900
U	f		
	2	3.99893800	49.92403500
	2	1.99884000	-24.67404200
	2	0.99564100	1.38948000

O nelec 2

O	s		
	2	10.44567000	50.77106900
O	p		
	2	18.04517400	-4.90355100
O	d		
	2	8.16479800	-3.31212400

END

SO

U	p		
	2	3.986181	1.816350
	2	2.000160	11.543940
	2	0.960841	0.794644
U	d		
	2	4.147972	0.353683
	2	2.234563	3.499282
	2	0.913695	0.514635
U	f		
	2	3.998938	4.744214
	2	1.998840	-5.211731
	2	0.995641	1.867860

END

dft

```
mult 1
xc hfexch
odft
grid fine
convergence energy 1.000000E-06
convergence density 1.000000E-05
convergence gradient 1E-05
iterations 100
mulliken
end

task sodft
```

Chapter 13

COSMO

COSMO is the continuum solvation ‘COnductor-like Screening MOdel’ of A. Klamt and G. Schüürmann to describe dielectric screening effects in solvents.

1. A. Klamt and G. Schüürmann, *J. Chem. Soc. Perkin Trans. 2*, 1993 799 (1993).

The NWChem COSMO module implements algorithm for calculation of the energy for the following methods:

1. Restricted Hartree-Fock (RHF),
2. Restricted open-shell Hartree-Fock (ROHF),
3. Restricted Kohn-Sham DFT (DFT),
4. Unrestricted Kohn-Sham DFT (ODFT),

by determining the solvent reaction field self-consistently with the solute charge distribution from the respective methods. Note that COSMO for unrestricted Hartree-Fock (UHF) method can also be performed by invoking the DFT module with appropriate keywords.

Correlation energy of solvent molecules may also be evaluated at

1. MP2,
2. CCSD,
3. CCSD+T(CCSD),
4. CCSD(T),

levels of theory. It is cautioned, however, that these correlated COSMO calculations determine the solvent reaction field using the HF charge distribution of the solute rather than the charge distribution of the correlation theory and are not entirely self consistent in that respect. In other words, these calculations assume that the correlation effect and solvation effect are largely additive, and the combination effect thereof is neglected. COSMO for MCSCF has not been implemented yet.

In the current implementation the code calculates the gas-phase energy of the system followed by the solution-phase energy, and returns the electrostatic contribution to the solvation free energy. At the present gradients are

calculated by finite difference of the energy. Known problems include that the code does not work with spherical basis functions. The code does not calculate the non-electrostatic contributions to the free energy, except for the cavitation/dispersion contribution to the solvation free energy, which is computed and printed. It should be noted that one must in general take into account the standard state correction besides the electrostatic and cavitation/dispersion contribution to the solvation free energy, when a comparison to experimental data is made.

Invoking the COSMO solvation model is done by specifying the input COSMO input block with the input options as:

```
cosmo
  [off]
  [dielec <real dielec default 78.4>]
  [radius <real atom1>
    <real atom2>
    . . .
    <real atomN>]
  [rsolv <real rsolv default 0.00>]
  [iscren <integer iscren default 0>]
  [minbem <integer minbem default 2>]
  [maxbem <integer maxbem default 3>]
  [ificos <integer ificos default 0>]
  [lineq <integer lineq default 1>]
end
```

followed by the task directive specifying the wavefunction and type of calculation, e.g., `task scf energy`, `task mp2 energy`, `task dft optimize`, etc.

`off` can be used to turn off COSMO in a compound (multiple task) run. By default, once the COSMO solvation model has been defined it will be used in subsequent calculations. Add the keyword `off` if COSMO is not needed in subsequent calculations.

`Dielec` is the value of the dielectric constant of the medium, with a default value of 78.4 (the dielectric constant for water).

`Radius` is an array that specifies the radius of the spheres associated with each atom and that make up the molecule-shaped cavity. Default values are Van der Waals radii. Values are in units of angstroms. The codes uses the following Van der Waals radii by default:

```
data vdwr(103) /
1  0.80,0.49,0.00,0.00,0.00,1.65,1.55,1.50,1.50,0.00,
2  2.30,1.70,2.05,2.10,1.85,1.80,1.80,0.00,2.80,2.75,
3  0.00,0.00,1.20,0.00,0.00,0.00,2.70,0.00,0.00,0.00,
4  0.00,0.00,0.00,1.90,1.90,0.00,0.00,0.00,0.00,1.55,
5  0.00,1.64,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,
6  0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,
7  0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,
8  0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,
9  0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,
1  0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,0.00,1.65,
2  0.00,0.00,0.00/
```

with 0.0 values replaced by 1.80. Other radii can be used as well. See for examples:

1. E. V. Stefanovich and T. N. Truong, Chem. Phys. Lett. 244, 65 (1995).

2. V. Barone, M. Cossi, and J. Tomasi, *J. Chem. Phys.* 107, 3210 (1997).

`rsolv` is a parameter used to define the solvent accessible surface. See the original reference of Klamt and Schuurmann for a description. The default value is 0.00 (in angstroms).

`iscren` is a flag to define the dielectric charge scaling option. “`iscren 1`” implies the original scaling from Klamt and Schuurmann, mainly “ $(\epsilon - 1)/(\epsilon + 1/2)$ ”, where ϵ is the dielectric constant. “`iscren 0`” implies the modified scaling suggested by Stefanovich and Truong, mainly “ $(\epsilon - 1)/\epsilon$ ”. Default is to use the modified scaling. For high dielectric the difference between the scaling is not significant.

The next three parameters define the tessellation of the unit sphere. The approach follows the original proposal by Klamt and Schuurmann. A very fine tessellation is generated from `maxbem` refining passes starting from either an octahedron or an icosahedron. The boundary elements created with the fine tessellation are condensed down to a coarser tessellation based on `minbem`. The induced point charges from the polarization of the medium are assigned to the centers of the coarser tessellation. Default values are “`minbem 2`” and “`maxbem 3`”. The flag `ificos` serves to select the original tessellation, “`ificos 0`” for an octahedron (default) and “`ificos 1`” for an icosahedron. Starting from an icosahedron yields a somewhat finer tessellation that converges somewhat faster. Solvation energies are not really sensitive to this choice for sufficiently fine tessellations.

The `lineq` parameter serves to select the numerical algorithm to solve the linear equations yielding the effective charges that represent the polarization of the medium. “`lineq 0`” selects an iterative method (default), “`lineq 1`” selects a dense matrix linear equation solver. For large molecules where the number of effective charges is large, the codes selects the iterative method.

The following example is for a water molecule in ‘water’, using the HF/6-31G** level of theory:

```
start
echo
  title "h2o"
geometry
o          .0000000000          .0000000000          -.0486020332
h          .7545655371          .0000000000          .5243010666
h          -.7545655371          .0000000000          .5243010666
end
basis segment cartesian
  o library 6-31g**
  h library 6-31g**
end
cosmo
  dielec 78.0
  radius 1.40
         1.16
         1.16
  rsolv 0.50
  lineq 0
end
task scf energy
```


Chapter 14

CIS, TDHF, and TDDFT

14.1 Overview

NWChem supports a spectrum of single excitation theories for vertical excitation energy calculations, namely, configuration interaction singles (CIS),¹ time-dependent Hartree–Fock (TDHF or also known as random-phase approximation RPA), time-dependent density functional theory (TDDFT),² and Tamm–Dancoff approximation to TDDFT.³ These methods are implemented in a single framework that invokes Davidson’s trial vector algorithm (or its modification for a non-Hermitian eigenvalue problem).⁴ The capabilities of the module are summarized as follows:

- Vertical excitation energies,
- Spin-restricted singlet and triplet excited states for closed-shell systems,
- Spin-unrestricted doublet, etc., excited states for open-shell systems,
- Tamm–Dancoff and full time-dependent linear response theories,
- Davidson’s trial vector algorithm,
- Symmetry (irreducible representation) characterization and specification,
- Spin multiplicity characterization and specification,
- Transition moments and oscillator strengths,
- Geometrical first and second derivatives of vertical excitation energies by numerical differentiation,
- Disk-based and fully incore algorithms,
- Multiple and single trial-vector processing algorithms,
- Frozen core and virtual approximation.

New capability added in the latest version (4.6) is:

¹J. B. Foreman, M. Head-Gordon, J. A. Pople, and M. J. Frisch, *J. Phys. Chem.* **96**, 135 (1992).

²C. Jamorski, M. E. Casida, and D. R. Salahub, *J. Chem. Phys.* **104**, 5134 (1996); R. Bauernschmitt and R. Ahlrichs, *Chem. Phys. Lett.* **256**, 454 (1996); R. Bauernschmitt, M. Häser, O. Treutler, and R. Ahlrichs, *Chem. Phys. Lett.* **264**, 573 (1997).

³S. Hirata and M. Head-Gordon, *Chem. Phys. Lett.* **314**, 291 (1999).

⁴E. R. Davidson, *J. Comput. Phys.* **17**, 87 (1975); J. Olsen, H. J. Aa. Jensen, and P. Jørgensen, *J. Comput. Phys.* **74**, 265 (1988).

- Asymptotically correct exchange-correlation potential by van Leeuwen and Baerends,⁵
- Asymptotic correction by Casida and Salahub,⁶
- Asymptotic correction by Hirata, Zhan, Aprà, Windus, and Dixon.⁷

These are very effective way to rectify the shortcomings of TDDFT when applied to Rydberg excited states (see below).

14.2 Performance of CIS, TDHF, and TDDFT methods

The accuracy of CIS and TDHF for excitation energies of closed-shell systems are comparable to each other, and are normally considered a zeroth-order description of the excitation process. These methods are particularly well balanced in describing Rydberg excited states, in contrast to TDDFT. However, for open-shell systems, the errors in the CIS and TDHF excitation energies are often excessive, primarily due to the multi-determinantal character of the ground and excited state wave functions of open-shell systems in a HF reference.⁸ The scaling of the computational cost of a CIS or TDHF calculation per state with respect to the system size is the same as that for a HF calculation for the ground state, since the critical step of the both methods are the Fock build, namely, the contraction of two-electron integrals with density matrices. It is usually necessary to include two sets of diffuse exponents in the basis set to properly account for the diffuse Rydberg excited states of neutral species.

The accuracy of TDDFT may vary depending on the exchange-correlation functional. In general, the exchange-correlation functionals that are widely used today and are implemented in NWChem work well for low-lying valence excited states. However, for high-lying diffuse excited states and Rydberg excited states in particular, TDDFT employing these conventional functionals breaks down and the excitation energies are substantially underestimated. This is because of the fact that the exchange-correlation potentials generated from these functionals decay too rapidly (exponentially) as opposed to the slow $-1/r$ asymptotic decay of the true potential. A rough but useful index is the negative of the highest occupied KS orbital energy; when the calculated excitation energies become close to this threshold, these numbers are most likely underestimated relative to experimental results.⁹ It appears that TDDFT provides a better-balanced description of radical excited states.¹⁰ This may be traced to the fact that, in DFT, the ground state wave function is represented well as a single KS determinant, with less multi-determinantal character and less spin contamination, and hence the excitation thereof is described well as a simple one electron transition. The computational cost per state of TDDFT calculations scales as the same as the ground state DFT calculations, although the prefactor of the scaling may be much greater in the former.

A very simple and effective way to rectify the TDDFT's failure for Rydberg excited states has been proposed by Tozer and Handy¹¹ and by Casida and Salahub (see previous reference). They proposed to splice a $-1/r$ asymptotic tail to an exchange-correlation potential that does not have the correct asymptotic behavior. Because the approximate exchange-correlation potentials are too shallow everywhere, a negative constant must be added to them before they can be spliced to the $-1/r$ tail seamlessly in a region that is not sensitive to chemical effects or to the long-range behavior. The negative constant or the shift is usually taken to be the difference of the HOMO energy from the true ionization potential, which can be obtained either from experiment or from a Δ SCF calculation. Recently, we proposed a new, expedient, and self-contained asymptotic correction that does not require an ionization potential (or shift) as an external parameter from a separate calculation.¹² In this scheme, the shift is computed by a semi-empirical

⁵R. van Leeuwen and E. J. Baerends, *Phys. Rev. A* **49**, 2421 (1994).

⁶M. E. Casida, C. Jamorski, K. C. Casida, and D. R. Salahub, *J. Chem. Phys.* **108**, 4439 (1998).

⁷S. Hirata, C.-G. Zhan, E. Aprà, T. L. Windus, and D. A. Dixon, *J. Phys. Chem. A* **107**, 10154 (2003).

⁸D. Maurice and M. Head-Gordon, *J. Phys. Chem.* **100**, 6131 (1996).

⁹M. E. Casida, C. Jamorski, K. C. Casida, and D. R. Salahub, *J. Chem. Phys.* **108**, 4439 (1998).

¹⁰S. Hirata and M. Head-Gordon, *Chem. Phys. Lett.* **302**, 375 (1999).

¹¹D. J. Tozer and N. C. Handy, *J. Chem. Phys.* **109**, 10180 (1998).

¹²S. Hirata, C.-G. Zhan, E. Aprà, T. L. Windus, and D. A. Dixon, *J. Phys. Chem. A* **107**, 10154 (2003).

formula proposed by Zhan, Nichols, and Dixon.¹³ Both Casida-Salahub scheme and this new asymptotic correction scheme give considerably improved (Koopmans type) ionization potentials and Rydberg excitation energies. The latter, however, supply the shift by itself unlike to former.

14.3 Input syntax

The module is called TDDFT as TDDFT employing a hybrid HF-DFT functional encompasses all of the above-mentioned methods implemented. To use this module, one needs to specify TDDFT on the task directive, e.g.,

```
TASK TDDFT ENERGY
```

for a single-point excitation energy calculation, and

```
TASK TDDFT OPTIMIZE
```

for an excited-state geometry optimization (and perhaps an adiabatic excitation energy calculation), and

```
TASK TDDFT FREQUENCIES
```

for an excited-state vibrational frequency calculation. The TDDFT module first invokes DFT module for a ground-state calculation (regardless of whether the calculations uses a HF reference as in CIS or TDHF or a DFT functional), and hence there is no need to perform a separate ground-state DFT calculation prior to calling a TDDFT task. When no second argument of the task directive is given, a single-point excitation energy calculation will be assumed. For geometry optimizations, it is usually necessary to specify the target excited state and its irreducible representation it belongs to. See the subsections TARGET and TARGETSYM for more detail.

Individual parameters and keywords may be supplied in the TDDFT input block. The syntax is:

```
TDDFT
  [(CIS||RPA) default RPA]
  [NROOTS <integer nroots default 1>]
  [MAXVECS <integer maxvecs default 1000>]
  [(SINGLET||NOSINGLET) default SINGLET]
  [(TRIPLET||NOTRIPLET) default TRIPLET]
  [THRESH <double thresh default 1e-4>]
  [MAXITER <integer maxiter default 100>]
  [TARGET <integer target default 1>]
  [TARGETSYM <character targetsym default 'none'>]
  [SYMMETRY]
  [ALGORITHM <integer algorithm default 0>]
  [FREEZE [[core] (atomic || <integer nfzc default 0>)] \
    [virtual <integer nfzv default 0>]]
  [PRINT (none||low||medium||high||debug)
    <string list_of_names ...>]
END
```

The user can also specify the reference wave function in the DFT input block (even when CIS and TDHF calculations are requested). See the section of Sample input and output for more details.

Since each keyword has a default value, a minimal input file will be

¹³C.-G. Zhan, J. A. Nichols, and D. A. Dixon, *J. Phys. Chem. A* **107**, 4184 (2003).

```
GEOMETRY
Be 0.0 0.0 0.0
END

BASIS
Be library 6-31G**
END

TASK TDDFT ENERGY
```

Note that the keyword for the asymptotic correction must be given in the DFT input block, since all the effects of the correction (and also changes in the computer program) occur in the SCF calculation stage. See Chapter 11 (keyword CS00 and LB94) for details.

14.4 Keywords of TDDFT input block

14.4.1 CIS and RPA — the Tamm–Dancoff approximation

These keywords toggle the Tamm–Dancoff approximation. CIS means that the Tamm–Dancoff approximation is used and the CIS or Tamm–Dancoff TDDFT calculation is requested. RPA, which is the default, requests TDHF (RPA) or TDDFT calculation.

The performance of CIS (Tamm–Dancoff TDDFT) and RPA (TDDFT) are comparable in accuracy. However, the computational cost is slightly greater in the latter due to the fact that the latter involves a non-Hermitian eigenvalue problem and requires left and right eigenvectors while the former needs just one set of eigenvectors of a Hermitian eigenvalue problem. The latter has much greater chance of aborting the calculation due to triplet near instability or other instability problems.

14.4.2 NROOTS — the number of excited states

One can specify the number of excited state roots to be determined. The default value is 1. It is advised that the users request several more roots than actually needed, since owing to the nature of the trial vector algorithm, some low-lying roots can be missed when they do not have sufficient overlap with the initial guess vectors.

14.4.3 MAXVECS — the subspace size

This keyword limits the subspace size of Davidson’s algorithm; in other words, it is the maximum number of trial vectors that the calculation is allowed to hold. Typically, 10 to 20 trial vectors are needed for each excited state root to be converged. However, it need not exceed the product of the number of occupied orbitals and the number of virtual orbitals. The default value is 1000.

14.4.4 SINGLET and NOSINGLET — singlet excited states

SINGLET (NOSINGLET) requests (suppresses) the calculation of singlet excited states when the reference wave function is closed shell. The default is SINGLET.

14.4.5 TRIPLET and NOTRIPLET — triplet excited states

TRIPLET (NOTRIPLET) requests (suppresses) the calculation of triplet excited states when the reference wave function is closed shell. The default is TRIPLET.

14.4.6 THRESH — the convergence threshold of Davidson iteration

This keyword specifies the convergence threshold of Davidson's iterative algorithm to solve a matrix eigenvalue problem. The threshold refers to the norm of residual, namely, the difference between the left-hand side and right-hand side of the matrix eigenvalue equation with the current solution vector. With the default value of $1e-4$, the excitation energies are usually converged to $1e-5$ hartree.

14.4.7 MAXITER — the maximum number of Davidson iteration

It typically takes 10–30 iterations for the Davidson algorithm to get converged results. The default value is 100.

14.4.8 TARGET and TARGETSYM — the target root and its symmetry

At the moment, the first and second geometrical derivatives of excitation energies that are needed in force, geometry, and frequency calculations are obtained by numerical differentiation. These keywords may be used to specify which excited state root is being used for the geometrical derivative calculation. For instance, when TARGET 3 and TARGETSYM alg are included in the input block, the total energy (ground state energy plus excitation energy) of the third lowest excited state root (excluding the ground state) transforming as the irreducible representation alg will be passed to the module which performs the derivative calculations. The default values of these keywords are 1 and none, respectively.

The keyword TARGETSYM is essential in excited state geometry optimization, since it is very common that the order of excited states changes due to the geometry changes in the course of optimization. Without specifying the TARGETSYM, the optimizer could (and would likely) be optimizing the geometry of an excited state that is different from the one the user had intended to optimize at the starting geometry. On the other hand, in the frequency calculations, TARGETSYM must be none, since the finite displacements given in the course of frequency calculations will lift the spatial symmetry of the equilibrium geometry. When these finite displacements can alter the order of excited states including the target state, the frequency calculation is not be feasible.

14.4.9 SYMMETRY — restricting the excited state symmetry

By adding this keyword to the input block, the user can request the module to generate the initial guess vectors transforming as the same irreducible representation as TARGETSYM. This causes the final excited state roots be (exclusively) dominated by those with the specified irreducible representation. This may be useful, when the user is interested in just the optically allowed transitions, or in the geometry optimization of an excited state root with a particular irreducible representation. By default, this option is not set. TARGETSYM must be specified when SYMMETRY is invoked.

14.4.10 ALGORITHM — algorithms for tensor contractions

There are four distinct algorithms to choose from, and the default value of 0 (optimal) means that the program makes an optimal choice from the four algorithms on the basis of available memory. In the order of decreasing memory

requirement, the four algorithms are:

- ALGORITHM 1 : Incore, multiple tensor contraction,
- ALGORITHM 2 : Incore, single tensor contraction,
- ALGORITHM 3 : Disk-based, multiple tensor contraction,
- ALGORITHM 4 : Disk-based, single tensor contraction.

The incore algorithm stores all the trial and product vectors in memory across different nodes with the GA, and often decreases the `MAXITER` value to accommodate them. The disk-based algorithm stores the vectors on disks across different nodes with the DRA, and retrieves each vector one at a time when it is needed. The multiple and single tensor contraction refers to whether just one or more than one trial vectors are contracted with integrals. The multiple tensor contraction algorithm is particularly effective (in terms of speed) for CIS and TDHF, since the number of the direct evaluations of two-electron integrals is diminished substantially.

14.4.11 FREEZE — the frozen core/virtual approximation

Some of the lowest-lying core orbitals and/or some of the highest-lying virtual orbitals may be excluded in the CIS, TDHF, and TDDFT calculations by this keyword (this does not affect the ground state HF or DFT calculation). No orbitals are frozen by default. To exclude the atom-like core regions altogether, one may request

```
FREEZE atomic
```

To specify the number of lowest-lying occupied orbitals be excluded, one may use

```
FREEZE 10
```

which causes 10 lowest-lying occupied orbitals excluded. This is equivalent to writing

```
FREEZE core 10
```

To freeze the highest virtual orbitals, use the `virtual` keyword. For instance, to freeze the top 5 virtuals

```
FREEZE virtual 5
```

14.4.12 PRINT — the verbosity

This keyword changes the level of output verbosity. One may also request some particular items in Table 14.1 printed.

14.5 Sample input

The following is a sample input for a spin-restricted TDDFT calculation of singlet excitation energies for the water molecule at the B3LYP/6-31G*.

Table 14.1: Printable items in the TDDFT modules and their default print levels.

Item	Print Level	Description
"timings"	high	CPU and wall times spent in each step
"trial vectors"	high	Trial CI vectors
"initial guess"	debug	Initial guess CI vectors
"general information"	default	General information
"xc information"	default	HF/DFT information
"memory information"	default	Memory information
"convergence"	debug	Convergence
"subspace"	debug	Subspace representation of CI matrices
"transform"	debug	MO to AO and AO to MO transformation of CI vectors
"diagonalization"	debug	Diagonalization of CI matrices
"iteration"	default	Davidson iteration update
"contract"	debug	Integral transition density contraction
"ground state"	default	Final result for ground state
"excited state"	low	Final result for target excited state

```
START h2o
```

```
TITLE "B3LYP/6-31G* H2O"
```

```
GEOMETRY
```

```
O    0.00000000    0.00000000    0.12982363
H    0.75933475    0.00000000   -0.46621158
H   -0.75933475    0.00000000   -0.46621158
END
```

```
BASIS
```

```
* library 6-31G*
END
```

```
DFT
```

```
XC B3LYP
END
```

```
TDDFT
```

```
RPA
NROOTS 20
END
```

```
TASK TDDFT ENERGY
```

To perform a spin-unrestricted TDHF/aug-cc-pVDZ calculation for the CO⁺ radical,

```
START co
```

```
TITLE "TDHF/aug-cc-pVDZ CO+"
```

```
CHARGE 1
```

```
GEOMETRY
C 0.0 0.0 0.0
O 1.5 0.0 0.0
END

BASIS
* library aug-cc-pVDZ
END

DFT
XC HFexch
MULT 2
END

TDDFT
RPA
NROOTS 5
END

TASK TDDFT ENERGY
```

A geometry optimization followed by a frequency calculation for an excited state is carried out for BF at the CIS/6-31G* level in the following sample input.

```
START bf

TITLE "CIS/6-31G* BF optimization frequencies"

GEOMETRY
B 0.0 0.0 0.0
F 0.0 0.0 1.2
END

BASIS
* library 6-31G*
END

DFT
XC HFexch
END

TDDFT
CIS
NROOTS 3
NOTRIPLET
TARGET 1
END

TASK TDDFT OPTIMIZE
```

```
TASK TDDFT FREQUENCIES
```

TDDFT with an asymptotically corrected SVWN exchange-correlation potential. Casida-Salahub scheme has been used with the shift value of 0.1837 a.u. supplied as an input parameter.

```
START tddft_ac_co
```

```
GEOMETRY
```

```
O 0.0 0.0 0.0000
```

```
C 0.0 0.0 1.1283
```

```
END
```

```
BASIS SPHERICAL
```

```
C library aug-cc-pVDZ
```

```
O library aug-cc-pVDZ
```

```
END
```

```
DFT
```

```
XC Slater VWN_5
```

```
CS00 0.1837
```

```
END
```

```
TDDFT
```

```
NROOTS 12
```

```
END
```

```
TASK TDDFT ENERGY
```

TDDFT with an asymptotically corrected B3LYP exchange-correlation potential. Hirata-Zhan-Apra-Windus-Dixon scheme has been used (this is only meaningful with B3LYP functional).

```
START tddft_ac_co
```

```
GEOMETRY
```

```
O 0.0 0.0 0.0000
```

```
C 0.0 0.0 1.1283
```

```
END
```

```
BASIS SPHERICAL
```

```
C library aug-cc-pVDZ
```

```
O library aug-cc-pVDZ
```

```
END
```

```
DFT
```

```
XC B3LYP
```

```
CS00
```

```
END
```

```
TDDFT
```

```
NROOTS 12
```

```
END
```

TASK TDDFT ENERGY

Chapter 15

Tensor Contraction Engine Module: CI, MBPT, and CC

15.1 Overview

The Tensor Contraction Engine (TCE) Module of NWChem implements a variety of approximations that converge at the exact solutions of Schrödinger equation. They include configuration interaction theory through singles, doubles, triples, and quadruples substitutions, coupled-cluster theory through connected singles, doubles, triples, and quadruples substitutions, and many-body perturbation theory through fourth order in its tensor formulation. Not only optimized parallel programs of some of these high-end correlation theories are new, but also the way in which they have been developed is unique. The working equations of all of these methods have been derived completely automatically by a symbolic manipulation program called a Tensor Contraction Engine (TCE), and the optimized parallel programs have also been computer-generated by the same program, which were interfaced to NWChem. The development of the TCE program and this portion of the NWChem program has been financially supported by the United States Department of Energy, Office of Science, Office of Basic Energy Science, through the SciDAC program.

The capabilities of the module include:

- Restricted Hartree–Fock, unrestricted Hartree–Fock, and restricted open-shell Hartree–Fock references,
- Restricted KS DFT and unrestricted KS DFT references,
- Unrestricted configuration interaction theory (CISD, CISDT, and CISDTQ),
- Unrestricted coupled-cluster theory (LCCD, CCD, LCCSD, CCSD, QCISD, CCSDT, CCSDTQ),
- Unrestricted iterative many-body perturbation theory [MBPT(2), MBPT(3), MBPT(4)] in its tensor formulation,

New capabilities added in the version 4.6 are:

- Unrestricted coupled-cluster singles and doubles with perturbative connected triples {CCSD(T), CCSD[T]},
- Unrestricted equation-of-motion coupled-cluster theory (EOM-CCSD, EOM-CCSDT, EOM-CCSDTQ) for excitation energies, transition moments and oscillator strengths, and excited-state dipole moments,
- Unrestricted coupled-cluster theory (CCSD, CCSDT, CCSDTQ) for dipole moments.

Version 4.6 and onwards the distributed binary executables do not contain CCSDTQ and its derivative methods, owing to their large volume. The source code includes them, so a user can reinstate them by `setenv CCSDTQ yes` and recompile TCE module. The following optimizations have been used in the module:

- Spin symmetry (spin integration is performed wherever possible within the unrestricted framework, making the present unrestricted program optimal for an open-shell system. The spin adaption was not performed, although in a restricted calculation for a closed-shell system, certain spin blocks of integrals and amplitudes are further omitted by symmetry, and consequently, the present unrestricted CCSD requires only twice as many operations as a spin-adapted restricted CCSD for a closed-shell system),
- Point-group symmetry,
- Index permutation symmetry,
- Runtime orbital range tiling for memory management,
- Dynamic load balancing (local index sort and matrix multiplications) parallelism,
- Multiple parallel I/O schemes including fully incore algorithm using Global Arrays,
- Frozen core and virtual approximation.
- DIIS extrapolation and Jacobi update of excitation amplitudes

This extensible module is designed such that an existing or new model of many-electron theory can be added and further optimization can be incorporated with ease by virtue of the TCE. This module is still being actively enhanced by the TCE and we hope to include more models and optimizations in future releases!

15.2 Performance of CI, MBPT, and CC methods

For reviews or tutorials of these highly-accurate correlation methods, the user is referred to:

- A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry: Introduction to Advanced Electronic Structure Theory*,
- R. J. Bartlett and J. F. Stanton, *Applications of Post-Hartree-Fock Methods: A Tutorial*, in *Reviews in Computational Chemistry*, Volume V,
- R. J. Bartlett, *Coupled-Cluster Theory: An Overview of Recent Developments*, in *Modern Electronic Structure Theory*, Part II,
- B. O. Roos (editor), *Lecture Notes in Quantum Chemistry I and II*.

For algorithms and applications of TCE, see:

- S. Hirata, *J. Phys. Chem. A* **107**, 9887 (2003).
- S. Hirata, T. Yanai, W. A. de Jong, T. Nakajima, and K. Hirao, *J. Chem. Phys.* **120**, 3297 (2004).
- S. Hirata, *J. Chem. Phys.* (in press) (2004).

15.3 Algorithms of CI, MBPT, and CC methods

15.3.1 Spin, spatial, and index permutation symmetry

The TCE thoroughly analyzes the working equation of many-electron theory models and automatically generates a program that takes full advantage of these symmetries at the same time. To do so, the TCE first recognizes the index permutation symmetries among the working equations, and perform strength reduction and factorization by carefully monitoring the index permutation symmetries of intermediate tensors. Accordingly, every input and output tensor (such as integrals, excitation amplitudes, residuals) has just two independent but strictly ordered index strings, and each intermediate tensor has just four independent but strictly ordered index strings. The operation cost and storage size of tensor contraction is minimized by using the index range restriction arising from these index permutation symmetries and also spin and spatial symmetry integration.

15.3.2 Runtime orbital range tiling

To maintain the peak local memory usage at a manageable level, in the beginning of the calculation, the orbitals are rearranged into tiles (blocks) that contains orbitals with the same spin and spatial symmetries. So the tensor contractions in these methods are carried out at the tile level; the spin, spatial, and index permutation symmetry is employed to reduce the operation and storage cost at the tile level also.

15.3.3 Dynamic load balancing parallelism

In a parallel execution, dynamic load balancing of tile-level local tensor index sorting and local tensor contraction (matrix multiplication) will be invoked.

15.3.4 Parallel I/O schemes

Each process is assigned a local tensor index sorting and tensor contraction dynamically. It must first retrieve the tiles of input tensors, and perform these local operations, and accumulate the output tensors to the storage. We have developed a uniform interface for these I/O operations to either (1) a global file on a global file system, (2) a global memory on a global or distributed memory system, and (3) semi-replicated files on a distributed file systems. Some of these operations depend on the ParSoft library.

15.4 Input syntax

The keyword to invoke the many-electron theories in the module is TCE. To perform a single-point energy calculation, include

```
TASK TCE ENERGY
```

in the input file, which may be preceded by the TCE input block that details the calculations:

```
TCE
  [(DFT||HF||SCF) default HF=SCF]
  [FREEZE [[core] (atomic || <integer nfzc default 0>)] \
    [virtual <integer nfzv default 0>]]
```

```
[ (LCCD| |CCD| |CCSD| |LCCSD| |CCSDT| |CCSDTQ| | \
  CCSD(T)| |CCSD[T]| |QCISD| |CISD| |CISDT| |CISDTQ| | \
  MBPT2| |MBPT3| |MBPT4| |MP2| |MP3| |MP4) default CCSD]
[THRESH <double thresh default 1e-6>]
[MAXITER <integer maxiter default 100>]
[PRINT (none| |low| |medium| |high| |debug)
  <string list_of_names ...>]
[IO (fortran| |eaf| |ga| |sf| |replicated| |dra| |ga_eaf) default ga]
[DIIS <integer diis default 5>]
[NROOTS <integer nroots default 0>]
[TARGET <integer target default 1>]
[TARGETSYM <character targetsym default 'none'>]
[SYMMETRY]
[DIPOLE]
[TILESIZE <no default (automatically adjusted)>]
[(NO)FOCK <logical recompf default .true.>]
[FRAGMENT <default -1 (off)>]
END
```

Also supported are energy gradient calculation, geometry optimization, and vibrational frequency (or hessian) calculation, on the basis of numerical differentiation. To perform these calculations, use

```
TASK TCE GRADIENT
```

or

```
TASK TCE OPTIMIZE
```

or

```
TASK TCE FREQUENCIES
```

Alternatively, more descriptive keywords for each individual method can be used. For instance, to perform a CCSDT energy, gradient, etc. calculation, use

```
TASK UCCSDT ENERGY
```

or

```
TASK UCCSDT GRADIENT
```

or

```
TASK UCCSDT OPTIMIZE
```

or

```
TASK UCCSDT FREQUENCIES
```

with an (optional) input block enclosed either by UCCSDT and END or by UCC and END. The keywords for individual methods of TCE module always start with letter U which stands for “unrestricted” to avoid confusion with other related methods (such as spin-restricted CCSD and various canonical MP2 implementation) already in place in NWChem.

```
(UCCSDT|UCC)
  [(DFT|HF|SCF) default HF=SCF]
  [FREEZE [[core] (atomic || <integer nfzc default 0>)] \
    [virtual <integer nfzv default 0>]]
  [THRESH <double thresh default 1e-6>]
  [MAXITER <integer maxiter default 100>]
  [PRINT (none||low|medium|high|debug)]
    <string list_of_names ...>]
  [IO (fortran|eaf|ga|sf|replicated|dra|ga_eaf) default ga]
  [DIIS <integer diis default 5>]
  [NROOTS <integer nroots default 0>]
  [TARGET <integer target default 1>]
  [TARGETSYM <character targetsym default 'none'>]
  [SYMMETRY]
  [DIPOLE]
  [TILESIZE <no default (automatically adjusted)>]
  [(NO)FOCK <logical recompf default .true.>]
  [FRAGMENT <default -1 (off)>]
END
```

When a method (CCSDT in this example) is specified in the task directive, a duplicate method specification is not necessary (indeed not allowed) in the corresponding (UCCSDT or UCC in this case) input block. The keywords of the other methods for task directive are:

```
TASK (UCCD|ULCCD|UCCSD|ULCCSD|UQCISD|UCCSDT|UCCSDTQ) ENERGY
```

or

```
TASK (UCISD|UCISDT|UCISDTQ) ENERGY
```

or

```
TASK (UMP2|UMP3|UMP4|UMBPT2|UMBPT3|UMBPT4) ENERGY
```

etc. The input block can be specified by the same name (UCISDT and END block for TASK UCISDT ENERGY) or UCC for the CC family, UCI for the CI family, and UMP or UMBPT for the MP family of methods.

The user may also specify the parameters of reference wave function calculation in a separate block for either HF (SCF) or DFT, depending on the first keyword in the above syntax.

Since each keyword has a default value, a minimal input file will be

```
GEOMETRY
Be 0.0 0.0 0.0
END
```

```
BASIS
```

```
Be library cc-pVDZ
END
```

```
TASK TCE ENERGY
```

which performs a CCSD/cc-pVDZ calculation of the Be atom in its singlet ground state with a spin-restricted HF reference.

15.5 Keywords of TCE input block

15.5.1 HF, SCF, or DFT — the reference wave function

This keyword tells the module which of the HF (SCF) or DFT module is going to be used for the calculation of a reference wave function. The keyword HF and SCF are one and the same keyword internally, and are default. When these are used, the details of the HF (SCF) calculation can be specified in the SCF input block, whereas if DFT is chosen, DFT input block may be provided.

For instance, RHF-RCCSDT calculation (R standing for spin-restricted) can be performed with the following input blocks:

```
SCF
SINGLET
RHF
END
```

```
TCE
SCF
CCSDT
END
```

```
TASK TCE ENERGY
```

or

```
SCF
SINGLET
RHF
END
```

```
UCCSDT
SCF
END
```

```
TASK UCCSDT ENERGY
```

or

```
SCF
SINGLET
```

```
RHF
END
```

```
UCC
SCF
END
```

```
TASK UCCSDT ENERGY
```

This calculation (and any correlation calculation in the TCE module using a RHF or RDFT reference for a closed-shell system) skips the storage and computation of all β spin blocks of integrals and excitation amplitudes. ROHF-UCCSDT (U standing for spin-unrestricted) for an open-shell doublet system can be requested by

```
SCF
DOUBLET
ROHF
END
```

```
TCE
SCF
CCSDT
END
```

```
TASK TCE ENERGY
```

and likewise, UHF-UCCSDT for an open-shell doublet system can be specified with

```
SCF
DOUBLET
UHF
END
```

```
TCE
SCF
CCSDT
END
```

```
TASK TCE ENERGY
```

The operation and storage costs of the last two calculations are identical. To use the KS DFT reference wave function for a UCCSD calculation of an open-shell doublet system,

```
DFT
ODFT
MULT 2
END
```

```
TCE
DFT
CCSD
END
```

TASK TCE ENERGY

Note that the default model of the DFT module is LDA.

15.5.2 CCSD,CCSDT,CCSDTQ,CISD,CISDT,CISDTQ,MBPT2,MBPT3,MBPT4, etc. — the correlation model

These keywords stand for the following models:

- LCCD: linearized coupled-cluster doubles,
- CCD: coupled-cluster doubles,
- LCCSD: linearized coupled-cluster singles & doubles,
- CCSD: coupled-cluster singles & doubles (also EOM-CCSD),
- CCSDT: coupled-cluster singles, doubles, & triples (also EOM-CCSDT),
- CCSDTQ: coupled-cluster singles, doubles, triples, & quadruples (also EOM-CCSDTQ),
- CCSD(T): CCSD and perturbative connected triples,
- CCSD[T]: CCSD and perturbative connected triples,
- QCISD: quadratic configuration interaction singles & doubles,
- CISD: configuration interaction singles & doubles,
- CISDT: configuration interaction singles, doubles, & triples,
- CISDTQ: configuration interaction singles, doubles, triples, & quadruples,
- MBPT2=MP2: iterative tensor second-order many-body or Møller–Plesset perturbation theory,
- MBPT3=MP3: iterative tensor third-order many-body or Møller–Plesset perturbation theory,
- MBPT4=MP4: iterative tensor fourth-order many-body or Møller–Plesset perturbation theory,

All of these models are based on spin-orbital expressions of the amplitude and energy equations, and designed primarily for spin-unrestricted reference wave functions. However, for a restricted reference wave function of a closed-shell system, some further reduction of operation and storage cost will be made. Within the unrestricted framework, all these methods take full advantage of spin, spatial, and index permutation symmetries to save operation and storage costs at every stage of the calculation. Consequently, these computer-generated programs will perform significantly faster than, for instance, a hand-written spin-adapted CCSD program in NWChem, although the nominal operation cost for a spin-adapted CCSD is just one half of that for spin-unrestricted CCSD (in spin-unrestricted CCSD there are three independent sets of excitation amplitudes, whereas in spin-adapted CCSD there is only one set, so the nominal operation cost for the latter is one third of that of the former. For a restricted reference wave function of a closed-shell system, all β spin block of the excitation amplitudes and integrals can be trivially mapped to the all α spin block, reducing the ratio to one half).

While the MBPT (MP) models implemented in the TCE module give identical correlation energies as conventional implementation for a canonical HF reference of a closed-shell system, the former are intrinsically more general and theoretically robust for other less standard reference wave functions and open-shell systems. This is because the zeroth

order of Hamiltonian is chosen to be the full Fock operator (not just the diagonal part), and no further approximation was invoked. So unlike the conventional implementation where the Fock matrix is assumed to be diagonal and a correlation energy is evaluated in a single analytical formula that involves orbital energies (or diagonal Fock matrix elements), the present tensor MBPT requires the iterative solution of amplitude equations and subsequent energy evaluation and is generally more expensive than the former. For example, the operation cost of many conventional implementation of MBPT(2) scales as the fourth power of the system size, but the cost of the present tensor MBPT(2) scales as the fifth power of the system size, as the latter permits non-canonical HF reference and the former does not (to reinstate the non-canonical HF reference in the former makes it also scale as the fifth power of the system size).

15.5.3 THRESH — the convergence threshold of iterative solutions of amplitude equations

This keyword specifies the convergence threshold of iterative solutions of amplitude equations, and applies to all of the CI, CC, and MBPT models. The threshold refers to the norm of residual, namely, the deviation from the amplitude equations. The default value is $1e-6$.

15.5.4 MAXITER — the maximum number of iterations

It sets the maximum allowed number iterations for the iterative solutions of amplitude equations. The default value is 100.

15.5.5 IO — parallel I/O scheme

There are five parallel I/O schemes implemented for all the models, which need to be wisely chosen for a particular problem and computer architecture.

- `fortran` : Fortran77 direct access,
- `eaf` : Exclusive Access File library,
- `ga` : Fully incore, Global Array virtual file,
- `sf` : Shared File library,
- `replicated` : Semi-replicated file on distributed file system with EAF library.
- `dra` : Distributed file on distributed file system with DRA library.
- `ga_eaf` : Semi-replicated file on distributed file system with EAF library. GA is used to speedup the file reconciliation.

The GA algorithm, which is default, stores all input (integrals and excitation amplitudes), output (residuals), and intermediate tensors in the shared memory area across all nodes by virtue of GA library. This fully incore algorithm replaces disk I/O by inter-process communications. This is a recommended algorithm whenever feasible. Note that the memory management through runtime orbital range tiling described above applies to local (unshared) memory of each node, which may be separately allocated from the shared memory space for GA. So when there is not enough shared memory space (either physically or due to software limitations, in particular, `shmmax` setting), the GA algorithm can crash due to an out-of-memory error. The replicated scheme is the currently the only disk-based algorithm for a genuinely distributed file system. This means that each node keeps an identical copy of input tensors and it holds non-identical overlapping segments of intermediate and output tensors in its local disk. Whenever data coherency is required, a file reconciliation process will take place to make the intermediate and output data identical throughout the

nodes. This algorithm, while requiring redundant data space on local disk, performs reasonably efficiently in parallel. For sequential execution, this reduces to the EAF scheme. For a global file system, the SF scheme is recommended. This together with the Fortran77 direct access scheme does not usually exhibit scalability unless shared files on the global file system also share the same I/O buffer. For sequential executions, the SF, EAF, and replicated schemes are interchangeable, while the Fortran77 scheme is appreciably slower.

Two new I/O algorithms `dra` and `ga_eaf` combines GA and DRA or EAF based replicated algorithm. In the former, arrays that are not active (e.g., prior T amplitudes used in DIIS or EOM-CC trial vectors) in GA algorithm will be moved to DRA. In the latter, the intermediates that are formed by tensor contractions are initially stored in GA, thereby avoiding the need to accumulate the fragments of the intermediate scattered in EAFs in the original EAF algorithm. Once the intermediate is formed completely, then it will be replicated as EAFs.

15.5.6 DIIS — the convergence acceleration

It sets the number iterations in which a DIIS extrapolation is performed to accelerate the convergence of excitation amplitudes. The default value is 5, which means in every five iteration, one DIIS extrapolation is performed (and in the rest of the iterations, Jacobi rotation is used). When zero or negative value is specified, the DIIS is turned off. It is not recommended to perform DIIS every iteration, whereas setting a large value for this parameter necessitates a large memory (disk) space to keep the excitation amplitudes of previous iterations.

15.5.7 FREEZE — the frozen core/virtual approximation

Some of the lowest-lying core orbitals and/or some of the highest-lying virtual orbitals may be excluded in the calculations by this keyword (this does not affect the ground state HF or DFT calculation). No orbitals are frozen by default. To exclude the atom-like core regions altogether, one may request

```
FREEZE atomic
```

To specify the number of lowest-lying occupied orbitals be excluded, one may use

```
FREEZE 10
```

which causes 10 lowest-lying occupied orbitals excluded. This is equivalent to writing

```
FREEZE core 10
```

To freeze the highest virtual orbitals, use the `virtual` keyword. For instance, to freeze the top 5 virtuals

```
FREEZE virtual 5
```

15.5.8 NROOTS — the number of excited states

One can specify the number of excited state roots to be determined. The default value is 1. It is advised that the users request several more roots than actually needed, since owing to the nature of the trial vector algorithm, some low-lying roots can be missed when they do not have sufficient overlap with the initial guess vectors.

15.5.9 TARGET and TARGETSYM — the target root and its symmetry

At the moment, the first and second geometrical derivatives of excitation energies that are needed in force, geometry, and frequency calculations are obtained by numerical differentiation. These keywords may be used to specify which excited state root is being used for the geometrical derivative calculation. For instance, when TARGET 3 and TARGETSYM a1g are included in the input block, the total energy (ground state energy plus excitation energy) of the third lowest excited state root (excluding the ground state) transforming as the irreducible representation a1g will be passed to the module which performs the derivative calculations. The default values of these keywords are 1 and none, respectively.

The keyword TARGETSYM is essential in excited state geometry optimization, since it is very common that the order of excited states changes due to the geometry changes in the course of optimization. Without specifying the TARGETSYM, the optimizer could (and would likely) be optimizing the geometry of an excited state that is different from the one the user had intended to optimize at the starting geometry. On the other hand, in the frequency calculations, TARGETSYM must be none, since the finite displacements given in the course of frequency calculations will lift the spatial symmetry of the equilibrium geometry. When these finite displacements can alter the order of excited states including the target state, the frequency calculation is not be feasible.

15.5.10 SYMMETRY — restricting the excited state symmetry

By adding this keyword to the input block, the user can request the module to seek just the roots of the specified irreducible representation as TARGETSYM. By default, this option is not set. TARGETSYM must be specified when SYMMETRY is invoked.

15.5.11 DIPOLE — the ground- and excited-state dipole moments

When this is set, the ground-state CC calculation will enter another round of iterative step for the so-called Λ equation to obtain the one-particle density matrix and dipole moments. Likewise, for excited-states (EOM-CC), the transition moments and dipole moments will be computed when (and only when) this option is set. In the latter case, EOM-CC left hand side solutions will be sought incurring approximately three times the computational cost of excitation energies alone (note that the EOM-CC effective Hamiltonian is not Hermitian and has distinct left and right eigenvectors).

15.5.12 (NO) FOCK — (not) recompute Fock matrix

The default is FOCK meaning that the Fock matrix will be reconstructed (as opposed to using the orbital energies as the diagonal part of Fock). This is essential in getting correct correlation energies with ROHF or DFT reference wave functions. However, currently, this module cannot reconstruct the Fock matrix when one-component relativistic effects are operative. So when a user wishes to run TCE's correlation methods with DK or other relativistic reference, NOFOCK must be set and orbital energies must be used for the Fock matrix.

15.5.13 PRINT — the verbosity

This keyword changes the level of output verbosity. One may also request some particular items in Table 15.1 printed.

Table 15.1: Printable items in the TCE modules and their default print levels.

Item	Print Level	Description
"time"	vary	CPU and wall times
"tile"	vary	Orbital range tiling information
"t1"	debug	T_1 excitation amplitude dumping
"t2"	debug	T_2 excitation amplitude dumping
"t3"	debug	T_3 excitation amplitude dumping
"t4"	debug	T_4 excitation amplitude dumping
"general information"	default	General information
"correlation information"	default	TCE information
"mbpt2"	debug	Caonical HF MBPT2 test
"get_block"	debug	I/O information
"put_block"	debug	I/O information
"add_block"	debug	I/O information
"files"	debug	File information
"offset"	debug	File offset information
"ao1e"	debug	AO one-electron integral evaluation
"ao2e"	debug	AO two-electron integral evaluation
"mo1e"	debug	One-electron integral transformation
"mo2e"	debug	Two-electron integral transformation

15.6 Sample input

The following is a sample input for a ROHF-UCCSD energy calculation of a water radical cation.

```

START h2o

TITLE "ROHF-UCCSD/cc-pVTZ H2O"

CHARGE 1

GEOMETRY
O      0.00000000    0.00000000    0.12982363
H      0.75933475    0.00000000   -0.46621158
H     -0.75933475    0.00000000   -0.46621158
END

BASIS
* library cc-pVTZ
END

SCF
ROHF
DOUBLET
THRESH 1.0e-10
TOL2E  1.0e-10
END

```

```
TCE
CCSD
END
```

```
TASK TCE ENERGY
```

The same result can be obtained by the following input:

```
START h2o
```

```
TITLE "ROHF-UCCSD/cc-pVTZ H2O"
```

```
CHARGE 1
```

```
GEOMETRY
```

```
O      0.00000000      0.00000000      0.12982363
H      0.75933475      0.00000000     -0.46621158
H     -0.75933475      0.00000000     -0.46621158
END
```

```
BASIS
```

```
* library cc-pVTZ
```

```
END
```

```
SCF
```

```
ROHF
```

```
DOUBLET
```

```
THRESH 1.0e-10
```

```
TOL2E 1.0e-10
```

```
END
```

```
TASK UCCSD ENERGY
```

EOM-CCSDT calculation for excitation energies, excited-state dipole, and transition moments.

```
START tce_h2o_eomcc
```

```
GEOMETRY UNITS BOHR
```

```
H      1.474611052297904      0.0000000000000000      0.863401706825835
O      0.0000000000000000      0.0000000000000000     -0.215850436155089
H     -1.474611052297904      0.0000000000000000      0.863401706825835
END
```

```
BASIS
```

```
* library sto-3g
```

```
END
```

```
SCF
```

```
SINGLET
```

```
RHF
```

```
END
```

```
TCE  
CCSDT  
DIPOLE  
FREEZE CORE ATOMIC  
NROOTS 1  
END
```

```
TASK TCE ENERGY
```

Chapter 16

MP2

There are (at least) three algorithms within NWChem that compute the Møller-Plesset (or many-body) perturbation theory second-order correction to the Hartree-Fock energy (MP2). They vary in capability, the size of system that can be treated and use of other approximations

- Semi-direct — this is recommended for most large applications (up to about 2800 basis functions), especially on the IBM SP and other machines with significant disk I/O capability. Partially transformed integrals are stored on disk, multi-passing as necessary. RHF and UHF references may be treated including computation of analytic derivatives. This is selected by specifying `mp2` on the task directive, e.g.

```
TASK MP2
```

- Fully-direct — this is of utility if only limited I/O resources are available (up to about 2800 functions). Only RHF references and energies are available. This is selected by specifying `direct_mp2` on the task directive, e.g.

```
TASK DIRECT_MP2
```

- Resolution of the identity (RI) approximation MP2 (RI-MP2) — this uses the RI approximation and is therefore only exact in the limit of a complete fitting basis. However, with some care, high accuracy may be obtained with relatively modest fitting basis sets. An RI-MP2 calculation can cost over 40 times less than the corresponding exact MP2 calculation. RHF and UHF references with only energies are available. This is selected by specifying `rmp2` on the task directive, e.g.,

```
TASK RIMP2
```

All three MP2 tasks share the same input block.

```
MP2
[FREEZE [[core] (atomic || <integer nfzc default 0>)] \
        [virtual <integer nfzv default 0>]]
[TIGHT]
[PRINT]
[NOPRINT]
[VECTORS <string filename default scf-output-vectors> \
        [swap [(alpha||beta)] <integer pair-list>] ]
```

```
[RIAPPROX <string riapprox default V>]
[FILE3C <string filename default $file_prefix$.mo3cint">]
[SCRATCHDISK <integer>]
END
```

16.1 FREEZE — Freezing orbitals

All MP2 modules support frozen core orbitals, however, only the direct MP2 and RI-MP2 modules support frozen virtual orbitals.

By default, no orbitals are frozen. The `atomic` keyword causes orbitals to be frozen according to the rules in Table 16.1. Note that *no* orbitals are frozen on atoms on which the nuclear charge has been modified either by the user or due to the presence of an ECP. The actual input would be

```
freeze atomic
```

For example, in a calculation on $\text{Si}(\text{OH})_2$, by default the lowest seven orbitals would be frozen (the oxygen 1s, and the silicon 1s, 2s and 2p).

Period	Elements	Core Orbitals	Number of Core
0	H – He	—	0
1	Li – Ne	1s	1
2	Na – Ar	1s2s2p	5
3	K – Kr	1s2s2p3s3p	9
4	Rb – Xe	1s2s2p3s3p4s3d4p	18
5	Cs – Rn	1s2s2p3s3p4s3d4p5s4d5p	27
6	Fr – Lr	1s2s2p3s3p4s3d4p5s4d5p6s4f5d6p	43

Table 16.1: Number of orbitals considered “core” in the “freeze by atoms” algorithm.

Caution: The rule for freezing orbitals “by atoms” are rather unsophisticated: the number of orbitals to be frozen is computed from the Table 16.1 by summing the number of core orbitals in each atom present. The corresponding number of lowest-energy orbitals are frozen — if for some reason the actual core orbitals are not the lowest lying, then correct results will not be obtained. From limited experience, it seems that special attention should be paid to systems including third- and higher- period atoms.

The user may also specify the number of orbitals to be frozen by atom. Following the $\text{Si}(\text{OH})_2$ example, the user could specify

```
freeze atomic 0 1 Si 3
```

In this case only the lowest four orbitals would be frozen. If the user does not specify the orbitals by atom, the rules default to Table 16.1.

Caution: The system does not check for a valid number of orbitals per atom. If the user specifies to freeze more orbitals than are available for the atom, the system will not catch the error. The user must specify a logical number of orbitals to be frozen for the atom.

The `FREEZE` directive may also be used to specify the number of core orbitals to freeze. For instance, to freeze the first 10 orbitals

```
freeze 10
```

or equivalently, using the optional keyword `core`

```
freeze core 10
```

Again, note that if the 10 orbitals to be frozen do not correspond to the first 10 orbitals, then the `swap` keyword of the `VECTORS` directive must be used to order the input orbitals correctly (Section 16.5).

To freeze the highest virtual orbitals, use the `virtual` keyword. For instance, to freeze the top 5 virtuals

```
freeze virtual 5
```

Again, note that this only works for the direct-MP2 and RI-MP2 energy codes.

16.2 TIGHT — Increased precision

The `TIGHT` directive can be used to increase the precision in the MP2 energy and gradients.

By default the MP2 gradient package should compute energies accurate to better than a micro-Hartree, and gradients accurate to about five decimal places (atomic units). However, if there is significant linear dependence in the basis set the precision might not be this good. Also, for computing very accurate geometries or numerical frequencies, greater precision may be desirable.

This option increases the precision to which both the SCF (from 10^{-6} to 10^{-8}) and CPHF (from 10^{-4} to 10^{-6}) are solved, and also tightens thresholds for computation of the AO and MO integrals (from 10^{-9} to 10^{-11}) within the MP2 code.

16.3 SCRATCHDISK — Limiting I/O usage

This directive – used only in the semi-direct algorithm – allows to limit the per process disk usage. Mandatory argument for this keyword is the maximum number of MBytes. For example, the following input line

```
scratchdisk 512
```

puts an upper limit of 512 MBytes to the semi-direct MP2 usage of disk (again, on a per process base).

16.4 PRINT and NOPRINT

The standard print control options are recognized. The list of recognized names are given in Table 16.2.

16.5 VECTORS — MO vectors

All of the (supported) MP2 modules require use of converged canonical SCF (RHF or UHF) orbitals for correct results. The vectors are by default obtained from the preceding SCF calculation, but it is possible to specify a different source using the `VECTORS` directive. For instance, to obtain vectors from the file `/tmp/h2o.movecs`, use the directive

```
vectors /tmp/h2o.movecs
```

Table 16.2: Printable items in the MP2 modules and their default print levels.

Item	Print Level	Description
RI-MP2		
"2/3 ints"	debug	Partial 3-center integrals
"3c ints"	debug	MO 3-center integrals
"4c ints b"	debug	"B" matrix with approx. 4c integrals
"4c ints"	debug	Approximate 4-center integrals
"amplitudes"	debug	"B" matrix with denominators
"basis"	high	
"fit xf"	debug	Transformation for fitting basis
"geombas"	debug	Detailed basis map info
"geometry"	high	
"information"	low	General information about calc.
"integral i/o"	high	File size information
"mo ints"	debug	
"pair energies"	debug	(working only in direct_mp2)
"partial pair energies"	debug	Pair energy matrix each time it is updated
"progress reports"	default	Report completion of time-consuming steps
"reference"	high	Details about reference wavefunction
"warnings"	low	Non-fatal warnings

As noted above (Section 16.1) if the SCF orbitals are not in the correct order, it is necessary to permute the input orbitals using the `swap` keyword of the `VECTORS` directive. For instance, if it is desired to freeze a total six orbitals corresponding to the SCF orbitals 1–5, and 7, it is necessary to swap orbital 7 into the 6th position. This is accomplished by

```
vectors swap 6 7
```

The swap capability is examined in more detail in Section 10.5.

16.6 RI-MP2 fitting basis

The RI-MP2 method requires a fitting basis, which must be specified with the name `"ri-mp2 basis"` (see Section 7). For instance,

```
basis "ri-mp2 basis"
  O s; 10000.0 1
  O s;  1000.0 1
  O s;   100.0 1
  ...
end
```

Alternatively, using a standard capability of basis sets (Section 7) another named basis may be associated with the fitting basis. For instance, the following input specifies a basis with the name `"small fitting basis"` and then defines this to be the `"ri-mp2 basis"`.

```

basis "small fitting basis"
  H s; 10    1
  H s;  3    1
  H s;  1    1
  H s;  0.1  1
  H s;  0.01 1
end

set "ri-mp2 basis" "small fitting basis"

```

16.7 FILE3C — RI-MP2 3-center integral filename

The default name for the file used to store the transformed 3-center integrals is "\$file_prefix\$.mo3cint" in the scratch directory. This may be overridden using the FILE3C directive. For instance, to specify the file /scratch/h2o.3c, use this directive

```
file3c /scratch/h2o.3c
```

16.8 RIAPPROX — RI-MP2 Approximation

The type of RI approximation used in the RI-MP2 calculation is controlled by means of the RIAPPROX directive. The two possible values are V and SVS (case sensitive), which correspond to the approximations with the same names described in O. Vahtras, J Almlöf, and M. W. Feyereisen, *Chem. Phys. Lett.* **213**, 514–518 (1993). The default is V.

16.9 Advanced options for RI-MP2

These options, which functioned at the time of writing, are not currently supported.

16.9.1 Control of linear dependence

Construction of the RI fit requires the inversion of a matrix of fitting basis integrals which is carried out via diagonalization. If the fitting basis includes near linear dependencies, there will be small eigenvalues which can ultimately lead to non-physical RI-MP2 correlation energies. Eigenvectors of the fitting matrix are discarded if the corresponding eigenvalue is less than \$mineval\$ which defaults to 10^{-8} . This parameter may be changed by setting the parameter in the database. For instance, to set it to 10^{-10}

```
set "mp2:fit min eval" 1e-10
```

16.9.2 Reference Spin Mapping for RI-MP2 Calculations

The user has the option of specifying that the RI-MP2 calculations are to be done with variations of the SCF reference wavefunction. This is accomplished with a SET directive of the form,

```
set "mp2:reference spin mapping" <integer array default 0>
```

Each element specified for `array` is the SCF spin case to be used for the corresponding spin case of the correlated calculation. The number of elements set determines the overall type of correlated calculation to be performed. The default is to use the unadulterated SCF reference wavefunction.

For example, to perform a spin-unrestricted calculation (two elements) using the alpha spin orbitals (spin case 1) from the reference for both of the correlated reference spin cases, the `SET` directive would be as follows,

```
set "mp2:reference spin mapping" 1 1
```

The SCF calculation to produce the reference wavefunction could be either RHF or UHF in this case.

The `SET` directive for a similar case, but this time using the beta-spin SCF orbitals for both correlated spin cases, is as follows,

```
set "mp2:reference spin mapping" 2 2
```

The SCF reference calculation must be UHF in this case.

The `SET` directive for a spin-restricted calculation (one element) from the beta-spin SCF orbitals using this option is as follows,

```
set "mp2:reference spin mapping" 2
```

The `SET` directive for a spin-unrestricted calculation with the spins flipped from the original SCF reference wavefunction is as follows,

```
set "mp2:reference spin mapping" 2 1
```

16.9.3 Batch Sizes for the RI-MP2 Calculation

The user can control the size of each batch in the transformation and energy evaluation in the MP2 calculation, and consequently the memory requirements and number of passes required. This is done using two `SET` directives of the following form,

```
set "mp2:transformation batch size" <integer size default -1>
set "mp2:energy batch size" <integer isize jsize default -1 -1>
```

The default is for the code to determine the batch size based on the available memory. Should there be problems with the program-determined batch sizes, these variables allow the user to override them. The program will always use the smaller of the user's value of these entries and the internally computed batch size.

The transformation batch size computed in the code is the number of occupied orbitals in the (*occ vir|fit*) three-center integrals to be produced at a time. If this entry is less than the number of occupied orbitals in the system, the transformation will require multiple passes through the two-electron integrals. The memory requirements of this stage are *two* global arrays of dimension $\langle batchsize \rangle \times vir \times fit$ with the "fit" dimension distributed across all processors (on shell-block boundaries). The compromise here is memory space versus multiple integral evaluations.

The energy evaluation batch sizes are computed in the code from the number of occupied orbitals in the two sets of three-center integrals to be multiplied together to produce a matrix of approximate four-center integrals. Two blocks of integrals of dimension $(\langle batchsize \rangle \times vir)$ and $(\langle batchjsize \rangle \times vir)$ by *fit* are read in from disk and multiplied together to produce $\langle batchsize \rangle \langle batchjsize \rangle vir^2$ approximate integrals. The compromise here is performance of the distributed matrix multiplication (which requires large matrices) versus memory space.

16.9.4 Energy Memory Allocation Mode: RI-MP2 Calculation

The user must choose a strategy for the memory allocation in the energy evaluation phase of the RI-MP2 calculation, either by minimizing the amount of I/O, or minimizing the amount of computation. This can be accomplished using a SET directive of the form,

```
set "mp2:energy mem minimize" <string mem_opt default I>
```

A value of I entered for the string mem_opt means that a strategy to minimize I/O will be employed. A value of C tells the code to use a strategy that minimizes computation.

When the option to minimize I/O is selected, the block sizes are made as large as possible so that the total number of passes through the integral files is as small as possible. When the option to minimize computation is selected, the blocks are chosen as close to square as possible so that permutational symmetry in the energy evaluation can be used most effectively.

16.9.5 Local Memory Usage in Three-Center Transformation

For most applications, the code will be able to size the blocks without help from the user. Therefore, it is unlikely that users will have any reason to specify values for these entries except when doing very particular performance measurements.

The size of `xf3ci:AO 1 batch size` is the most important of the three, in terms of the effect on performance.

Local memory usage in the first two steps of the transformation is controlled in the RI-MP2 calculation using the following SET directives,

```
set "xf3ci:AO 1 batch size" <integer max>
set "xf3ci:AO 2 batch size" <integer max>
set "xf3ci:fit batch size" <integer max>
```

The size of the local arrays determines the sizes of the two matrix multiplications. These entries set limits on the size of blocks to be used in each index. The listing above is in order of importance of the parameters to performance, with `xf3ci:AO 1 batch size` being most important.

Note that these entries are only upper bounds and that the program will size the blocks according to what it determines as the best usage of the available local memory. The absolute maximum for a block size is the number of functions in the AO basis, or the number of fitting basis functions on a node. The absolute minimum value for block size is the size of the largest shell in the appropriate basis. Batch size entries specified for max that are larger than these limits are automatically reset to an appropriate value.

16.10 One-electron properties and natural orbitals

If an MP2 energy gradient is computed, all contributions are available to form the MP2 linear-response density. This is the density that when contracted with any spin-free, one-electron operator yields the associated property defined as the derivative of the energy. Thus, the reported MP2 dipole moment is the derivative of the energy w.r.t. an external magnetic field and is *not* the expectation value of the operator over the wavefunction. Only dipole moments are printed by the MP2 gradient code, but natural orbitals are produced and stored in the permanent directory with a file extension of ".mp2nos". These may be fed into the property package (see Section 28) to compute more general properties.

Note that the MP2 linear response density matrix is not necessarily positive definite so it is not unusual to see a few small negative natural orbital occupation numbers.

Chapter 17

Multiconfiguration SCF

The NWChem multiconfiguration SCF (MCSCF) module can currently perform complete active space SCF (CASSCF) calculations with at most 20 active orbitals and about 500 basis functions. It is planned to extend it to handle 1000+ basis functions.

```
MCSCF
  STATE <string state>
  ACTIVE <integer nactive>
  ACTELEC <integer nactelec>
  MULTIPLICITY <integer multiplicity>
  [SYMMETRY <integer symmetry default 1>]
  [VECTORS [[input] <string input_file default $file_prefix$.movecs>]
           [swap <integer vec1 vec2> ...] \
           [output <string output_file default input_file>] \
           [lock]
  [HESSIAN (exact||onel)]
  [MAXITER <integer maxiter default 20>]
  [THRESH <real thresh default 1.0e-4>]
  [TOL2E <real tol2e default 1.0e-9>]
  [LEVEL <real shift default 0.1d0>]
END
```

Note that the ACTIVE, ACTELEC, and MULTIPLICITY directives are *required*. The symmetry and multiplicity may alternatively be entered using the STATE directive.

17.1 ACTIVE — Number of active orbitals

The number of orbitals in the CASSCF active space must be specified using the ACTIVE directive.

E.g.,

```
active 10
```

The input molecular orbitals (see the vectors directive, Sections 17.6 and 10.5) must be arranged in order

1. doubly occupied orbitals,
2. active orbitals, and
3. unoccupied orbitals.

17.2 ACTELEC — Number of active electrons

The number of electrons in the CASSCF active space must be specified using the the ACTELEC directive. An error is reported if the number of active electrons and the multiplicity are inconsistent.

The number of closed shells is determined by subtracting the number of active electrons from the total number of electrons (which in turn is derived from the sum of the nuclear charges minus the total system charge).

17.3 MULTIPLICITY

The spin multiplicity must be specified and is enforced by projection of the determinant wavefunction.

E.g., to obtain a triplet state

```
multiplicity 3
```

17.4 SYMMETRY — Spatial symmetry of the wavefunction

This species the irreducible representation of the wavefunction as an integer in the range 1—8 using the same numbering of representations as output by the SCF program. Note that only Abelian point groups are supported.

E.g., to specify a B_1 state when using the C_{2v} group

```
symmetry 3
```

17.5 STATE — Symmetry and multiplicity

The electronic state (spatial symmetry and multiplicity) may alternatively be specified using the conventional notation for an electronic state, such as 3B_2 for a triplet state of B_2 symmetry. This would be accomplished with the input

```
state 3b2
```

which is equivalent to

```
symmetry 4
multiplicity 3
```

17.6 VECTORS — Input/output of MO vectors

Calculations are best started from RHF/ROHF molecular orbitals (see Section 10), and by default vectors are taken from the previous MCSCF or SCF calculation. To specify another input file use the `VECTORS` directive. Vectors are by default output to the input file, and may be redirected using the `output` keyword. The `swap` keyword of the `VECTORS` directive may be used to reorder orbitals to obtain the correct active space. See Section 10.5 for an example.

The `LOCK` keyword allows the user to specify that the ordering of orbitals will be locked to that of the initial vectors, insofar as possible. The default is to order by ascending orbital energies within each orbital space. One application where locking might be desirable is a calculation where it is necessary to preserve the ordering of a previous geometry, despite flipping of the orbital energies. For such a case, the `LOCK` directive can be used to prevent the SCF calculation from changing the ordering, even if the orbital energies change.

Output orbitals of a converged MCSCF calculation are canonicalized as follows:

- Doubly occupied and unoccupied orbitals diagonalize the corresponding blocks of an effective Fock operator. Note that in the case of degenerate orbital energies this does not fully determine the orbitals.
- Active-space orbitals are chosen as natural orbitals by diagonalization of the active space 1-particle density matrix. Note that in the case of degenerate occupations that this does not fully determine the orbitals.

17.7 HESSIAN — Select preconditioner

The MCSCF will use a one-electron approximation to the orbital-orbital Hessian until some degree of convergence is obtained, whereupon it will attempt to use the exact orbital-orbital Hessian which makes the micro iterations more expensive but potentially reduces the total number of macro iterations. Either choice may be forced throughout the calculation by specifying the appropriate keyword on the `HESSIAN` directive.

E.g., to specify the one-electron approximation throughout

```
hessian onel
```

17.8 LEVEL — Level shift for convergence

The Hessian (Section 17.7) used in the MCSCF optimization is by default level shifted by 0.1 until the orbital gradient norm falls below 0.01, at which point the level shift is reduced to zero. The initial value of 0.1 may be changed using the `LEVEL` directive. Increasing the level shift may make convergence more stable in some instances.

E.g., to set the initial level shift to 0.5

```
level 0.5
```

17.9 PRINT and NOPRINT

Specific output items can be selectively enabled or disabled using the `print` control mechanism (5.6) with the available print options listed in table(17.9).

Option	Class	Synopsis
ci energy	default	CI energy eigenvalue
fock energy	default	Energy derived from Fock matrices
gradient norm	default	Gradient norm
movecs	default	Converged occupied MO vectors
trace energy	high	Trace Energy
converge info	high	Convergence data and monitoring
precondition	high	Orbital preconditioner iterations
microci	high	CI iterations in line search
canonical	high	Canonicalization information
new movecs	debug	MO vectors at each macro-iteration
ci guess	debug	Initial guess CI vector
density matrix	debug	One- and Two-particle density matrices

Table 17.1: MCSCF Print Options

Chapter 18

Selected CI

The selected CI module is integrated into NWChem but as yet no input module has been written. The input thus consists of setting the appropriate variables in the database.

It is assumed that an initial SCF/MCSCF calculation has completed, and that MO vectors are available. These will be used to perform a four-index transformation, if this has not already been performed.

18.1 Background

This is a general spin-adapted, configuration-driven CI program which can perform arbitrary CI calculations, the only restriction being that all spin functions are present for each orbital occupation. CI wavefunctions may be specified using a simple configuration generation program, but the prime usage is intended to be in combination with perturbation correction and selection of new configurations. The second-order correction (Epstein-Nesbet) to the CI energy may be computed, and at the same time configurations that interact greater than a certain threshold with the current CI wavefunction may be chosen for inclusion in subsequent calculations. By repeating this process (typically twice is adequate) with the same threshold until no new configurations are added, the CI expansion may be made consistent with the selection threshold, enabling tentative extrapolation to the full-CI limit.

A typical sequence of calculations is as follows:

1. Pick as an initial CI reference the previously executed SCF/MCSCF.
2. Define an initial selection threshold.
3. Determine the roots of interest in the current reference space.
4. Compute the perturbation correction and select additional configurations that interact greater than the current threshold.
5. Repeat steps 3 and 4.
6. Lower the threshold (a factor of 10 is common) and repeat steps 3, 4, and 5. The *first* pass through step 4 will yield the approximately self-consistent CI and CI+PT energies from the *previous* selection threshold.

To illustrate this, below is some abbreviated output from a calculation on water in an augmented cc-PVDZ basis set with one frozen core orbital. The SCF was converged to high precision in C_{2v} symmetry with the following input

```

start h2o
geometry; symmetry c2v
  O 0 0 0; H 0 1.43042809 -1.10715266
end
basis
  H library aug-cc-pvdz; O library aug-cc-pvdz
end
task scf
scf; thresh 1d-8; end

```

The following input restarts from the SCF to perform a sequence of selected CI calculations with the specified tolerances, starting with the SCF reference.

```

restart h2o
set fourindex:occ_frozen 1
set selci:mode select
set "selci:selection thresholds" \
  0.001 0.001 0.0001 0.0001 0.00001 0.00001 0.000001
task selci

```

Table 18.1 summarizes the output from each of the major computational steps that were performed.

Step	Description	CI dimension	Energy
1	Four-index, one frozen-core		
2	Config. generator, SCF default	1	
3+4	CI diagonalization	1	$E_{CI} = -76.041983$
5	PT selection $T=0.001$	1	$E_{CI+PT} = -76.304797$
6+7	CI diagonalization	75	$E_{CI} = -76.110894$
8	PT selection $T=0.001$	75	$E_{CI+PT} = -76.277912$
9+10	CI diagonalization	75	$E_{CI}(T = 0.001) = -76.110894$
11	PT selection $T=0.0001$	75	$E_{CI+PT}(T = 0.001) = -76.277912$
12+13	CI diagonalization	823	$E_{CI} = -76.228419$
14	PT selection $T=0.0001$	823	$E_{CI+PT} = -76.273751$
15+16	CI diagonalization	841	$E_{CI}(T = 0.0001) = -76.2300544$
17	PT selection $T=0.00001$	841	$E_{CI+PT}(T = 0.0001) = -76.274073$
18+19	CI diagonalization	2180	$E_{CI} = -76.259285$
20	PT selection $T=0.00001$	2180	$E_{CI+PT} = -76.276418$
21+22	CI diagonalization	2235	$E_{CI}(T = 0.00001) = -76.259818$
23	PT selection $T=0.000001$	2235	$E_{CI+PT}(T = 0.00001) = -76.276478$
24	CI diagonalization	11489	

Table 18.1: Summary of steps performed in a selected CI calculation on water.

18.2 Files

Currently, no direct control is provided over filenames. All files are prefixed with the standard file-prefix, and any files generated by all nodes are also postfixed with the processor number. Thus, for example the molecular integrals file,

used only by process zero, might be called `h2o.moints` whereas the off-diagonal Hamiltonian matrix element file used by process number eight would be called `h2o.hamil.8`.

`ciconf` — the CI configuration file, which holds information about the current CI expansion, indexing vectors, etc.

This is the most important file and is required for all restarts. Note that the CI configuration generator is only run if this file does not exist. Referenced only by process zero.

`moints` — the molecular integrals, generated by the four-index transformation. As noted above these must currently be manually deleted, or the database entry `selci:moints:force` set, to force regeneration. Referenced only by process zero.

`civecs` — the CI vectors. Referenced only by process zero.

`wmatrx` — temporary file used to hold coupling coefficients. Deleted at calculation end. Referenced only by process zero.

`rtname`, `roname` — restart information for the PT selection. Should be automatically deleted if no restart is necessary. Referenced only by process zero.

`hamdg` — diagonal elements of the Hamiltonian. Deleted at calculation end. Referenced only by process zero.

`hamil` — off-diagonal Hamiltonian matrix elements. All processes generate a file containing a subset of these elements. These files can become very large. Deleted at calculation end.

18.3 Configuration Generation

If no configuration is explicitly specified then the previous SCF/MCSCF wavefunction is used, adjusting for any orbitals frozen in the four-index transformation. The four-index transformation must have completed successfully before this can execute. Orbital configurations for use as reference functions may also be explicitly specified.

Once the default/user-input reference configurations have been determined additional reference functions may be generated by applying multiple sets of creation-annihilation operators, permitting for instance, the ready specification of complete or restricted active spaces.

Finally, a uniform level of excitation from the current set of configurations into all orbitals may be applied, enabling, for instance, the simple creation of single or single+double excitation spaces from an MCSCF reference.

18.3.1 Specifying the reference occupation

A single orbital configuration or occupation is specified by

```
ns (socc(i),i=1,ns) (docc(i),i=1,nd)
```

where `ns` specifies the number of singly occupied orbitals, `socc()` is the list of singly occupied orbitals, and `docc()` is the list of doubly occupied orbitals (the number of doubly occupied orbitals, `nd`, is inferred from `ns` and the total number of electrons). All occupations may be strung together and inserted into the database as a single integer array with name `"selci:conf"`. For example, the input

```
set "selci:conf" \
  0          1  2  3  4 \
  0          1  2  3 27 \
```

```

0          1 3 4 19 \
2  11 19   1 3 4 \
2    8 27   1 2 3 \
0          1 2 4 25 \
4    3 4 25 27 1 2 \
4    2 3 19 20 1 4 \
4    2 4 20 23 1 3

```

specifies the following nine orbital configurations

```

1(2) 2(2) 3(2) 4(2)
1(2) 2(2) 3(2) 27(2)
1(2) 3(2) 4(2) 19(2)
1(2) 3(2) 4(2) 11(1) 19(1)
1(2) 2(2) 3(2) 8(1) 27(1)
1(2) 2(2) 4(2) 25(2)
1(2) 2(2) 3(1) 4(1) 25(1) 27(1)
1(2) 2(1) 3(1) 4(2) 19(1) 20(1)
1(2) 2(1) 3(2) 4(1) 20(1) 23(1)

```

The optional formatting of the input is just to make this arcane notation easier to read. Relatively few configurations can be currently specified in this fashion because of the input line limit of 1024 characters.

18.3.2 Applying creation-annihilation operators

Up to 10 sets of creation-annihilation operator pairs may be specified, each set containing up to 255 pairs. This suffices to specify complete active spaces with up to ten electrons.

The number of sets is specified as follows,

```
set selci:ngen 4
```

which indicates that there will be four sets. Each set is then specified as a separate integer array

```

set "selci:refgen 1" 5 4    6 4    5 3    6 3
set "selci:refgen 2" 5 4    6 4    5 3    6 3
set "selci:refgen 3" 5 4    6 4    5 3    6 3
set "selci:refgen 4" 5 4    6 4    5 3    6 3

```

In the absence of friendly, input note that the names "selci:refgen n" must be formatted with n in I2 format. Each set specifies a list of creation-annihilation operator pairs (in that order). So for instance, in the above example each set is the same and causes the excitations

```
4->5    4->6    3->5    3->6
```

If orbitals 3 and 4 were initially doubly occupied, and orbitals 5 and 6 initially unoccupied, then the application of this set of operators four times in succession is sufficient to generate the four electron in four orbital complete active space.

The precise sequence in which operators are applied is

1. loop through sets of operators

2. loop through reference configurations
3. loop through operators in the set
4. apply the operator to the configuration, if the result is new add it to the new list
5. end the loop over operators
6. end the loop over reference configurations
7. augment the configuration list with the new list
8. end the loop over sets of operators

18.3.3 Uniform excitation level

By default no excitation is applied to the reference configurations. If, for instance, you wanted to generate a single excitation CI space from the current configuration list, specify

```
set selci:exci 1
```

Any excitation level may be applied, but since the list of configurations is explicitly generated, as is the CI Hamiltonian matrix, you will run out of disk space if you attempt to use more than a few tens of thousands of configurations.

18.4 Number of roots

By default, only one root is generated in the CI diagonalization or perturbation selection. The following requests that 2 roots be generated

```
set selci:nroot 2
```

There is no imposed upper limit. If many roots are required, then, to minimize root skipping problems, it helps to perform an initial approximate diagonalization with several more roots than required, and then resetting this parameter once satisfied that the desired states are obtained.

18.5 Accuracy of diagonalization

By default, the CI wavefunctions are converged to a residual norm of 10^{-6} which provides similar accuracy in the perturbation corrections to the energy, and much higher accuracy in the CI eigenvalues. This may be adjusted with

```
set "selci:diag tol" 1d-3
```

the example setting much lower precision, appropriate for the approximate diagonalization discussed in the preceding section.

18.6 Selection thresholds

When running in the selected-CI mode the program will loop through a list of selection thresholds (T), performing the CI diagonalization, computing the perturbation correction, and augmenting the CI expansion with configurations that make an energy lowering to any root greater than T . The list of selection thresholds is specified as follows

```
set "selci:selection thresholds" \  
    0.001 0.001 0.0001 0.0001 0.00001 0.00001 0.000001
```

There is no default for this parameter.

18.7 Mode

By default the program runs in "ci+davids" mode and just determines the CI eigenvectors/values in the current configuration space. To perform a selected-CI with perturbation correction use the following

```
set selci:mode select
```

and remember to define the selection thresholds.

18.8 Memory requirements

No global arrays are used inside the selected-CI, though the four-index transformation can be automatically invoked and it does use GAs. The selected CI replicates inside each process

- all unique two-electron integrals in the MO basis that are non-zero by symmetry, and
- all CI information, including the CI vectors.

These large data structures are allocated on the local stack. A fatal error will result if insufficient memory is available.

18.9 Forcing regeneration of the MO integrals

When scanning a potential energy surface or optimizing a geometry the MO integrals need to be regenerated each time. Specify

```
set selci:moints:force logical .true.
```

to accomplish this.

18.10 Disabling update of the configuration list

When computing CI+PT energy the reference configuration list is normally updated to reflect all configurations that interact more than the specified threshold. This is usually desirable. But when scanning a potential energy surface or

optimizing a geometry the reference list must be kept fixed to keep the potential energy surface continuous and well defined. To do this specify

```
set selci:update logical .false.
```

18.11 Orbital locking in CI geometry optimization

The selected CI wavefunction is not invariant to orbital rotations or to swapping two or more orbitals. Orbitals could be swapped or rotated when the geometry is changed in a geometry optimization step. The keyword `lock` has to be set in the SCF/MCSCF (vectors) input block to keep the orbitals in the same order throughout the geometry optimization.

Chapter 19

Coupled Cluster Calculations

The NWChem coupled cluster energy module is primarily the work of Alistair Rendell and Rika Kobayashi, with contributions from David Bernholdt.

The coupled cluster code can perform calculations with full iterative treatment of single and double excitations and non-iterative inclusion of triple excitation effects. It is presently limited to closed-shell (RHF) references.

Note that symmetry is not used within most of the CCSD(T) code. This can have a profound impact on performance since the speed-up from symmetry is roughly the square of the number of irreducible representations. In the absence of symmetry, the performance of this code is competitive with other programs.

The operation of the coupled cluster code is controlled by the input block

```
CCSD
  [MAXITER <integer maxiter default 20>]
  [THRESH <real thresh default 10e-6>]
  [TOL2E <real tol2e default min(10e-12 , 0.01*$thresh$)>]
  [DIISBAS <integer diisbas default 5>]
  [FREEZE [[core] (atomic || <integer nfzc default 0>)] \
    [virtual <integer nfzv default 0>]]
  [IPRT <integer IPRT default 0>]
  [PRINT ...]
  [NOPRINT ...]
END
```

Note that the keyword CCSD is used for the input block regardless of the actual level of theory desired (specified with the TASK directive). The following directives are recognized within the CCSD group.

19.1 MAXITER — Maximum number of iterations

The maximum number of iterations is set to 20 by default. This should be quite enough for most calculations, although particularly troublesome cases may require more.

```
MAXITER <integer maxiter default 20>
```

19.2 THRESH — Convergence threshold

Controls the convergence threshold for the iterative part of the calculation. Both the RMS error in the amplitudes *and* the change in energy must be less than `thresh`.

```
THRESH <real thresh default 10e-6>
```

19.3 TOL2E — integral screening threshold

```
TOL2E <real tol2e default min(10e-12 , 0.01*$thresh$)>
```

The variable `tol2e` is used in determining the integral screening threshold for the evaluation of the energy and related quantities.

CAUTION! At the present time, the `tol2e` parameter only affects the three- and four-virtual contributions, and the triples, all of which are done “on the fly”. The transformations used for the other parts of the code currently have a hard-wired threshold of 10^{-12} . The default for `tol2e` is set to match this, and since user input can only make the threshold smaller, setting this parameter can only make calculations take longer.

19.4 DIISBAS — DIIS subspace dimension

Specifies the maximum size of the subspace used in DIIS convergence acceleration. Note that DIIS requires the amplitudes and errors be stored for each iteration in the subspace. Obviously this can significantly increase memory requirements, and could force the user to reduce `DIISBAS` for large calculations.

Measures to alleviate this problem, including more compact storage of the quantities involved, and the possibility of disk storage are being considered, but have not yet been implemented.

```
DIISBAS <integer diisbas default 5>
```

19.5 FREEZE — Freezing orbitals

```
[FREEZE [[core] (atomic || <integer nfzc default 0>)] \
  [virtual <integer nfzv default 0>]]
```

This directive is identical to that used in the MP2 module, Section 16.1.

19.6 IPRT — Debug printing

This directive controls the level of output from the code, mostly to facilitate debugging and the like. The larger the value, the more output printed. From looking at the source code, the interesting values seem to be `IPRT > 5`, 10, and 50.

```
IPRT <integer IPRT default 0>
```

19.7 PRINT and NOPRINT

The coupled cluster module supports the standard NWChem print control keywords, although very little in the code is actually hooked into this mechanism yet.

Item	Print Level	Description
"reference"	high	Wavefunction information
"guess pair energies"	debug	MP2 pair energies
"byproduct energies"	default	Intermediate energies
"term debugging switches"	debug	Switches for individual terms

19.8 Methods (Tasks) Recognized

Currently available methods are

- CCSD – Full iterative inclusion of single and double excitations
- CCSD+T (CCSD) – The fourth order triples contribution computed with converged singles and doubles amplitudes
- CCSD (T) – The linearized triples approximation due to Raghavachari.

The calculation is invoked using the the TASK directive, so to perform a CCSD+T(CCSD) calculation, for example, the input file should include the directive

```
TASK CCSD+T(CCSD)
```

Lower-level results which come as by-products (such as MP3/MP4) of the requested calculation are generally also printed in the output file and stored on the run-time database, but the method specified in the TASK directive is considered the primary result.

19.9 Debugging and Development Aids

The information in this section is intended for use by experts (both with the methodology and with the code), primarily for debugging and development work. Messing with stuff in listed in this section will probably make your calculation quantitatively **wrong!** Consider yourself warned!

19.9.1 Switching On and Off Terms

The /DEBUG/ common block contains a number of arrays which control the calculation of particular terms in the program. These are 15-element integer arrays (although from the code only a few elements actually effect anything) which can be set from the input deck. See the code for details of how the arrays are interpreted.

Printing of this data at run-time is controlled by the "term debugging switches" print option. The values are checked against the defaults at run-time and a warning is printed to draw attention to the fact that the calculation does not correspond precisely to the requested method.

```
DOA <integer array default 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2>
DOB <integer array default 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2>
DOG <integer array default 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1>
DOH <integer array default 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1>
DOJK <integer array default 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2 2>
DOS <integer array default 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1>
DOD <integer array default 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1>
```

Chapter 20

Geometry Optimization with DRIVER

The DRIVER module is one of two drivers (see Section 21 for documentation on STEPPER) to perform a geometry optimization function on the molecule defined by input using the GEOMETRY directive (see Section 6). Geometry optimization is either an energy minimization or a transition state optimization. The algorithm programmed in DRIVER is a quasi-newton optimization with line searches and approximate energy Hessian updates.

DRIVER is selected by default out of the two available modules to perform geometry optimization. In order to force use of DRIVER (e.g., because a previous optimization used STEPPER) provide a DRIVER input block (below) — even an empty block will force use of DRIVER.

Optional input for this module is specified within the compound directive,

```
DRIVER
  (LOOSE || DEFAULT || TIGHT)
  GMAX <real value>
  GRMS <real value>
  XMAX <real value>
  XRMS <real value>

  EPREC <real eprec default 1e-7>

  TRUST <real trust default 0.3>
  SADSTP <real sadstp default 0.1>

  CLEAR
  REDOAUTOZ

  INHESS <integer in Hess default 0>

  (MODDIR || VARDIR) <integer dir default 0>
  (FIRSTNEG || NOFIRSTNEG)

  MAXITER <integer maxiter default 20>

  BSCALE <real BSCALE default 1.0>
  ASCALE <real ASCALE default 0.25>
  TSCALE <real TSCALE default 0.1>
```

```

HSCALE <real HSCALE default 1.0>

PRINT ...

XYZ [<string xyz default $file_prefix$>]
NOXYZ

END

```

On each optimization step a line search is performed. To speed up calculations (up to two times), it may be beneficial to turn off the line search using following directive:

```
set driver:linopt 0
```

20.1 Convergence criteria

```

(LOOSE || DEFAULT || TIGHT)
GMAX <real value>
GRMS <real value>
XMAX <real value>
XRMS <real value>

```

In version 3.3 Gaussian-style convergence criteria have been adopted. The defaults may be used, or the directives LOOSE, DEFAULT, or TIGHT specified to use standard sets of values, or the individual criteria adjusted. All criteria are in atomic units. GMAX and GRMS control the maximum and root mean square gradient in the coordinates being used (Z-matrix, redundant internals, or Cartesian). XMAX and XRMS control the maximum and root mean square of the Cartesian step.

	LOOSE	DEFAULT	TIGHT
GMAX	0.0045d0	0.00045	0.000015
GRMS	0.0030d0	0.00030	0.00001
XMAX	0.0054d0	0.00180	0.00006
XRMS	0.0036d0	0.00120	0.00004

Note that GMAX and GRMS used for convergence of geometry may significantly vary in different coordinate systems such as Z-matrix, redundant internals, or Cartesian. The coordinate system is defined in the input file (default is Z-matrix). Therefore the choice of coordinate system may slightly affect converged energy. Although in most cases XMAX and XRMS are last to converge which are always done in Cartesian coordinates, which insures convergence to the same geometry in different coordinate systems.

The old criterion may be recovered with the input

```
gmax 0.0008; grms 1; xrms 1; xmax 1
```

20.2 Available precision

```
EPREC <real eprec default 1e-7>
```

In performing a line search the optimizer must know the precision of the energy (this has nothing to do with convergence criteria). The default value of $1e-7$ should be adjusted if less, or more, precision is available. Note that the default EPREC for DFT calculations is $5e-6$ instead of $1e-7$.

20.3 Controlling the step length

```
TRUST <real trust default 0.3>  
SADSTP <real sadstp default 0.1>
```

A fixed trust radius (`trust`) is used to control the step during minimizations, and is also used for modes being minimized during saddle-point searches. It defaults to 0.3 for minimizations and 0.1 for saddle-point searches. The parameter `sadstp` is the trust radius used for the mode being maximized during a saddle-point search and defaults to 0.1.

20.4 Maximum number of steps

```
MAXITER <integer maxiter default 20>
```

By default at most 20 geometry optimization steps will be taken, but this may be modified with this directive.

20.5 Discard restart information

```
CLEAR
```

By default Driver reuses Hessian information from a previous optimization, and, to facilitate a restart also stores which mode is being followed for a saddle-point search. This option deletes all restart data.

20.6 Regenerate internal coordinates

```
REDOAUTOZ
```

Deletes Hessian data and regenerates internal coordinates at the current geometry. Useful if there has been a large change in the geometry that has rendered the current set of coordinates invalid or non-optimal.

20.7 Initial Hessian

```
INHESS <integer inhess default 0>
```

- 0 = Default ... use restart data if available, otherwise use diagonal guess.
- 1 = Use diagonal initial guess.

- 2 = Use restart data if available, otherwise transform Cartesian Hessian from previous frequency calculation.

In addition, the diagonal elements of the initial Hessian for internal coordinates may be scaled using separate factors for bonds, angles and torsions with the following

```
BSCALE <real bscale default 1.0>
ASCALE <real ascale default 0.25>
TSCALE <real tscale default 0.1>
```

These values typically give a two-fold speedup over unit values, based on about 100 test cases up to 15 atoms using 3-21g and 6-31g* SCF. However, if doing many optimizations on physically similar systems it may be worth fine tuning these parameters.

Finally, the entire Hessian from any source may be scaled by a factor using the directive

```
HSCALE <real hscale default 1.0>
```

It might be of utility, for instance, when computing an initial Hessian using SCF to start a large MP2 optimization. The SCF vibrational modes are expected to be stiffer than the MP2, so scaling the initial Hessian by a number less than one might be beneficial.

20.8 Mode or variable to follow to saddle point

```
(MODDIR || VARDIR) <integer dir default 0>
(FIRSTNEG || NOFIRSTNEG)
```

When searching for a transition state the program, by default, will take an initial step uphill and then do mode following using a fuzzy maximum overlap (the lowest eigen-mode with an overlap with the previous search direction of 0.7 times the maximum overlap is selected). Once a negative eigen-value is found, that mode is followed regardless of overlap.

The initial uphill step is appropriate if the gradient points roughly in the direction of the saddle point, such as might be the case if a constrained optimization was performed at the starting geometry. Alternatively, the initial search direction may be chosen to be along a specific internal variable (using the directive VARDIR) or along a specific eigen-mode (using MODDIR). Following a variable might be valuable if the initial gradient is either very small or very large. Note that the eigen-modes in the optimizer have next-to-nothing to do with the output from a frequency calculation. You can examine the eigen-modes used by the optimizer with

```
driver; print hvecs; end
```

The selection of the first negative mode is usually a good choice if the search is started in the vicinity of the transition state and the initial search direction is satisfactory. However, sometimes the first negative mode might not be the one of interest (e.g., transverse to the reaction direction). If NOFIRSTNEG is specified, the code will not take the first negative direction and will continue doing mode-following until that mode goes negative.

20.9 Optimization history as XYZ files

```
XYZ [<string xyz default $fileprefix>]
NOXYZ
```

The XYZ directive causes the geometry at each step (but not intermediate points of a line search) to be output into separate files in the permanent directory in XYZ format. The optional string will prefix the filename. The NOXYZ directive turns this off.

For example, the input

```
driver; xyz test; end
```

will cause files test-000.xyz, test-001.xyz, ... to be created in the permanent directory.

The script rasmolmovie in the NWChem contrib directory can be used to turn these into an animated GIF movie.

20.10 Print options

The UNIX command "egrep '^@' < output" will extract a pretty table summarizing the optimization.

If you specify the NWChem input

```
scf; print none; end
driver; print low; end
task scf optimize
```

you'll obtain a pleasantly terse output.

For more control, these options for the standard print directive are recognized

- debug - prints a large amount of data. Don't use in parallel.
- high - print the search direction in internals
- default - prints geometry for each major step (not during the line search), gradient in internals (before and after application of constraints)
- low - prints convergence and energy information. At convergence prints final geometry, change in internals from initial geometry

and these specific print options

- finish (low) - print geometry data at end of calculation
- bonds (default) - print bonds at end of calculation
- angles (default) - print angles at end of calculation
- hvecs (never) - print eigen-values/vectors of the Hessian
- searchdir (high) - print the search direction in internals
- 'internal gradient' (default) - print the gradient in internals
- sadmode (default) - print the mode being followed to the saddle point

Chapter 21

Geometry Optimization with STEPPER

The STEPPER module performs a search for critical points on the potential energy surface of the molecule defined by input using the GEOMETRY directive (see Section 6). Since STEPPER is **not** the primary geometry optimization module in NWChem the compound directive is required; the DRIVER module is the default (see Section 20). Input for this module is specified within the compound directive,

```
STEPPER
...
END
```

The presence of the STEPPER compound directive automatically turns off the default geometry optimization tool DRIVER. Input specified for the STEPPER module must appear in the input file *after* the GEOMETRY directive, since it must know the number of atoms that are to be used in the geometry optimization. In the current version of NWChem, STEPPER can be used only with geometries that are defined in Cartesian coordinates. STEPPER removes translational and rotational components before determining the step direction (5 components for linear systems and 6 for others) using a standard Eckart algorithm. The default initial guess nuclear Hessian is the identity matrix.

The default in STEPPER is to minimize the energy as a function of the geometry with a maximum of 20 geometry optimization iterations. When this is the desired calculation, no input is required other than the STEPPER compound directive. However, the user also has the option of defining different tasks for the STEPPER module, and can vary the number of iterations and the convergence criteria from the default values. The input for these options is described in the following sections.

21.1 MIN and TS — Minimum or transition state search

The default is for STEPPER to minimize the energy with respect to the geometry of the system. This default behavior may be forced with the directive

```
MIN
```

STEPPER can also be used to find the transition state by following the lowest eigenvector of the nuclear Hessian. This is usually invoked by using the `saddle` keyword on the TASK directive (Section 5.10), but it may also be selected by specifying the directive

```
TS
```

in the STEPPER input.

21.2 TRACK — Mode selection

STEPPER has the ability to “track” a specific mode during an optimization for a transition state search, the user can also have the module track the eigenvector corresponding to a specific mode. This is done by specifying the directive

```
TRACK [nmode <integer nmode default 1>]
```

The keyword TRACK tells STEPPER to track the eigenvector corresponding to the integer value of <nmode> during a transition state walk. (Note: this input is invalid for a minimization walk since following a specific eigenvector will not necessarily give the desired local minimum.) The step is constructed to go up in energy along the nmode eigenvector and down in all other degrees of freedom.

21.3 MAXITER — Maximum number of steps

In most applications, 20 stepper iterations will be sufficient to obtain the energy minimization. However, the user has the option of specifying the maximum number of iterations allowed, using the input line,

```
MAXITER <integer maxiter default 20>
```

The value specified for the integer <maxiter> defines the maximum number of geometry optimization steps. The geometry optimization will restart automatically.

21.4 TRUST — Trust radius

The size of steps that can be taken in STEPPER is controlled by the trust radius which has a default value of 0.1. Steps are constrained to be no larger than the trust radius. The user has the option of overriding this default using the keyword TRUST, with the following input line,

```
TRUST <real radius default 0.1>
```

The larger the value specified for the variable radius, the larger the steps that can be taken by STEPPER. Experience has shown that for larger systems (i.e., those with 20 or more atoms), a value of 0.5, or greater, usually should be entered for <radius>.

21.5 CONVGGM, CONVGG and CONVGE — Convergence criteria

Three convergence criteria can be specified explicitly for the STEPPER calculations. The keyword CONVGGM allows the user to specify the convergence tolerance for the largest component of the gradient. This is the primary convergence criterion, as per the default settings, although all three criteria are in effect. this default setting is consistent with the other optimizer module DRIVER. The input line for CONVGGM has the following form,

```
CONVGGM <real convggm default 8.0d-04>
```

The keyword `CONVGG` allows the user to specify the convergence tolerance for the gradient norm for all degrees of freedom. The input line is of the following form,

```
CONVGG <real convgg default 1.0d-02>
```

The entry for the real variable `<convgg>` should be approximately equal to the square root of the energy convergence tolerance.

The energy convergence tolerance is the convergence criterion for the energy difference in the geometry optimization in STEPPER. It can be specified by input using a line of the following form,

```
CONVGE <real convge default 1.0d-04>
```

21.6 Backstepping in STEPPER

If a step taken during the optimization is too large (e.g., the step causes the energy to go up for a minimization or down for a transition state search), the STEPPER optimizer will automatically “backstep” and correct the step based on information prior to the faulty step. If you have an optimization that “backsteps” frequently then the initial trust radius should most likely be decreased.

21.7 Initial Nuclear Hessian Options

Stepper uses a modified Fletcher-Powell algorithm to find the transition state or energy minimum on the potential energy hypersurface. There are two files left in the user’s permanent directory that are used to provide an initial hessian to the critical point search algorithm. If these files do not exist then the default is to use a unit matrix as the initial hessian. Once Stepper executes it generates a binary dump file by the name `name.stpr41` which will be used on all subsequent stepper runs and modified with the current updated hessian. The default file prefix is the “name” that is used (c.f., 5.1). It also stores the information for the last valid step in case the algorithm must take a “backstep” (c.f., 21.6). This file is the working data store for all stepper-based optimizations. This file is never deleted by default and is the *first* source of an initial hessian. The second source of an initial hessian is an ascii file that contains the lower triangular values of the initial hessian. This is stored in file `name.hess`, where “name” is again the default file prefix. This is the *second* source of an initial hessian and is the method used to incorporate an initial hessian from any other source (e.g., another *ab initio* code, a molecular mechanics code, etc.). To get a decent starting hessian at a given point you can use the task specification `task scf hessian`, with a smaller basis set, which will by default generate the `name.hess` file. Then you may define your basis set of choice and proceed with the optimization you desire.¹

¹If you have done a geometry optimization and hessian generation in the same input deck using a small basis set, you must make sure you delete the `name.stpr41` file since stepper will by default use that hessian and not the one in the `name.hess` file

Chapter 22

Constraints for Geometry Optimization

The constraints directive allows the user to specify which constraints should be imposed on the system during the geometry optimization. Currently such constraints are limited to fixed atom positions and harmonic restraints (springs) on the distance between the two atoms. The general form of constraints block is presented below:

```
CONSTRAINTS [string name ] \  
    [clear] \  
    [enable||disable] \  
    [fix atom <integer list>] \  
    [spring bond <integer atom1> <integer atom2> <real k> <real r0> ]  
END
```

The keywords are described below

- *name* – optional keyword that associates a name with a given set of constraints. Any unnamed set of constraints will be given a name "default" and will be automatically loaded prior to a calculation. Any constraints with the name other than "default" will have to be loaded manually using SET directive. For example,

```
CONSTRAINTS one  
    spring bond 1 3 5.0 1.3  
    fix atom 1  
END
```

```
#the above constraints can be loaded using set directive  
set constraints one  
....  
task ....
```

- *clear* – destroys any prior constraint information. This may be useful when the same constraints have to be redefined or completely removed from the runtime database.
- *enable||disable* – enables or disables particular set of constraints without actually removing the information from the runtime database.
- *fix atom* – fixes atom positions during geometry optimization. This directive requires an integer list that specifies which atoms are to be fixed. This directive can be repeated within a given constraints block. To illustrate the use

"*fix atom*" directive let us consider a situation where we would like to fix atoms 1, 3, 4, 5, 6 while performing an optimization on some hypothetical system. There are actually several ways to enter this particular constraint. There is a straightforward option which requires the most typing

```
constraints
  fix atom 1 3 4 5 6
end
```

Second method uses list input

```
constraints
  fix atom 1 3:6
end
```

Third approach illustrates the use of multiple *fix atom* directives

```
constraints
  fix atom 1
  fix atom 3:6
end
```

- *spring bond* $\langle ijkr_0 \rangle$ – places a spring with a spring constant k and equilibrium length r_0 between atoms i and j (all in atomic units). Please note that this type of constraint adds an additional term to the total energy expression

$$E = E_{total} + \frac{1}{2}k(r_{ij} - r_0)^2$$

This additional term forces the distance between atoms i and j to be in the vicinity of r_0 but never exactly that. In general the spring energy term will always have some nonzero residual value, and this has to be accounted for when comparing total energies. The "*spring bond*" directive can be repeated within a given constraints block. If the spring between the same pair of atoms is defined more than once, it will be replaced by the latest specification in the order it appears in the input block.

Chapter 23

Hybrid Calculations with ONIOM

ONIOM is the hybrid method of Morokuma and co-workers that enables different levels of theory to be applied to different parts of a molecule/system and combined to produce a consistent energy expression. The objective is to perform a high-level calculation on just a small part of the system and to include the effects of the remainder at lower levels of theory, with the end result being of similar accuracy to a high-level calculation on the full system.

1. M. Svensson, S. Humbel, R.D.J. Froese, T. Mastubara, S. Sieber, and K. Morokuma, *J. Phys. Chem.*, 100, 19357 (1996).
2. S. Dapprich, I. Komaromi, K.S. Byun, K. Morokuma, and M.J. Frisch, *J. Mol. Struct. (Theochem)*, 461-462, 1 (1999).
3. R.D.J. Froese and K. Morokuma in “Encyclopedia of Computational Chemistry,” volume 2, pp.1244-1257, (ed. P. von Rague Schleyer, John Wiley and Sons, Chichester, Sussex, 1998).

The NWChem ONIOM module implements two- and three-layer ONIOM models for use in energy, gradient, geometry optimization, and vibrational frequency calculations with any of the pure quantum mechanical methods within NWChem. At the present time, it is not possible to perform ONIOM calculations with either solvation models or classical force fields. Nor is it yet possible to compute properties except as derivatives of the total energy.

Using the terminology of Morokuma et al., the full molecular geometry including all atoms is referred to as the “real” geometry and it is treated using a “low”-level of theory. A subset of these atoms (referred to as the “model” geometry) are treated using both the “low”-level and a “high”-level of theory. A three-layer model also introduces an “intermediate” model geometry and a “medium” level of theory.

The two-layer model requires a high and low level of theory and a real and model molecular geometry. The energy at the high-level of theory for the real geometry is estimated as

$$E(\text{High,Real}) = E(\text{Low,Real}) + [E(\text{High,Model}) - E(\text{Low,Model})].$$

The three-layer model requires high, medium and low levels of theory, and real, intermediate and model geometries and the corresponding energy estimate is

$$E(\text{High,Real}) = E(\text{Low,Real}) + [E(\text{High,Model}) - E(\text{Medium,Model})] \\ + [E(\text{Medium,Inter}) - E(\text{Low,Inter})].$$

When does ONIOM work well? The approximation for a two-layer model will be good if

- the model system includes the interactions that dominate the energy difference being computed and the high-level of theory describes these to the required precision, and
- the interactions between the model and the rest of the real system (substitution effects) are described to sufficient accuracy at the lower level of theory.

ONIOM is used to compute energy differences and the absolute energies are not all that meaningful even though they are well defined. Due to cancellation of errors, ONIOM actually works better than you might expect, but a poorly designed calculation can yield very bad results. Please read and heed the caution at the end of the article by Dapprich et al.

The input options are as follows

```

ONIOM
HIGH <string theory> [basis <string basis default "ao basis">] \
    [ecp <string ecp>] [input <string input>]
[MEDIUM <string theory> [basis <string basis default "ao basis">] \
    [ecp <string ecp>] [input <string input>]]
LOW <string theory> [basis <string basis default "ao basis">] \
    [ecp <string ecp>] [input <string input>]
MODEL <integer natoms> [charge <double charge>] \
    [<integer i1 j1> <real g1> [<string tag1>] ...]
[INTER <integer natoms> [charge <double charge>] \
    [<integer i1 j1> <real g1> [<string tag1>] ...]]
[VECTORS [low-real <string mofile>] [low-model <string mofile>] \
    [high-model <string mofile>] [medium-model <string mofile>] \
    [medium-inter <string mofile>] [low-inter <string mofile>]]
[PRINT ...]
[NOPRINT ...]
END

```

which are described in detail below.

For better validation of user input, the HIGH, LOW and MODEL directives must always be specified. If the one of the MEDIUM or INTER directives are specified, then so must the other.

23.1 Real, model and intermediate geometries

The geometry and total charge of the full or real system should be specified as normal using the geometry directive (see Section 6). If N_{model} of the atoms are to be included in the model system, then these should be specified first in the geometry. Similarly, in a three-layer calculation, if there are N_{inter} atoms to be included in the intermediate system, then these should also be arranged together at the beginning of the geometry. The implicit assumption is that the model system is a subset of the intermediate system which is a subset of the real system. The number of atoms to be included in the model and intermediate systems are specified using the MODEL and INTER directives. Optionally, the total charge of the model and intermediate systems may be adjusted. The default is that all three systems have the same total charge.

Example 1. A two-layer calculation on $K^+(H_2O)$ taking the potassium ion as the model system. Note that no bonds are broken so no link atoms are introduced. The real geometry would be specified with potassium (the model) first.

```

geometry autosym
  K  0   0.00   1.37
  O  0   0.00  -1.07
  H  0  -0.76  -1.68
  H  0   0.76  -1.68
end

```

and the following directive in the ONIOM input block indicates that one atom (implicitly the first in the geometry) is in the model system

```

model 1

```

23.1.1 Link atoms

Link atoms for bonds spanning two regions are automatically generated from the bond information. The additional parameters on the MODEL and INTER directives describe the broken bonds including scale factors for placement of the link atom and, optionally, the type of link atom. The type of link atom defaults to hydrogen, but any type may be specified (actually here you are specifying a geometry tag which is used to associate a geometrical center with an atom type and basis sets, etc.. See section 6.3). For each broken bond specify the numbers of the two atoms (i and j), the scale factor (g) and optionally the tag of the link atom. Link atoms are placed along the vector connecting the first to the second atom of the bond according to the equation

$$\underline{R}_{link} = (1 - g)\underline{R}_1 + g * \underline{R}_2$$

where g is the scale factor. If the scale factor is one, then the link atom is placed where the second atom was. More usually, the scale factor is less than one, in which case the link atom is placed between the original two atoms. The scale factor should be chosen so that the link atom (usually hydrogen) is placed near its equilibrium bond length from the model atom. E.g., when breaking a single carbon-carbon bond (typical length 1.528 Angströms) using a hydrogen link atom we will want a carbon-hydrogen bond length of about 1.084 Angströms, so the scale factor should be chosen as $1.084/1.528 \approx 0.709$.

Example 2. A calculation on acetaldehyde ($H_3C - CHO$) using aldehyde ($H - CHO$) as the model system. The covalent bond between the two carbon atoms is broken and a link atom must be introduced to replace the methyl group. The link atom is automatically generated — all you need to do is specify the atoms in the model system that are also in the real system (here CHO) and the broken bonds. Here is the geometry of acetaldehyde with the CHO of aldehyde first

```

geometry
  C   -0.383   0.288   0.021
  H   -1.425   0.381   0.376
  O    0.259   1.263  -0.321

  H    0.115  -1.570   1.007
  H   -0.465  -1.768  -0.642
  H    1.176  -1.171  -0.352
  C    0.152  -1.150   0.005
end

```

There are three atoms (the first three) of the real geometry included in the model geometry, and we are breaking the bond between atoms 1 and 7, replacing atom 7 with a hydrogen link atom. This is all accomplished by the directive

```
model 3 1 7 0.709 H
```

Since the default link atom is hydrogen there is actually no need to specify the “H”.

See also Section 23.6.3 for a more complex example.

23.1.2 Numbering of the link atoms

The link atoms are appended to the atoms of the model or intermediate systems in the order that the broken bonds are specified in the input. This is of importance only if manually constructing an initial guess.

23.2 High, medium and low theories

The two-layer model requires both the high-level and low-level theories be specified. The three-layer model also requires the medium-level theory. Each of these includes a theory (such as SCF, MP2, DFT, CCSD, CCSD(T), etc.), an optional basis set, an optional ECP, and an optional string containing general NWChem input.

23.2.1 Basis specification

The basis name on the theory directive (high, medium, or low) is that specified on a basis set directive (see Section 7) and *not* the name of a standard basis in the library. If not specified, the basis set for the high-level theory defaults to the standard "ao basis". That for the medium level defaults to the high-level basis, and the low-level basis defaults to the medium-level basis. Other wavefunction parameters are obtained from the standard wavefunction input blocks. See 23.6.2 for an example.

23.2.2 Effective core potentials

If an effective core potential is specified in the usual fashion (see Section 8) outside of the ONIOM input then this will be used in all calculations. If an alternative ECP name (the name specified on the ECP directive in the same manner as done for basis sets) is specified on one of the theory directives, then this ECP will be used in preference for that level of theory. See Section 23.6.2 for sample input.

23.2.3 General input strings

For many purposes, the ability to specify the theory, basis and effective core potential is adequate. All of the options for each theory are determined from their independent input blocks. However, if the same theory (e.g., DFT) is to be used with different options for the ONIOM theoretical models, then the general input strings must be used. These strings are processed as NWChem input each time the theoretical model is invoked. The strings may contain any NWChem input, except for options pertaining to ONIOM and the task directive. The intent that the strings be used just to control the options pertaining to the theory being used.

A word of caution. Be sure to check that the options are producing the desired results. Since the NWChem database is persistent and the ONIOM calculations happen in an undefined order, the input strings should fully define the calculation you wish to have happen.

For instance, if the high model is DFT/B3LYP/6-311g** and the low model is DFT/LDA/3-21g, the ONIOM input might look like this

```
oniom
  model 3
  low dft basis 3-21g   input "dft\; xc\; end"
  high dft basis 6-311g** input "dft\; xc b3lyp\; end"
end
```

The empty XC directive restores the default LDA exchange-correlation option (see Section 11.3). Note that semicolons and other quotation marks inside the input string must be preceded by a backslash to avoid special interpretation.

See Section 23.6.4 for another example.

23.3 Use of symmetry

Symmetry should work just fine as long as the model and intermediate regions respect the symmetry — i.e., symmetry equivalent atoms need to be treated equivalently. If symmetry equivalent atoms must be treated in separate regions then the symmetry must be lowered (or completely switched off).

23.4 Molecular orbital files

The VECTORS directive in the ONIOM block is different to that elsewhere in NWChem. For each of the necessary combinations of theory and geometry you can specify a different file for the molecular orbitals. By default each combination will store the MO vectors in the permanent directory using a file name created by appending to the name of the calculation the following string

- low-real — ".lrmos"
- low-inter — ".limos"
- low-model — ".lmmos"
- medium-inter — ".mimos"
- medium-model — ".mmmms"
- high-model — ".hmmos"

Each calculation will utilize the appropriate vectors which is more efficient during geometry optimizations and frequency calculations, and is also useful for the initial calculation. In the absence of existing MO vectors files, the default atomic guess is used (see Section 10.5).

If special measures must be taken to converge the initial SCF, DFT or MCSCF calculation for one or more of the systems, then initial vectors may be saved in a file with the default name, or another name may be specified using the VECTORS directive. Note that subsequent vectors (e.g., from a geometry optimization) will be written back to this file, so take a copy if you wish to preserve it. To generate the initial guess for the model or intermediate systems it is necessary to generate the geometries which is most readily done, if there are link atoms, by just running NWChem on the input for the ONIOM calculation on your workstation. It will print these geometries before starting any calculations which you can then terminate.

E.g., in a calculation on Fe(III) surrounded by some ligands, it is hard to converge the full (real) system from the atomic guess so as to obtain a d^5 configuration for the iron atom since the d orbitals are often nominally lower in energy than some of the ligand orbitals. The most effective mechanism is to converge the isolated Fe(III) and then to

use the fragment guess (see Section 10.5.1) as a starting guess for the real system. The resulting converged molecular orbitals can be saved either with the default name (as described above in this section), in which case no additional input is necessary. If an alternative name is desired, then the `VECTORS` directive may be used as follows

```
vectors low-real /u/rjh/jobs/fe_ether_water.mos
```

23.5 Restarting

Restart of ONIOM calculations does not currently work as smoothly as we would like. For geometry optimizations that terminated gracefully by running out of iterations, the restart will work as normal. Otherwise, specify in the input of the restart job the last geometry of the optimization. The Hessian information will be reused and the calculation should proceed losing at most the cost of one ONIOM gradient evaluation. For energy or frequency calculations, restart may not currently be possible.

23.6 Examples

23.6.1 Hydrocarbon bond energy

A simple two-layer model changing just the wavefunction with one link atom.

This reproduces the two-layer ONIOM (MP2:HF) result from Dapprich et al. for the reaction $R-CH_3 = R-CH_2 + H$ with $R = CH_3$ using CH_4 as the model. The geometries of $R-CH_3$ and $R-CH_2$ are optimized at the DFT-B3LYP/6-311++G** level of theory, and then ONIOM is used to compute the binding energy using UMP2 for the model system and HF for the real system. The results, including MP2 calculations on the full system for comparison, are as given in Table 23.1

Theory	Me-CH2	Me-Me	H	De(Hartree)	De(kcal/mol)
B3LYP	-79.185062	-79.856575	-0.502256	0.169257	106.2
HF	-78.620141	-79.251701	-0.499817	0.131741	82.7
MP2	-78.904716	-79.571654	-0.499817	0.167120	104.9
MP2:HF	-78.755223	-79.422559	-0.499817	0.167518	105.1

Table 23.1: Energies for ONIOM example 1, hydrocarbon bond energy using MP2:HF two-layer model.

The following input first performs a calculation on CH_3-CH_2 , and then on CH_3-CH_3 . Note that in the second calculation we cannot use the full symmetry since we are breaking the C-C bond in forming the model system (the non-equivalence of the methyl groups is perhaps more apparent if we write $R-CH_3$).

```
start

basis spherical
  H library 6-311++G**; C library 6-311++G**
end

title "ONIOM Me-CH2"

geometry autosym
```

```

H    -0.23429328    1.32498565    0.92634814
H    -0.23429328    1.32498565   -0.92634814
C    -0.13064265    0.77330370    0.00000000
H    -1.01618703   -1.19260361    0.00000000
H     0.49856072   -1.08196901   -0.88665533
H     0.49856072   -1.08196901    0.88665533
C    -0.02434414   -0.71063687    0.00000000
end

scf; uhf; doublet; thresh 1e-6; end
mp2; freeze atomic; end

oniom
  high mp2
  low  scf
  model 3   3 7 0.724
end

task oniom

title "ONIOM Me-Me"

geometry # Note cannot use full D3D symmetry here
H    -0.72023641    0.72023641   -1.16373235
H     0.98386124    0.26362482   -1.16373235
H    -0.26362482   -0.98386124   -1.16373235
C     0.00000000    0.00000000   -0.76537515
H     0.72023641   -0.72023641    1.16373235
H    -0.98386124   -0.26362482    1.16373235
H     0.26362482    0.98386124    1.16373235
C     0.00000000    0.00000000    0.76537515
end

scf; rhf; singlet; end

oniom
  high mp2
  low  scf
  model 4   4 8 0.724
end

task oniom

```

23.6.2 Optimization and frequencies

A two-layer model including modification of theory, basis, ECP and total charge and no link atoms.

This input reproduces the ONIOM optimization and vibrational frequency calculation of $Rh(CO)_2Cp$ of Dapprich et al. The model system is $Rh(CO)_2^+$. The low theory is the Gaussian LANL2MB model (Hay-Wadt n+1 ECP with minimal basis on Rh, STO-3G on others) with SCF. The high theory is the Gaussian LANL2DZ model (another Hay-

Wadt ECP with a DZ basis set on Rh, Dunning split valence on the other atoms) with DFT/B3LYP. Note that different names should be used for the basis set and ECP since the same mechanism is used to store them in the database.

```

start

ecp LANL2DZ_ECP
  rh library LANL2DZ_ECP
end

basis LANL2DZ spherical
  rh library LANL2DZ_ECP
  o library SV_(Dunning-Hay); c library SV_(Dunning-Hay); h library SV_(Dunning-Hay)
end

ecp Hay-Wadt_MB_(n+1)_ECP
  rh library Hay-Wadt_MB_(n+1)_ECP
end

# This is the minimal basis used by Gaussian. It is not the same
# as the one in the EMSL basis set library for this ECP.
basis Hay-Wadt_MB_(n+1) spherical
  Rh s; .264600D+01 -.135541D+01; .175100D+01 .161122D+01; .571300D+00 .589381D+00
  Rh s; .264600D+01 .456934D+00; .175100D+01 -.595199D+00; .571300D+00 -.342127D+00
    .143800D+00 .410138D+00; .428000D-01 .780486D+00
  Rh p; .544000D+01 -.987699D-01; .132900D+01 .743359D+00; .484500D+00 .366846D+00
  Rh p; .659500D+00 -.370046D-01; .869000D-01 .452364D+00; .257000D-01 .653822D+00
  Rh d; .366900D+01 .670480D-01; .142300D+01 .455084D+00; .509100D+00 .479584D+00
    .161000D+00 .233826D+00
  o library sto-3g; c library sto-3g; h library sto-3g
end

charge 0
geometry autosym
  rh      0.00445705    -0.15119674    0.00000000
  c      -0.01380554   -1.45254070    1.35171818
  c      -0.01380554   -1.45254070   -1.35171818
  o      -0.01805883   -2.26420212    2.20818932
  o      -0.01805883   -2.26420212   -2.20818932
  c      1.23209566    1.89314720    0.00000000
  c      0.37739392    1.84262319   -1.15286640
  c     -1.01479160    1.93086461   -0.70666350
  c     -1.01479160    1.93086461    0.70666350
  c      0.37739392    1.84262319    1.15286640
  h      2.31251453    1.89903673    0.00000000
  h      0.70378132    1.86131979   -2.18414218
  h     -1.88154273    1.96919306   -1.35203550
  h     -1.88154273    1.96919306    1.35203550
  h      0.70378132    1.86131979    2.18414218
end

dft; grid fine; convergence gradient 1e-6 density 1e-6; xc b3lyp; end

```

```

scf; thresh 1e-6; end

oniom
  low scf basis Hay-Wadt_MB_(n+1) ecp Hay-Wadt_MB_(n+1)_ECP
  high dft basis LANL2DZ ecp LANL2DZ_ECP
  model 5 charge 1
  print low
end

task oniom optimize
task oniom freq

```

23.6.3 A three-layer example

A three layer example combining CCSD(T), and MP2 with two different quality basis sets, and using multiple link atoms.

The full system is tetra-dimethyl-amino-ethylene (TAME) or $(\text{N}(\text{Me})_2)_2\text{C}=\text{C}-(\text{N}(\text{Me})_2)_2$. The intermediate system is $(\text{NH}_2)_2\text{C}=\text{C}-(\text{NH}_2)_2$ and $\text{H}_2\text{C}=\text{CH}_2$ is the model system. CCSD(T)+aug-cc-pvtz is used for the model region, MP2+aug-cc-pvtz for the intermediate region, and MP2+aug-cc-pvdz for everything.

In the real geometry the first two atoms (C, C) are the model system (link atoms will be added automatically). The first six atoms (C, C, N, N, N, N) describe the intermediate system (again with link atoms to be added automatically). The atoms have been numbered using comments to make the bonding input easier to generate.

To make the model system, four C-N bonds are broken between the ethylene fragment and the dimethyl-amino groups and replaced with C-H bonds. To make the intermediate system, eight C-N bonds are broken between the nitrogens and the methyl groups and replaced with N-H bonds. The scaling factor could be chosen differently for each of the bonds.

```

start

geometry
  C  0.40337795 -0.17516305 -0.51505208 # 1
  C -0.40328664  0.17555927  0.51466084 # 2
  N  1.87154979 -0.17516305 -0.51505208 # 3
  N -0.18694782 -0.60488524 -1.79258692 # 4
  N  0.18692927  0.60488318  1.79247594 # 5
  N -1.87148219  0.17564718  0.51496494 # 6
  C  2.46636552  1.18039452 -0.51505208 # 7
  C  2.48067731 -1.10425355  0.46161675 # 8
  C -2.46642715 -1.17982091  0.51473105 # 9
  C -2.48054940  1.10495864 -0.46156202 # 10
  C  0.30027136  0.14582197 -2.97072148 # 11
  C -0.14245927 -2.07576980 -1.96730852 # 12
  C -0.29948109 -0.14689874  2.97021079 # 13
  C  0.14140463  2.07558249  1.96815181 # 14
  H  0.78955302  2.52533887  1.19760764
  H -0.86543435  2.50958894  1.88075113
  ... and 22 other hydrogen atoms on the methyl groups
end

```

```

basis aug-cc-pvtz spherical
  C library aug-cc-pvtz; H library aug-cc-pvtz
end

basis aug-cc-pvdz spherical
  C library aug-cc-pvtz; H library aug-cc-pvtz
end

oniom
  high ccSD(t) basis aug-cc-pvtz
  medium mp2 basis aug-cc-pvtz
  low mp2 basis aug-cc-pvdz
  model 2   1 3 0.87   1 4 0.87   2 5 0.87   2 6 0.87

  inter 6   3 7 0.69   3 8 0.69   4 11 0.69   4 12 0.69 \
           5 13 0.69   5 14 0.69   6 9 0.69    6 10 0.69
end

task oniom

```

23.6.4 DFT with and without charge fitting

Demonstrates use of general input strings.

A two-layer model for anthracene (a linear chain of three fused benzene rings) using benzene as the model system. The high-level theory is DFT/B3LYP/TZVP with exact Coulomb. The low level is DFT/LDA/DZVP2 with charge fitting.

Note the following.

1. The semi-colons and quotation marks inside the input string must be quoted with backslash.
2. The low level of theory sets the fitting basis set and the high level of theory unsets it.

```

start
geometry
  symmetry d2h
  C   0.71237329   -1.21458940   0.0
  C  -0.71237329   -1.21458940   0.0
  C   0.71237329    1.21458940   0.0
  C  -0.71237329    1.21458940   0.0
  C  -1.39414269    0.00000000   0.0
  C   1.39414269    0.00000000   0.0
  H  -2.47680865    0.00000000   0.0
  H   2.47680865    0.00000000   0.0
  C   1.40340535   -2.48997027   0.0
  C  -1.40340535   -2.48997027   0.0
  C   1.40340535    2.48997027   0.0
  C  -1.40340535    2.48997027   0.0
  C   0.72211503    3.64518615   0.0

```

```
C   -0.72211503    3.64518615    0.0
C    0.72211503   -3.64518615    0.0
C   -0.72211503   -3.64518615    0.0
H    2.48612947    2.48094825    0.0
H    1.24157357    4.59507342    0.0
H   -1.24157357    4.59507342    0.0
H   -2.48612947    2.48094825    0.0
H    2.48612947   -2.48094825    0.0
H    1.24157357   -4.59507342    0.0
H   -1.24157357   -4.59507342    0.0
H   -2.48612947   -2.48094825    0.0
end

basis small
  h library DZVP_(DFT_Orbital)
  c library DZVP_(DFT_Orbital)
end

basis fitting
  h library DGauss_A1_DFT_Coulomb_Fitting
  c library DGauss_A1_DFT_Coulomb_Fitting
end

basis big
  h library TZVP_(DFT_Orbital)
  c library TZVP_(DFT_Orbital)
end

oniom
  model 8   1 9 0.75   2 10 0.75   3 11 0.75   4 12 0.75
  high dft basis big   input "unset \"cd basis\""; dft\; xc b3lyp\; end"
  low  dft basis small input "set \"cd basis\" fitting\; dft\; xc\; end"
end

task oniom
```


Chapter 24

Hessians

This section relates to the computation of analytic Hessians which are available for open and closed shell SCF, except ROHF and for closed shell DFT. Analytic Hessians are not currently available for SCF or DFT calculations relativistic all-electron methodologies or for charge fitting with DFT. The current algorithm is fully in-core and does not use symmetry. This will be changed in the next release.

There is no required input for the Hessian module. This module only impacts the Hessian calculation. For options for calculating the frequencies, please see Section 25, the Vibrational module.

24.1 Hessian Module Input

All input for the Hessian Module is optional since the default definitions are usually correct for most purposes. The generic module input begins with `hessian` and has the form:

```
hessian
  thresh <real tol default 1d-6>
  print ...
  profile
end
```

24.1.1 Defining the wavefunction threshold

You may modify the default threshold for the wavefunction. This keyword is identical to `THRESH` in the SCF, Section 10.7, and the `CONVERGENCE` gradient in the DFT, Section 11.7. The usual defaults for the convergence of the wavefunction for single point and gradient calculations is generally not tight enough for analytic Hessians. Therefore, the Hessian, by default, tightens these up to 1d-6 and runs an additional energy point if needed. If, during an analytic Hessian calculation, you encounter an error:

```
cphf_solve:the available MOs do not satisfy the SCF equations
```

the convergence criteria of the wavefunction generally needs to be tightened.

24.1.2 Profile

The PROFILE keyword provides additional information concerning the computation times of different sections of the hessian code. Summary information is given about the maximum, minimum and average times that a particular section of the code took to complete. This is normally only useful for developers.

24.1.3 Print Control

Known controllable print options are:

Name	Print Level	Description
"hess_follow"	high	more information about where the calculation is
"cphf_cont"	debug	detailed CPHF information
"nucdd_cont"	debug	detailed nuclear contribution information
"onedd_cont"	debug	detailed one electron contribution information
"twodd_cont"	debug	detailed two electron contribution information
"fock_xc"	debug	detailed XC information during the fock builds

Table 24.1: Hessian Print Control Specifications

Chapter 25

Vibrational frequencies

The nuclear hessian which is used to compute the vibrational frequencies can be computed by finite difference for any ab initio wave-function that has analytic gradients or by analytic methods for SCF and DFT (see Section 24 for details). The appropriate nuclear hessian generation algorithm is chosen based on the user input when `TASK <theory> frequencies` is the task directive.

The vibrational package was integrated from the Utah Messkit and can use any nuclear hessian generated from the driver routines, finite difference routines or any analytic hessian modules. There is no required input for the “VIB” package. VIB computes the Infra Red frequencies and intensities¹ for the computed nuclear hessian and the “projected” nuclear hessian. The VIB module projects out the translations and rotations of the nuclear hessian using the standard Eckart projection algorithm. It also computes the zero point energy for the molecular system based on the frequencies obtained from the projected hessian.

The default mass of each atom is used unless an alternative mass is provided via the geometry input, (c.f., 6) or redefined using the vibrational module input. The default mass is the mass of the most abundant isotope of each element.² If the abundance was roughly equal, the mass of the isotope with the longest half life was used.

25.1 Vibrational Module Input

All input for the Vibrational Module is optional since the default definitions will compute the frequencies and IR intensities³. The generic module input can begin with `vib`, `freq`, `frequency` and has the form:

```
{freq || vib || frequency}
  reuse [<string> hessian_filename]
  mass <integer> lexical_index <real> new_mass
  mass <string> tag_identifier <real> new_mass
  animate [<real> step_size_for_animation]
end
```

¹Intensities are only computed if the dipole derivatives are available; these are computed by default for most methods that use the finite difference driver routines

²c.f., “The Elements” by John Emsley, Oxford University Press, (C) 1989, ISBN 0-19-855237-8.

³The geometry specification at the point where the hessian is computed must be the default “geometry” on the current run-time-data-base for the projection to work properly.

25.1.1 Hessian File Reuse

By default the `task <theory> frequencies` directive will recompute the hessian. To reuse the previously computed hessian you need only specify `reuse` in the module input block. If you have stored the hessian in an alternate place you may redirect the `reuse` directive to that file by specifying the path to that file.

```
reuse /path_to_hessian_file
```

This will reuse your saved Hessian data but one caveat is that the geometry specification at the point where the hessian is computed must be the default “geometry” on the current run-time-data-base for the projection to work properly.

25.1.2 Redefining Masses of Elements

You may also modify the mass of a specific center or a group of centers via the input.

To modify the mass of a specific center you can simply use:

```
mass 3 4.00260324
```

which will set the mass of center 3 to 4.00260324 AMUs. The lexical index of centers is determined by the geometry object.

To modify all Hydrogen atoms in a molecule you may use the tag based mechanism:

```
mass hydrogen 2.014101779
```

The mass redefinitions always start with the default masses and change the masses in the order given in the input. Care must be taken to change the masses properly. For example, if you want all hydrogens to have the mass of Deuterium and the third hydrogen (which is the 6th atomic center) to have the mass of Tritium you must set the Deuterium masses first with the tag based mechanism and then set the 6th center’s mass to that of Tritium using the lexical center index mechanism.

The mass redefinitions are not fully persistent on the run-time-data-base. Each input block that redefines masses will invalidate the mass definitions of the previous input block. For example,

```
freq
  reuse
  mass hydrogen 2.014101779
end
task scf frequencies
freq
  reuse
  mass oxygen 17.9991603
end
task scf frequencies
```

will use the new mass for all hydrogens in the first frequency analysis. The mass of the oxygen atoms will be redefined in the second frequency analysis but the hydrogen atoms will use the default mass. To get a modified oxygen and hydrogen analysis you would have to use:

```
freq
```

```
reuse
mass hydrogen 2.014101779
end
task scf frequencies
freq
reuse
mass hydrogen 2.014101779
mass oxygen 17.9991603
end
task scf frequencies
```

25.1.3 Animation

The “VIB” module also can generate mode animation input files in the standard xyz file format for graphics packages like RasMol or XMol. There are scripts to automate this for RasMol in `$NWCHEM_TOP/contrib/rasmolmovie`. Each mode will have 20 xyz files generated that cycle from the equilibrium geometry to 5 steps in the positive direction of the mode vector, back to 5 steps in the negative direction of the mode vector, and finally back to the equilibrium geometry. By default these files are **not** generated. To activate this mechanism simply use the following input directive

```
animate
```

anywhere in the frequency/vib input block.

Controlling the Step Size Along the Mode Vector

By default, the step size used is 0.15 a.u. which will give reliable animations for most systems. This can be changed via the input directive

```
animate real <step_size>
```

where `<step_size>` is the real number that is the magnitude of each step along the eigenvector of each nuclear hessian mode in atomic units.

25.1.4 An Example Input Deck

This example input deck will optimize the geometry for the given basis set, compute the frequencies for H₂O, D₂O, HDO, and TDO.

```
start h2o
title Water
geometry units au autosym
  O      0.00000000    0.00000000    0.00000000
  H      0.00000000    1.93042809   -1.10715266
  H      0.00000000   -1.93042809   -1.10715266
end
basis noprint
  H library sto-3g
  O library sto-3g
end
scf; thresh 1e-6; end
driver; tight; end
task scf optimize

scf; thresh 1e-8; print none; end
task scf freq

freq
  reuse; mass H 2.014101779
end
task scf freq

freq
  reuse; mass 2 2.014101779
end
task scf freq

freq
  reuse; mass 2 2.014101779 ; mass 3 3.01604927
end
task scf freq
```

Chapter 26

DPLOT

```
DPLOT
...
END
```

This directive is used to obtain the plots of various types of electron densities (or orbitals) of the molecule. The electron density is calculated on a specified set of grid points using the molecular orbitals from SCF or DFT calculation. The output file is either in MSI Insight II contour format (default) or in the Gaussian Cube format. DPLOT is not executed until the “task dplot” directive is given. Different sub-directives are described below.

26.1 GAUSSIAN — Gaussian Cube format

```
GAUSSIAN
```

A outputfile is generate in Gaussian Cube format. You can visualize this file using gOpenMol (after converting the Gaussian Cube file with gcube2plt), Molden or Molekel.

26.2 TITLE — Title directive

```
TITLE <string Title default Unknown Title>
```

This sub-directive specifies a title line for the generated input to the *Insight* program or for the Gaussian cube file. Only one line is allowed.

26.3 LIMITXYZ — Plot limits

```
LIMITXYZ [units <string Units default angstroms>]
<real X_From> <real X_To> <integer No_Of_Spacings_X>
<real Y_From> <real Y_To> <integer No_Of_Spacings_Y>
<real Z_From> <real Z_To> <integer No_Of_Spacings_Z>
```

This sub-directive specifies the limits of the cell to be plotted. The grid is generated using `No_Of_Spacings + 1` points along each direction. The known names for `Units` are `angstroms`, `au` and `bohr`.

26.4 SPIN — Density to be plotted

```
SPIN <string Spin default total>
```

This sub-directive specifies, what kind of density is to be computed. The known names for `Spin` are `total`, `alpha`, `beta` and `spindens`, the last being computed as the difference between α and β electron densities.

26.5 OUTPUT — Filename

```
OUTPUT <string File_Name default dplot>
```

This sub-directive specifies the name of the generated input to the *Insight* program or the generated Gaussian cube file. The name `OUTPUT` is reserved for the standard NWChem output.

26.6 VECTORS — MO vector file name

```
VECTORS <string File_Name default movecs> [<string File_Name2>]
```

This sub-directive specifies the name of the molecular orbital file. If the second file is optionally given the density is computed as the difference between the corresponding electron densities. The vector files have to match.

26.7 WHERE — Density evaluation

```
WHERE <string Where default grid>
```

This sub-directive specifies where the density is to be computed. The known names for `Where` are `grid` (the calculation of the density is performed on the set of a grid points specified by the sub-directive `LimitXYZ` and the file specified by the sub-directive `Output` is generated), `nuclei` (the density is computed at the position of the nuclei and written to the NWChem output) and `g+n` (both).

26.8 ORBITAL — Orbital sub-space

```
ORBITALS [<string Option default density>]
<integer No_Of_Orbitals>
<integer Orb_No_1 Orb_No_2 ...>
```

This sub-directive specifies the subset of the orbital space for the calculation of the electron density. The density is computed using the occupation numbers from the orbital file modified according to the `Spin` directive. If the contours

of the orbitals are to be plotted `Option` should be set to `view`. Note, that in this case `No_Of_Orbitals` should be set to 1 and sub-directive `Where` is automatically set to `grid`. Also specification of two orbital files conflicts with the `view` option. α orbitals are always plotted unless `Spin` is set to `beta`.

26.9 Examples

Charge Density

Example of charge density plot (with Gaussian Cube output):

```
start n2
geometry
  n 0 0 0.53879155
  n 0 0 -0.53879155
end
basis; n library cc-pvdz;end
scf
vectors output n2.movecs
end
dplot
  TITLE HOMO
  vectors n2.movecs
  LimitXYZ
  -3.0 3.0 10
  -3.0 3.0 10
  -3.0 3.0 10
  spin total
  gaussian
  output lumo.cube
end
task scf
task dplot
```

Molecular Orbital

Example of orbital plot (with Insight II contour output):

```
start n2
geometry
  n 0 0 0.53879155
  n 0 0 -0.53879155
end
basis; n library cc-pvdz;end
scf
vectors output n2.movecs
end
dplot
  TITLE HOMO
```

```
vectors n2.movecs
  LimitXYZ
-3.0 3.0 10
-3.0 3.0 10
-3.0 3.0 10
  spin total
  orbitals view; 1; 7
  output homo.grd
end
task scf
task dplot
```

Chapter 27

Electron Transfer Calculations with ET

The NWChem electron transfer (ET) module calculates the electronic coupling energy (also called the electron transfer matrix element) between ET reactant and product states. The electronic coupling (V_{RP}), activation energy (ΔG^*), and nuclear reorganization energy (λ) are all components of the electron transfer rate defined by Marcus' theory, which also depends on the temperature (reference 1):

$$k_{ET} = \frac{2\pi}{\hbar} V_{RP}^2 \frac{1}{\sqrt{4\pi\lambda k_B T}} \exp\left(\frac{-\Delta G^*}{k_B T}\right) \quad (27.1)$$

The ET module utilizes the method of *Corresponding Orbital Transformation* to calculate V_{RP} . The only input required are the names of the files containing the open-shell (UHF or ODFT) MO vectors for the ET reactant and product states (R and P).

Although open-shell DFT orbitals can be used as input, the current implementation of the ET module uses a Hartree-Fock formalism (ref.3). Therefore, for consistency, UHF orbitals should be used for the calculation of V_{RP} , although this is not required.

The basis set used in the calculation of V_{RP} must be the same as the basis set used to calculate the MO vectors of R and P . The magnitude of V_{RP} depends on the amount of overlap between R and P , which is important to consider when choosing the basis set. Diffuse functions may be necessary to fill in the overlap, particularly when the ET distance is long.

The MO's of R and P must correspond to localized states. for instance, in the reaction $A^- B \rightarrow A B^-$ the transferring electron is localized on A in the reactant state and is localized on B in the product state. To verify the localization of the electron in the calculation of the vectors, carefully examine the Mulliken population analysis. In order to determine which orbitals are involved in the electron transfer, use the print keyword "mulliken ao" which prints the Mulliken population of each basis function.

An effective core potential (ECP) basis can be used to replace core electrons. However, there is one caveat: the orbitals involved in electron transfer must not be replaced with ECP's. Since the ET orbitals are valence orbitals, this is not usually a problem, but the user should use ECP's with care.

Suggested references are listed below. The first two references gives a good description of Marcus' two-state ET model, and the appendix of the third reference details the method used in the ET module.

1. R.A. Marcus, N. Sutin, *Biochimica Biophysica Acta* 35, 437, (1985).
2. J.R. Bolton, N. Mataga, and G. McLendon in "Electron Transfer in Inorganic, Organic and Biological Systems" (American Chemical Society, Washington, D.C., 1991)

3. A. Farazdel, M. Dupuis, E. Clementi, and A. Aviram, J. Am. Chem. Soc., 112, 4206 (1990).

27.1 VECTORS — input of MO vectors for ET reactant and product states

```
VECTORS [reactants] <string reactants_filename>
VECTORS [products ] <string products_filename>
```

In the VECTORS directive the user specifies the source of the molecular orbital vectors for the ET reactant and product states. This is required input, as no default filename will be set by the program. In fact, this is the only required input in the ET module, although there are other optional keywords described below.

27.2 FOCK/NOFOCK — method for calculating the two-electron contribution to V_{RP}

```
<string (FOCK|NOFOCK) default FOCK>
```

This directive enables/disables the use of the NWChem's Fock matrix routine in the calculation of the two-electron portion of the ET Hamiltonian. Since the Fock matrix routine has been optimized for speed, accuracy and parallel performance, it is the most efficient choice.

Alternatively, the user can calculate the two-electron contribution to the ET Hamiltonian with another subroutine which may be more accurate for systems with a small number of basis functions, although it is slower.

27.3 TOL2E — integral screening threshold

```
TOL2E <real tol2e default max(10e-12,min(10e-7, S(RP)*10e-7 )>
```

The variable `tol2e` is used in determining the integral screening threshold for the evaluation of the two-electron contribution to the Hamiltonian between the electron transfer reactant and product states. As a default, `tol2e` is set depending on the magnitude of the overlap between the ET reactant and product states (S_{RP}), and is not less than $1.0d-12$ or greater than $1.0d-7$.

The input to specify the threshold explicitly within the ET directive is, for example:

```
tol2e 1e-9
```

27.4 Example

The following example is for a simple electron transfer reaction, $He \rightarrow He^+$. The ET calculation is easy to execute, but it is crucial that ET reactant and product wavefunctions reflect *localized states*. This can be accomplished using either a fragment guess (shown in the example, see 10.5.1), or a charged atomic density guess (see 10.5.2). For self-exchange ET reactions such as this one, you can use the REORDER keyword to move the electron from the first helium to the second (see 10.5).

Example input :

```
#ET reactants:
charge 1
scf
  doublet; uhf; vectors input fragment HeP.mo He.mo output HeA.mo
# HeP.mo are the vectors for He(+),
# He.mo are the vectors for neutral He.
end
task scf

#ET products:
charge 1
scf
  doublet; uhf; vectors input HeA.mo reorder 2 1 output HeB.mo
end
task scf

et
  vectors reactants HeA.mo
  vectors products HeB.mo
end
task scf et
```

Here is what the output looks like for this example:

```

                                Electron Transfer Calculation
                                -----

MO vectors for reactants: HeA.mo
MO vectors for products : HeB.mo

Electronic energy of reactants      H(RR)      -5.3402392824
Electronic energy of products      H(PP)      -5.3402392824

Reactants/Products overlap        S(RP)      -0.0006033839

Reactants/Products interaction energy:
-----
One-electron contribution          H1(RP)      0.0040314092

Beginning calculation of 2e contribution
Two-electron integral screening (tol2e) : 6.03E-11

Two-electron contribution          H2(RP)     -0.0007837138
Total interaction energy            H(RP)      0.0032476955

Electron Transfer Coupling Energy |V(RP)|      0.0000254810
                                         5.592 cm-1
```

0.000693 eV
0.016 kcal/mol

The overlap between the ET reactant and product states (S_{RP}) is small, so the magnitude of the coupling between the states is also small. If the fragment guess or charged atomic density guess were not used, the Mulliken spin population would be 0.5 on both He atoms, the overlap between the ET reactant and product states would be 100 % and an infinite V_{RP} would result.

Chapter 28

Properties

```
PROPERTY
  [property name]
  [CENTER ((com || coc || origin || arb <real x y z>) default coc)]
  [VECTORS ...]
END
```

Calculation of properties is accomplished with `TASK PROPERTY` after the completion of an energy (or MP2 gradient) calculation. The following properties can be computed for all wavefunctions that produce orbitals, including Hartree-Fock (closed-shell RHF, open-shell ROHF, and open-shell UHF), DFT (closed-shell and open-shell spin unrestricted), MCSCF (complete active space), and MP2 (closed-shell RHF and open-shell UHF).

- natural bond analysis
- dipole moment
- quadrupole moment
- octupole moment
- Mulliken population analysis and bond order analysis
- electrostatic potential (diamagnetic shielding) at nuclei
- electric field at nuclei
- electric field gradient at nuclei
- electron and spin density at nuclei
- NMR chemical shifts (GIAO method) only for closed-shell RHF

The default molecular orbital file `$file_prefix$.movecs` is used unless a `vectors` directive (Section 10.5) is provided. It is therefore only necessary to include a `vectors` directive if the MO vectors to be analyzed are not coming from the default file, e.g., if they have been previously redirected, or if MP2 natural orbitals (file extension `".mp2nos"`) are being analyzed. The MP2 natural orbitals **MUST** be used if the user wants MP2 properties.

28.1 Subdirectives

Note that presenting any property input causes all previous property input to be “forgotten”, unlike other NWChem modules.

Each property can be requested by means of a subdirective among the subdirectives provided :

- NBOFILE
- DIPOLE
- QUADRUPOLE
- OCTUPOLE
- MULLIKEN
- ESP
- EFIELD
- EFIELDGRAD
- ELECTRONDENSITY
- GIAO
- ALL

The “ALL” keyword generates all currently available properties.

The request NBOFILE does not execute the Natural Bond Analysis code, but simply creates an input file to be used as input to the stand-alone NBO code. To execute the NBO analysis directly, see Section 37.1. All other properties are calculated upon request.

An additional subdirective is provided to specify the origin of the molecular orbitals used in the calculation of the molecular properties. This is the ‘VECTORS’ subdirective, also used in the SCF and DFT tasks. For a full description of this subdirective the user is referred to the description found in the SCF description. By default, the input file used for the calculation of the properties has the `.movecs` name extension.

The user also has the option to choose the center of expansion for the dipole, quadrupole, and octupole calculations.

```
[CENTER ((com || coc || origin || arb <real x y z>) default coc)]
```

`com` is the center of mass, `coc` is the center of charge, `origin` is (0.0, 0.0, 0.0) and `arb` is any arbitrary point which must be accompanied by the coordinates to be used. Currently the x, y, and z coordinates must be given in the same units as UNITS in GEOMETRY (See Section 6.1).

28.1.1 Nbofile

Following the successful completion of an electronic structure calculation, a Natural Bond Orbital (NBO) analysis may be carried out in the following way. On restart specify the TASK as PROPERTY and supply the sub-directive NBOFILE to the PROPERTY directive. NWChem will query the rtdb and construct an ASCII file, `<file_prefix>.gen`, that may be used as input to the stand alone version of the NBO program, `gennbo`. `<file_prefix>` is equal to string following the RESTART directive. The input deck may be edited to provide additional options to the NBO calculation, (see the NBO user’s manual for details.) The other option in to directly run the NBO analysis (See Section 37.1 for more information).

Chapter 29

Electrostatic potentials

The NWChem Electrostatic Potential (ESP) module derives partial atomic charges that fit the quantum mechanical electrostatic potential on selected grid points.

The ESP module is specified by the NWChem task directive

```
task esp
```

The input for the module is taken from the ESP input block

```
ESP  
...  
END
```

29.1 Grid specification

The grid points for which the quantum mechanical electrostatic potential is evaluated and used in the fitting procedure of the partial atomic charges all lie outside the van der Waals radius of the atoms and within a cutoff distance from the atomic centers. The following input parameters determine the selection of grid points.

- If a grid file is found, the grid will be read from that file. If no grid file is found, or the keyword

```
recalculate
```

is given, the grid and the electrostatic potential is recalculated.

- The extent of the grid is determined by

```
range <real rcut>
```

where *rcut* is the maximum distance in *nm* between a grid point and any of the atomic centers. When omitted, a default value for *rcut* of 0.3 *nm* is used.

- The grid spacing is specified by

```
spacing <real spac>
```

where `spac` is the grid spacing in *nm* for the regularly spaced grid points. If not specified, a default spacing of 0.05 *nm* is used.

- The van der Waals radius of an element can be specified by

```
radius <integer iatnum> <real atrad>
```

where `iatnum` is the atomic number for which a van der Waals radius of `atrad` in *nm* will be used in the grid point determination. Default values will be used for atoms not specified.

- The probe radius in *nm* determining the envelope around the molecule is specified by

```
probe <real probe default 0.07>
```

- The distance between atomic center and probe center can be multiplied by a constant factor specified by

```
factor <real factor default 1.0>
```

All grid points are discarded that lie within a distance `factor*(radius(i)+probe)` from any atom *i*.

- Schwarz screening is applied using

```
screen [<real scrtol default 1.0D-5>]
```

29.2 Constraints

Additional constraints to the partial atomic charges can be imposed during the fitting procedure.

- The net charge of a subset of atoms can be constrained using

```
constrain <real charge> {<integer iatom>}
```

where `charge` is the net charge of the set of atoms `{iatom}`. A negative atom number `iatom` can be used to specify that the partial charge of that atom is subtracted in the sum for the set.

- The net charge of a sequence of atoms can be constrained using

```
constrain <real charge> <integer iatom> through <integer jatom>
```

where `charge` is the net charge of the set of atoms `{[iatom:jatom]}`.

- A group of atoms can be constrained to have the same charge with

```
constrain equal {<integer iatom>}
```

- The individual charge of a group of atoms can be constrained to be equal to those of a second group of atoms with

```
constrain group <integer iatom> <integer jatom> to <integer katom> <integer latom>
```

resulting in the same charge for atoms `iatom` and `katom`, for atoms `iatom+1` and `katom+1`, ... for atoms `jatom` and `latom`.

- A special constraint

```
constrain xhn <integer iatom> {<integer jatom>}
```

can be used to constrain the set $\{i\text{atom}, \{j\text{atom}\}\}$ to zero charge, and constrain all atoms in $\{j\text{atom}\}$ to have the same charge. This can be used, for example, to restrain a methyl group to zero charge, and have all hydrogen carrying identical charges.

29.3 Restraints

Restraints can be applied to each partial charge using the RESP charge fitting procedure.

- The directive for charge restraining is

```
restrain [hfree] (harmonic [<real scale>] | \  
hyperbolic [<real scale> [<real tight>]] \  
[maxiter <integer maxit>] [tolerance <real toler>])
```

where *hfree* can be specified to exclude hydrogen atoms from the restraining procedure. Variable *scale* is the strength of the restraint potential, with a default of $0.005au$ for the harmonic restraint and a default value of $0.001au$ for the hyperbolic restraint. For the hyperbolic restraints the tightness *tight* can be specified to change the default value of $0.1e$. The iteration count that needs to be carried out for the hyperbolic restraint is determined by the maximum number of allowed iterations *maxiter*, with a default value of 25, and the tolerance in the convergence of the partial charges *toler*, with a default of $0.001e$.

Chapter 30

Prepare

The **prepare** module is used to set up the necessary files for a molecular dynamics simulation with **NWChem**. User supplied coordinates can be used to generate topology and restart files. The topology file contains all static information about a molecular system, such as lists of atoms, bonded interactions and force field parameters. The restart file contains all dynamic information about a molecular system, such as coordinates, velocities and properties.

Without any input, the prepare module checks the existence of a topology and restart file for the molecular systems. If these files exist, the module returns to the main task level without action. The module will generate these files when they do not exist. Without any input to the module, the generated system will be for a non-solvated isolated solute system.

To update existing files, including solvation, the module requires input directives read from an input deck,

```
prepare
...
end
```

The prepare module performs three sub-tasks:

* **sequence generation**

This sub-task analyzes the supplied coordinates from a PDB-formatted file or from the input geometry, and generates a sequence file, containing the description of the system in terms of basic building blocks found as fragment or segment files in the database directories for the force field used. If these files do not exist, they are generated based on the supplied coordinates. This process consists of generating a fragment file with the list of atoms with their force field dependent atom types, partial atomic charges calculated from a Hartree Fock calculation for the fragment, followed by a restrained electrostatic potential fit, and a connectivity list. From the information on this fragment file the lists of all bonded interactions are generated, and the complete lists are written to a segment file.

* **topology generation**

Based on the generated or user-supplied sequence file and the force field specific segment database files, this sub-task compiles the lists of atoms, bonded interactions, excluded pairs, and substitutes the force field parameters. Special commands may be given to specify interaction parameters that will be changing in a free energy evaluation.

* **restart generation**

Using the user supplied coordinates and the topology file for the chemical system, this sub-task generates a

restart file for the system with coordinates, velocities and other dynamic information. This step may include solvation of the chemical system and specifying periodic boundary conditions.

Files involved in the preparation phase exist in the following hierarchy:

* **standards**

The standard database files contain the original force field information. These files are to reside in a directory that is specified in the file \$HOME/.nwchemrc. There will be such a directory for each supported force field. These directories contain fragment files (with extension frg), segment files (with extension sgm) and a parameter file (with the name of the force field and with extension par).

* **extensions**

These database files contain generally accepted extensions to the original force field and are to reside in a separate directory that is specified in the file \$HOME/.nwchemrc. There will be such a directory for each supported force field. These directories contain fragment files (with extension frg), segment files (with extension sgm) and a parameter file (with the name of the force field and with extension par).

* **contributed**

These database files contain contributed definitions, also required for the quality assurance tests and are to reside in a separate directory that is specified in the file \$HOME/.nwchemrc. There will be such a directory for each supported force field. These directories contain fragment files (with extension frg), segment files (with extension sgm) and a parameter file (with the name of the force field and with extension par).

* **user preferences**

These database files contain user preferred extensions to the original force field and are to reside in a separate directory that is specified in the file \$HOME/.nwchemrc. Separate directories of this type should be defined for each supported force field. This directory may contain fragment files (with extension frg), segment files (with extension sgm) and a parameter file (with the name of the force field and with extension par).

* **temporary files**

Temporary database files contain user preferred extensions to the original force field and are to reside in a separate directory that is specified in the file \$HOME/.nwchemrc. There may be such a directory for each supported force field. This directory may contain fragment files (with extension frg), segment files (with extension sgm) and a parameter file (with the name of the force field and with extension par).

* **current files**

Database files that contain user preferred extensions to the original force field and are to reside in a separate directory that is specified in the file \$HOME/.nwchemrc. Typically this will be the current working directory, although it may be defined as a specific directory. This directory may contain fragment files (with extension frg), segment files (with extension sgm) and a parameter file (with the name of the force field and with extension par). If not specified, files will be taken from the current directory.

Data is taken from the database files searched in the above order. If data is specified more than once, the last found values are used. For example, if some standard segment is redefined in a temporary file, the latter one will be used. This allows the user to redefine standards or extensions without having to modify those database files, which may reside in a generally available, non-modifiable directory. If a filename is specified rather than a directory, the filename indicates the parameter file definition. All other files (frg and sgm files) will be taken from the specified directory.

The most common problems with the **prepare** module are

The format of the pdb file does not conform to the pdb standard. In particular, atom names need to correspond with definitions in the fragment and segment database files, and should adhere to IUPAC recommendations as adopted by the pdb standard. If this problem occurs, the pdb file will need to be corrected.

Non-standard segments may contain atoms that could not be atom typed with the existing typing rules in the force field parameter files. When this happens, additional typing rules can be included in the parameter file, or the fragment file may be manually typed.

Parameters for atom types or bonded interactions do not exist in the force field. When this happens, additional parameters may be defined in the parameter files, or the segment file may be edited to include explicit parameters.

30.1 Default database directories

The file \$HOME/.nwchemrc may contain the following entries that determine which files are used by the prepare module.

```
ffield <string ffname>
```

This entry specifies the default force field. Database files supplied with **NWChem** currently support values for `ffname` of **amber**, referring to AMBER95, and **charmm**, referring to the academic CHARMM22 force field.

```
<string ffname>_(s | x | q | u | t | c) <string ffdir>[<string parfile>]
```

Entries of this type specify the directory `ffdir` in which force field database files can be found. Optionally the parameterfile in this directory may be specified as `parfile`. The prepare module will only use files in directories specified here. One exception is that files in the current work directory will be used if no directory with current files is specified.

```
<string solvnam> <string solvfil>
```

This entry may be used to identify a pure solvent restart file `solvfil` by a name `solvnam`

An example file \$HOME/.nwchemrc is:

```
ffield amber
amber_s /msrc/proj/nwchem/share/amber/amber_s/amber99.par
amber_x /msrc/proj/nwchem/share/amber/amber_x/
amber_u /usr/people/username/data/amber/amber_u/
spce /msrc/proj/nwchem/share/solvents/spce.rst
charmm_s /msrc/proj/nwchem/share/charmm/charmm_s/
charmm_x /msrc/proj/nwchem/share/charmm/charmm_x/
```

30.2 System name and coordinate source

```
system <string sys_calc>
```

The system name can be explicitly specified for the **prepare** module. If not specified, the system name will be taken from a specification in a previous **md** input block, or derived from the run time database name.

```
source ( pdb | rtdb )
```

The source of the coordinates can be explicitly specified to be from a PDB formatted file `sys.pdb`, or from a geometry object in the run time database. If not specified, a `pdb` file will be used when it exists in the current directory or the `rtdb` geometry otherwise.

```
model <integer modpdb default 0>
```

If a PDB formatted source file contains different MODELS, the `model` keyword can be used to specify which MODEL will be used to generate the topology and restart file. If not specified, the first MODEL found on the PDB file will be read.

```
altloc <character locpdb default ' '>
```

The `altloc` keyword may be used to specify the use of alternate location coordinates on a PDB file.

```
chain <character chnpdb default ' '>
```

The `chain` keyword may be used to specify the chain identifier for coordinates on a PDB file.

```
histidine ( hid | hie | hip )
```

specifies the default protonation state of histidine.

```
sscyx
```

Keyword `sscyx` may be used to rename cysteine residues that form sulphur bridges to `CYX`.

```
hbuild
```

Keyword `hbuild` may be used to add hydrogen atoms to the unknown segments of the structure found on the `pdb` file.

30.3 Sequence file generation

If no existing sequence file is present in the current directory, or if the `new_seq` keyword was specified in the **prepare** input deck, a new sequence file is generated from information from the `pdb` file, and the following input directives.

```
maxscf <integer maxscf default 20>
```

Variable `maxscf` specifies the maximum number of atoms in a segment for which partial atomic charges will be determined from an SCF calculation followed by RESP charge fitting. For larger segments a crude partial charge guesstimation will be done.

```
qscale <real qscale default 1.0>
```

Variable `qscale` specifies the factor with which SCF/RESP determined charges will be multiplied.

```
modify sequence { <integer sgmnum>:<string sgmnam> }
```

This command specifies that segment **sgmnam** should be used for segment with number *sgmnum*. This command can be used to specify a particular protonation state. For example, the following command specifies that residue 114 is a histidine protonated at the N_ε site and residue 202 is a histidine protonated at the N_δ site:

```
modify sequence 114:HIE 202:HID
```

Links between atoms can be enforced with

```
link <string atomname> <string atomname>
```

For example, to link atom SG in segment 20 with atom FE in segment 55, use:

```
link 20:_SG 55:FE
```

The format of the sequence file is given in Table 34.8. In addition to the list of segments this file also includes links between non-standard segments or other non-standard links. These links are generated based on distances found between atoms on the pdb file. When atoms are involved in such non-standard links that have not been identified in the fragment of segment files as a non-chain link atom, the prepare module will ignore these links and report them as skipped. If one or more of these links are required, the user has to include them with explicit link directives in the sequence file, making them forced links. Alternatively, these links can be made forced-links by changing `link` into `LINK` in the sequence file.

```
fraction { <integer imol> }
```

Directive `fraction` can be used to separate solute molecules into fractions for which energies will be separately reported during molecular dynamics simulations. The listed molecules will be the last molecule in a fraction. Up to 10 molecules may be specified in this directive.

```
counter <integer num> <string ion>
```

Directive `counter` adds `num` counter ions of type `ion` to the sequence file. Up to 10 `counter` directives may appear in the input block.

```
counter <real factor>
```

This directive scales the counter ion charge by the specified factor in the determination of counter ions positions.

30.4 Topology file generation

```
new_top [ new_seq ]
```

Keyword `new_top` is used to force the generation of a new topology file. An existing topology file for the system in the current directory will be overwritten. If keyword `new_seq` is also specified, an existing sequence file will also be overwritten with a newly generated file.

amber | charmm

The prepare module generates force field specific fragment, segment and topology files. The force field may be explicitly specified in the prepare input block by specifying its name. Currently **AMBER** and **CHARMM** are the supported force fields. A default force field may be specified in the file \$HOME/.nwchemrc.

```
standard <string dir_s>[<string par_s>]
extensions <string dir_x>[<string par_x>]
contributed <string dir_q>[<string par_q>]
user <string dir_u>[<string par_u>]
temporary <string dir_t>[<string par_t>]
current <string dir_c>[<string par_c>]
```

The user can explicitly specify the directories where force field specific databases can be found. These include force field standards, extensions, quality assurance tests, user preferences, temporary, and current database files. Defaults for the directories where database files reside may be specified in the file \$HOME/.nwchemrc for each of the supported force fields. Fragment, segment and sequence files generated by the **prepare** module are written in the temporary directory. When not specified, the current directory will be used. Topology and restart files are always created in the current directory.

The following directives control the modifications of a topology file. These directives are executed in the order in which they appear in the **prepare** input deck. The topology modifying commands are not stored on the run-time database and are, therefore, not persistent.

```
modify atom <string atomname> [set <integer mset> | initial | final] \
( type <string atomtyp> | charge <real atomcharge> | \
  polar <real atompolar> | dummy | self | quantum | quantum_high )
```

These modify commands change the atom type, partial atomic charge, atomic polarizability, specify a dummy, self-interaction and quantum atom, respectively. If mset is specified, the modification will only apply to the specified set, which has to be 1, 2 or 3. If not specified, the modification will be applied to all three sets. The quantum region in QM/MM simulations is defined by specifying atoms with the quantum or quantum_high label. For atoms defined quantum_high basis sets labeled X_H will be used. The atomnam should be specified as <integer isgm>:<string name>, where isgm is the segment number, and name is the atom name. A leading blank in an atom name should be substituted with an underscore. The modify commands may be combined. For example, the following directive changes for the specified atom the charge and atom type in set 2 and specifies the atom to be a dummy in set 3.

```
modify atom 12:_C1 set 2 charge 0.12 type CA set 3 dummy
```

With the following directives modifications can be made for entire segments.

```
modify segment <integer isgm> [set <integer mset> | initial | final] \
( dummy | self | uncharged | quantum | quantum_high )
```

Modifications to bonded interaction parameters can be made with the following modify commands.

```
modify ( bond <string atomtyp> <string atomtyp> | \
  angle <string atomtyp> <string atomtyp> <string atomtyp> | \
```

```
torsion <string atomtyp> <string atomtyp> <string atomtyp> \
<string atomtyp> [ multiplicity <integer multip> ] | \
plane <string atomtyp> <string atomtyp> <string atomtyp> \
<string atomtyp> ) [set <integer mset> | initial | final] \
<real value> <real forcon>
```

where *atomtyp* and *mset* are defined as above, *multip* is the torsion multiplicity for which the modification is to be applied, *value* is the reference bond, angle, torsion angle of out-of-plane angle value respectively, and *forcon* is the force constant for bond, angle, torsion angle of out-of-plane angle. When *multip* or *mset* are not defined the modification will be applied to all multiplicities and sets, respectively, for the identified bonded interaction.

After modifying atoms to quantum atoms the bonded interactions in which only quantum atoms are involved are removed from the bonded lists using

```
update lists
```

Error messages resulting from parameters not being defined for bonded interaction in which only quantum atoms are involved are ignored using

```
ignore
```

30.5 Appending to an existing topology file

```
noe <string atom1> <string atom3> \
<real dist1> <real dist2> <real dist3> <real forc1> <real forc2>
```

This directive specifies a distance restraint potential between atoms *atom1* and *atom2*, with a harmonic function with force constant *forc1* between *dist1* and *dist2*, and a harmonic function with force constant *forc2* between *dist2* and *dist3*. For distances shorter than *dist1* or larger than *dist3*, a constant force is applied such that force and energy are continuous at *dist1* and *dist3*, respectively. Distances are given in nm, force constants in $\text{kJ mol}^{-1} \text{nm}^{-2}$.

```
select <integer isel> { <string atoms> }
```

Directive *select* specifies a group of atoms used in the definition of potential of mean force potentials.

The selected atoms are specified by the string *atoms* which takes the form

```
[{isgm [ - jsgm ] [,]} [:] [{aname[,]}]
```

For example, all carbon and oxygen atoms in segments 3 and 6 through 12 are selected for group 1 by

```
3,6-12:_C????,_O????
```

```
pmf [all] [bias] zalign <integer isel> <real forcon1> <real forcon2>
pmf [combine] [bias] xyplane <integer isel> <real forcon1> <real forcon2>
pmf [constraint] [bias] (distance | zdistance) <integer isel> <integer jsel> \
<real dist1> <real dist2> <real forcon1> <real forcon2>
```

```

pmf [bias] angle <integer isel> <integer jsel> <integer ksel> \
    <real angle1> <real angle2> <real forcon1> <real forcon2>
pmf [bias] torsion <integer isel> <integer jsel> <integer ksel> <integer lsel> \
    <real angle1> <real angle2> <real forcon1> <real forcon2>
pmf [bias] basepair <integer isel> <integer jsel> \
    <real dist1> <real dist2> <real forcon1> <real forcon2>
pmf [bias] (zaxis | zaxis-cog) <integer isel> <integer jsel> <integer ksel> \
    <real dist1> <real dist2> <real forcon1> <real forcon2>

```

Directive `pmf` specifies a potential of mean force potential in terms of the specified atom selection. Option `zalign` specifies the atoms in the selection to be restrained to a line parallel to the z-axis. Option `xyplane` specifies the atoms in the selection to be restrained to a plane perpendicular to the z-axis. Options `distance`, `angle` and `torsion`, are defined in terms of the center of geometry of the specified atom selections. Keyword `basepair` is used to specify a harmonic potential between residues `isel` and `jsel`. Keywords `zaxis` and `zaxis-cog` can be used to pull atoms toward the z-axis. Option `all` may be specified to apply an equivalent pmf to each of the equivalent solute molecules in the system. Option `combine` may be specified to apply the specified pmf to the atoms in all of the equivalent solute molecules. Option `constraint` may be specified to a distance pmf to treat the distance as a constraint. Option `bias` may be specified to indicate that this function should be treated as a biasing potential. Appropriate corrections to free energy results will be evaluated.

30.6 Generating a restart file

`new_rst`

Keyword `new_rst` will cause an existing restart file to be overwritten with a new file.

The following directives control the manipulation of restart files, and are executed in the order in which they appear in the **prepare** input deck.

```

solvent name <string*3 slvnam default ``HOH``> \
    model <string slvmdl default ``spce``>

```

The solvent keyword can be used to specify the three letter solvent name as expected on the PDB formatted file, and the name of the solvent model for which solvent coordinates will be used.

```

solvate [ < real rshell default 1.2 > ] \
    ( [ cube [ <real edge> ] ] | \
    [ box [ <real xedge> [ <real xedge> [ <real xedge> ]]] ] | \
    [ sphere <real radius> ] |
    [ troct <real edge> ] )

```

Solvation can be specified to be in a cubic box with specified edge, rectangular box with specified edges, or in a sphere with specified radius. Solvation in a cube or rectangular box will automatically also set periodic boundary conditions. Solvation in a sphere will only allow simulations without periodic boundary conditions. The size of the cubic and rectangular boxes will be expanded by a length specified by the `expand` variable. If no shape is specified, solvation will be done for a cubic box with an edge that leaves `rshell` nm between any solute atom and a periodic image of any solute atom after the solute has been centered. An explicit `write` is not needed to write the restart file. The `solvate` will write out a file `sys_calc.rst`. If not specified, the dimension of the solvation cell will be as large as to have at least a distance of `rshell` nm between any solute atom and the edge of the cell. The experimental `troct` directive generates a truncated octrahedral box.

```
touch <real touch default 0.23>
```

The variable `touch` specifies the minimum distance between a solvent and solute atom for which a solvent molecule will be accepted for solvation.

```
envelope <real xpndw default 0.0>
```

sets the expand values to be used in `solvate` operations.

```
expand <real xpndw default 0.1>
```

The variable `xpndw` specifies the size in nm with which the simulation volume will be increased after solvation.

```
read [rst | rst_old | pdb] <string filename>
write [rst | [solute [<integer nsolvent>]] ( [large] pdb | xyz)] <string filename>
```

These directives read and write the file `filename` in the specified format. The `solute` option instructs to write out the coordinates for solute and all, or if specified the first `nsolvent`, crystal solvent molecules only. If no format is specified, it will be derived from the extension of the filename. Recognized extensions are `rst`, `rst_old` (read only), `pdb`, `xyz` (write only) and `pov` (write only). Reading and then writing the same restart file will cause the sub-block size information to be lost. If this information needs to be retained a shell copy command needs to be used. The `large` keyword allows PDB files to be written with more than 9999 residues. Since the PDB file will not conform to the PDB convention, this option should only be used if required. NWChem will be able to read the resulting PDB file, but other codes may not.

```
scale <real scale default -1.0>
```

This directive scales the volume and coordinates written to `povray` files. A negative value of `scale` (default) scales the coordinates to lie in [-1:1].

```
cpk [<real cpk default 1.0>]
```

This directive causes `povray` files to contain `cpk` model output. The optional value is used to scale the atomic radii. A negative value of `cpk` resets the rendering to stick.

```
center | centerx | centery | centerz
```

These directives center the solute center of geometry at the origin, in the y-z plane, in the x-z plane or in the x-y plane, respectively.

```
orient
```

This directive orients the solute principal axes.

```
translate [atom | segment | molecule] \
<integer itrans> <integer itrans> <real xtrans(3)>
```

This directive translates solute atoms in the indicated range by `xtran`, without checking for bad contacts in the resulting structure.

```
rotate [atom | segment | molecule] \  
  <integer itrans> <integer itrans> <real angle> <real xrot(3)>
```

This directive rotates solute atoms in the indicated range by angle around the vector given by `xrot`, without checking for bad contacts in the resulting structure.

```
remove solvent [inside | outside] [x <real xmin> <real xmax>] \  
 [y <real ymin> <real ymax>] [z <real zmin> <real zmax>]
```

This directive removes solvent molecules inside or outside the specified coordinate range.

`periodic`

This directive enables periodic boundary conditions.

`vacuo`

This directive disables periodic boundary conditions.

```
grid <integer mgrid default 24> <real rgrid default 0.2>
```

This directive specifies the grid size of trial counter-ion positions and minimum distance between an atom in the system and a counter-ion.

`crop`

prints minimum and maximum solute coordinates.

`boxsize`

specifies to redetermine the box size.

`cube`

specifies to redetermine the smallest cubic box size.

```
box <real xsize> <real ysize> <real zsize>
```

The `box` directive resets the box size.

```
align <string atomi> <string atomj> <string atomk>
```

The `align` directive orients the system such that `atomi` and `atomj` are on the z-axis, and `atomk` in the x=y plane.

```
repeat [randomx | randomy | randomz] [chains | molecules | fractions ] \
  <integer nx> <integer ny> <integer nz> [<real dist>] [<real zdist>]
```

The `repeat` directive causes a subsequent `write pdb` directive to write out multiple copies of the system, with `nx` copies in the x, `ny` copies in the y, and `nz` copies in the z-direction, with a minimum distance of `dist` between any pair of atoms from different copies. If `nz` is -2, an inverted copy is placed in the z direction, with a separation of `zdist` nm. If `dist` is negative, the box dimensions will be used. For systems with solvent, this directive should be used with a negative `dist`. Optional keywords `chains`, `molecules` and `fractions` specify to write each repeating solute unit as a chain, to repeat each solute molecule, or each solute fraction separately. Optional keywords `randomx`, `randomy`, and `randomz` can be used to apply random rotations for each repeat unit around a vector through the center of geometry of the solute in the x, y or z direction.

```
skip <integer ix> <integer iy> <integer iz>
```

The `skip` directive can be used to skip single repeat unit from the `repeat` directive. Up to 100 `skip` directives may be specified, and will only apply to the previously specified `repeat` directive.

```
(collapsexy | collapsez) [ <integer nmoves>]
```

specifies to move all solute molecules toward the z-axis or x=y-plane, respectively, to within a distance of `touch` nm between any pair of atoms from different solute molecules. Parameter `nmoves` specifies the number of collapse moves that will be made. Monatomic ions will move with the nearest multi-atom molecule.

```
collapse_group <integer imol> <integer jmol>
```

specifies that molecule `jmol` will move together with molecule `imol` in collapse operations.

```
merge <real xtran(3)> <string pdbfile>
```

specifies to merge the coordinates found on the specified `pdb` file into the current structure after translation by `xtran(3)`.

Chapter 31

Molecular dynamics

31.1 Introduction

31.1.1 Spatial decomposition

The molecular dynamics module of **NWChem** uses a distribution of data based on a spacial decomposition of the molecular system, offering an efficient parallel implementation in terms of both memory requirements and communication costs, especially for simulations of large molecular systems.

Inter-processor communication using the global array tools and the design of a data structure allowing distribution based on spacial decomposition are the key elements in taking advantage of the distribution of memory requirements and computational work with minimal communication.

In the spacial decomposition approach, the physical simulation volume is divided into rectangular cells, each of which is assigned to a processor. Depending on the conditions of the calculation and the number of available processors, each processor contains one or more of these spacially grouped cells. The most important aspects of this decomposition are the dependence of the cell sizes and communication cost on the number of processors and the shape of the cells, the frequent reassignment of atoms to cells leading to a fluctuating number of atoms per cell, and the locality of communication which is the main reason for the efficiency of this approach for very large molecular systems.

To improve efficiency, molecular systems are broken up into separately treated solvent and solute parts. Solvent molecules are assigned to the domains according to their center of geometry and are always owned by a one node. This avoids solvent–solvent bonded interactions crossing node boundaries. Solute molecules are broken up into segments, with each segment assigned to a processor based on its center of geometry. This limits the number of solute bonded interactions that cross node boundaries. The processor to which a particular cell is assigned is responsible for the calculation of all interactions between atoms within that cell. For the calculation of forces and energies in which atoms in cells assigned to different processors are involved, data are exchanged between processors. The number of neighboring cells is determined by the size and shape of the cells and the range of interaction. The data exchange that takes place every simulation time step represents the main communication requirements. Consequently, one of the main efforts is to design algorithms and data structures to minimize the cost of this communication. However, for very large molecular systems, memory requirements also need to be taken into account.

To compromise between these requirements exchange of data is performed in successive point to point communications rather than using the shift algorithm which reduces the number of communication calls for the same amount of communicated data.

For inhomogeneous systems, the computational load of evaluating atomic interactions will generally differ between cell pairs. This will lead to load imbalance between processors. Two algorithms have been implemented that allow for dynamically balancing the workload of each processor. One method is the dynamic resizing of cells such that cells gradually become smaller on the busiest node, thereby reducing the computational load of that node. Disadvantages of this method are that the efficiency depends on the solute distribution in the simulation volume and the redistribution of work depends on the number of nodes which could lead to results that depend on the number of nodes used. The second method is based on the dynamic redistribution of intra-node cell-cell interactions. This method represents a more coarse load balancing scheme, but does not have the disadvantages of the cell resizing algorithm. For most molecular systems the cell pair redistribution is the more efficient and preferred method.

The description of a molecular system consists of static and dynamic information. The static information does not change during a simulation and includes items such as connectivity, excluded and third neighbor lists, equilibrium values and force constants for all bonded and non-bonded interactions. The static information is called the topology of the molecular system, and is kept on a separate topology file. The dynamic information includes coordinates and velocities for all atoms in the molecular system, and is kept in a so-called restart file.

31.1.2 Topology

The static information about a molecular system that is needed for a molecular simulation is provided to the simulation module in a topology file. Items in this file include, among many other things, a list of atoms, their non-bonded parameters for van der Waals and electrostatic interactions, and the complete connectivity in terms of bonds, angles and dihedrals.

In molecular systems, a distinction is made between *solvent* and *solute*, which are treated separately. A solvent molecule is defined only once in the topology file, even though many solvent molecules usually are included in the actual molecular system. In the current implementation only one solvent can be defined. Everything that is not solvent in the molecular system is solute. Each solute atom in the system must be explicitly defined in the topology.

Molecules are defined in terms of one or more *segments*. Typically, repetitive parts of a molecule are each defined as a single segment, such as the amino acid residues in a protein. Segments can be quite complicated to define and are, therefore, collected in a set of database files. The definition of a molecular system in terms of segments is a *sequence*.

Topology files are created using the **prepare** module.

31.1.3 Files

File names used have the form `$system$_$calc$. ext`, with exception of the topology file (Section 31.1.2), which is named `$system$.top`. Anything that refers to the definition of the chemical system can be used for `$system$`, as long as no periods or underlines are used. The identifier `$calc$` can be anything that refers to the type of calculation to be performed for the system with the topology defined. This file naming convention allows for the creation of a single topology file `$system$.top` that can be used for a number of different calculations, each identified with a different `$calc$`. For example, if `crown.top` is the name of the topology file for a crown ether, `crown_em`, `crown_md`, `crown_ti` could be used with appropriate extensions for the filenames for energy minimization, molecular dynamics simulation and multi-configuration thermodynamic integration, respectively. All of these calculations would use the same topology file `crown.top`.

The extensions `<ext>` identify the kind of information on a file, and are pre-determined.

dbg	debug file
frg	fragment file
gib	free energy data file
mri	free energy multiple run input file
mro	free energy multiple run output file
nw	NWChem input file
nwout	NWChem output file
out	molecular dynamics output file
pdb	PDB formatted coordinate file
prp	property file
qrs	quenched restart file, resulting from an energy minimization
rst	restart file, used to start or restart a simulation
seq	sequence file, describing the system in segments
sgm	segment file, describing segments
syn	synchronization time file
tst	test file
tim	timing analysis file
top	topology file, contains the static description of a system
trj	trajectory file

Table 31.1: List of file extensions for nwchem chemical system files.

31.1.4 Databases

Database file used by the **prepare** module are found in directories with name `$ffield$_$level$`, where `$ffield$` is any of the supported force fields (Section 31.1.5). The source of the data is identified by `$level$`, and can be

level	Description	Availability
s	original published data	public
x	additional published data	public
q	contributed data	public
u	user preferred data	private
t	user defined temporary data	private
c	current working data	private

Typically, only the level **s**, **x** and **q** databases are publicly available. The user is responsible for the private level **u**, **t** and **c** database files. When the **prepare** module scans the databases, the priority is **c>t>u>q>x>s**.

The extension `<ext>` defines the type of database file within each database directory.

frg	fragments
par	parameters
seq	sequences
sgm	segments

Table 31.2: List of database file extensions.

The paths of the different database directories should be defined in a file `.nwchemrc` in a user's home directory, and provides the user the option to select which database files are scanned.

31.1.5 Force fields

Force fields recognized are

Keyword	Force field	Status
amber	AMBER99	AMBER95, GLYCAM also available
charmm	CHARMM	incomplete

31.2 Format of fragment files

Fragment files contain the basic information needed to specify all interactions that need to be considered in a molecular simulation. The format of the fragment files is described in Table 34.1. Normally these files are created by the **prepare** module. Manual editing is needed when, for example, the **prepare** module could not complete atom typing, or when modified charges are required.

31.3 Creating segment files

The **prepare** module is used to generate segment files from corresponding fragment files. A segment file contains all information for the calculation of bonded and non-bonded interactions for a given chemical system using a specific force field.

Which atoms form a fragment is specified in the coordinate file, currently only in PDB format. The restriction is that bonded interactions may only involve atoms on at most two segments does no longer exist as of **NWChem** release 3.2.1. The segment entries define three sets of parameters for each interaction.

Free energy perturbations can be performed using set 1 for the generation of the ensemble while using sets 2 and/or 3 as perturbations. Free energy multiconfiguration thermodynamic integration and multistep thermodynamic perturbation calculations are performed by gradually changing the interactions in the system from parameter set 2 to parameter set 3. These modifications can be edited into the segment files manually, or introduced directly into the topology file using the `modify` commands in the input for the **prepare** module.

The format of a segment is described in Tables 34.2–34.7.

31.4 Creating sequence files

A sequence file describes a molecular system in terms of segments. This file is generated by the **prepare** module for the molecular system provided on a PDB-formatted coordinate file. The file format is given in Table 34.8

31.5 Creating topology files

The topology (Section 31.1.2) describes all static information that describes a molecular system. This includes the connectivity in terms of bond-stretching, angle-bending and torsional interactions, as well as the non-bonded van der Waals and Coulombic interactions.

The topology of a molecular system is generated by the **prepare** module from the sequence in terms of segments as specified on the PDB file. For each unique segment specified in this file the segment database directories are searched

for the segment definition. For segments not found in one of the database directories a segment definition is generated in the temporary directory if a fragment file was found. If a fragment file could not be found, it is generated by the **prepare** module base on what is found on the PDB file.

When all segments are found or created, the parameter substitutions are performed, using force field parameters taken from the parameter databases. After all lists have been generated the topology is written to a local topology file `$system$.top`.

31.6 Creating restart files

Restart files contain all dynamical information about a molecular system and are created by the **prepare** module if a topology file is available. The **prepare** module will automatically generate coordinates for hydrogen atoms and monatomic counter ions not found on the PDB formatted coordinate file, if no fragment or segment files were generated using that PDB file.

The **prepare** module has a number of other optional input command, including solvation.

31.7 Molecular simulations

The type of molecular dynamics simulation is specified by the **NWChem** task directive.

```
task md [ energy | optimize | dynamics | thermodynamics ]
```

where the theory keyword `md` specifies use of the molecular dynamics module, and the operation keyword is one of

`energy` for single configuration energy evaluation

`optimize` for energy minimization

`dynamics` for molecular dynamics simulations and single step thermodynamic perturbation free energy molecular dynamics simulations

`thermodynamics` for combined multi-configuration thermodynamic integration and multiple step thermodynamic perturbation free energy molecular dynamics simulations.

31.8 System specification

The chemical system for a calculation is specified in the topology and restart files. These files should be created using the utilities **nwtop** and **nwrst** before a simulation can be performed. The names of these files are determined from the required `system` directive.

```
system <string systemid>_<string calcid>
```

where the strings `systemid` and `calcid` are user defined names for the chemical system and the type of calculation to be performed, respectively. These names are used to derive the filenames used for the calculation. The topology file used will be `systemid.top`, while all other files are named `systemid_calcid.ext`.

31.9 Restarting and continuing simulations

`finish`

specifies that the current job will finish a previous, incomplete simulation, using the input data that have been recorded by that previous run in the restart file. Most of the input in the current md input block will be ignored.

`resume`

specifies that the current job will be an extension of a previous simulation, using most of the input data that have been recorded by that previous run in the restart file. Typically the input in the current md input block defines a larger number of steps than the previous job.

31.10 Parameter set

`set <integer iset>`

specifies the use of parameter set `<iset>` for the molecular dynamics simulation. The topology file contains three separate parameters sets that can be used. The default for `<iset>` is 1.

`lambda <integer ilambda> <integer ilambda>`

specifies the use of parameter set for the `ilambda`-th of `mlambda` steps.

`pset <integer isetp1> [<integer isetp2>]`

specifies the parameter sets to be used as perturbation potentials in single step thermodynamic perturbation free energy evaluations, where `<isetp1>` specifies the first perturbation parameter set and `<isetp2>` specifies the second perturbation parameter set. Legal values for `<isetp1>` are 2 and 3. Legal value for `<isetp2>` is 3, in which case `<isetp1>` can only be 2. If specified, `<iset>` is automatically set to 1.

`pmf [equilharm <integer npmfc> | scale <real facpmf>]`

specifies that any potential of mean force functions defined in the topology files are to be used. If `equilharm` is specified, the first `npmfc` dynamics steps will use a harmonic potential in stead of any pmf constraint. If `scale` is specified, all pmf force constants are scaled by a factor `facpmf`.

`distar [draver [<integer ndaver default 1>]]`
`[scale <real drsscl>]`
`[after <integer nfdrss>]`

specifies that any distance restraint functions defined in the topology files are to be used.

31.11 Energy minimization algorithms

The energy minimization of the system as found in the restart file is performed with the following directives. If both are specified, steepest descent energy minimization precedes conjugate gradient minimization.

`sd <integer msdit> [init <real dx0sd>] [min <real dxsdmx>] \`
`[max <real dxmsd>]`

specifies the variables for steepest descent energy minimizations, where `<msdit>` is the maximum number of steepest descent steps taken, for which the default is 100, `<dx0sd>` is the initial step size in nm for which the default is 0.001, `<dxsdmx>` is the threshold for the step size in nm for which the default is 0.0001, and `<dxmsd>` is the maximum allowed step size in nm for which the default is 0.05.

```
cg <integer mcgit> [init <real dx0cg>] [min <real dxcgmx>] \
    [cy <integer ncgcy>]
```

specifies the variables for conjugate gradient energy minimizations, where `<mcgit>` is the maximum number of conjugate gradient steps taken, for which the default is 100, `<dx0cg>` is the initial search interval size in nm for which the default is 0.001, `<dxcgmx>` is the threshold for the step size in nm for which the default is 0.0001, and `<ncgcy>` is the number of conjugate gradient steps after which the gradient history is discarded for which the default is 10. If conjugate gradient energy minimization is preceded by steepest descent energy minimization, the search interval is set to twice the final step of the steepest descent energy minimization.

31.12 Multi-configuration thermodynamic integration

The following keywords control free energy difference simulations. Multi-configuration thermodynamic integrations are always combined with multiple step thermodynamic perturbations.

```
(forward | reverse) [[<integer mrun> of] <integer maxlam>]
```

specifies the direction and number of integration steps in free energy evaluations, with `forward` being the default direction. `<mrun>` is the number of ensembles that will be generated in this calculation, and `<maxlam>` is the total number of ensembles to complete the thermodynamic integration. The default value for `<maxlam>` is 21. The default value of `<mrun>` is the value of `<maxlam>`.

```
error <real edacq>
```

specifies the maximum allowed statistical error in each generated ensemble, where `<edacq>` is the maximum error allowed in the ensemble average derivative of the Hamiltonian with respect to λ with a default of 5.0 kJ mol⁻¹.

```
drift <real ddacq>
```

specifies the maximum allowed drift in the free energy result, where `<ddacq>` is the maximum drift allowed in the ensemble average derivative of the Hamiltonian with respect to λ with a default of 5.0 kJ mol⁻¹ps⁻¹.

```
factor <real fdacq>
```

specifies the maximum allowed change in ensemble size where `<fdacq>` is the minimum size of an ensemble relative to the previous ensemble in the calculation with a default value of 0.75.

```
decomp
```

specifies that a free energy decomposition is to be carried out. Since free energy contributions are path dependent, results from a decomposition analysis can no be unambiguously interpreted, and the default is not to perform this decomposition.

```
sss [delta <real delta>]
```

specifies that atomic non-bonded interactions describe a dummy atom in either the initial or final state of the thermodynamic calculation will be calculated using separation-shifted scaling, where `<delta>` is the separation-shifted scaling factor with a default of 0.075 nm². This scaling method prevents problems associated with singularities in the interaction potentials.

```
new | renew | extend
```

specifies the initial conditions for thermodynamic calculations. `new` indicates that this is an initial mcti calculation, which is the default. `renew` instructs to obtain the initial conditions for each λ from the **mro**-file from a previous mcti calculation, which has to be renamed to an **mri**-file. The keyword `extend` will extend a previous mcti calculation from the data read from an **mri**-file.

31.13 Time and integration algorithm directives

Following directives control the integration of the equations of motion.

`leapfrog | leapfrog_bc`

specifies the integration algorithm, where `leapfrog` specifies the default leap frog integration, and `leapfrog_bc` specifies the Brown-Clarke leap frog integrator.

`guided [<real fguide default 0.2> [<real tguide default 0.2>]]`

specifies the use of the guided molecular dynamics simulation technique. Variable *fguide* defines the fraction of the averaged forces *g* to be added to the forces f^f evaluated using the force field functions to obtain the forces *f* used to advance the coordinates.

$$f_i = f_i^f + fguide * g_{i-1} \quad (31.1)$$

Variable *tguide* defines the length of the averaging relative to the timestep Δt .

$$g_i = \frac{\Delta t}{tguide} f_i + \left(1 - \frac{\Delta t}{tguide}\right) g_{i-1} \quad (31.2)$$

The current implementation is still under development.

`equil <integer mequi>`

specifies the number of equilibration steps `<mequi>`, with a default of 100.

`data <integer mdacq> [over <integer ldacq>]]`

specifies the number of data gathering steps `<mdacq>` with a default of 500. In multi-configuration thermodynamic integrations `<mequi>` and `<mdacq>` are for each of the ensembles, and variable `<ldacq>` specifies the minimum number of data gathering steps in each ensemble. In regular molecular dynamics simulations `<ldacq>` is not used. The default value for `<ldacq>` is the value of `<mdacq>`.

`time <real stime>`

specifies the initial time `<stime>` of a molecular simulation in ps, with a default of 0.0.

`step <real tstep>`

specifies the time step `<tstep>` in ps, with 0.001 as the default value.

31.14 Ensemble selection

Following directives control the ensemble type.

`isotherm [<real tmpext> [<real tmpext2>]] [trelax <real tmprlx> [<real tmsrlx>]] \`
`[anneal [<real tann1>] <real tann2>]`

specifies a constant temperature ensemble using Berendsen's thermostat, where `<tmpext>` is the external temperature with a default of 298.15 K, and `<tmprlx>` and `<tmsrlx>` are temperature relaxation times in ps with a default of 0.1. If only `<tmprlx>` is given the complete system is coupled to the heat bath with relaxation time `<tmprlx>`. If both relaxation times are supplied, solvent and solute are independently coupled to the heat bath with relaxation times `<tmprlx>` and `<tmsrlx>`, respectively. If keyword `anneal` is specified, the external temperature will change from `tmpext` to `tmpext2` between simulation time `tann1` and `tann2`

```
isobar [<real prsext>] [trelax <real prsrlx> ] \
      [compress <real compr>] [anisotropic] [xy | z | xy-z]
```

specifies a constant pressure ensemble using Berendsen's piston, where $\langle \text{prsext} \rangle$ is the external pressure with a default of $1.025 \cdot 10^5$ Pa, $\langle \text{prsrlx} \rangle$ is the pressure relaxation time in ps with a default of 0.5, and $\langle \text{compr} \rangle$ is the system compressibility in m^2N^{-1} with a default of $4.53\text{E-}10$. Optional keywords *xy*, *z* and *xy-z* may be used to specify that pressure scaling is to be applied in the *x* and *y* dimension only, the *z* dimension only, or, in all three dimensions with identical scaling in the *x* and *y* dimension. The last option requires that *anisotropic* is also specified.

31.15 Velocity reassignments

Velocities can be periodically reassigned to reflect a certain temperature.

```
vreass <integer nfgaus> <real tgauss>
      [fraction [<real frgaus default 0.5>]]
      [once]
      [(first | initial)] [(last | final)]
```

specifies that velocities will be reassigned every $\langle \text{nfgaus} \rangle$ molecular dynamics steps, reflecting a temperature of $\langle \text{tgauss} \rangle$ K. The default is not to reassign velocities, i.e. $\langle \text{nfgaus} \rangle$ is 0. Keyword *fraction* allows the specification of the fraction of the new velocities are random. Keyword *once* specifies that velocity reassignment only should be done in the first step. Keywords *first* or *initial* and *last* or *final* specify that velocity reassignment should only be applied in the first and last window of multiple run simulations.

31.16 Cutoff radii

Cutoff radii can be specified for short range and long range interactions.

```
cutoff [short] <real rshort> [long <real rlong>] \
      [qmmm <real rqmmm>]
```

specifies the short range cutoff radius $\langle \text{rshort} \rangle$, and the long range cutoff radius $\langle \text{rlong} \rangle$ in nm. If the long range cutoff radius is larger than the short range cutoff radius the twin range method will be used, in which short range forces and energies are evaluated every molecular dynamics step, and long range forces and energies with a frequency of $\langle \text{nflong} \rangle$ molecular dynamics steps. Keyword *qmmm* specifies the radius of the zone around quantum atoms defining the QM/MM bare charges. The default value for $\langle \text{rshort} \rangle$, $\langle \text{rlong} \rangle$ and $\langle \text{rqmmm} \rangle$ is 0.9 nm.

31.17 Polarization

First order and self consistent electronic polarization models have been implemented.

```
polar (first | scf [[<integer mpolit>] <real ptol>])
```

specifies the use of polarization potentials, where the keyword *first* specifies the first order polarization model, and *scf* specifies the self consistent polarization field model, iteratively determined with a maximum of $\langle \text{mpolit} \rangle$ iterations to within a tolerance of $\langle \text{ptol} \rangle$ D in the generated induced dipoles. The default is not to use polarization models.

31.18 External electrostatic field

```
field <real xfield> [freq <real xffreq>] [vector <real xfvect(1:3)>]
```

specifies an external electrostatic field, where `<xfield>` is the field strength, `<xffreq>` is the frequency in MHz and `<xfvect>` is the external field vector.

31.19 Constraints

Constraints are satisfied using the SHAKE coordinate resetting procedure.

```
shake [<integer mshitw> [<integer mshits>]] \
      [<real tlwsha> [<real tlssha>]]
```

specifies the use of SHAKE constraints, where `<mshitw>` is the maximum number of solvent SHAKE iterations, and `<mshits>` is the maximum number of solute SHAKE iterations. If only `<mshitw>` is specified, the value will also be used for `<mshits>`. The default maximum number of iterations is 100 for both. `<tlwsha>` is the solvent SHAKE tolerance in nm, and `<tlssha>` is the solute SHAKE tolerance in nm. If only `<tlwsha>` is specified, the value given will also be used for `<tlssha>`. The default tolerance is 0.001 nm for both.

```
noshake (solvent | solute)
```

disables SHAKE and treats the bonded interaction according to the force field.

31.20 Long range interaction corrections

Long range electrostatic interactions are implemented using the smooth particle mesh Ewald technique, for neutral periodic cubic systems in the constant volume ensemble, using pair interaction potentials. Particle-mesh Ewald long range interactions can only be used in molecular dynamics simulations using effective pair potentials, and not in free energy simulations, QMD or QM/MM simulations.

```
pme [grid <integer ng>] [alpha <real ealpha>] \
     [order <integer morder>] [fft <integer imfft>]\
     [procs <integer nprocs>] [solvent]
```

specifies the use of smooth particle-mesh Ewald long range interaction treatment, where `ng` is the number of grid points per dimension, `ealpha` is the Ewald coefficient in nm^{-1} , with a default that leads to a tolerance of 10^{-4} at the short range cutoff radius, and `morder` is order of the Cardinal B-spline interpolation which must be an even number and at least 4 (default value). A platform specific 3D fast Fourier transform is used, if available, when `imfft` is set to 2. `nprocs` can be used to define a subset of processors to be used to do the FFT calculations. If `solvent` is specified, the charge grid will be calculated from the solvent charges only.

```
react [<real dielec default 80.0>]
```

specifies that a simple reaction field correction is used with a dielectric constant `dielec`. This is an experimental option that has not been well tested.

31.21 Fixing coordinates

Solvent or solute may be fixed using the following keywords.

```
( fix | free )
  solvent ( [<integer idfirst> [<integer idlast>]] |
           ( within | beyond) <real rfix> <string atomname> ) | \
  solute  ( [<integer idfirst> [<integer idlast>]] [ heavy | {<string atomname>}] |
           ( within | beyond) <real rfix> <string atomname> )
  [permanent]
```

For solvent the molecule numbers `idfirst` and `idlast` may be specified to be the first and last molecule to which the directive applies. If omitted, the directive applies to all molecules. For solute, the segment numbers `idfirst` and `idlast` may be specified to be the first and last segment to which the directive applies. If omitted, the directive applies to all segments. In addition, the keyword `heavy` may be specified to apply to all non hydrogen atoms in the solute, or a set of atom names may be specified in which a wildcard character `?` may be used. Keyword `permanent` is used to keep the specification on the restart file for subsequent simulations.

31.22 Special options

```
import [<integer impfr default 1> [<integer impto default impfr> \
      [<integer nftri default 1>]]
```

specifies the import of frames `impfr` to `impto` with frequency `nftri` from a trajectory file with extension `tri` for which energies and forces are to be recalculated. This option only applied to `task md energy`.

`detail`

specifies that moments of inertia and radii of gyration will be part of the recorded properties.

`profile`

specifies that execution time profiling data will be part of the recorded properties.

`scale <real scaleq>`

specifies that all charges will be scaled by the factor `scaleq`.

`collapse [<real fcoll default 10.0> [z | xy]`

specifies that additional forces directed to the origin of the simulation cell with strength `fcoll` will be applied to all solute molecules. If `z` or `xy` is specified, these forces will only apply in the specified dimension(s).

`include fixed`

specifies that energies will be evaluated between fixed atoms. Normally these interactions are excluded from the pairlists.

`eqm <real eqm>`

specifies the zero point of energy in QMD simulations.

`atomlist`

specifies that pairlists will be atom based. Normally pairlist are charge group based.

31.23 Autocorrelation function

For the evaluation of the statistical error of multi-configuration thermodynamic integration free energy results a correlated data analysis is carried out, involving the calculation of the autocorrelation function of the derivative of the Hamiltonian with respect to the control variable λ .

```
auto <integer lacf> [fit <integer nfit>] [weight <real weight>]
```

controls the calculation of the autocorrelation, where `<lacf>` is the length of the autocorrelation function, with a default of 1000, `<nfit>` is the number of functions used in the fit of the autocorrelation function, with a default of 15, and `<weight>` is the weight factor for the autocorrelation function, with a default value of 0.0.

31.24 Print options

Keywords that control print to the output file, with extension **out**. Print directives may be combined to a single directive.

```
print [topol [nonbond] [solvent] [solute]] \
      [step <integer nfoutp> [extra] [energy]] \
      [stat <integer nfstat>] \
      [energies [<integer nfener>]] \
      [forces [<integer nfforce>]] \
      [matrix] \
      [expect <integer npxpct>] \
      [timing] \
      [pmf [<integer iprpfm>]] \
      [out6] \
      [dayout]
```

Keyword `topol` specifies printing the topology information, where `nonbond` refers to the non-bonded interaction parameters, `solvent` to the solvent bonded parameters, and `solute` to the solute bonded parameters. If only `topol` is specified, all topology information will be printed to the output file.

Keyword `step` specifies the frequency `nfoutp` of printing molecular dynamics step information to the output file. If the keyword `extra` is specified additional energetic data are printed for solvent and solute separately. If the keyword `energy` is specified, information is printed for all bonded solute interactions. The default for `nfoutp` is 0. For molecular dynamics simulations this frequency is in time steps, and for multi-configuration thermodynamic integration in λ -steps.

Keyword `stat` specifies the frequency `nfstat` of printing statistical information of properties that are calculated during the simulation. For molecular dynamics simulation this frequency is in time steps, for multi-configuration thermodynamic integration in λ -steps.

Keyword `energies` specifies the frequency `nfener` of printing solute bonded energies the output file for energy/import calculations. The default for `nfener` is 0.

Keyword `forces` specifies the frequency `nfforc` of printing solute forces the output file for energy/import calculations. The default for `nfforc` is 0.

Keyword `matrix` specifies that a solute distance matrix is to be printed.

Keyword `expect` is obsolete.

Keyword `timing` specifies that timing data is printed.

Keyword `pmf` specifies that `pmf` data is printed every `iprpf` steps. Keyword `out6` specifies that output is written to standard out in stead of the output file with extension `out`.

Keyword `dayout` is obsolete.

31.25 Periodic updates

Following keywords control periodic events during a molecular dynamics or thermodynamic integration simulation. Update directives may be combined to a single directive.

```
update [pairs <integer nfpair default 1>] \
      [long <integer nflong default 1>] \
      [center <integer nfcntnr default 0> [zonly | xyonly] \
        [fraction <integer idscb(1:5)>] \
      [motion <integer nslow default 0>] \
      [analysis <integer nfanal default 0>] \
      [rdf <integer nfrdf default 0> \
        [range <real rrdf>] [bins <integer ngl>] \
```

Keyword `pairs` specifies the frequency `<nfpair>` in molecular dynamics steps of updating the pair lists. The default for the frequency is 1. In addition, pair lists are also updated after each step in which recording of the restart or trajectory files is performed. Updating the pair lists includes the redistribution of atoms that changed domain and load balancing, if specified.

Keyword `long` specifies the frequency `<nflong>` in molecular dynamics steps of updating the long range forces. The default frequency is 1. The distinction of short range and long range forces is only made if the long range cutoff radius was specified to be larger than the short range cutoff radius. Updating the long range forces is also done in every molecular dynamics step in which the pair lists are regenerated.

Keyword `center` specifies the frequency `<nfcntnr>` in molecular dynamics steps in which the center of geometry of the solute(s) is translated to the center of the simulation volume. Optional keyword `zonly` or `xyonly` can be used to specify that centering will take place in the z-direction or in the xy-plane only. The solute fractions determining the solutes that will be centered are specified by the keyword `fraction` and the vector `<idscb>`, with a maximum of 5 entries. This translation is implemented such that it has no effect on any aspect of the simulation. The default is not to center, i.e. `nfcntnr` is 0. The default fraction used to center solute is 1.

Keyword `motion` specifies the frequency `<nslow>` in molecular dynamics steps of removing the center of mass motion.

Keyword `analysis` specifies the frequency `<nfanal>` in molecular dynamics steps of invoking the analysis module. This option is obsolete.

Keyword `rdf` specifies the frequency `<nfrdf>` in molecular dynamics steps of calculating contributions to the radial distribution functions. The default is 0. The range of the radial distribution functions is given by `<rrdf>` in nm, with a default of the short range cutoff radius. Note that radial distribution functions are not evaluated beyond the short range cutoff radius. The number of bins in each radial distribution function is given by `<ngl>`, with a default of 1000. This option is no longer supported. If radial distribution function are to be calculated, a **rdf** files needs to be available in which the contributions are specified as follows.

Card	Format	Description
I-1	i	Type, 1=solvent-solvent, 2=solvent-solute, 3-solute-solute
I-2	i	Number of the rdf for this contribution
I-3	i	First atom number
I-4	i	Second atom number

31.26 Recording

The following keywords control recording data to file. Record directives may be combined to a single directive.

```
record [rest <integer nfrest> [keep]] \
      [coord <integer nfcoord default 0>] \
      [wcoord <integer nfwcoord default 0>] \
      [scoord <integer nfscoord default 0>] \
      [veloc <integer nfvelo default 0>] \
      [wvelo <integer nfwvelo default 0>] \
      [svelo <integer nfsvelo default 0>] \
      [force <integer nfvelo default 0>] \
      [wforc <integer nfwvelo default 0>] \
      [sforc <integer nfsvelo default 0>] \
      [(prop | prop_average) <integer nfprop default 0>] \
      [free <integer nffree default 1>] \
      [sync <integer nfsync default 0>] \
      [times <integer nftime default 0>] \
      [acf] [cnv] [fet]
      [binary] [ascii] [ecce] [argos]
```

Keyword `rest` specifies the frequency `<nfrest>` in molecular dynamics steps of rewriting the restart file, with extension `rst`. For multi-configuration thermodynamic integration simulations the frequency is in steps in λ . The default is not to record. The restart file is used to start or restart simulations. The keyword `keep` causes all restart files written to be kept on disk, rather than to be overwritten.

Keyword `coord` specifies the frequency `<nfcoord>` in molecular dynamics steps of writing coordinates to the trajectory file. This directive redefines previous `coord`, `wcoord` and `scoord` directives. The default is not to record.

Keyword `wcoord` specifies the frequency `<nfwcoord>` in molecular dynamics steps of writing solvent coordinates to the trajectory file. This keyword takes precedent over `coord`. This directive redefines previous `coord`, `wcoord` and `scoord` directives. The default is not to record.

Keyword `scoord` specifies the frequency `<nfscoord>` in molecular dynamics steps of writing solute coordinates to the trajectory file. This keyword takes precedent over `coord`. This directive redefines previous `coord`, `wcoord` and `scoord` directives. The default is not to record.

Keyword `veloc` specifies the frequency `<nfvelo>` in molecular dynamics steps of writing velocities to the trajectory file. This directive redefines previous `veloc`, `wvelo` and `svelo` directives. The default is not to record.

Keyword `wvelo` specifies the frequency `<nfwvelo>` in molecular dynamics steps of writing solvent velocities to the trajectory file. This keyword takes precedent over `veloc`. This directive redefines previous `veloc`, `wvelo` and `svelo` directives. The default is not to record.

Keyword `svelo` specifies the frequency `<nfsvelo>` in molecular dynamics steps of writing solute velocities to the trajectory file. This keyword takes precedent over `veloc`. This directive redefines previous `veloc`, `wvelo` and `svelo` directives. The default is not to record.

Keyword `force` specifies the frequency `<nfvelo>` in molecular dynamics steps of writing forces to the trajectory file. This directive redefines previous `vforce`, `wforc` and `sforc` directives. The default is not to record.

Keyword `wforc` specifies the frequency `<nfwvelo>` in molecular dynamics steps of writing solvent forces to the trajectory file. This keyword takes precedent over `force`. This directive redefines previous `vforce`, `wforc` and `sforc` directives. The default is not to record.

Keyword `sforc` specifies the frequency `<nfsvel>` in molecular dynamics steps of writing solute forces to the trajectory file. This keyword takes precedent over `force`. This directive redefines previous `vforce`, `wforc` and `sforc` directives. The default is not to record.

Keyword `prop` specifies the frequency `<nfprop>` in molecular dynamics steps of writing information to the property file, with extension `prp`. The default is not to record.

Keyword `prop_average` specifies the frequency `<nfprop>` in molecular dynamics steps of writing average information to the property file, with extension `prp`. The default is not to record.

Keyword `free` specifies the frequency `<nffree>` in multi-configuration thermodynamic integration steps to record data to the free energy data file, with extension `gib`. The default is 1, i.e. to record at every λ . This option is obsolete. All data are required to do the final analysis.

Keyword `sync` specifies the frequency `<nfsync>` in molecular dynamics steps of writing information to the synchronization file, with extension `syn`. The default is not to record. The information written is the simulation time, the wall clock time of the previous MD step, the wall clock time of the previous force evaluation, the total synchronization time, the largest synchronization time and the node on which the largest synchronization time was found. The recording of synchronization times is part of the load balancing algorithm. Since load balancing is only performed when pair-lists are updated, the frequency `<nfsync>` is correlated with the frequency of pair-list updates `<nfpair>`. This directive is only needed for analysis of the load balancing performance. For normal use this directive is not used.

Keyword `times` specifies the frequency `<nfsync>` in molecular dynamics steps of writing information to the timings file, with extension `tim`. The default is not to record. The information written is wall clock time used by each of the processors for the different components in the force evaluation. This directive is only needed for analysis of the wall clock time distribution. For normal use this directive is not used.

Keywords `acf`, `cnv` and `fet` are obsolete.

Keywords `binary`, `ascii`, `ecce` and `argos` are obsolete.

31.27 Program control options

```
load [reset]
  ( none |
    size [<real factld>] |
    sizez [<real factld>] | pairs |
    (pairs [<integer ldpair>] size [<real factld>]) )
  [last]
  [minimum]
  [average]
  [combination]
  [iotime]
  [experimental]
```

determines the type of dynamic load balancing performed, where the default is `none`. Load balancing option `size` is resizing cells on a node, and `pairs` redistributes the cell-cell interactions over nodes. Keyword `reset` will reset the load balancing read from the restart file. The level of cell resizing can be influenced with `factld`. The cells on the busiest node are resized with a factor

$$\left(1 - factld * \frac{T_{sync} - t_{sync}^{min}}{n_p t_{wall}} \right)^{\frac{1}{3}} \quad (31.3)$$

where T_{sync} is the accumulated synchronization time of all nodes, n_p is the total number of nodes, t_{sync}^{min} is the synchronization time of the busiest node, and t_{wall} is the wall clock time of the molecular dynamics step. For the combined load balancing, `ldpair` is the number of successive pair redistribution load balancing steps in which the accumulated synchronization time increases, before a resizing load balancing step will be attempted. Load balancing is only performed in molecular dynamics steps in which the pair-list is updated. The default load balancing is equivalent to specifying

```
load pairs 10 size 0.75
```

Keyword `last` specifies that the load balancing is based on the synchronization times of the last step. This is the default. Keyword `average` specifies that the load balancing is based on the average synchronization times since the last load balancing step. Keyword `minimum` specifies that the load balancing is based on the minimum synchronization times since the last load balancing step. Keywords `combination`, `iotime` and `experimental` are experimental load balancing options that should not be used in production runs.

```
(pack | nopack)
```

specifies if data are communicated in packed or unpacked form. The default is `pack`.

```
procs <integer npx> <integer npy> <integer npz>
```

specifies the distribution of the available processors over the three Cartesian dimensions. The default distribution is chosen such that, $\langle npx \rangle * \langle npy \rangle * \langle npz \rangle = \langle np \rangle$ and $\langle npx \rangle \leq \langle npy \rangle \leq \langle npz \rangle$, where $\langle npx \rangle$, $\langle npy \rangle$ and $\langle npz \rangle$ are the processors in the x, y and z dimension respectively, and $\langle np \rangle$ is the number of processors allocated for the calculation. Where more than one combination of $\langle npx \rangle$, $\langle npy \rangle$ and $\langle npz \rangle$ are possible, the combination is chosen with the minimum value of $\langle npx \rangle + \langle npy \rangle + \langle npz \rangle$. To change the default setting the following optional input option is provided.

```
cells <integer nbx> <integer nby> <integer nbz>
```

specifies the distribution of cells, where $\langle nbx \rangle$, $\langle nby \rangle$ and $\langle nbz \rangle$ are the number of cells in x, y and z direction, respectively. The molecular system is decomposed into cells that form the smallest unit for communication of atomic data between nodes. The size of the cells is per default set to the short-range cutoff radius. If long-range cutoff radii are used the cell size is set to half the long-range cutoff radius if it is larger than the short-range cutoff. If the number of cells in a dimension is less than the number of processors in that dimension, the number of cells is set to the number of processors.

```
extra <integer madbox>
```

sets the number of additional cells for which memory is allocated. In rare events the amount of memory set aside per node is insufficient to hold all atomic coordinates assigned to that node. This leads to execution which aborts with the message that `mwm` or `msa` is too small. Jobs may be restarted with additional space allocated by where $\langle madbox \rangle$ is the number of additional cells that are allocated on each node. The default for $\langle madbox \rangle$ is 6. In some cases $\langle madbox \rangle$ can be reduced to 4 if memory usage is a concern. Values of 2 or less will almost certainly result in memory shortage.

```
mwm <integer mwmreq>
```

sets the maximum number of solvent molecules $\langle mwmreq \rangle$ per node, allowing increased memory to be allocated for solvent molecules. This option can be used if execution aborted because `mwm` was too small.

```
msa <integer msareq>
```

sets the maximum number of solute atoms $\langle msareq \rangle$ per node, allowing increased memory to be allocated for solute atoms. This option can be used if execution aborted because `msa` was too small.

`mcells <integer mbbreq>`

sets the maximum number of cell pairs `<mbbreq>` per node, allowing increased memory to be allocated for the cell pair lists. This option can be used if execution aborted because `mdbl` was too small.

`boxmin <real rbox>`

sets the minimum size of a cell. This directive is obsolete. The use of `mcells` is preferred.

`segmentsize <real rsgm>`

sets the maximum size of a segment. This value is used to determine which segments at the boundary of the cutoff radius should be considered in the generation of the pairlists. This value is also determined by the prepare module and written to the restart file. Use of this directive is not needed for simulations that use the current prepare module to generate the restart file.

`memory <integer memlim>`

sets a limit `<memlim>` in kB on the allocated amount of memory used by the molecular dynamics module. Per default all available memory is allocated. Use of this command is required for QM/MM simulations only.

`expert`

enables the use of certain combinations of features that are considered unsafe. This directive should not be used for production runs.

`develop <integer idevel>`

enables the use of certain development options specified by the integer `idevel`. This option is for development purposes only, and should not be used for production runs.

`control <integer icntrl>`

enables the use of certain development options specified by the integer `icntrl`. This option is for development purposes only, and should not be used for production runs.

`numerical`

writes out analytical and finite difference forces for test purposes.

`server <string servername> <integer serverport>`

allows monitoring over a socket connection to the specified port on the named server of basic data as a simulation is running.

For development purposes debug information can be written to the debug file with extension **dbg** with

`debug <integer idebug>`

where *idebug* specifies the type of debug information being written.

For testing purposes test information can be written to the test file with extension **tst** with

`test <integer itest>`

where *itest* specifies the number of steps test information is written.

Chapter 32

Analysis

The **analysis** module is used to analyze molecular trajectories generated by the **NWChem** molecular dynamics module, or partial charges generated by the **NWChem** electrostatic potential fit module. This module should not be run in parallel mode.

Directives for the **analysis** module are read from an input deck,

```
analysis
...
end
```

The analysis is performed as post-analysis of trajectory files through using the task directive

```
task analysis
```

or

```
task analyze
```

32.1 System specification

```
system <string systemid>_<string calcid>
```

where the strings `systemid` and `calcid` are user defined names for the chemical system and the type of calculation to be performed, respectively. These names are used to derive the filenames used for the calculation. The topology file used will be `systemid.top`, while all other files are named `systemid_calcid.ext`.

32.2 Reference coordinates

Most analyses require a set of reference coordinates. These coordinates are read from a **NWChem** restart file by the directive,

```
reference <string filename>
```

where `filename` is the name of an existing restart file. This input directive is required.

32.3 File specification

The trajectory file(s) to be analyzed are specified with

```
file <string filename> [<integer firstfile> <integer lastfile>]
```

where `filename` is an existing `trj` trajectory file. If `firstfile` and `lastfile` are specified, the specified filename needs to have a `?` wild card character that will be substituted by the 3-character integer number from `firstfile` to `lastfile`, and the analysis will be performed on the series of files. For example,

```
file tr_md?.trj 3 6
```

will instruct the analysis to be performed on files `tr_md003.trj`, `tr_md004.trj`, `tr_md005.trj` and `tr_md006.trj`.

From the specified files the subset of frames to be analyzed is specified by

```
frames [<integer firstframe default 1>] <integer lastframe> \  
      [<integer frequency default 1>]
```

For example, to analyze the first 100 frames from the specified trajectory files, use

```
frames 100
```

To analyze every 10-th frame between frames 200 and 400 recorded on the specified trajectory files, use

```
frames 200 400 10
```

Solute coordinates of the reference set and each subsequent frame read from a trajectory file are translated to have the center of geometry of the specified solute molecule at the center of the simulation box. After this translation all molecules are folded back into the box according to the periodic boundary conditions. The directive for this operation is

```
center <integer imol> [<integer jmol default imol>]
```

Coordinates of each frame read from a trajectory file can be rotated using

```
rotate ( off | x | y | z ) <real angle units degrees>
```

If `center` was defined, rotation takes place after the system has been centered. The `rotate` directives only apply to frames read from the trajectory files, and not to the reference coordinates. Upto 100 `rotate` directives can be specified, which will be carried out in the order in which they appear in the input deck. `rotate off` cancels all previously defined `rotate` directives.

To perform a hydrogen bond analysis:

```

hbond [distance [[<real rhbmin default 0.0>] <real rhbmin>]] \
      [angle [<real hbdmin> [ <real hbdmax default pi>]]] \
      [solvent [<integer numwhb>]]

```

32.4 Selection

Analyses can be applied to a selection of solute atoms and solvent molecules. The selection is determined by

```

select ( [ super ] [ { <string atomlist> } ] |
solvent <real range> | save <string filename> | read <string filename> )

```

where {atomlist} is the set of atom names selected from the specified residues. By default all solute atoms are selected. When keyword `super` is specified the selection applies to the superimposition option.

The selected atoms are specified by the string `atomlist` which takes the form

```
[{isgm [ - jsgm ] [,]} [:] [{aname[,]}]
```

where `isgm` and `jsgm` are the first and last residue numbers, and `aname` is an atom name. In the atomname a question mark may be used as a wildcard character.

For example, all protein backbone atoms are selected by

```
select _N,_CA,_C
```

To select the backbone atoms in residues 20 to 80 and 90 to 100 only, use

```
select 20-80,90-100:_N,_CA,_C
```

This selection is reset to apply to all atoms after each file directive.

Solvent molecules within range nm from any selected solute atom are selected by

```
select solvent <real range>
```

After solvent selection, the solute atom selection is reset to being all selected.

The current selection can be saved to, or read from a file using the `save` and `read` keywords, respectively.

Some analysis are performed on groups of atoms. These groups of atoms are defined by

```
define <integer igroup> [<real rsel>] [solvent] { <string atomlist> }
```

The string atom in this definitions again takes the form

```
[{isgm [ - jsgm ] [,]} [:] [{aname[,]}]
```

where `isgm` and `jsgm` are the first and last residue numbers, and `aname` is an atom name. In the atomname a question mark may be used as a wildcard character.

Multiple define directive can be used to define a single set of atoms.

32.5 Coordinate analysis

To analyze the root mean square deviation from the specified reference coordinates:

```
rmsd
```

To analyze protein ϕ - ψ and backbone hydrogen bonding:

```
ramachandran
```

To define a distance:

```
distance <integer ibond> <string atomi> <string atomj>
```

To define an angle:

```
angle <integer iangle> <string atomi> <string atomj> <string atomk>
```

To define a torsion:

```
torsion <integer itorsion> <string atomi> <string atomj> \  
      <string atomk> <string atoml>
```

To define a vector:

```
vector <integer ivector> <string atomi> <string atomj>
```

The atom string in these definitions takes the form

```
<integer segment>:<string atomname> | w<integer molecule>:<string atomname>
```

for solute and solvent atom specification, respectively.

To define charge distribution in z-direction:

```
charge_distribution <integer bins>
```

Analyses on atoms in a predefined group are specified by

```
group [<integer igroup> [periodic <integer ipbc>] \  
      ( local [<real rsel default 0.0>] [<real rval default rsel>] \  
        <string function> )
```

where `igroup` specifies the group of atoms defined with a `define` directive. Keyword `periodic` can be used to specify the periodicity, `ipbc=1` for periodicity in `z`, `ipbc=2` for periodicity in `x` and `y`, and `ipbc=3` for periodicity in `x`, `y` and `z`. Currently the only option is `local` which prints all selected solute atom with a distance between `rsel` and `rval` from the atoms defined in `igroup`. The actual analysis is done by the `scan` directive. A formatted report is printed from `group` analyses using

```
report <string filename> local
```

Analyses on pairs of atoms in predefined groups are specified by

```
groups [<integer igroup> [<integer jgroup>]] [periodic [<integer ipbc default 3>]] \
  <string function> [<real value1> [<real value2>]] [<string filename>]
```

where *igroup* and *jgroup* are groups of atoms defined with a `define` directive. Keyword `periodic` specifies that periodic boundary conditions need to be applied in *ipbc* dimensions. The type of analysis is defined by *function*, *value1* and *value2*. If *filename* is specified, the analysis is applied to the reference coordinates and written to the specified file. If no *filename* is given, the analysis is applied to the specified trajectory and performed as part of the `scan` directive. Implemented analyses defined by `<string function> [<real value1> [<real value2>]]` include

`distance` to calculate the distance between the centers of geometry of the two specified groups of atoms, and `distances` to calculate all atomic distances between atoms in the specified groups that lie between *value1* and *value2*.

Coordinate histograms are specified by

```
histogram <integer ideo> [<integer length>] zcoordinate <string filename>
```

where *ideo* is the atom group definition number, *length* is the size of the histogram, `zcoordinate` is the currently only histogram option, and *filename* is the filename to which the histogram is written.

Order parameters are evaluated using

```
order <integer isel> <integer jsel> <string atomi> <string atomj>
```

This is an experimental feature.

To write the average coordinates of a trajectory

```
average [super] <string filename>
```

To perform the coordinate analysis:

```
scan [ super ] <string filename>
```

which will create, depending on the specified analysis options files *filename.rms* and *filename.ana*. After the `scan` directive previously defined coordinate analysis options are all reset. Optional keyword `super` specifies that frames read from the trajectory file(s) are superimposed to the reference structure before the analysis is performed.

32.6 Essential dynamics analysis

Essential dynamics analysis is performed by

```
essential
```

This can be followed by one or more

```
project <integer vector> <string filename>
```

to project the trajectory onto the specified vector. This will create files filename with extensions frm or trj, val, vec, _min.pdb and _max.pdb, with the projected trajectory, the projection value, the eigenvector, and the minimum and maximum projection structure.

For example, an essential dynamics analysis with projection onto the first vector generating files firstvec.{trj, val, vec, _min.pdb, _max.pdb} is generated by

```
essential
project 1 firstvec
```

32.7 Trajectory format conversion

To write a single frame in PDB or XYZ format, use

```
write [<integer number default 1>] [super] [solute] <string filename>
```

To copy the selected frames from the specified trajectory file(s), onto a new file, use

```
copy [solute] [rotate <real tangle>] <string filename>
```

To superimpose the selected atoms for each specified frame to the reference coordinates before copying onto a new file, use

```
super [solute] [rotate <real tangle>] <string filename>
```

The rotate directive specifies that the structure will make a full rotation every tangle ps. This directive only has effect when writing povray files.

The format of the new file is determined from the extension, which can be one of

```
amb  AMBER formatted trajectory file (obsolete)
arc  DISCOVER archive file
bam  AMBER unformatted trajectory file
crd  AMBER formatted trajectory file
dcd  CHARMM formatted trajectory file
esp  gOpenMol formatted electrostatic potential files
frm  ecce frames file (obsolete)
pov  povray input files
trj  NWChem trajectory file
```

If no extension is specified, a trj formatted file will be written.

A special tag can be added to frm and pov formatted files using

```
label <integer itag> <string tag> [ <real rval default 1.0> ] \\
  [ <integer iatag> [ <integer jatag default iatag> ] [ <real rtag default 0.0> ] ]
  [ <string anam> ]
```

where tag number *itag* is set to the string *tag* for all atoms *anam* within a distance *rtag* from segments *iatag* through *jatag*. A question mark can be used in *anam* as a wild card character.

Atom rendering is specified using

```
render ( cpk | stick ) [ <real rval default 1.0> ] \\
  [ <integer iatag> [ <integer jatag default iatag> ] [ <real rtag default 0.0> ] ]
  [ <string anam> ]
```

for all atoms *anam* within a distance *rtag* from segments *iatag* through *jatag*, and a scaling factor of *rval*. A question mark can be used in *anam* as a wild card character.

Atom color is specified using

```
color ( <string color> | atom ) \\
  [ <integer iatag> [ <integer jatag default iatag> ] [ <real rtag default 0.0> ] ]
  [ <string anam> ]
```

for all atoms *anam* within a distance *rtag* from segments *iatag* through *jatag*. A question mark can be used in *anam* as a wild card character.

For example, to display all carbon atoms in segments 34 through 45 in green and rendered cpk in povray files can be specified with

```
render cpk 34 45 _C??
color green 34 45 _C??
```

Coordinates written to a pov file can be scaled using

```
scale <real factor>
```

A zero or negative scaling factor will scale the coordinates to lie within [-1,1] in all dimensions.

The cpk rendering in povray files can be scaled by

```
cpk <real factor default 1.0>
```

The stick rendering in povray files can be scaled by

```
stick <real factor default 1.0>
```

The initial sequence number of esp related files is defined by

```
index <integer index default 1>
```

32.8 Electrostatic potentials

A file in plt format of the electrostatic potential resulting from partial charges generated by the ESP module is generated by the command

```
esp [ <integer spacing default 10> ] \  
    [ <real rcut default 1.0> ] [periodic [<integer iper default 3>]] \  
    [ <string xfile> [ <string pltfile> ] ]
```

The input coordinates are taken from the `xyzq` file that can be generated from a `rst` by the `prepare` module. Parameter `spacing` specifies the number of gridpoints per nm, `rcut` specifies extent of the charge grid beyond the molecule. Periodic boundaries will be used if `periodic` is specified. If `iper` is set to 2, periodic boundary conditions are applied in `x` and `y` dimensions only. If `periodic` is specified, a negative value of `rcut` will extend the grid in the periodic dimensions by `abs(rcut)`, otherwise this value will be ignored in the periodic dimensions. The resulting `plt` formatted file `pltfile` can be viewed with the `gOpenMol` program. The resulting electrostatic potential grid is in units of $\text{kJ mol}^{-1}\text{e}^{-1}$. If no files are specified, only the parameters are set. This analysis applies to solute(s) only.

Chapter 33

Combined quantum and molecular mechanics

Combined or hybrid Quantum Mechanics and Molecular Mechanics (QM/MM) is a simulation methodology that is about 15 years old but in all the literature there are cautions that calibration computations must be done to validate the model for each particular chemical system studied. This is not a black box style computation and the NWChem users are advised that without calibration QM/MM may not give the appropriate results¹.

The QM/MM module in NWChem is driven by the molecular dynamics module. This module currently works for any QM method that has analytic gradients². The input for this requires the definition of chemical system via the same interface that is used by the MD module (c.f. Section 31). The extensions to this interface include the definition of “Quantum” atoms and “Link” where appropriate. The QM information must be present in the traditional NWChem input deck except for the geometry³. The geometrical information will be constructed automatically by nwmd. For dynamics and free energy simulations the input is again identical to that for nwmd with limitations on the kinds of simulations that can be done. Link atoms are defined by the molecular dynamics module, based on the specification of the quantum atoms in the prepare module. For the link atoms, basis sets `X_L` need to be defined in the basis set input block.

The QM/MM module is invoked with the task directive where the “theory” is QMMM. The recognized operations on the QM/MM theory directive are energy, optimize, and dynamics.

```
task qmmm (energy | optimize | dynamics)
```

Tasks `gradient`, `saddle`, `frequencies` and `thermodynamics` are currently not available in the QM/MM mode.

The QM/MM input consists of the standard NWChem input block:

```
qmmm  
...
```

¹c.f., Singh and Kollman, *J. Comp. Chem.* **7**, 718 (1986); M. J. Field, P. A. Bash and M. Karplus, *J. Comp. Chem.* **11**, 700, (1990); J. Gao, “Methods and Applications of Combined Quantum Mechanical and Molecular Mechanical Potentials.” In *Reviews in Computational Chemistry*; K. B. Lipkowitz, D. B. Boyd, Eds.; VCH Publishers: New York; Vol. 7, pp 119-185 (1995); and M. A. Thompson and G. K. Schenter, *J. Phys. Chem* **99** 6374 (1995)

²The QM/MM method will work with numerical gradients available in NWChem, but it is expected that the performance will not allow any substantive simulations

³Any geometry information in the traditional form will be ignored

end

The `qmmm` has the following the additional sub-directive that the user may specify for the particular simulation. These options currently are:

- `eatoms < real eatoms >`

There is one compound input directive that must exist for the QM/MM simulation to proceed. This sets the relative zero of energy for the QM component of the system. It is not incorrect to leave this value as zero but the energetics of the QM system will likely over shadow the MM component of the system. Properties based on energy fluctuations of the system will be overly sensitive to the energy of the QM component of the system. The zero of energy for the MM system is by definition of most parameterized force fields the separated atom energy. The zero of energy for QM systems by definition of most QM methods is the vacuum. The *a priori* determination of the separated atom energy for a particular QM method is not well defined and thus leads to a number of assumptions or guess work depending upon the particular QM method being utilized. Therefore, the determination of the QM separated atom energy (“eatoms”) is left to the user. There is no default for this and the input **must** be present for a QM/MM simulation.

- `qatoms < real qatoms >`

This input directive is used to specify the total charge on the set of quantum atoms.

- `link hydrogen`

This directive specified to use hydrogens as link atoms. The default is to use the next group VII atom.

- `nobq (hydrogen | all | none)`

This directive specifies that hydrogen atoms, all atoms, or no atoms, respectively, bonded to link atoms will carry a zero charge in the QM part of the calculation.

All other parameters that control the QM/MM simulation are set via the input to `nwmd` (see chapter 31).

Chapter 34

File formats

Card	Format	Description
I-1-1	a1	\$ to identify the start of a fragment
I-1-2	a10	name of the fragment, the tenth character N: identifies beginning of a chain C: identifies end of a chain blank: identifies chain fragment M: identifies an integral molecule
I-2-1	i5	number of atoms in the fragment
For each atom one deck II		
II-1-1	i5	atom sequence number
II-1-2	a6	atom name
II-1-3	a5	atom type
II-1-4	a1	dynamics type : normal D : dummy atom S : solute interactions only Q : quantum atom other : intramolecular solute interactions only
II-1-5	i5	link number 0: no link 1: first atom in chain 2: second atom in chain 3 and up: other links
II-1-6	i5	environment type 0: no special identifier 1: planar, using improper torsion 2: tetrahedral, using improper torsion 3: tetrahedral, using improper torsion 4: atom in aromatic ring
II-1-7	i5	charge group
II-1-8	i5	polarization group
II-1-9	f12.6	atomic partial charge
II-1-10	f12.6	atomic polarizability
Any number of cards in deck III to specify complete connectivity		
III-1-1	16i5	connectivity, duplication allowed

Table 34.1: The format of fragment files.

Deck	Format	Description
I-1-1	a1	\$ to identify the start of a segment
I-1-2	a10	name of the segment, the tenth character N: identifies beginning of a chain C: identifies end of a chain blank: identifies chain fragment M: identifies an integral molecule
I-2-1	i5	number of atoms in the segment
I-2-2	i5	number of bonds in the segment
I-2-3	i5	number of angles in the segment
I-2-4	i5	number of proper dihedrals in the segment
I-2-5	i5	number of improper dihedrals in the segment

Table 34.2: Segment file format, table 1 of 6.

Deck	Format	Description
For each atom one deck II		
II-1-1	i5	atom sequence number
II-1-2	a6	atom name
II-1-3	a5	atom type, generic set 1
II-1-4	a1	dynamics type : normal D : dummy atom S : solute interactions only Q : quantum atom other : intramolecular solute interactions only
II-1-4	a5	atom type, generic set 2
II-1-5	a1	dynamics type : normal D : dummy atom S : solute interactions only Q : quantum atom other : intramolecular solute interactions only
II-1-6	a5	atom type, generic set 3
II-1-7	a1	dynamics type : normal D : dummy atom S : solute interactions only Q : quantum atom other : intramolecular solute interactions only
II-1-8	i5	charge group
II-1-9	i5	polarization group
II-1-10	i5	link number
II-1-11	i5	environment type 0: no special identifier 1: planar, using improper torsion 2: tetrahedral, using improper torsion 3: tetrahedral, using improper torsion 4: atom in aromatic ring
II-2-1	f12.6	atomic partial charge in e, set 1
II-2-2	f12.6	atomic polarizability/ $4\pi\epsilon_0$ in nm ³ , set 1
II-2-3	f12.6	atomic partial charge in e, set 2
II-2-4	f12.6	atomic polarizability/ $4\pi\epsilon_0$ in nm ³ , set 2
II-2-5	f12.6	atomic partial charge in e, set 3
II-2-6	f12.6	atomic polarizability/ $4\pi\epsilon_0$ in nm ³ , set 3

Table 34.3: Segment file format, table 2 of 6.

Deck	Format	Description
For each bond a deck III		
III-1-1	i5	bond sequence number
III-1-2	i5	bond atom i
III-1-3	i5	bond atom j
III-1-4	i5	bond type 0: harmonic 1: constrained bond
III-1-5	i5	bond parameter origin 0: from database, next card ignored 1: from next card
III-2-1	f12.6	bond length in nm, set 1
III-2-2	e12.5	bond force constant in $\text{kJ nm}^2 \text{mol}^{-1}$, set 1
III-2-3	f12.6	bond length in nm, set 2
III-2-4	e12.5	bond force constant in $\text{kJ nm}^2 \text{mol}^{-1}$, set 2
III-2-5	f12.6	bond length in nm, set 3
III-2-6	e12.5	bond force constant in $\text{kJ nm}^2 \text{mol}^{-1}$, set 3

Table 34.4: Segment file format, table 3 of 6.

Deck	Format	Description
For each angle a deck IV		
IV-1-1	i5	angle sequence number
IV-1-2	i5	angle atom i
IV-1-3	i5	angle atom j
IV-1-4	i5	angle atom k
IV-1-5	i5	angle type 0: harmonic
IV-1-6	i5	angle parameter origin 0: from database, next card ignored 1: from next card
IV-2-1	f10.6	angle in radians, set 1
IV-2-2	e12.5	angle force constant in kJ mol^{-1} , set 1
IV-2-3	f10.6	angle in radians, set 2
IV-2-4	e12.5	angle force constant in kJ mol^{-1} , set 2
IV-2-5	f10.6	angle in radians, set 3
IV-2-6	e12.5	angle force constant in kJ mol^{-1} , set 3

Table 34.5: Segment file format, table 4 of 6.

Deck	Format	Description
For each proper dihedral a deck V		
V-1-1	i5	proper dihedral sequence number
V-1-2	i5	proper dihedral atom i
V-1-3	i5	proper dihedral atom j
V-1-4	i5	proper dihedral atom k
V-1-5	i5	proper dihedral atom l
V-1-6	i5	proper dihedral type 0: $C \cos(m\phi - \delta)$
V-1-7	i5	proper dihedral parameter origin 0: from database, next card ignored 1: from next card
V-2-1	i3	multiplicity, set 1
V-2-2	f10.6	proper dihedral in radians, set 1
V-2-3	e12.5	proper dihedral force constant in kJ mol^{-1} , set 1
V-2-4	i3	multiplicity, set 2
V-2-5	f10.6	proper dihedral in radians, set 2
V-2-6	e12.5	proper dihedral force constant in kJ mol^{-1} , set 2
V-2-7	i3	multiplicity, set 3
V-2-8	f10.6	proper dihedral in radians, set 3
V-2-9	e12.5	proper dihedral force constant in kJ mol^{-1} , set 3

Table 34.6: Segment file format, table 5 of 6.

Deck	Format	Description
For each improper dihedral a deck VI		
VI-1-1	i5	improper dihedral sequence number
VI-1-2	i5	improper dihedral atom i
VI-1-3	i5	improper dihedral atom j
VI-1-4	i5	improper dihedral atom k
VI-1-5	i5	improper dihedral atom l
VI-1-6	i5	improper dihedral type 0: harmonic
VI-1-7	i5	improper dihedral parameter origin 0: from database, next card ignored 1: from next card
VI-2-1	3x,f10.6	improper dihedral in radians, set 1
VI-2-2	e12.5	improper dihedral force constant in kJ mol^{-1} , set 1
VI-2-3	3x,f10.6	improper dihedral in radians, set 2
VI-2-4	e12.5	improper dihedral force constant in kJ mol^{-1} , set 2
VI-2-5	3x,f10.6	improper dihedral in radians, set 3
VI-2-6	e12.5	improper dihedral force constant in kJ mol^{-1} , set 3

Table 34.7: Segment file format, table 6 of 6.

Card	Format	Description
I-1-1	a1	\$ to identify the start of a sequence
I-1-2	a10	name of the sequence
Any number of cards 1 and 2 in deck II to specify the system		
II-1-1	i5	segment number
II-1-2	a10	segment name, last character will be determined from chain
II-2-1	a	break to identify a break in the molecule chain
II-2-1	a	molecule to identify the end of a solute molecule
II-2-1	a	fraction to identify the end of a solute fraction
II-2-1	a5	link to specify a link
II-2-2	i5	segment number of first link atom
II-2-3	a4	name of first link atom
II-2-4	i5	segment number of second link atom
II-2-5	a4	name of second link atom
II-2-1	a	solvent to identify solvent definition on next card
II-2-1	a	stop to identify the end of the sequence
II-2-1	a6	repeat to repeat next <i>ncard</i> cards <i>ncount</i> times
II-2-2	i5	number of cards to repeat (<i>ncards</i>)
II-2-3	i5	number of times to repeat cards (<i>ncount</i>)
Any number of cards in deck II to specify the system		

Table 34.8: Sequence file format.

Card	Format	Description
I-1-1	a6	keyword header
I-2-1	i10	number of atoms per solvent molecule
I-2-2	i10	number of solute atoms
I-2-3	i10	number of solute bonds
I-2-4	i10	number of solvent bonds
For each atoms per solvent molecule one card I-3		
I-3-1	a5	solvent name
I-3-2	a5	atom name
I-3-3	6x,i10	solvent atom counter
For each solute atom one card I-4		
I-4-1	a5	segment name
I-4-2	a5	atom name
I-4-3	i6	segment number
I-4-4	i10	solute atom counter
For each solvent bond one card I-5		
I-5-1	i8	atom index i for bond between i and j
I-5-2	i8	atom index j for bond between i and j
For each solute bond one card I-6		
I-6-1	i8	atom index i for bond between i and j
I-6-2	i8	atom index j for bond between i and j
For each frame one deck II		
II-1-1	a5	keyword frame
II-2-1	f12.6	time of frame in ps
II-2-2	f12.6	temperature of frame in K
II-2-3	e12.5	pressure of frame in Pa
II-3-1	f12.6	box dimension x
II-3-2	12x,f12.6	box dimension y
II-3-3	24x,f12.6	box dimension z
II-4-1	l1	logical lxw for solvent coordinates
II-4-2	l1	logical lvw for solvent velocities
II-4-3	l1	logical lxs for solute coordinates
II-4-4	l1	logical lvs for solute velocities
II-4-5	i10	number of solvent molecules
II-4-6	i10	number of solvent atoms
II-4-7	i10	number of solute atoms
For each solvent molecule one card II-5 for each atom		
II-5-1	3f8.3	solvent atom coordinates, if lxw or lvw
II-5-4	3f8.3	solvent atom velocities, if lvw
For each solute atom one card II-6 for each atom		
II-6-1	3f8.3	solute atom coordinates, if lxs or lvs
II-6-4	3f8.3	solute atom velocities, if lvs

Table 34.9: Trajectory file format.

Card	Format	Description
For each step in λ one deck I		
I-1-1	i7	number n_{deriv} of data summed in derivative decomposition array $deriv$
I-1-2	i7	length n_{data} of total derivative array dfr
I-1-3	f12.6	current value of λ
I-1-4	f12.6	step size of λ
I-2-1	4e12.12	derivative decomposition array $deriv(1 : 24)$
I-3-1	4e12.12	total derivative array $dfr(1 : n_{da})$
I-4-1	i10	size of ensemble at current λ
I-4-2	e20.12	average temperature at current λ
I-4-3	e20.12	average exponent reverse perturbation energy at current λ
I-4-4	e20.12	average exponent forward perturbation energy at current λ

Table 34.10: Free energy file format.

Card	Format	Description
For each analyzed time step one card I-1		
I-1-1	f12.6	time in ps
I-1-2	f12.6	total rms deviation of the selected atoms before superimposition
I-1-3	f12.6	total rms deviation of the selected atoms after superimposition
II-1-1	a8	keyword <i>analysis</i>
For each solute atom one card II-2		
II-2-1	a5	segment name
II-2-2	a5	atom name
II-2-3	i6	segment number
II-2-4	i10	atom number
II-2-5	i5	selected if 1
II-2-6	f12.6	average atom rms deviation after superimposition
III-1-1	a8	keyword <i>analysis</i>
For each solute segment one card III-2		
III-2-1	a5	segment name
III-2-2	i6	segment number
III-2-3	f12.6	average segment rms deviation after superimposition

Table 34.11: Root mean square deviation file format.

Card	Format	Description
I-1-1	i7	number n_{prop} of recorded properties
I-1-2	1x,2a10	date and time
For each of the n_{prop} properties one card I-2		
I-2-1	a50	description of recorded property
For each recorded step one deck II		
II-1-1	4(1pe12.5)	value of property

Table 34.12: Property file format.

Chapter 35

Pseudopotential plane-wave density functional theory (NWPW)

The NWChem plane-wave (NWPW) module uses pseudopotentials and plane-wave basis sets to perform Density Functional Theory calculations. This module complements the capabilities of the more traditional Gaussian function based approaches by having an accuracy at least as good for many applications, yet is still fast enough to treat systems containing hundreds of atoms. Another significant advantage is its ability to simulate dynamics on a ground state potential surface directly at run-time using the Car-Parrinello algorithm. This method's efficiency and accuracy make it a desirable first principles method of simulation in the study of complex molecular, liquid, and solid state systems. Applications for this first principles method include the calculation of free energies, search for global minima, explicit simulation of solvated molecules, and simulations of complex vibrational modes that cannot be described within the harmonic approximation.

The NWPW module is a collection of three modules.

- PSPW - (PSeudopotential Plane-Wave) A gamma point code for calculating molecules, liquids, crystals, and surfaces.
- Band - A band structure code for calculating crystals and surfaces with small band gaps (e.g. semi-conductors and metals).
- PAW - a prototype (gamma point) projector augmented plane-wave code for calculating molecules, crystals, and surfaces

The PSPW, Band, and PAW modules can be used to compute the energy and optimize the geometry. Both the PSPW and Band modules can also be used to find saddle points, and compute numerical second derivatives. In addition the PSPW module can also be used to perform Car-Parrinello molecular dynamics.

Section 35.1 describes the tasks contained within the PSPW module, section 35.2 describes the tasks contained within the Band module, section 35.3 describes the tasks contained within the PAW module, and section 35.4 describes the pseudopotential library included with NWChem. The datafiles used by the PSPW module are described in section 35.5. Car-Parrinello output data files are described in section 35.5.7, and the minimization and Car-Parrinello algorithms are described in section 35.6. Examples of how to setup and run a PSPW geometry optimization, a Car-Parrinello simulation, a band structure minimization, and a PAW geometry optimization are presented in sections 35.7, 35.8, and 35.10, and 35.11. Finally in section 35.12 the capabilities and limitations of the NWPW module are discussed.

If you are a first time user of this module it is recommended that you skip the next five sections and proceed directly to the tutorials in sections 35.7-35.11.

35.1 PSPW Tasks

All input to the PSPW Tasks is contained within the compound PSPW block,

```
PSPW
  ...
END
```

To perform an actual calculation a TASK PSPW directive is used (Section 5.10).

```
TASK PSPW
```

In addition to the directives listed in Section 5.10, i.e.

```
TASK pspw energy
TASK pspw gradient
TASK pspw optimize
TASK pspw saddle
TASK pspw frequencies
TASK pspw vib
```

there are additional directives that are specific to the PSPW module, which are:

```
TASK PSPW [Car-Parrinello      ||
           pspw_dplot          ||
           wannier              ||
           psp_generator        ||
           steepest_descent     ||
           psp_formatter        ||
           wavefunction_initializer ||
           v_wavefunction_initializer ||
           wavefunction_expander  ]
```

Once a user has specified a geometry, the PSPW module can be invoked with no input directives (defaults invoked throughout). However, the user will probably always specify the simulation cell used in the computation, since the default simulation cell is not well suited for most systems. There are sub-directives which allow for customized application; those currently provided as options for the PSPW module are:

```
PSPW
  CELL_NAME <string cell_name default 'cell_default'>
  INPUT_WAVEFUNCTION_FILENAME <string input_wavefunctions default input_movecs>
  OUTPUT_WAVEFUNCTION_FILENAME <string output_wavefunctions default input_movecs>
  FAKE_MASS <real fake_mass default 400000.0>
  TIME_STEP <real time_step default 5.8>
  LOOP <integer inner_iteration outer_iteration default 10 100>
```

```

TOLERANCES <real tole tolc default 1.0e-7 1.0e-7>
ENERGY_CUTOFF <real ecut default (see input description)>
WAVEFUNCTION_CUTOFF <real wcut default (see input description)>
EWALD_NCUT <integer ncut default 1>]
EWALD_RCUT <real rcut default (see input description)>
XC (Vosko || PBE96 default Vosko)
DFT || ODFT || RESTRICTED || UNRESTRICTED
MULT <integer mult default 1>
MULLIKEN
ALLOW_TRANSLATION

SIMULATION_CELL ... (see input description) END
DPLOT ... (see input description) END
WANNIER ... (see input description) END
CAR-PARRINELLO ... (see input description) END
PSP_GENERATOR ... (see input description) END
WAVEFUNCTION_INITIALIZER ... (see input description) END
V_WAVEFUNCTION_INITIATIZER ... (see input description) END
WAVEFUNCTION_EXPANDER ... (see input description) END
STEEPEST_DESCENT ... (see input description) END
END

```

The following list describes the keywords contained in the PSPW input block.

- <cell_name> - name of the simulation_cell named <cell_name>. See section 35.1.1.
- <input_wavefunctions> - name of the file containing one-electron orbitals
- <output_wavefunctions> - name of the file that will contain the one-electron orbitals at the end of the run.
- <fake_mass> - value for the electronic fake mass (μ). This parameter is not presently used in a conjugate gradient simulation
- <time_step> - value for the time step (Δt). This parameter is not presently used in a conjugate gradient simulation.
- <inner_iteration> - number of iterations between the printing out of energies and tolerances
- <outer_iteration> - number of outer iterations
- <tole> - value for the energy tolerance.
- <tolc> - value for the one-electron orbital tolerance.
- <edit> - value for the cutoff energy used to define the density. Default is set to be the maximum value that will fit within the simulation_cell <cell_name>.
- <wcut> - value for the cutoff energy used to define the one-electron orbitals. Default is set to be the maximum value that will fit within the simulation_cell <cell_name>.
- <ncut> - value for the number of unit cells to sum over (in each direction) for the real space part of the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
- <rcut> - value for the cutoff radius used in the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
Default set to be $\frac{MIN(\vec{a}_i)}{\pi}, i = 1, 2, 3$.

- (Vosko || PBE96) - Choose between Vosko et al's LDA parameterization or the Perdew, Burke, and Ernzerhof GGA functional.
- MULT - optional keyword which if specified allows the user to define the spin multiplicity of the system
- MULLIKEN - optional keyword which if specified causes a Mulliken analysis to be performed at the end of the simulation.
- ALLOW_TRANSLATION - By default the the center of mass forces are projected out of the computed forces. This optional keyword if specified allows the center of mass forces to not be zero.
- SIMULATION_CELL (see section 35.1.1)
- DPLOT (see section 35.1.3)
- WANNIER (see section 35.1.4)
- CAR-PARRINELLO(see section 35.1.7)
- PSP_GENERATOR (see section 35.1.9)
- WAVEFUNCTION_INITIALIZER (see section 35.1.10)
- V_WAVEFUNCTION_INITIALIZER (see section 35.1.11)
- WAVEFUNCTION_EXPANDER (see section 35.1.12).
- STEEPEST_DESCENT (see section 35.1.13)

A prototype limited memory BFGS (LMBFGS) minimizer can be used to minimize the energy. To use this new optimizer the following SET directive needs to be specified:

```
set nwpw:mimimizer 1 # Default - Grassman conjugate gradient minimizer is used to minimize
set nwpw:mimimizer 2 # Grassman LMBFGS minimizer is used to minimize the energy.
set nwpw:minimizer 4 # Stiefel conjugate gradient minimizer is used to minimize the energy
set nwpw:minimizer 5 # Band-by-band minimizer is used to minimize the energy.
```

Limited testing suggests that the Grassman LMBFGS minimizer is about twice as fast as the conjugate gradient minimizer. However, there are several known cases where this optimizer fails, so it is currently not a default option, and should be used with caution.

In addition the following SET directives can be specified:

```
set nwpw:lcao_skip .false. # Default - initial wavefunctions generated using an LCAO guess
set nwpw:lcao_skip .true. # Initial wavefunctions generated using a random plane-wave guess

set nwpw:lcao_print .false. # Default - Ouput not produced during the generation of the LCAO guess
set nwpw:lcao_print .true. # Output produced during the generation of the LCAO guess.

set nwpw:lcao_iterations 2 #specifies the number of LCAO iterations
```

35.1.1 Simulation Cell

The simulation cell parameters are entered by defining a `simulation_cell` sub-block within the PSPW block. Listed below is the format of a `simulation_cell` sub-block.

```
PSPW
...
  SIMULATION_CELL
    CELL_NAME <string name default 'cell_default'>
    BOUNDARY_CONDITIONS (periodic || aperiodic default periodic)
    LATTICE_VECTORS
      <real a1.x a1.y a1.z default 20.0 0.0 0.0>
      <real a2.x a2.y a2.z default 0.0 20.0 0.0>
      <real a3.x a3.y a3.z default 0.0 0.0 20.0>
    NGRID <integer na1 na2 na3 default 32 32 32>
  END
...
END
```

Basically, the user needs to enter the dimensions, gridding and boundary conditions of the simulation cell. The following list describes the input in detail.

- `<name>` - user-supplied name for the simulation block.
- `periodic` - keyword specifying that the simulation cell has periodic boundary conditions.
- `aperiodic` - keyword specifying that the simulation cell has free-space boundary conditions.
- `<a1.x a1.y a1.z>` - user-supplied values for the first lattice vector
- `<a2.x a2.y a2.z>` - user-supplied values for the second lattice vector
- `<a3.x a3.y a3.z>` - user-supplied values for the third lattice vector
- `<na1 na2 na3>` - user-supplied values for discretization along lattice vector directions.

Alternatively, instead of explicitly entering lattice vectors, users can enter the unit cell using the standard cell parameters, a , b , c , α , β , and γ , by using the LATTICE block. The format for input is as follows:

```
PSPW
...
  SIMULATION_CELL
    ...
    LATTICE
      [lat_a <real a default 20.0>]
      [lat_b <real b default 20.0>]
      [lat_c <real c default 20.0>]
      [alpha <real alpha default 90.0>]
      [beta <real beta default 90.0>]
      [gamma <real gamma default 90.0>]
    END
    ...
  END
```

```

END
...
END

```

The user can also enter the lattice vectors of standard unit cells using the keywords SC, FCC, BCC, for simple cubic, face-centered cubic, and body-centered cubic respectively. Listed below is an example of the format of this type of input.

```

PSPW
...
SIMULATION_CELL
  SC 20.0
  ....
END
...
END

```

Finally, the lattice vectors from the unit cell can also be defined using the fractional coordinate input in the GEOMETRY input (see section 6.7). Listed below is an example of the format of this type of input for an 8 atom silicon carbide unit cell.

```

geometry units au
system crystal
  lat_a 8.277d0
  lat_b 8.277d0
  lat_c 8.277d0
  alpha 90.0d0
  beta 90.0d0
  gamma 90.0d0
end
Si -0.50000d0 -0.50000d0 -0.50000d0
Si 0.00000d0 0.00000d0 -0.50000d0
Si 0.00000d0 -0.50000d0 0.00000d0
Si -0.50000d0 0.00000d0 0.00000d0
C -0.25000d0 -0.25000d0 -0.25000d0
C 0.25000d0 0.25000d0 -0.25000d0
C 0.25000d0 -0.25000d0 0.25000d0
C -0.25000d0 0.25000d0 0.25000d0
end

```

35.1.2 Unit Cell Optimization

The PSPW module using the DRIVER geometry optimizer can optimize a crystal unit cell. Currently this type of optimization works only if the geometry is specified in fractional coordinates. The following SET directive is used to tell the DRIVER geometry optimizer to optimize the crystal unit cell in addition to the geometry.

```

set includestress .true.

```

35.1.3 DPLOT

The pspw dplot task is used to generate plots of various types of electron densities (or orbitals) of a molecule. The electron density is calculated on the specified set of grid points from a PSPW calculation. The output file generated is in the Gaussian Cube format. Input to the DPLOT task is contained within the DPLOT sub-block.

```
PSPW
  ...
  DPLOT
    ...
  END
  ...
END
```

To run a DPLOT calculation the following directive is used:

```
TASK PSPW PSPW_DPLOT
```

Listed below is the format of a DPLOT sub-block.

```
PSPW
  ...
  DPLOT
    VECTORS <string input_wavefunctions default input_movecs>
    DENSITY [total||difference||alpha||beta||laplacian||potential default total] <string orbital_name no default>
    ELF [restricted|alpha|beta] <string elf_name no default>
    ORBITAL <integer orbital_number no default> <string orbital_name no default>

    [LIMITXYZ [units <string Units default angstroms>]
    <real X_From> <real X_To> <integer No_Of_Spacings_X>
    <real Y_From> <real Y_To> <integer No_Of_Spacings_Y>
    <real Z_From> <real Z_To> <integer No_Of_Spacings_Z>]

  END

  ...
END
```

The following list describes the input for the DPLOT sub-block.

```
VECTORS <string input_wavefunctions default input_movecs>
```

This sub-directive specifies the name of the molecular orbital file. If the second file is optionally given the density is computed as the difference between the corresponding electron densities. The vector files have to match.

```
DENSITY [total||difference||alpha||beta||laplacian||potential default total] <string orbital_name no default>
```

This sub-directive specifies, what kind of density is to be plotted. The known names for total, difference, alpha, beta, laplacian, and potential.

```
ELF [restricted|alpha|beta] <string elf_name no default>
```

This sub-directive specifies that an electron localization function (ELF) is to be plotted.

```
ORBITAL <integer orbital_number no default> <string orbital_name no default>
```

This sub-directive specifies the molecular orbital number that is to be plotted.

```
LIMITXYZ [units <string Units default angstroms>]
<real X_From> <real X_To> <integer No_Of_Spacings_X>
<real Y_From> <real Y_To> <integer No_Of_Spacings_Y>
<real Z_From> <real Z_To> <integer No_Of_Spacings_Z>
```

By default the grid spacing and the limits of the cell to be plotted are defined by the input wavefunctions. Alternatively the user can use the LIMITXYZ sub-directive to specify other limits. The grid is generated using No_Of_Spacings + 1 points along each direction. The known names for Units are angstroms, au and bohr.

35.1.4 Wannier

The pspw wannier task is generate maximally localized (Wannier) molecular orbitals. The algorithm proposed by Silvestrelli et al is use to generate the Wannier orbitals. The current version of this code works only for cubic cells.

Input to the Wannier task is contained within the Wannier sub-block.

```
PSPW
...
Wannier
...
END
...
END
```

To run a Wannier calculation the following directive is used:

```
TASK PSPW Wannier
```

Listed below is the format of a Wannier sub-block.

```
PSPW
...
Wannier
  OLD_WAVEFUNCTION_FILENAME <string input_wavefunctions default input_movecs>
  NEW_WAVEFUNCTION_FILENAME <string output_wavefunctions default input_movecs>
END
...
END
```

The following list describes the input for the Wannier sub-block.

- <input_wavefunctions> - name of pspw wavefunction file.
- <output_wavefunctions> - name of pspw wavefunction file that will contain the Wannier orbitals.

35.1.5 Self-Interaction Corrections

The SET directive is used to specify the molecular orbitals contribute to the self-interaction-correction (SIC) term.

```
set pspw:SIC_orbitals <integer list_of_molecular_orbital_numbers>
```

This defines only the molecular orbitals in the list as SIC active. All other molecular orbitals will not contribute to the SIC term.

For example the following directive specifies that the molecular orbitals numbered 1,5,6,7,8, and 15 are SIC active.

```
set pspw:SIC_orbitals 1 5:8 15
```

or equivalently

```
set pspw:SIC_orbitals 1 5 6 7 8 15
```

The following directive turns on self-consistent SIC.

```
set pspw:SIC_relax .false. # Default - Perturbative SIC calculation
set pspw:SIC_relax .true.  # Self-consistent SIC calculation
```

Two types of solvers can be used and they are specified using the following SET directive

```
set pspw:SIC_solver_type 1 # Default - cutoff coulomb kernel
set pspw:SIC_solver_type 2 # Free-space boundary condition kernel
```

The parameters for the cutoff coulomb kernel are defined by the following SET directives:

```
set pspw:SIC_screening_radius <real rcut>
set pspw:SIC_screening_power <real rpower>
```

35.1.6 Point Charge Analysis

The MULLIKEN option can be used to generate derived atomic point charges from a plane-wave density. This analysis is based on a strategy suggested in the work of P.E. Blochl, J. Chem. Phys. vol. 103, page 7422 (1995). In this strategy the low-frequency components a plane-wave density are fit to a linear combination of atom centered Gaussian functions.

The following SET directives are used to define the fitting.

```
set pspw_APC:Gc <real Gc_cutoff> # specifies the maximum frequency component of the density
```

```
set pspw_APC:nga <integer number_gauss> # specifies the the number of Gaussian functions per atom.
```

```
set pspw_APC:gamma <real gamma_list> # specifies the decay lengths of each atom centered Gaussian
```

We suggest using the following parameters.

```
set pspw_APC:Gc 2.5
set pspw_APC:nga 3
set pspw_APC:gamma 0.6 0.9 1.35
```

35.1.7 Car-Parrinello

The Car-Parrinello task is used to perform ab initio molecular dynamics using the scheme developed by Car and Parrinello. In this unified ab initio molecular dynamics scheme the motion of the ion cores is coupled to a fictitious motion for the Kohn-Sham orbitals of density functional theory. Constant energy or constant temperature simulations can be performed. A detailed description of this method is described in section 35.6.

Input to the Car-Parrinello simulation is contained within the Car-Parrinello sub-block.

```
PSPW
...
Car-Parrinello
...
END
...
END
```

To run a Car-Parrinello calculation the following directive is used:

```
TASK PSPW Car-Parrinello
```

The Car-Parrinello sub-block contains a great deal of input, including pointers to data, as well as parameter input. Listed below is the format of a Car-Parrinello sub-block.

```
PSPW
...
Car-Parrinello
  CELL_NAME <string cell_name default 'cell_default'>
  INPUT_WAVEFUNCTION_FILENAME <string input_wavefunctions default input_movecs>
  OUTPUT_WAVEFUNCTION_FILENAME <string output_wavefunctions default input_movecs>
  INPUT_V_WAVEFUNCTION_FILENAME <string input_v_wavefunctions default input_vmovecs>
  OUTPUT_V_WAVEFUNCTION_FILENAME <string output_v_wavefunctions default input_vmovecs>
  FAKE_MASS <real fake_mass default default 1000.0>
  TIME_STEP <real time_step default 5.0>
  LOOP <integer inner_iteration outer_iteration default 10 1>
  SCALING <real scale_c scale_r default 1.0 1.0>
  ENERGY_CUTOFF <real ecut default (see input description)>
  WAVEFUNCTION_CUTOFF <real wcut default (see input description)>
  EWALD_NCUR <integer ncut default 1>
  EWALD_RCUR <real rcut default (see input description)>
  XC (Vosko || PBE96 default Vosko)
  [Nose-Hoover <real Period_electron Temperature_electrion Period_ion Temperature_ion
    default 100.0 298.15 100.0 298.15>]
  [SA_decay <real sa_scale_c sa_scale_r default 1.0 1.0>]
  XYZ_FILENAME <string xyz_filename default XYZ>
  EMOTION_FILENAME <string emotion_filename default EMOTION>
  HMOTION_FILENAME <string hmotion_filename default HMOTION>
  OMOTION_FILENAME <string omotion_filename default OMOTION>
  EIGMOTION_FILENAME <string eigmotion_filename default EIGMOTION>
  ION_MOTION_FILENAME <string ion_motion_filename default MOTION>
```

```

END
...
END

```

The following list describes the input for the Car-Parrinello sub-block.

- `<cell_name>` - name of the the simulation_cell named `<cell_name>`. See section 35.1.1.
- `<input_wavefunctions>` - name of the file containing one-electron orbitals
- `<output_wavefunctions>` - name of the file that will contain the one-electron orbitals at the end of the run.
- `<input_v_wavefunctions>` - name of the file containing one-electron orbital velocities.
- `<output_v_wavefunctions>` - name of the file that will contain the one-electron orbital velocities at the end of the run.
- `<fake_mass>` - value for the electronic fake mass (μ).
- `<time_step>` - value for the Verlet integration time step (Δt).
- `<inner_iteration>` - number of iterations between the printing out of energies.
- `<outer_iteration>` - number of outer iterations
- `<scale_c>` - value for the initial velocity scaling of the one-electron orbital velocities.
- `<scale_r>` - value for the initial velocity scaling of the ion velocities.
- `<ecut>` - value for the cutoff energy used to define the density. Default is set to be the maximum value that will fit within the simulation_cell `<cell_name>`.
- `<wcut>` - value for the cutoff energy used to define the one-electron orbitals. Default is set to be the maximum value that will fit within the simulation_cell `<cell_name>`.
- `<ncut>` - value for the number of unit cells to sum over (in each direction) for the real space part of the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
- `<rcut>` - value for the cutoff radius used in the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
Default set to be $\frac{MIN(|\vec{a}_i|)}{\pi}, i = 1, 2, 3$.
- (Vosko || PBE96) - Choose between Vosko et al's LDA parameterization or the Perdew, Burke, and Erzhherhoff GGA functional.
- Nose-Hoover - optional subblock which if specified causes the simulation to perform Nose-Hoover dynamics. If this subblock is not specified the simulation performs constant energy dynamics. See section 35.6.2 for a description of the parameters.
 - `<Period_electron>` $\equiv P_{electron}$ - estimated period for fictitious electron thermostat.
 - `<Temperature_electron>` $\equiv T_{electron}$ - temperature for fictitious electron motion
 - `<Period_ion>` $\equiv P_{ion}$ - estimated period for ionic thermostat
 - `<Temperature_ion>` $\equiv T_{ion}$ - temperature for ion motion

- SA_decay - optional subblock which if specified causes the simulation to run a simulated annealing simulation. For simulated annealing to work the Nose-Hoover subblock needs to be specified. The initial temperature are taken from the Nose-Hoover subblock. See section 35.6.2 for a description of the parameters.
 - <sa_scale_c> $\equiv \tau_{electron}$ - decay rate in atomic units for electronic temperature.
 - <sa_scale_r> $\equiv \tau_{ionic}$ - decay rate in atomic units for the ionic temperature.
- <xyz_filename> - name of the XYZ motion file generated
- <emotion_filename> - name of the emotion motion file. See section 35.5.7 for a description of the datafile.
- <hmotion_filename> - name of the hmotion motion file. See section 35.5.7 for a description of the datafile.
- <eigmotion_filename> - name of the eigmotion motion file. See section 35.5.7 for a description of the datafile.
- <ion_motion_filename> - name of the ion_motion motion file. See section 35.5.7 for a description of the datafile.
- MULLIKEN - optional keyword which if specified causes an omotion motion file to be created.
- <omotion_filename> - name of the omotion motion file. See section 35.5.7 for a description of the datafile.

When a DPLOT sub-block is specified the following SET directive can be used to output dplot data during a Car-Parrinello simulation:

```
set pspw_dplot:iteration_list <integer list_of_iteration_numbers>
```

The Gaussian cube files specified in the DPLOT sub-block are appended with the specified iteration number.

For example, the following directive specifies that at the 3,10,11,12,13,14,15, and 50 iterations Gaussian cube files are to be produced.

```
set pspw_dplot:iteration_list 3,10:15,50
```

35.1.8 Adding Geometry Constraints To A Car-Parrinello Simulation

The Car-Parrinello module allows users to freeze the cartesian coordinates in a simulation (Note - the Car-Parrinello code recognizes Cartesian constraints, but it does not recognize internal coordinate constraints). The SET directive (Section 6.6) is used to freeze atoms, by specifying a directive of the form:

```
set geometry:actlist <integer list_of_center_numbers>
```

This defines only the centers in the list as active. All other centers will have zero force assigned to them, and will remain frozen at their starting coordinates during a Car-Parrinello simulation.

For example, the following directive specifies that atoms numbered 1, 5, 6, 7, 8, and 15 are active and all other atoms are frozen:

```
set geometry:actlist 1 5:8 15
```

or equivalently,

```
set geometry:actlist 1 5 6 7 8 15
```

If this option is not specified by entering a SET directive, the default behavior in the code is to treat all atoms as active. To revert to this default behavior after the option to define frozen atoms has been invoked, the UNSET directive must be used (since the database is persistent, see Section 3.2). The form of the UNSET directive is as follows:

```
unset geometry:actlist
```

In addition, the Car-Parrinello module allows users to freeze bond lengths via a Shake algorithm. The following SET directive shows how to do this.

```
set nwpw:shake_constraint "2 6 L 6.9334"
```

This input fixes the bond length between atoms 2 and 6 to be 6.9334 bohrs. Note that this input only recognizes bohrs.

When using constraints it is usually necessary to turn off center of mass shifting. This can be done by the following SET directive.

```
set nwpw:com_shift .false.
```

35.1.9 PSP_GENERATOR

A one-dimensional pseudopotential code has been integrated into NWChem. This code allows the user to modify and develop pseudopotentials. Currently, only the Hamann and Troullier-Martins norm-conserving pseudopotentials can be generated. In future releases, the pseudopotential library (section 35.4) will be more complete, so that the user will not have explicitly generate pseudopotentials using this module.

Input to the PSP_GENERATOR task is contained within the PSP_GENERATOR sub-block.

```
PSPW
...
PSP_GENERATOR
...
END
...
END
```

To run a PSP_GENERATOR calculation the following directive is used:

```
TASK PSPW PSP_GENERATOR
```

Listed below is the format of a PSP_GENERATOR sub-block.

```
PSPW
...
PSP_GENERATOR
  PSEUDOPOTENTIAL_FILENAME: <string psp_name>
  ELEMENT: <string element>
  CHARGE: <real charge>
  MASS_NUMBER: <real mass_number>
  ATOMIC_FILLING: <integer ncore nvalence>
  ( (1|2|...) (s|p|d|f|...) <real filling> \
```

```

... )

[CUTOFF: <integer lmax>
  ( (s||p||d||f||g) <real rcut>\
  ... )
]
PSEUDOPOTENTIAL_TYPE: (TROULLIER-MARTINS || HAMANN default HAMANN)
SOLVER_TYPE: (PAULI || SCRHODINGER default PAULI)
EXCHANGE_TYPE: (dirac || PBE96 default DIRAC)
CORRELATION_TYPE: (VOSKO || PBE96 default VOSKO)
[SEMICORE_RADIUS: <real rcore>]

end
...
END

```

The following list describes the input for the PSP_GENERATOR sub-block.

- <psp_name> - name that points to a.
- <element> - Atomic symbol.
- <charge> - charge of the atom
- <mass> - mass number for the atom
- <ncore> - number of core states
- <nvalence> - number of valence states.
- ATOMIC_FILLING:.....(see below)
- <filling> - occupation of atomic state
- CUTOFF:.....(see below)
- <rcore> - value for the semicore radius (see below)

ATOMIC_FILLING Block

This required block is used to define the reference atom which is used to define the pseudopotential. After the ATOMIC_FILLING: <ncore> <nvalence> line, the core states are listed (one per line), and then the valence states are listed (one per line). Each state contains two integer and a value. The first integer specifies the radial quantum number, n , The second integer specifies the angular momentum quantum number, l , and the third value specifies the occupation of the state.

For example to define a pseudopotential for the Neon atom in the $1s^2 2s^2 2p^6$ state could have the block

```

ATOMIC_FILLING: 1 2
  1 s 2.0 #core state - 1s^2
  2 s 2.0 #valence state - 2s^2
  2 p 6.0 #valence state - 2p^6

```

for a pseudopotential with a $2s$ and $2p$ valence electrons or the block

```

ATOMIC_FILLING: 3 0
    1 s 2.0 #core state
    2 s 2.0 #core state
    2 p 6.0 #core state

```

could be used for a pseudopotential with no valence electrons.

CUTOFF **Block**

This optional block specifies the cutoff distances used to match the all-electron atom to the pseudopotential atom. For Hamann pseudopotentials $r_{cut}(l)$ defines the distance where the all-electron potential is matched to the pseudopotential, and for Troullier-Martins pseudopotentials $r_{cut}(l)$ defines the distance where the all-electron orbital is matched to the pseudowavefunctions. Thus the definition of the radii depends on the type of pseudopotential. The cutoff radii used in Hamann pseudopotentials will be smaller than the cutoff radii used in Troullier-Martins pseudopotentials.

For example to define a softened Hamann pseudopotential for Carbon would be

```

ATOMIC_FILLING: 1 2
    1 s 2.0
    2 s 2.0
    2 p 2.0
CUTOFF: 2
    s 0.8
    p 0.85
    d 0.85

```

while a similarly softened Troullier-Marting pseudopotential for Carbon would be

```

ATOMIC_FILLING: 1 2
    1 s 2.0
    2 s 2.0
    2 p 2.0
CUTOFF: 2
    s 1.200
    p 1.275
    d 1.275

```

SEMICORE_RADIUS **Option**

Specifying the SEMICORE_RADIUS option turns on the semicore correction approximation proposed by Louie et al (S.G. Louie, S. Froyen, and M.L. Cohen, Phys. Rev. B, **26**, 1738, (1982)). This approximation is known to dramatically improve results for systems containing alkali and transition metal atoms.

The implementation in the PSPW module defines the semi-core density, $\rho_{semicore}$ in terms of the core density, ρ_{core} , by using the sixth-order polynomial

$$\rho_{semicore}(r) = \begin{cases} \rho_{core} & \text{if } r \geq r_{semicore} \\ c_0 + c_3 r^3 + c_4 r^4 + c_5 r^5 + c_6 r^6 & \text{if } r < r_{semicore} \end{cases} \quad (35.1)$$

This expansion was suggested by Fuchs and Scheffler (M. Fuchs, and M. Scheffler, Comp. Phys. Comm., **119**, 67 (1999)), and is better behaved for taking derivatives (i.e. calculating ionic forces) than the expansion suggested by Louie et al.

35.1.10 WAVEFUNCTION_INITIALIZER

The functionality of this task is now performed automatically. For backward compatibility, we provide a description of the input to this task.

The `wavefunction_initializer` task is used to generate an initial wavefunction datafile. Input to the `WAVEFUNCTION_INITIALIZER` task is contained within the `WAVEFUNCTION_INITIALIZER` sub-block.

```
PSPW
...
WAVEFUNCTION_INITIALIZER
...
END
...
END
```

To run a `WAVEFUNCTION_INITIALIZER` calculation the following directive is used:

```
TASK PSPW WAVEFUNCTION_INITIALIZER
```

Listed below is the format of a `WAVEFUNCTION_INITIALIZER` sub-block.

```
PSPW
...
WAVEFUNCTION_INITIALIZER
  CELL_NAME: <string cell_name>
  WAVEFUNCTION_FILENAME: <string wavefunction_name default input_movecs>
  (RESTRICTED|UNRESTRICTED)
  if (RESTRICTED)
    RESTRICTED_ELECTRONS: <integer restricted electrons>
  if (UNRESTRICTED)
    UP_ELECTRONS: <integer up_electrons>
    DOWN_ELECTRONS: <integer down_electrons>
  END
...
END
```

The following list describes the input for the `WAVEFUNCTION_INITIALIZER` sub-block.

- `<cell_name>` - name of the simulation_cell named `<cell_name>`. See section 35.1.1.
- `<wavefunction_name>` - name that will point to a wavefunction file.
- `RESTRICTED` - keyword specifying that the calculation is restricted.
- `UNRESTRICTED` - keyword specifying that the calculation is unrestricted.
- `<restricted_electrons>` - number of restricted electrons. Not used if an `UNRESTRICTED` calculation.
- `<up_electrons>` - number of spin-up electrons. Not used if a `RESTRICTED` calculation.
- `<down_electrons>` - number of spin-down electrons. Not used if a `RESTRICTED` calculation.

Old Style Input (version 3.3) to WAVEFUNCTION_INITIALIZER

For backward compatibility, the input to the WAVEFUNCTION_INITIALIZER sub-block can also be of the form

```
PSPW
...
  WAVEFUNCTION_INITIALIZER
    CELL_NAME: <string cell_name>
    WAVEFUNCTION_FILENAME: <string wavefunction_name default input_movecs>
    (RESTRICTED | UNRESTRICTED)

    [UP_FILLING: <integer up_filling>
      [0 0 0 0]
      {<integer kx ky kz> (-2||-1||1||2)}]
    [DOWN_FILLING: <integer down_filling>
      [0 0 0 0]
      {<integer kx ky kz> (-2||-1||1||2)}]
  END
...
END
```

where

- <cell_name> - name of the simulation_cell named <cell_name>. See section 35.1.1.
- <wavefunction_name> - name that will point to a wavefunction file.
- RESTRICTED - keyword specifying that the calculation is restricted.
- UNRESTRICTED - keyword specifying that the calculation is unrestricted.
- <up_filling> - number of restricted molecular orbitals if RESTRICTED and number of spin-up molecular orbitals if UNRESTRICTED.
- <down_filling> - number of spin-down molecular orbitals if UNRESTRICTED. Not used if a RESTRICTED calculation.
- <kx ky kz> - specifies which planewave is to be filled.

The values for the planewave $(-2||-1||1||2)$ are used to represent whether the specified planewave is a cosine or a sine function, in addition random noise can be added to these base functions. That is +1 represents a cosine function, and -1 represents a sine function. The +2 and -2 values are used to represent a cosine function with random components added and a sine function with random components added respectively.

35.1.11 V_WAVEFUNCTION_INITIALIZER

The functionality of this task is now performed automatically. For backward compatibility, we provide a description of the input to this task.

The v_wavefunction_initializer task is used to generate an initial velocity wavefunction datafile. Input to the V_WAVEFUNCTION_INITIALIZER task is contained within the V_WAVEFUNCTION_INITIALIZER sub-block.

```

PSPW
  ...
  V_WAVEFUNCTION_INITIALIZER
    ...
  END
  ...
END

```

To run a V_WAVEFUNCTION_INITIALIZER calculation the following directive is used:

```
TASK PSPW WAVEFUNCTION_INITIALIZER
```

Listed below is the format of a V_WAVEFUNCTION_INITIALIZER sub-block.

```

PSPW
  ...
  V_WAVEFUNCTION_INITIALIZER
    V_WAVEFUNCTION_FILENAME: <string v_wavefunction_name default input_vmoves>
    CELL_NAME: <string cell_name>
    (RESTRICTED|UNRESTRICTED)
    UP_FILLING: <integer up_filling>
    DOWN_FILLING: <integer down_filling>
  END
  ...
END

```

The following list describes the input for the V_WAVEFUNCTION_INITIALIZER sub-block.

- <cell_name> - name of the simulation_cell named <cell_name>. See section 35.1.1.
- <wavefunction_name> - name that will point to a velocity wavefunction file.
- RESTRICTED - keyword specifying that the calculation is restricted.
- UNRESTRICTED - keyword specifying that the calculation is unrestricted.
- <up_filling> - number of restricted velocity molecular orbitals if RESTRICTED and number of spin-up velocity molecular orbitals if UNRESTRICTED.
- <down_filling> - number of spin-down velocity molecular orbitals if UNRESTRICTED. Not used if a RESTRICTED calculation.

35.1.12 WAVEFUNCTION_EXPANDER

The functionality of this task is now performed automatically. For backward compatibility, we provide a description of the input to this task.

The wavefunction_expander task is used to convert a new wavefunction file that spans a larger grid space from an old wavefunction file. Input to the WAVEFUNCTION_EXPANDER task is contained within the WAVEFUNCTION_EXPANDER sub-block.

```
PSPW
...
  WAVEFUNCTION_EXPANDER
...
  END
...
END
```

To run a WAVEFUNCTION_EXPANDER calculation the following directive is used:

```
TASK PSPW WAVEFUNCTION_EXPANDER
```

Listed below is the format of a WAVEFUNCTION_EXPANDER sub-block.

```
PSPW
...
  WAVEFUNCTION_EXPANDER
    OLD_WAVEFUNCTION_FILENAME: <string old_wavefunction_name default input_moves>
    NEW_WAVEFUNCTION_FILENAME: <string new_wavefunction_name default input_moves>
    NEW_NGRID: <integer na1 na2 na3>

  END
...
END
```

The following list describes the input for the WAVEFUNCTION_EXPANDER sub-block.

- <old_wavefunction_name> - name of the wavefunction file.
- <new_wavefunction_name> - name that will point to a wavefunction file.
- <na1 na2 na3> - number of grid points in each dimension for the new wavefunction file.

35.1.13 STEEPEST_DESCENT

The functionality of this task is now performed automatically by the PSPW minimizer. For backward compatibility, we provide a description of the input to this task.

The steepest_descent task is used to optimize the one-electron orbitals with respect to the total energy. In addition it can also be used to optimize geometries. This method is meant to be used for coarse optimization of the one-electron orbitals.

Input to the steepest_descent simulation is contained within the steepest_descent sub-block.

```
PSPW
...
  STEEPEST_DESCENT
...
  END
...
END
```

To run a `steepest_descent` calculation the following directive is used:

```
TASK PSPW steepest_descent
```

The `steepest_descent` sub-block contains a great deal of input, including pointers to data, as well as parameter input. Listed below is the format of a `STEEPEST_DESCENT` sub-block.

```
PSPW
...
  STEEPEST_DESCENT
    CELL_NAME <string cell_name>
    [GEOMETRY_OPTIMIZE]
    INPUT_WAVEFUNCTION_FILENAME <string input_wavefunctions default input_moves>
    OUTPUT_WAVEFUNCTION_FILENAME <string output_wavefunctions default input_moves>
    FAKE_MASS <real fake_mass default 400000.0>
    TIME_STEP <real time_step default 5.8>
    LOOP <integer inner_iteration outer_iteration default 10 1>
    TOLERANCES <real tole tolc tolr default 1.0d-9 1.0d-9 1.0d-4>
    ENERGY_CUTOFF <real ecut default (see input description)>
    WAVEFUNCTION_CUTOFF <real wcut default (see input description)>
    EWALD_NCUT <integer ncut default 1>
    EWALD_RCUT <real rcut default (see input description)>
    XC (Vosko || PBE96 default Vosko)
    [MULLIKEN]

  END
...
END
```

The following list describes the input for the `STEEPEST_DESCENT` sub-block.

- `<cell_name>` - name of the simulation_cell named `<cell_name>`. See section 35.1.1.
- `GEOMETRY_OPTIMIZE` - optional keyword which if specified turns on geometry optimization.
- `<input_wavefunctions>` - name of the file containing one-electron orbitals
- `<output_wavefunctions>` - name of the file that will contain the one-electron orbitals at the end of the run.
- `<fake_mass>` - value for the electronic fake mass (μ).
- `<time_step>` - value for the time step (Δt).
- `<inner_iteration>` - number of iterations between the printing out of energies and tolerances
- `<outer_iteration>` - number of outer iterations
- `<tole>` - value for the energy tolerance.
- `<tolc>` - value for the one-electron orbital tolerance.
- `<tolr>` - value for the ion position tolerance.

- `<ecut>` - value for the cutoff energy used to define the density. Default is set to be the maximum value that will fit within the simulation_cell `<cell_name>`.
- `<wcut>` - value for the cutoff energy used to define the one-electron orbitals. Default is set to be the maximum value that will fit within the simulation_cell `<cell_name>`.
- `<ncut>` - value for the number of unit cells to sum over (in each direction) for the real space part of the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
- `<rcut>` - value for the cutoff radius used in the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
Default set to be $\frac{MIN(|\vec{a}_i|)}{\pi}, i = 1, 2, 3$.
- (Vosko || PBE96) - Choose between Vosko et al's LDA parameterization or the Perdew, Burke, and Ernzerhof GGA functional.
- MULLIKEN - optional keyword which if specified causes a Mulliken analysis to be performed at the end of the simulation.

35.2 Band Tasks

All input to the Band Tasks is contained within the compound NWPW block,

```
NWPW
...
END
```

To perform an actual calculation a TASK Band directive is used (Section 5.10).

```
TASK Band
```

Once a user has specified a geometry, the Band module can be invoked with no input directives (defaults invoked throughout). There are sub-directives which allow for customized application; those currently provided as options for the Band module are:

```
NWPW
CELL_NAME <string cell_name default 'cell_default'>
ZONE_NAME <string zone_name default 'zone_default'>
INPUT_WAVEFUNCTION_FILENAME <string input_wavefunctions default input_movecs>
OUTPUT_WAVEFUNCTION_FILENAME <string output_wavefunctions default input_movecs>
FAKE_MASS <real fake_mass default 400000.0>
TIME_STEP <real time_step default 5.8>
LOOP <integer inner_iteration outer_iteration default 10 100>
TOLERANCES <real tole tolc default 1.0e-7 1.0e-7>
ENERGY_CUTOFF <real ecut default (see input description)>
WAVEFUNCTION_CUTOFF <real wcut default (see input description)>
EWALD_NCUR <integer ncut default 1>]
EWALD_RCUR <real rcut default (see input description)>
EXCHANGE_CORRELATION: (Vosko || PBE96 default Vosko)
DFT | |ODFT | |RESTRICTED | |UNRESTRICTED
```

```

MULT <integer mult default 1>

SIMULATION_CELL ... (see input description) END
BRILLOUIN_ZONE ... (see input description) END
MONKHORST-PACK <real n1 n2 n3 default 1 1 1>

END

```

The following list describes these keywords.

- <cell_name> - name of the simulation_cell named <cell_name>. See section 35.1.1.
- <input_wavelfunctions> - name of the file containing one-electron orbitals
- <output_wavelfunctions> - name that will point to file containing the one-electron orbitals at the end of the run.
- <fake_mass> - value for the electronic fake mass (μ). This parameter is not presently used in a conjugate gradient simulation
- <time_step> - value for the time step (Δt). This parameter is not presently used in a conjugate gradient simulation.
- <inner_iteration> - number of iterations between the printing out of energies and tolerances
- <outer_iteration> - number of outer iterations
- <tole> - value for the energy tolerance.
- <tolc> - value for the one-electron orbital tolerance.
- <ecut> - value for the cutoff energy used to define the density. Default is set to be the maximum value that will fit within the simulation_cell <cell_name>.
- <wcut> - value for the cutoff energy used to define the one-electron orbitals. Default is set to be the maximum value that will fit within the simulation_cell <cell_name>.
- <ncut> - value for the number of unit cells to sum over (in each direction) for the real space part of the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
- <rcut> - value for the cutoff radius used in the Ewald summation. Note Ewald summation is only used if the simulation_cell is periodic.
Default set to be $\frac{MIN(\vec{a}_i)}{\pi}, i = 1, 2, 3$.
- (Vosko || PBE96) - Choose between Vosko et al's LDA parameterization or the Perdew, Burke,
- SIMULATION_CELL (see section 35.1.1)
- BRILLOUIN_ZONE (see section 35.2.1)
- MONKHORST-PACK - Alternatively, the MONKHORST-PACK keyword can be used to enter a MONKHORST-PACK sampling of the Brillouin zone.

35.2.1 Brillouin Zone

To supply the special points of the Brillouin zone, the user defines a `brillouin_zone` sub-block within the `NWPW` block. Listed below is the format of a `brillouin_zone` sub-block.

```
NWPW
...
  BRILLOUIN_ZONE
    ZONE_NAME <string name default 'zone_default'>
    (KVECTOR <real k1 k2 k3 no default> <real weight default (see input description)>
      ...)
  END
...
END
```

The user enters the special points and weights of the Brillouin zone. The following list describes the input in detail.

- `<name>` - user-supplied name for the simulation block.
- `<k1 k2 k3>` - user-supplied values for a special point in the Brillouin zone.
- `<weight>` - user-supplied weight. Default is to set the weight so that the sum of all the weights for the entered special points adds up to unity.

35.3 PAW Tasks

All input to the PAW Tasks is contained within the compound `NWPW` block,

```
NWPW
...
END
```

To perform an actual calculation the following is used (Section 5.10).

```
TASK PAW steepest\_descent
```

Once a user has specified a geometry, the PAW module can be invoked with no input directives (defaults invoked throughout). There are sub-directives which allow for customized application; those currently provided as options for the PAW module are:

```
NWPW
  CELL_NAME <string cell_name default 'cell_default'>
  [GEOMETRY_OPTIMIZE]
  INPUT_WAVEFUNCTION_FILENAME <string input_wavefunctions default input_movecs>
  OUTPUT_WAVEFUNCTION_FILENAME <string output_wavefunctions default input_movecs>
  FAKE_MASS <real fake_mass default 400000.0>
  TIME_STEP <real time_step default 5.8>
  LOOP <integer inner_iteration outer_iteration default 10 100>
  TOLERANCES <real tole tol c default 1.0e-7 1.0e-7>
```

```

ENERGY_CUTOFF      <real ecut default (see input description)>
WAVEFUNCTION_CUTOFF <real wcut default (see input description)>
EWALD_NCUT <integer ncut default 1>]
EWALD_RCUT <real rcut default (see input description)>
EXCHANGE_CORRELATION: (Vosko || PBE96 default Vosko)
DFT || ODFT || RESTRICTED || UNRESTRICTED
MULT <integer mult default 1>

SIMULATION_CELL ... (see input description) END

```

END

The following list describes these keywords.

- <cell_name> - name of the the simulation_cell named <cell_name>. The current version of PAW only accepts periodic unit cells. See section 35.1.1.
- GEOMETRY_OPTIMIZE - optional keyword which if specified turns on geometry optimization.
- <input_wavefunctions> - name of the file containing one-electron orbitals
- <output_wavefunctions> - name of the file that will contain the one-electron orbitals at the end of the run.
- <fake_mass> - value for the electronic fake mass (μ). This parameter is not presently used in a conjugate gradient simulation
- <time_step> - value for the time step (Δt). This parameter is not presently used in a conjugate gradient simulation.
- <inner_iteration> - number of iterations between the printing out of energies and tolerances
- <outer_iteration> - number of outer iterations
- <tole> - value for the energy tolerance.
- <tolc> - value for the one-electron orbital tolerance.
- <ecut> - value for the cutoff energy used to define the density. Default is set to be the maximum value that will fit within the simulation_cell <cell_name>.
- <wcut> - value for the cutoff energy used to define the one-electron orbitals. Default is set to be the maximum value that will fix within the simulation_cell <cell_name>.
- <ncut> - value for the number of unit cells to sum over (in each direction) for the real space part of the smooth compensation summation.
- <rcut> - value for the cutoff radius used in the smooth compensation summation.
Default set to be $\frac{MIN(|\vec{a}_i|)}{\pi}, i = 1, 2, 3$.
- (Vosko || PBE96) - Choose between Vosko et al's LDA parameterization or the Perdew, Burke,
- SIMULATION_CELL (see section 35.1.1)

35.4 Pseudopotential and PAW basis Libraries

A library of pseudopotentials used by PSPW and BAND is currently available in the directory

```
$NWCHEM_TOP/src/nwpw/libraryyps/pspw_default
```

The elements listed in the following table are present:

H									
Li Be				B C N O					
Na Mg				Al Si P S Cl					
K		Ti		Fe		Cu Zn Ga Ge As			
Sr									
				Pt Au		Pb			

The pseudopotential libraries are continually being tested and added. Also, a stand-alone code is available that allows users to convert certain pseudopotentials generated with the OPIUM package of Rappe et al. The user can request additional pseudopotentials as well as the conversion code from Eric J. Bylaska at (Eric.Bylaska@pnl.gov).

Similarly, a library of PAW basis used by PAW is currently available in the directory

```
$NWCHEM_TOP/src/nwpw/libraryyps/paw_default
```

Currently there are not very many elements available for PAW. However, the user can request additional basis sets from Eric J. Bylaska at (Eric.Bylaska@pnl.gov). The elements available are: H, O Al, Sc, V, and Fe.

A preliminary implementation of the HGH pseudopotentials (Hartwigsen, Goedecker, and Hutter) has been implemented into the PSPW module. To access the pseudopotentials the pseudopotentials input block is used. For example, to redirect the code to use HGH pseudopotentials for carbon and hydrogen, the following input would be used.

```
nwpw
...
pseudopotentials
  C library HGH_LDA
  H library HGH_LDA
end
...
end
```

The implementation of HGH pseudopotentials is rather limited in this release. HGH pseudopotentials cannot be used to optimize unit cells, and they do not work with the MULLIKEN option. They also have not yet been implemented into the BAND structure code.

If you wish to redirect the code to a different directory other than the default one, you need to set the environmental variable `NWCHEM_NWPW_LIBRARY` to the new location of the `libraryyps` directory.

35.5 NWPW RTDB Entries and DataFiles

Input to the PSPW and Band modules are contained in both the RTDB and datafiles. The RTDB is used to store input that the user will need to directly specify. Input of this kind includes ion positions, ion velocities, and simulation cell parameters. The datafiles are used to store input, such the one-electron orbitals, one-electron orbital velocities, formatted pseudopotentials, and one-dimensional pseudopotentials, that the user will in most cases run a program to generate.

35.5.1 Ion Positions

The positions of the ions are stored in the default geometry structure in the RTDB and must be specified using the GEOMETRY directive.

35.5.2 Ion Velocities

The velocities of the ions are stored in the default geometry structure in the RTDB, and must be specified using the GEOMETRY directive.

35.5.3 Wavefunction Datafile

The one-electron orbitals are stored in a wavefunction datafile. This is a binary file and cannot be directly edited. This datafile is used by steepest_descent and Car-Parrinello tasks and can be generated using the wavefunction_initializer or wavefunction_expander tasks.

35.5.4 Velocity Wavefunction Datafile

The one-electron orbital velocities are stored in a velocity wavefunction datafile. This is a binary file and cannot be directly edited. This datafile is used by the Car-Parrinello task and can be generated using the v_wavefunction_initializer task.

35.5.5 Formatted Pseudopotential Datafile

The pseudopotentials in Kleinman-Bylander form expanded on a simulation cell (3d grid) are stored in a formatted pseudopotential datafile. This is a binary file and cannot be directly edited. This datafile is used by steepest_descent and Car-Parrinello tasks and can be generated using the pseudopotential_formatter task.

35.5.6 One-Dimensional Pseudopotential Datafile

The one-dimensional pseudopotentials are stored in a one-dimensional pseudopotential file. This is an ASCII file and can be directly edited with a text editor. However, the user will usually use the psp_generator task to generate this datafile.

The data stored in the one-dimensional pseudopotential file is

```
character*2 element      :: element name
```

```

integer    charge      :: valence charge of ion
real       mass        :: mass of ion
integer    lmax        :: maximum angular component
real       rcut(lmax)  :: cutoff radii used to define pseudopotentials
integer    nr          :: number of points in the radial grid
real       dr          :: linear spacing of the radial grid
real       r(nr)       :: one-dimensional radial grid
real       Vpsp(nr,lmax) :: one-dimensional pseudopotentials
real       psi(nr,lmax) :: one-dimensional pseudowavefunctions
real       r_semicore  :: semicore radius
real       rho_semicore(nr) :: semicore density

```

and the format of it is:

```

[line 1:      ] element
[line 2:      ] charge mass lmax
[line 3:      ] (rcut(1), l=1,lmax)
[line 4:      ] nr dr
[line 5:      ]   r(1) (Vpsp(1,1), l=1,lmax)
[line 6:      ] ....
[line nr+4:   ] r(nr) (Vpsp(nr,1), l=1,lmax)
[line nr+5:   ] r(1) (psi(1,1), l=1,lmax)
[line nr+6:   ] ....
[line 2*nr+4:] r(nr) (psi(nr,1), l=1,lmax)
[line 2*nr+5:] r_semicore
if (r_semicore read) then
[line 2*nr+6:] r(1) rho_semicore(1)
[line 2*nr+7:] ....
[line 3*nr+5:] r(nr) rho_semicore(nr)
end if

```

35.5.7 PSPW Car-Parrinello Output Datafiles

XYZ motion file

Data file that stores ion positions and velocities as a function of time in XYZ format.

```

[line 1:      ] n_ion
[line 2:      ]
do ii=1,n_ion
[line 2+ii:    ] atom_name(ii), x(ii),y(ii),z(ii),vx(ii),vy(ii),vz(ii)
end do
[line n_ion+3 ] n_nion

do ii=1,n_ion
[line n_ion+3+ii:] atom_name(ii), x(ii),y(ii),z(ii), vx(ii),vy(ii),vz(ii)
end do
[line 2*n_ion+4:] ....

```

ION_MOTION motion file

Datafile that stores ion positions and velocities as a function of time

```
[line 1:          ] it_out, n_ion, omega
[line 2:          ] time
do ii=1,n_ion
[line 2+ii:       ] x(ii),y(ii),z(ii), vx(ii),vy(ii),vz(ii)
end do
[line n_ion+3     ] time
do
do ii=1,n_ion
[line n_ion+3+ii:] x(ii),y(ii),z(ii), vx(ii),vy(ii),vz(ii)
end do
[line 2*n_ion+4:  ] ....
```

EMOTION motion file

Datafile that store energies as a function of time

```
[line 1:          ] time, E1,E2,E3,E4,E5,E6,E7,E8, (E9,E10, if Nose-Hoover)
[line 2:          ] ...
```

HMOTION motion file

Datafile that stores the rotation matrix as a function of time.

```
[line 1:          ] time
[line 2:          ] ms,ne(ms),ne(ms)
do i=1,ne(ms)
[line 2+i:        ] (hml(i,j), j=1,ne(ms)
end do
[line 3+ne(ms):   ] time
[line 4+ne(ms):   ] ....
```

EIGMOTION motion file

Datafile that stores the eigenvalues for the one-electron orbitals as a function of time.

```
[line 1:          ] time, (eig(i), i=1,number_orbitals)
[line 2:          ] ...
```

OMOTION motion file

Datafile that stores a reduced representation of the one-electron orbitals. To be used with a molecular orbital viewer that will be ported to NWChem in the near future.

35.6 Car-Parrinello Scheme for Ab Initio Molecular Dynamics

Car and Parrinello developed a unified scheme for doing *ab initio* molecular dynamics by combining the motion of the ion cores and a fictitious motion for the Kohn-Sham orbitals of density-functional theory (R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471, (1985)). At the heart of this method they introduced a fictitious kinetic energy functional for the Kohn-Sham orbitals.

$$KE(\{\psi_{i,\sigma}(\vec{r})\}) = \sum_{i,\sigma}^{occ} \int d\vec{r} \mu |\dot{\psi}_{i,\sigma}(\vec{r})|^2 \quad (35.2)$$

Given this kinetic energy the constrained equations of motion are found by taking the first variation of the auxiliary Lagrangian.

$$\begin{aligned} L = & \sum_{i,\sigma}^{occ} \int d\vec{r} \mu |\dot{\psi}_{i,\sigma}(\vec{r})|^2 + \frac{1}{2} \sum_I M_I |\dot{\vec{R}}_I|^2 - E[\{\psi_{i,\sigma}(\vec{r})\}, \{\vec{R}_I\}] \\ & + \sum_{ij,\sigma} \Lambda_{ij,\sigma} \left(\int d\vec{r} \psi_{i,\sigma}^*(\vec{r}) \psi_{j,\sigma}(\vec{r}) - \delta_{ij,\sigma} \right) \end{aligned} \quad (35.3)$$

Which generates a dynamics for the wavefunctions $\psi_{i,\sigma}(\vec{r})$ and atoms positions \vec{R}_I through the constrained equations of motion:

$$\mu \ddot{\psi}_{i,\sigma}(\vec{r}, t) = -\frac{\delta E}{\delta \psi_{i,\sigma}^*(\vec{r}, t)} + \sum_j \Lambda_{ij,\sigma} \psi_{j,\sigma}(\vec{r}, t) \quad (35.4)$$

$$M_I \ddot{\vec{R}}_I = -\frac{\partial E}{\partial \vec{R}_I} \quad (35.5)$$

where μ is the fictitious mass for the electronic degrees of freedom and M_I are the ionic masses. The adjustable parameter μ is used to describe the relative rate at which the wavefunctions change with time. $\Lambda_{ij,\sigma}$ are the Lagrangian multipliers for the orthonormalization of the single-particle orbitals $\psi_{i,\sigma}(\vec{r})$. They are defined by the orthonormalization constraint conditions and can be rigorously found. However, the equations of motion for the Lagrange multipliers depend on the specific algorithm used to integrate Eqs. 35.4-35.5.

For this method to give ionic motions that are physically meaningful the kinetic energy of the Kohn-Sham orbitals must be relatively small when compared to the kinetic energy of the ions. There are two ways where this criterion can fail. First, the numerical integrations for the Car-Parrinello equations of motion can often lead to large relative values of the kinetic energy of the Kohn-Sham orbitals relative to the kinetic energy of the ions. This kind of failure is easily fixed by requiring a more accurate numerical integration, i.e. use a smaller time step for the numerical integration. Second, during the motion of the system a the ions can be in locations where there is a Kohn-Sham orbital level crossing, i.e. the density-functional energy can have two states that are nearly degenerate. This kind of failure often occurs in the study of chemical reactions. This kind of failure is not easily fixed and requires the use of a more sophisticated density-functional energy that accounts for low-lying excited electronic states.

35.6.1 Verlet Algorithm for Integration

Eqs. 35.4-35.5 integrated using the Verlet algorithm results in

$$\Psi_{i,\sigma}^{t+\Delta t} \leftarrow 2\Psi_{i,\sigma}^t - \Psi_{i,\sigma}^{t-\Delta t} + \frac{(\Delta t)^2}{\mu} \left[\frac{\delta E}{\delta \Psi_{i,\sigma}^*} + \sum_j \Psi_{j,\sigma} \Lambda_{ji,\sigma} \right]_t \quad (35.6)$$

$$\vec{R}_I^{t+\Delta t} \leftarrow 2\vec{R}_I^t - \vec{R}_I^{t-\Delta t} + \frac{(\Delta t)^2}{M_I} \frac{\partial E}{\partial \vec{R}_I} \quad (35.7)$$

In this molecular dynamic procedure we have to know variational derivative $\frac{\delta E}{\delta \Psi_{i,\sigma}^*}$ and the matrix $\Lambda_{ij,\sigma}$. The variational derivative $\frac{\delta E}{\delta \Psi_{i,\sigma}^*}$ can be analytically found and is

$$\begin{aligned} \frac{\delta E}{\delta \Psi_{i,\sigma}^*} &= -\frac{1}{2} \nabla^2 \Psi_{i,\sigma}(\vec{r}) \\ &+ \int d\vec{r}' W_{ext}(\vec{r}, \vec{r}') \Psi_{i,\sigma}(\vec{r}') \\ &+ \int d\vec{r}' \frac{n(\vec{r}')}{|\vec{r} - \vec{r}'|} \Psi_{i,\sigma}(\vec{r}') \\ &+ \mu_{xc}^\sigma(\vec{r}) \Psi_{i,\sigma}(\vec{r}) \\ &\equiv \hat{H} \Psi_{i,\sigma} \end{aligned} \quad (35.8)$$

To find the matrix $\Lambda_{ij,\sigma}$ we impose the orthonormality constraint on $\Psi_{i,\sigma}^{t+\Delta t}$ to obtain a matrix Riccati equation, and then Riccati equation is solved by an iterative solution

35.6.2 Constant Temperature Simulations: Nose-Hoover Thermostats

Nose-Hoover Thermostats for the electrons and ions can also be added to the Car-Parrinello simulation. In this type of simulation thermostats variables x_e and x_R are added to the simulation by adding the auxiliary energy functionals to the total energy.

$$ION_THERMOSTAT(x_R) = \frac{1}{2} Q_R \dot{x}_R^2 + E_{R0} x_R \quad (35.9)$$

$$ELECTRON_THERMOSTAT(x_e) = \frac{1}{2} Q_e \dot{x}_e^2 + E_{e0} x_e \quad (35.10)$$

In these equations, the average kinetic energy for the ions is

$$E_{R0} = \frac{1}{2} f k_B T \quad (35.11)$$

where f is the number of atomic degrees of freedom, k_B is Boltzmann's constant, and T is the desired temperature. Defining the average fictitious kinetic energy of the electrons is not as straightforward. Blöchl and Parrinello (P.E.

Blöchl and M. Parrinello, Phys. Rev. B, **45**, 9413, (1992)) have suggested the following formula for determining the average fictitious kinetic energy

$$E_{e0} = 4k_B T \frac{\mu}{M} \sum_i \langle \psi_i | -\frac{1}{2} \nabla^2 | \psi_i \rangle \quad (35.12)$$

where μ is the fictitious electronic mass, M is average mass of one atom, and $\sum_i \langle \psi_i | -\frac{1}{2} \nabla^2 | \psi_i \rangle$ is the kinetic energy of the electrons.

Blöchl and Parrinello suggested that the choice of mass parameters, Q_e , and Q_R should be made such that the period of oscillating thermostats should be chosen larger than the typical time scale for the dynamical events of interest but shorter than the simulation time.

$$P_{ion} = 2\pi \sqrt{\frac{Q_R}{4E_{R0}}} \quad (35.13)$$

$$P_{electron} = 2\pi \sqrt{\frac{Q_e}{4E_{e0}}} \quad (35.14)$$

where P_{ion} and $P_{electron}$ are the periods of oscillation for the ionic and fictitious electronic thermostats.

In simulated annealing simulations the electronic and ionic Temperatures are scaled according to an exponential cooling schedule,

$$T_e(t) = T_e^0 \exp^{-\frac{t}{\tau_e}} \quad (35.15)$$

$$T_{ionic}(t) = T_{ionic}^0 \exp^{-\frac{t}{\tau_{ionic}}} \quad (35.16)$$

where T_e^0 and T_{ionic}^0 are the initial temperatures, and τ_e and τ_{ionic} are the cooling rates in atomic units.

35.7 PSPW Tutorial 1: Minimizing the geometry for a C₂ molecule

In this section we show how use the PSPW module to optimize the geometry for a C₂ molecule at the PBE96 levels.

In the following example we show the input needed to optimize the geometry for a C₂ molecule at the LDA level. In this example, default pseudopotentials from the pseudopotential library are used for C, the boundary condition is free-space, the exchange correlation functional is PBE96, The boundary condition is free-space, and the simulation cell is aperiodic and cubic with a side length of 10.0 Angstroms and has 40 grid points in each direction (cutoff energy is 44 Ry).

```
start c2_pspw_pbe96
title "C2 restricted singlet dimer optimization - PBE96/44Ry"

geometry
C    -0.62 0.0 0.0
C     0.62 0.0 0.0
end

pspw
simulation_cell units angstroms
```

```

        boundary_conditions aperiodic
        SC 10.0
        ngrid 40 40 40
    end
    xc pbe96
end
set nwpw:minimizer 2
task pspw optimize

```

35.8 PSPW Tutorial 2: Running a Car-Parrinello Simulation

In this section we show how use the PSPW module to perform a Car-Parrinello molecular dynamic simulation for a C_2 molecule at the LDA level. Before running a PSPW Car-Parrinello simulation the system should be on the Born-Oppenheimer surface, i.e. the one-electron orbitals should be minimized with respect to the total energy (i.e. task pspw energy). The input needed is basically the same as for optimizing the geometry of a C_2 molecule at the LDA level, except that an additional Car-Parrinello sub-block is added.

In the following example we show the input needed to run a Car-Parrinello simulation for a C_2 molecule at the LDA level. In this example, default pseudopotentials from the pseudopotential library are used for C, the boundary condition is free-space, the exchange correlation functional is LDA, The boundary condition is free-space, and the simulation cell is aperiodic and cubic with a side length of 10.0 Angstroms and has 40 grid points in each direction (cutoff energy is 44 Ry). The time step and fake mass for the Car-Parrinello run are specified to be 5.0 au and 600.0 au, respectively.

```

start c2_pspw_lda_md
title "C2 restricted singlet dimer, LDA/44Ry - constant energy Car-Parrinello simulation"

geometry
C    -0.62 0.0 0.0
C     0.62 0.0 0.0
end

pspw
    simulation_cell units angstroms
    boundary_conditions aperiodic
    lattice
        lat_a 10.00d0
        lat_b 10.00d0
        lat_c 10.00d0
    end
    ngrid 40 40 40
end
Car-Parrinello
    fake_mass 600.0
    time_step 5.0
    loop 10 10
end
end
set nwpw:minimizer 2

```

```
task pspw energy
task pspw Car-Parrinello
```

35.9 PSPW Tutorial 3: optimizing a unit cell and geometry for Silicon-Carbide

The following example demonstrates how to use the PSPW module to optimize the unit cell and geometry for a silicon-carbide crystal.

```
title "SiC 8 atom cubic cell - geometry and unit cell optimization"

start SiC

#**** Enter the geometry using fractional coordinates ****
geometry units au center noautosym noautoz print
  system crystal
    lat_a 8.277d0
    lat_b 8.277d0
    lat_c 8.277d0
    alpha 90.0d0
    beta 90.0d0
    gamma 90.0d0
  end
Si   -0.50000d0  -0.50000d0  -0.50000d0
Si    0.00000d0   0.00000d0  -0.50000d0
Si    0.00000d0  -0.50000d0   0.00000d0
Si   -0.50000d0   0.00000d0   0.00000d0
C    -0.25000d0  -0.25000d0  -0.25000d0
C     0.25000d0   0.25000d0  -0.25000d0
C     0.25000d0  -0.25000d0   0.25000d0
C    -0.25000d0   0.25000d0   0.25000d0
end

#***** setup the nwpw gamma point code *****
nwpw
  simulation_cell
    ngrid 16 16 16
  end
  ewald_ncut 8
end
set nwpw:minimizer 2
set nwpw:psi_nolattice .true. # turns off unit cell checking for wavefunctions

driver
  clear
  maxiter 40
end
set includestress .true. # this option tells driver to optimize the unit cell
```

```
task pspw optimize
```

35.10 Band Tutorial 1: Minimizing the energy of a silicon-carbide crystal by running a PSPW and Band simulation in tandem

The following input deck performs a PSPW energy calculation followed by a Band energy calculation at the Γ -point for a cubic (8-atom) silicon-carbide crystal. Since the geometry is entered using fractional coordinates the unit cell parameters do not have to be re-specified in the `simulation_cell nwpw` sub-block. In this example, default pseudopotential from the pseudopotential library are used for C and Si. The advantage of running these calculations in tandem is that the Band code uses the wavefunctions generated from the faster PSPW calculation for its initial guess. The PSPW energy is -38.353570, and the Band energy is -38.353570.

```
start SiC_band
title "SiC 8 atom cubic cell"

#**** geometry entered using fractional coordinates ****
geometry units au center noautosym noautoz print
  system crystal
    lat_a 8.277d0
    lat_b 8.277d0
    lat_c 8.277d0
    alpha 90.0d0
    beta 90.0d0
    gamma 90.0d0
  end
Si   -0.50000d0  -0.50000d0  -0.50000d0
Si    0.00000d0   0.00000d0  -0.50000d0
Si    0.00000d0  -0.50000d0   0.00000d0
Si   -0.50000d0   0.00000d0   0.00000d0
C    -0.25000d0  -0.25000d0  -0.25000d0
C     0.25000d0   0.25000d0  -0.25000d0
C     0.25000d0  -0.25000d0   0.25000d0
C    -0.25000d0   0.25000d0   0.25000d0
end

#***** setup the nwpw gamma point code ****
nwpw
  simulation_cell
    ngrid 16 16 16
  end
  brillouin_zone
    kvector 0.0 0.0 0.0
  end
  ewald_ncut 8
end
set nwpw:minimizer 2
set nwpw:psi_brillioun_check .false.
task pspw energy
```

task band energy

35.11 PAW Tutorial

The following input deck performs for a water molecule a PSPW energy calculation followed by a PAW energy calculation and a PAW geometry optimization calculation. The default unit cell parameters are used (SC=20.0, ngrid 32 32 32). In this simulation, the first PAW run optimizes the wavefunction and the second PAW run optimizes the wavefunction and geometry in tandem.

```

title "paw steepest descent test"

start paw_test

charge 0

geometry units au nocenter noautoz noautosym
O      0.00000    0.00000    0.01390
H     -1.49490    0.00000   -1.18710
H      1.49490    0.00000   -1.18710
end

nwpw
  time_step 15.8
  ewald_rcut 1.50
  tolerances 1.0d-8 1.0d-8
end
set nwpw:lcao_iterations 1
set nwpw:minimizer 2
task pspw energy

task paw steepest_descent

nwpw
  time_step 5.8
  geometry_optimize
  ewald_rcut 1.50
  tolerances 1.0d-7 1.0d-7 1.0d-4
end
task paw steepest_descent

```

35.12 NWPW Capabilities and Limitations

- You cannot use more processors than the size of the third dimension (e.g. a 64x64x64 FFT grid can use at most 64 processors).
- The second and third dimensions of the FFT grid must be the same (i.e. the parameters na2 and na3 must be the same for each simulation cell).

- Wannier orbital generation only works with cubic unit cells ($\alpha = \beta = \gamma = 90^\circ$)
- PAW is not interfaced to driver, stepper, and vib modules.

35.13 Questions and Difficulties

Questions and encountered problems should be reported to nwchem-users@emsl.pnl.gov or to Eric J. Bylaska, Eric.Bylaska@pnl.gov

Chapter 36

Controlling NWChem with Python

Python (version 1.5.1) programs may be embedded into the NWChem input and used to control the execution of NWChem. Python is a very powerful and widely used scripting language that provides useful things such as variables, conditional branches and loops, and is also readily extended. Example applications include scanning potential energy surfaces, computing properties in a variety of basis sets, optimizing the energy w.r.t. parameters in the basis set, computing polarizabilities with finite field, and simple molecular dynamics.

Look in the NWChem `contrib` directory for useful scripts and examples. Visit the Python web-site <http://www.python.org> for a full manual and lots of useful code and resources.

36.1 How to input and run a Python program inside NWChem

A Python program is input into NWChem inside a Python compound directive.

```
python [print|noprint]
    ...
end
```

The END directive must be flush against the left margin (see the Troubleshooting section for the reason why).

The program is by default printed to standard output when read, but this may be disabled with the `noprint` keyword. Python uses indentation to indicate scope (and the initial level of indentation must be zero), whereas NWChem uses optional indentation only to make the input more readable. For example, in Python, the contents of a loop, or conditionally-executed block of code must be indented further than the surrounding code. Also, Python attaches special meaning to several symbols also used by NWChem. For these reasons, the input inside a `PYTHON` compound directive is read verbatim except that if the first line of the Python program is indented, the same amount of indentation is removed from all subsequent lines. This is so that a program may be indented inside the `PYTHON` input block for improved readability of the NWChem input, while satisfying the constraint that when given to Python the first line has zero indentation.

E.g., the following two sets of input specify the same Python program.

```
python
    print 'Hello'
    print 'Goodbye'
```

```

end

python
print 'Hello'
print 'Goodbye'
end

```

whereas this program is in error since the indentation of the second line is less than that of the first.

```

python
  print 'Hello'
print 'Goodbye'
end

```

The Python program is not executed until the following directive is encountered

```
task python
```

which is to maintain consistency with the behavior of NWChem in general. *The program is executed by all nodes.* This enables the full functionality and speed of NWChem to be accessible from Python, but there are some gotchas

- Print statements and other output will be executed by all nodes so you will get a lot more output than probably desired unless the output is restricted to just one node (by convention node zero).
- The calls to NWChem functions are all collective (i.e., all nodes must execute them). If these calls are not made collectively your program may deadlock (i.e., cease to make progress).
- When writing to the database (`rtdb_put ()`) it is the data from node zero that is written.
- NWChem overrides certain default signal handlers so care must be taken when creating processes (see Section 36.3.11).

36.2 NWChem extensions

Since we have little experience using Python, the NWChem-Python interface might change in a non-backwardly compatible fashion as we discover better ways of providing useful functionality. We would appreciate suggestions about useful things that can be added to the NWChem-Python interface. In principle, nearly any Fortran or C routine within NWChem can be extended to Python, but we are also interested in ideas that will enable users to build completely new things. For instance, how about being able to define your own energy functions that can be used with the existing optimizers or dynamics package?

Python has been extended with a module named "nwchem" which is automatically imported and contains the following NWChem-specific commands. They all handle NWChem-related errors by raising the exception "NWChemError", which may be handled in the standard Python manner (see Section 36.3.9).

- `input_parse(string)` — invokes the standard NWChem input parser with the data in `string` as input. Note that the usual behavior of NWChem will apply — the parser only reads input up to either end of input or until a TASK directive is encountered (the task directive is *not* executed by the parser).
- `task_energy(theory)` — returns the energy as if computed with the NWChem directive `TASK ENERGY <THEORY>`.

- `task_gradient(theory)` — returns a tuple `(energy, gradient)` as if computed with the NWChem directive `TASK GRADIENT <THEORY>`.
- `task_optimize(theory)` — returns a tuple `(energy, gradient)` as if computed with the NWChem directive `TASK OPTIMIZE <THEORY>`. The energy and gradient will be those at the last point in the optimization and consistent with the current geometry in the database.
- `ga_nodeid()` — returns the number of the parallel process.
- `rtdb_print(print_values)` — prints the contents of the RTDB. If `print_values` is 0, only the keys are printed, if it is 1 then the values are also printed.
- `rtdb_put(name, values)` or `rtdb_put(name, values, type)` — puts the values into the database with the given name. In the first form, the type is inferred from the first value, and in the second form the type is specified using the last argument as one of INT, DBL, LOGICAL, or CHAR.
- `rtdb_get(name)` — returns the data from the database associated with the given name.

An example below (Section 36.3.10) explains, in lieu of a Python wrapper for the geometry object, how to obtain the Cartesian molecular coordinates directly from the database.

36.3 Examples

Several examples will provide the best explanation of how the extensions are used, and how Python might prove useful.

36.3.1 Hello world

```
python
  print 'Hello world from process ', ga_nodeid()
end

task python
```

This input prints the traditional greeting from each parallel process.

36.3.2 Scanning a basis exponent

```
geometry units au
  O 0 0 0; H 0 1.430 -1.107; H 0 -1.430 -1.107
end

python
  exponent = 0.1
  while (exponent <= 2.01):
    input_parse(''
      basis noprint
      H library 3-21g; O library 3-21g; O d; %f 1.0
    end
    '' % (exponent))
```

```

    print ' exponent = ', exponent, ' energy = ', task_energy('scf')
    exponent = exponent + 0.1
end

print none

task python

```

This program augments a 3-21g basis for water with a d-function on oxygen and varies the exponent from 0.1 to 2.0 in steps of 0.1, printing the exponent and energy at each step.

The geometry is input as usual, but the basis set input is embedded inside a call to `input_parse()` in the Python program. The standard Python string substitution is used to put the current value of the exponent into the basis set (replacing the `%f`) before being parsed by NWChem. The energy is returned by `task_energy('scf')` and printed out. The `print none` in the NWChem input switches off all NWChem output so all you will see is the output from your Python program.

Note that execution in parallel may produce unwanted output since all process execute the print statement inside the Python program.

Look in the NWChem `contrib` directory for a routine that makes the above task easier.

36.3.3 Scanning a basis exponent revisited.

```

geometry units au
  O 0 0 0; H 0 1.430 -1.107; H 0 -1.430 -1.107
end

print none

python
  if (ga_nodeid() == 0): plotdata = open("plotdata",'w')

  def energy_at_exponent(exponent):
    input_parse('''
      basis noprint
      H library 3-21g; O library 3-21g; O d; %f 1.0
    end
    ''' % (exponent))
    return task_energy('scf')

  exponent = 0.1
  while exponent <= 2.01:
    energy = energy_at_exponent(exponent)
    if (ga_nodeid() == 0):
      print ' exponent = ', exponent, ' energy = ', energy
      plotdata.write('%f %f\n' % (exponent , energy))
      exponent = exponent + 0.1

  if (ga_nodeid() == 0): plotdata.close()
end

```

```
task python
```

This input performs exactly the same calculation as the previous one, but uses a slightly more sophisticated Python program, also writes the data out to a file for easy visualization with a package such as `gnuplot`, and protects write statements to prevent duplicate output in a parallel job. The only significant differences are in the Python program. A file called "plotdata" is opened, and then a procedure is defined which given an exponent returns the energy. Next comes the main loop that scans the exponent through the desired range and prints the results to standard output and to the file. When the loop is finished the additional output file is closed.

36.3.4 Scanning a geometric variable

```
python
  geometry = '''
    geometry noprint; symmetry d2h
      C 0 0 %f; H 0 0.916 1.224
    end
  '''
  x = 0.6
  while (x < 0.721):
    input_parse(geometry % x)
    energy = task_energy('scf')
    print ' x = %5.2f   energy = %10.6f' % (x, energy)
    x = x + 0.01
  end

basis; C library 6-31g*; H library 6-31g*; end

print none

task python
```

This scans the bond length in ethene from 1.2 to 1.44 in steps of 0.2 computing the energy at each geometry. Since it is using D_{2h} symmetry the program actually uses a variable (x) that is half the bond length.

Look in the NWChem `contrib` directory for a routine that makes the above task easier.

36.3.5 Scan using the BSSE counterpoise corrected energy

```
basis spherical
  Ne library cc-pvdz; BqNe library Ne cc-pvdz
  He library cc-pvdz; BqHe library He cc-pvdz
end

mp2; tight; freeze core atomic; end

print none

python noprint
```

```

supermolecule = 'geometry noprint;   Ne 0 0 0;   He 0 0 %f; end\n'
fragment1      = 'geometry noprint;   Ne 0 0 0; BqHe 0 0 %f; end\n'
fragment2      = 'geometry noprint; BqNe 0 0 0;   He 0 0 %f; end\n'

def energy(geometry):
    input_parse(geometry + 'scf; vectors atomic; end\n')
    return task_energy('mp2')

def bsse_energy(z):
    return energy(supermolecule % z) - \
           energy(fragment1 % z) - \
           energy(fragment2 % z)

z = 3.3
while (z < 4.301):
    e = bsse_energy(z)
    if (ga_nodeid() == 0):
        print ' z = %5.2f   energy = %10.7f ' % (z, e)
        z = z + 0.1
end

task python

```

This example scans the He—Ne bond-length from 3.3 to 4.3 and prints out the BSSE counterpoise corrected MP2 energy.

The basis set is specified as usual, noting that we will need functions on ghost centers to do the counterpoise correction. The Python program commences by defining strings containing the geometry of the super-molecule and two fragments, each having one variable to be substituted. Next, a function is defined to compute the energy given a geometry, and then a function is defined to compute the counterpoise corrected energy at a given bond length. Finally, the bond length is scanned and the energy printed. When computing the energy, the atomic guess has to be forced in the SCF since by default it will attempt to use orbitals from the previous calculation which is not appropriate here.

Since the counterpoise corrected energy is a linear combination of other standard energies, it is possible to compute the analytic derivatives term by term. Thus, combining this example and the next could yield the foundation of a BSSE corrected geometry optimization package.

36.3.6 Scan the geometry and compute the energy and gradient

```

basis noprint; H library sto-3g; O library sto-3g; end

python noprint
print '   y       z       energy                gradient'
print ' -----'
y = 1.2
while y <= 1.61:
    z = 1.0
    while z <= 1.21:
        input_parse(''
            geometry noprint units atomic
            O 0   0   0
            H 0   %f  -%f

```

```

        H 0 -%f -%f
    end
    ''' % (y, z, y, z))

    (energy,gradient) = task_gradient('scf')

    print ' %5.2f %5.2f %9.6f' % (y, z, energy),
    i = 0
    while (i < len(gradient)):
        print '%5.2f' % gradient[i],
        i = i + 1
    print ''
    z = z + 0.1
    y = y + 0.1
end

print none

task python

```

This program illustrates evaluating the energy and gradient by calling `task_gradient()`. A water molecule is scanned through several C_{2v} geometries by varying the y and z coordinates of the two hydrogen atoms. At each geometry the coordinates, energy and gradient are printed.

The basis set (sto-3g) is input as usual. The two while loops vary the y and z coordinates. These are then substituted into a geometry which is parsed by NWChem using `input_parse()`. The energy and gradient are then evaluated by calling `task_gradient()` which returns a tuple containing the energy (a scalar) and the gradient (a vector or list). These are printed out exploiting the Python convention that a print statement ending in a comma does not print end-of-line.

36.3.7 Reaction energies varying the basis set

```

mp2; freeze atomic; end

print none

python
    energies = {}
    c2h4 = 'geometry noprint; symmetry d2h; \
          C 0 0 0.672; H 0 0.935 1.238; end\n'
    ch4  = 'geometry noprint; symmetry td; \
          C 0 0 0; H 0.634 0.634 0.634; end\n'
    h2   = 'geometry noprint; H 0 0 0.378; H 0 0 -0.378; end\n'

    def energy(basis, geometry):
        input_parse('''
            basis spherical noprint
            c library %s ; h library %s
            end
            ''' % (basis, basis))

```

```

input_parse(geometry)
return task_energy('mp2')

for basis in ('sto-3g', '6-31g', '6-31g*', 'cc-pvdz', 'cc-pvtz'):
    energies[basis] = 2*energy(basis, ch4) \
        - 2*energy(basis, h2) - energy(basis, c2h4)
    if (ga_nodeid() == 0): print basis, ' %8.6f' % energies[basis]
end

task python

```

In this example the reaction energy for $2H_2 + C_2H_4 \rightarrow 2CH_4$ is evaluated using MP2 in several basis sets. The geometries are fixed, but could be re-optimized in each basis. To illustrate the useful associative arrays in Python, the reaction energies are put into the associative array `energies` — note its declaration at the top of the program.

36.3.8 Using the database

```

python
    rtdb_put("test_int2", 22)
    rtdb_put("test_int", [22, 10, 3], INT)
    rtdb_put("test_dbl", [22.9, 12.4, 23.908], DBL)
    rtdb_put("test_str", "hello", CHAR)
    rtdb_put("test_logic", [0,1,0,1,0,1], LOGICAL)
    rtdb_put("test_logic2", 0, LOGICAL)

    rtdb_print(1)

    print "test_str      = ", rtdb_get("test_str")
    print "test_int      = ", rtdb_get("test_int")
    print "test_int2     = ", rtdb_get("test_int2")
    print "test_dbl      = ", rtdb_get("test_dbl")
    print "test_logic    = ", rtdb_get("test_logic")
    print "test_logic2   = ", rtdb_get("test_logic2")
end

task python

```

This example illustrates how to access the database from Python.

36.3.9 Handling exceptions from NWChem

```

geometry; he 0 0 0; he 0 0 2; end
basis; he library 3-21g; end
scf; maxiter 1; end

python
    try:
        task_energy('scf')
    except NWChemError, message:

```

```

    print 'Error from NWChem ... ', message
end

task python

```

The above test program shows how to handle exceptions generated by NWChem by forcing an SCF calculation on He_2 to fail due to insufficient iterations.

If an NWChem command fails it will raise the exception "NWChemError" (case sensitive) unless the error was fatal. If the exception is not caught, then it will cause the entire Python program to terminate with an error. This Python program catches the exception, prints out the message, and then continues as if all was well since the exception has been handled.

If your Python program detects an error, raise an unhandled exception. Do not call `exit(1)` since this may circumvent necessary clean-up of the NWChem execution environment.

36.3.10 Accessing geometry information — a temporary hack

In an ideal world the geometry and basis set objects would have full Python wrappers, but until then a back-door solution will have to suffice. We've already seen how to use `input_parse()` to put geometry (and basis) data into NWChem, so it only remains to get the geometry data back after it has been updated by a geometry optimization or some other operation.

The following Python procedure retrieves the coordinates in the same units as initially input for a geometry of a given name. Its full source is included in the NWChem `contrib` directory.

```

def geom_get_coords(name):
    try:
        actualname = rtdb_get(name)
    except NWChemError:
        actualname = name
    coords = rtdb_get('geometry:' + actualname + ':coords')
    units = rtdb_get('geometry:' + actualname + ':user units')
    if (units == 'a.u.'):
        factor = 1.0
    elif (units == 'angstroms'):
        factor = rtdb_get('geometry:'+actualname+':angstrom_to_au')
    else:
        raise NWChemError,'unknown units'
    i = 0
    while (i < len(coords)):
        coords[i] = coords[i] / factor
        i = i + 1
    return coords

```

A geometry (see Section 6) with name NAME has its coordinates (in atomic units) stored in the database entry `geometry:NAME:coords`. A minor wrinkle here is that indirection is possible (and used by the optimizers) so that we must first check if NAME actually points to another name. In the program this is done in the `try...except` sequence. With the actual name of the geometry, we can get the coordinates. Any exceptions are passed up to the caller. The rest of the code is just to convert back into the initial input units — only atomic units or Angströms are handled in this simple example. Returned is a list of the atomic coordinates in the same units as your initial input.

The routine is used as follows

```
coords = geom_get_coords('geometry')
```

or, if you want better error handling

```
try:
    coords = geom_get_coords('geometry')
except NWChemError,message:
    print 'Coordinates for geometry not found ', message
else:
    print coords
```

This is very dirty and definitely not supported from one release to another, but, browsing the output of `rtodb_print()` at the end of a calculation is a good way to find stuff. To be on safer ground, look in the programmers manual since some of the high-level routines do pass data via the database in a well-defined and supported manner. *Be warned* — you must be very careful if you try to modify data in the database. The input parser does many important things that are not immediately apparent (e.g., ensure the geometry is consistent with the point group, mark the SCF as not converged if the SCF options are changed, ...). Where at all possible your Python program should generate standard NWChem input and pass it to `input_parse()` rather than setting parameters directly in the database.

36.3.11 Scanning a basis exponent yet again — plotting and handling child processes

```
geometry units au
  O 0 0 0; H 0 1.430 -1.107; H 0 -1.430 -1.107
end

print none

python
  import Gnuplot, time, signal

  def energy_at_exponent(exponent):
    input_parse('''
      basis noprint
      H library 3-21g; O library 3-21g; O d; %f 1.0
    end
    ''' % (exponent))
    return task_energy('scf')

  data = []
  exponent = 0.5
  while exponent <= 0.6:
    energy = energy_at_exponent(exponent)
    print ' exponent = ', exponent, ' energy = ', energy
    data = data + [[exponent,energy]]
    exponent = exponent + 0.02

  if (ga_nodeid() == 0):
    signal.signal(signal.SIGCHLD, signal.SIG_DFL)
```

```
g = Gnuplot.Gnuplot()
g('set data style linespoints')
g.plot(data)
time.sleep(30) # 30s to look at the plot

end

task python
```

This illustrates how to handle signals from terminating child processes and how to generate simple plots on UNIX systems. The example from Section 36.3.3 is modified so that instead of writing the data to a file for subsequent visualization, it is saved for subsequent visualization with Gnuplot (you'll need both Gnuplot and the corresponding package for Python in your PYTHONPATH. Look at <http://monsoon.harvard.edu/mhagger/download>).

The issue is that NWChem traps various signals from the O/S that usually indicate bad news in order to provide better error handling and reliable clean-up of shared, parallel resources. One of these signals is SIGCHLD which is generated whenever a child process terminates. If you want to create child processes within Python, then the NWChem handler for SIGCHLD must be replaced with the default handler. There seems to be no easy way to restore the NWChem handler after the child has completed, but this should have no serious side effect.

36.4 Troubleshooting

Common problems with Python programs inside NWChem.

1. You get the message

```
0:python_input: indentation must be >= that of first line: 4
```

This indicates that NWChem thinks that a line is less indented than the first line. If this is not the case then perhaps there is a tab in your input which NWChem treats as a single space character but appears to you as more spaces. Try running `untabify` in Emacs. The problem could also be the `END` directive that terminates the `PYTHON` compound directive — since Python also has an `end` statement. To avoid confusion the `END` directive for NWChem *must* be at the start of the line.

2. Your program hangs or deadlocks — most likely you have a piece of code that is restricted to executing on a subset of the processors (perhaps just node 0) but is calling (perhaps indirectly) a function that must execute on all nodes.

Chapter 37

Interfaces to Other Programs

NWChem has interfaces to several different packages which are listed below. In general, the NWChem authors work with the authors of the other packages to make sure that the interface works. However, any problems with the interface should be reported to the `nwchem-users@emsl.pnl.gov` e-mail list.

37.1 NBO — Natural Bond Orbital Analysis

```
NBO
...
END
```

This directive is used to run the NBO package within NWChem. (To only print out an input file for NBO, see Section 28.1.1.) The current version of NBO in NWChem is version 5.0.

Inside the NBO block are the typical commands that would be needed for NBO (Please see an NBO user's manual for more information.). The following directive is needed to execute NBO.

```
task nbo
```

As an example:

```
title "Methylamine...rhf/3-21g//Pople-Gordon standard geometry"

start methylamine

echo

memory 8 mw

basis
  C library 3-21g
  N library 3-21g
  H library 3-21g
end
```

```
geometry
C          .052902      .711852      .000000
N          .052902     -.758148      .000000
H         -.974760      1.075185      .000000
H          .566733      1.075185      .889981
H          .566733      1.075185     -.889981
H         -.423217     -1.094815      .824662
H         -.423217     -1.094815     -.824662
symmetry c1
end

task SCF energy

nbo
$nbo cmo $end
end

task nbo
```

37.2 DIRDYVTST — **DI**Rect **D**ynamics for **V**ariational **T**ransition **S**tate **T**heory

by Bruce C. Garrett,
Environmental Molecular Sciences Laboratory,
Pacific Northwest Laboratory, Richland, Washington

Yao-Yuan Chuang and Donald G. Truhlar,
Department of Chemistry and Super Computer Institute,
University of Minnesota, MN 55455-0431

and interfaced to NWChem by

Ricky A. Kendall,
Scalable Computing Laboratory,
Ames Laboratory and Iowa State University, Ames, IA 50011

Theresa L. Windus,
Environmental Molecular Sciences Laboratory,
Pacific Northwest Laboratory, Richland, Washington

If you use the DIRDYVTST portion of NWChem, please use following citation in addition to the usual NWChem citation from Section 1:

DIRDYVTST, Yao-Yuan Chuang and Donald G. Truhlar, Department of Chemistry and Super Computer Institute, University of Minnesota; Ricky A. Kendall, Scalable Computing Laboratory, Ames Laboratory and Iowa State University; Bruce C. Garrett and Theresa L. Windus, Environmental Molecular Sciences Laboratory, Pacific Northwest Laboratory.

37.2.1 Introduction

By using DIRDYVTST, a user can carry out electronic structure calculations with NWChem and use the resulting energies, gradients, and Hessians for direct dynamics calculations with POLYRATE. This program prepares the file30 input for POLYRATE from NWChem electronic structure calculations of energies, gradients and Hessians at the reactant, product, and saddle point geometries and along the minimum energy path. Cartesian geometries for the reactants, products, and saddle points need to be input to this program; optimization of geometries is not performed in this program. Note that DIRDYVTST is based on the DIRDYGAUSS program and is similar to two other programs: DDUTILITIES and GAUSSRATE. Users of this module are encouraged to read the POLYRATE manual since they will need to create the file fu5 input to run calculations with POLYRATE.

Notes about the code:

Input. The code has been written to parallel, as much as possible, the POLYRATE code.

Output. There is one default output file for each DIRDYVTST run - .file30.

Integrators for following the reaction path. Currently the Euler and three Page-McIver (PM) methods are implemented. The PM methods are the local quadratic approximation (LQA), the corrected LQA (CLQA), and the cubic (CUBE) algorithm. The PM methods are implemented so that the Hessian can be reused at intermediate steps at which only the gradient is updated.

37.2.2 Files

Test runs are located in directories in $\$NWCHEM_TOP/QA/tests$. Test runs are available for two systems: $H + H_2$ and $OH + H_2$.

The $H + H_2$ test uses the Euler integration method at the SCF/3-21G level of theory to calculate points along the reaction path. This test is located in the $\$NWCHEM_TOP/QA/tests/h3tr1$ directory.

The $OH + H_2$ test uses the Page-McIver CUBE algorithm to calculate points on the SCF/3-21G surface and does additional single point calculations at the SCF/6-31G* level of theory. This test is located in the $\$NWCHEM_TOP/QA/tests/oh3tr3$ directory.

Note: These tests are set up with SCF, however, other levels of theory can be used. The initial hessian calculations at the reactants, products and saddle point can cause some problems when numerical Hessians are required (especially when there is symmetry breaking in the wavefunction).

37.2.3 Detailed description of the input

The input consists of keywords for NWChem and keywords related to POLYRATE input. The first set of inputs are for NWChem with the general input block of the form:

```
DIRDYVTST [autosym [real tol default 1d-2] | noautosym]
  [THEORY <string theory> [basis <string basis default "ao basis">] \
    [ecp <string ecp>] [input <string input>]]
  [SPTHEORY <string theory> [basis <string basis default "ao basis">] \
    [ecp <string ecp>] [input <string input>]]
  ...
END
```

Use of symmetry

The use of symmetry in the calculation is controlled by the keyword `autosym` | `noautosym` which is used as described in the geometry directive (see Section 6). `Autosym` is on by default. A couple words of warning here. The tolerance related to `autosym` can cause problems when taking the initial step off of the transition state. If the tolerance is too large and the initial step relatively small, the resulting geometry will be close to a higher symmetry than is really wanted and the molecule will be symmetrized into the higher symmetry. To check this, the code prints out the symmetry at each geometry along the path. It is up to the user to check the symmetry and make sure that it is the required one. In preverse cases, the user may need to turn `autosym` off (`noautosym`) if changing the tolerance doesn't produce the desired results. In the case that `autosym` is used, the user does not need to worry about the different alignment of the molecule between NWChem and POLYRATE, this is taken care of internally in the DIRDYVTST module.

Basis specification

The basis name on the theory or sptheory directive is that specified on a basis set directive (see Section 7) and *not* the name of a standard basis in the library. If not specified, the basis set for the sptheory defaults to the theory basis which defaults to "ao basis".

Effective core potentials

If an effective core potential is specified in the usual fashion (see Section 8) outside of the DIRDYVTST input then this will be used in all calculations. If an alternative ECP name (the name specified on the ECP directive in the same manner as done for basis sets) is specified on one of the theory directives, then this ECP will be used in preference for that level of theory.

General input strings

For many purposes, the ability to specify the theory, basis and effective core potential is adequate. All of the options for each theory are determined from their independent input blocks. However, if the same theory (e.g., DFT) is to be used with different options for theory and sptheory, then the general input strings must be used. These strings are processed as NWChem input each time the theoretical calculation is invoked. The strings may contain any NWChem input, except for options pertaining to DIRDYVTST and the task directive. The intent is that the strings be used just to control the options pertaining to the theory being used.

A word of caution. Be sure to check that the options are producing the desired results. Since the NWChem database is persistent, the input strings should fully define the calculation you wish to have happen.

For instance, if the theory model is DFT/LDA/3-21g and the sptheory model is DFT/B3LYP/6-311g**, the DIRDYVTST input might look like this

```
dirdyvtst
  theory dft basis 3-21g   input "dft\; xc\; end"
  sptheory dft basis 6-311g** input "dft\; xc b3lyp\; end"
  . . . .
end
```

The empty XC directive restores the default LDA exchange-correlation option (see Section 11.3). Note that semicolons and other quotation marks inside the input string must be preceded by a backslash to avoid special interpretation.

POLYRATE related options

These keyword options are similar to the POLYRATE input format, except there are no ENERGETICS, OPTIMIZATION, SECOND, TUNNELING, and RATE sections.

GENERAL section The GENERAL section has the following format:

```
*GENERAL
  [TITLE <string title>]
  ATOMS
    <integer num> <string tag> [<real mass>]
    ...
  END
  [SINGLEPOINT]
  [SAVEFILE (vecs || hess || spc)]
```

Descriptions

TITLE is a keyword that allows the user to input a description of the calculation. In this version, the user can only have a single-line description.

For example: TITLE Calculation of D + HCl reaction

ATOMS is a list keyword that is used to input a list of the atoms. It is similar to POLYRATE in that the order of the atom and the atomic symbol are required in a single line. If isotope of the element is considered then the atomic mass is required in units of amu.

For example:

```
          ATOMS
            1      H      2.014
            2      H
            3      Cl
          END
```

SINGLEPOINT is a keyword that specifies that a single point calculation is to be performed at the reactants, products and saddle point geometries. The type of single point calculation is specified in the `sptheory` line.

SAVEFILE is a keyword that specifies that NWChem files are to be saved. Allowed values of variable input to SAVEFILE are `vecs`, `hess`, and `spc` for saving the files base theory movecs, base theory hessian and singlepoint calculation movecs.

REACT1, REACT2, PROD1, PROD2, and START sections These sections have the following format:

```
*(REACT1 || REACT2 || PROD1 || PROD2 || START)
  GEOM
    <integer num> <real x y z>
    ...
  END
  SPECIES (ATOMIC || LINRP || NONLINRP || LINTS || NONLINTS default NONLINRP)
```

REACT1 and REACT2 are input for each of the reactants and PROD1 and PROD2 are input for each of the products. REACT1 and PROD1 are required. START is the input for the transition state if one exists, or starting point to follow downhill the MEP.

Descriptions

GEOM is a list keyword that indicates the geometry of the molecule in Cartesian coordinates with atomic unit.

For example:

```

GEOM
  1      0.0    0.0    0.0
  2      0.0    0.0    1.5
END

```

SPECIES is a variable keyword that indicates the type of the molecule. Options are: ATOMIC (atomic reactant or product), LINRP (linear reactant or product), NONLINRP (nonlinear reactant or product), LINTS (linear transition state), and NONLINTS (nonlinear transition state).

For example: SPECIES atomic

PATH section The Path section has the format:

*PATH

```

[SCALEMASS <real scalemass default 1.0>]
[SSTEP <real sstep default 0.01>]
[SSAVE <real ssave default 0.1>]
[SHESS <real shess default SSAVE>]
[SLP <real slp default 1.0>]
[SLM <real slm default -1.0>]
[SIGN (REACTANT || PRODUCT default REACTANT)]
[INTEGRA (EULER || LQA || CLQA || CUBE default EULER)]
[PRINTFREQ (on || off default off)]

```

Descriptions

SCALEMASS is a variable keyword that indicates the arbitrary mass (in amu) used for mass-scaled Cartesian coordinates. This is the variable called mu in published papers. Normally, this is taken as either 1.0 amu or, for bimolecular reactions, as the reduced mass of relative translation of the reactants.

SSTEP is a variable keyword that indicates the numerical step size (in bohrs) for the gradient grid. This is the step size for following the minimum energy path.

SSAVE is a variable keyword that indicates the numerical step size (in bohrs) for saving the Hessian grid. At each save point the potential and its first and second derivatives are recalculated and written to the .file30 file. For example, if SSTEP=0.01 and SSAVE=0.1, then the potential information is written to .file30 every 10 steps along the gradient grid.

SHESS is a variable keyword that indicates the numerical step size (in bohrs) for recomputing the Hessian when using a Page-McIver integrator (e.g., LQA, CLQA, or CUBE). For Euler integration SHESS = SSAVE. For intermediate points along the gradient grid, the Hessian matrix from the last Hessian calculation is reused. For example, if SSTEP=0.01 and SHESS=0.05, then the Hessian matrix is recomputed every 5 steps along the gradient grid.

SLP is a variable keyword that indicates the positive limit of the reaction coordinate (in bohrs).

SLM is a variable keyword that indicates the negative limit of the reaction coordinate (in bohrs).

SIGN is a variable keyword used to ensure the conventional definition of the sign of s , $s < 0$ for the reactant side and $s > 0$ for the product side, is followed. PRODUCT should be used if the eigenvector at the saddle point points toward the product side and REACTANT if the eigenvector points toward the reactant side.

INTEGRA is a variable keyword that indicates the integration method used to follow the reaction path. Options are: EULER, LQA, CLQA, and CUBE.

PRINTFREQ is a variable keyword that indicates that projected frequencies and eigenvectors will be printed along the MEP.

Restart

DIRDYVTST calculations should be restarted through the normal NWChem mechanism (See Section 5.1). The user needs to change the `start` directive to a `restart` directive and get rid of any information that will overwrite important information in the RTDB. The `file.db` and `file.file30` need to be available for the calculation to restart properly.

Example

This is an example that creates the file30 file for POLYRATE for $H + H_2$. Note that the multiplicity is that of the entire supermolecule, a doublet. In this example, the initial energies, gradients, and Hessians are calculated at the UHF/3-21G level of theory and the singlepoint calculations are calculated at the MP2/cc-pVDZ level of theory with a tighter convergence threshold than the first SCF.

```
start h3test

basis
  h library 3-21G
end

basis singlepoint
  h library cc-pVDZ
end

scf
  uhf
  doublet
  thresh 1.0e-6
end

dirtyvtst autosym 0.001
  theory scf input "scf; uhf; doublet; thresh 1.0e-06; end"
  spttheory mp2 basis singlepoint input \
    "scf; uhf; doublet; thresh 1.0e-07; end"
*GENERAL
  TITLE
    Test run: H+H2 reaction, Page-McIver CLQA algorithm, no restart

  ATOMS
```

```
    1    H
    2    H
    3    H
END

SINGLEPOINT

*REACT1
  GEOM
    1  0.0  0.0  0.0
    2  0.0  0.0  1.3886144
  END

  SPECIES  LINRP

*REACT2
  GEOM
    3  0.0  0.0  190.3612132
  END

  SPECIES  ATOMIC

*PROD2
  GEOM
    1  0.0  0.0  190.3612132
  END

  SPECIES  ATOMIC

*PROD1
  GEOM
    2  0.0  0.0  1.3886144
    3  0.0  0.0  0.0
  END

  SPECIES  LINRP

*START
  GEOM
    1  0.0  0.0  -1.76531973
    2  0.0  0.0  0.0
    3  0.0  0.0  1.76531973
  END

  SPECIES  LINTS

*PATH
  SSTEP  0.05
  SSAVE  0.05
```

```
SLP      0.50
SLM      -0.50
SCALEMASS 0.6718993
INTEGRA  CLQA
end

task dirtyvtst
```


Chapter 38

Acknowledgments

This work was supported by funds from the Environmental and Molecular Sciences Laboratory Construction Project at Pacific Northwest National Laboratory. Development of some of the parallel programming tools and algorithms employed by NWChem was performed under the auspices of the High Performance Computing and Communications Program of the Mathematical, Information, and Computational Sciences Division, U.S. Department of Energy. Pacific Northwest National Laboratory is operated by Battelle Memorial Institute for the U.S. Department of Energy under Contract DE-AC06-76RLO 1830.

Appendix A

Standard Basis Sets

Basis sets and effective core potentials were obtained (1/1/2002) from the Extensible Computational Chemistry Environment (ECCE) Basis Set Database, as developed and distributed by the Molecular Science Computing Facility, Environmental and Molecular Sciences Laboratory which is part of the Pacific Northwest National Laboratory, P.O. Box 999, Richland, Washington 99352, USA, and is funded by the U.S. Department of Energy. The Pacific Northwest National Laboratory is a multi-program laboratory operated by Battelle Memorial Institute for the U.S. Department of Energy under contract DE-AC06-76RLO. Contact David Feller (df_feller@pnl.gov) or Deborah Gracio (gracio@pnl.gov) for further information.

The names in the NWChem library are consistent with those in the ECCE database and thus may include spaces. The standard NWChem input routines require that strings including spaces are enclosed in quotation marks (" . . . ") or that blanks are escaped with a backslash. As a convenience, basis set names may also have the blanks replaced with underscores. Thus, the following all yield the same basis set for oxygen

```
oxygen library "DZP + Diffuse (Dunning)"
oxygen library DZP\ +\ Diffuse\ (Dunning)
oxygen library DZP+_Diffuse_(Dunning)
```

Case may be ignored when specifying basis set names, but otherwise names should be specified exactly as provided below. A good method is just to cut/paste from the WWW pages since they were generated electronically from the library source.

Errors found in the basis set library of NWChem version 4.4 have been corrected in the current library of NWChem version 4.6. The changes are listed below:

1. Basis Set "Stuttgart RLC ECP"
Element: I
Effective core potential has been changed and now conforms to the published from by Bergner et al. (Mol. Phys. 1993).

Errors found in the basis set library of NWChem version 4.0 have been corrected in the current library of NWChem version 4.1. The changes are listed below:

1. Basis Set "aug-cc-pVTZ"
Element: Br
Correction of the diffuse s-exponent from 0.041601 to 0.045593.

2. Basis Set "cc-pV5Z"
Element: Ar
Questionable contraction coefficients were replaced with HF AO's.
3. Basis Set "Stuttgart RSC 1997 ECP"
Element: Ba
This basis was called Stuttgart RSC ECP in version 4.0. The 4th s- and p-primitive are now uncontracted.
4. ECP "LANL2DZ ECP"
Element: K and Ca
The ECP was defined with a core of 18 electrons, inconsistent with the other elements in this series (3d has 10 electron core). The ECP has been replaced with the Hay-Wadt ECP used in Hay-Wadt VDZ.

IMPORTANT NOTE: The Stuttgart basis set and ECP developers have recently developed a new set of Stuttgart RSC ECP's for the Lanthanide series of elements. This new set is now called "Stuttgart RSC ECP", while the previous, and most widely used, set has been renamed to "Stuttgart RSC 1997 ECP" in both the EMSL Basis Set library and in the NWChem basis set library.

Relativistic contractions of standard basis sets for use in the Douglas-Kroll and Dyall-modified-Dirac method have also been included in the library. These are identified by tags following the standard basis set name.

For the Douglas-Kroll method the tag `dk` should be specified:

```
oxygen library cc-pVDZ_DK
```

For the Dyall-modified-Dirac (DMD) method three tags should be specified. The first is a tag for the nuclear model, which can be `pt` or `fi` for a point or a finite Gaussian model (see Section 6). The second is a tag for the relativistic Hamiltonian: `sf` is for the spin-free modified Dirac Hamiltonian. The third tags the component type, `fw` for the atomic FW transformed large component, `lc` for the large component and `sc` for the small component:

```
oxygen library cc-pvdz_pt_sf_fw
oxygen library cc-pvdz_pt_sf_lc
```

Basis sets which are available with either the DmD or DK contractions are indicated in the list below. Additional DmD basis sets are available in the directory `nwchem/contrib/basissets/Dyall_DMD`.

Here is a list of known all-electron non-relativistic, DK and DmD basis sets, effective core potentials with their respective basis sets, fitting basis sets, and Polarization, Diffuse and Core-valence sets of functions, along with the elements included for each. Additional information about each basis set in the NWChem library can be obtained from the online EMSL Gaussian Basis Set library.

Standard all-electron basis sets:

1. Basis Set "STO-2G" (number of atoms 21)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sr
2. Basis Set "STO-3G" (number of atoms 53)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I
3. Basis Set "STO-6G" (number of atoms 36)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr

4. Basis Set "STO-3G*" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
5. Basis Set "3-21G" (number of atoms 55)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs
6. Basis Set "3-21++G" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
7. Basis Set "3-21G*" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
8. Basis Set "3-21++G*" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
9. Basis Set "3-21GSP" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
10. Basis Set "4-22GSP" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
11. Basis Set "4-31G" (number of atoms 13)
H He Li Be B C N O F Ne P S Cl
12. Basis Set "6-31G" (number of atoms 30)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn
13. Basis Set "6-31G*" (number of atoms 30)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn
14. Basis Set "6-31G**" (number of atoms 30)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn
15. Basis Set "6-31++G" (number of atoms 20)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca
16. Basis Set "6-31++G*" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
17. Basis Set "6-31++G**" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
18. Basis Set "6-31+G*" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
19. Basis Set "6-31G(3df, 3pd)" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
20. Basis Set "6-31G-Blaudeau" (number of atoms 2)
K Ca
21. Basis Set "6-31G*-Blaudeau" (number of atoms 2)
K Ca
22. Basis Set "6-311G" (number of atoms 25)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr I

23. Basis Set "6-311G*" (number of atoms 25)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr I
24. Basis Set "6-311G**" (number of atoms 25)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr I
25. Basis Set "6-311++G**" (number of atoms 10)
H He Li Be B C N O F Ne
26. Basis Set "6-311++G(2d, 2p)" (number of atoms 20)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca
27. Basis Set "6-311G(2df, 2pd)" (number of atoms 12)
H He Li Be B C N O F Ne K Ca
28. Basis Set "6-311+G*" (number of atoms 10)
H He Li Be B C N O F Ne
29. Basis Set "6-311++G(3df, 3pd)" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
30. Basis Set "MINI (Huzinaga)" (number of atoms 20)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca
31. Basis Set "MINI (Scaled)" (number of atoms 20)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca
32. Basis Set "MIDI (Huzinaga)" (number of atoms 19)
H He Li Be B C N O F Ne Na Al Si P S Cl Ar K Cs
33. Basis Set "MIDI!" (number of atoms 11)
H C N O F Si P S Cl Br I
34. Basis Set "SV (Dunning-Hay)" (number of atoms 9)
H Li Be B C N O F Ne
35. Basis Set "SVP (Dunning-Hay)" (number of atoms 9)
H Li Be B C N O F Ne
36. Basis Set "SVP + Diffuse (Dunning-Hay)" (number of atoms 9)
H Li Be B C N O F Ne
37. Basis Set "TZ (Dunning)" (number of atoms 8)
H Li B C N O F Ne
38. Basis Set "Chipman DZP + Diffuse" (number of atoms 6)
H B C N O F
39. Basis Set "DZ (Dunning)" (number of atoms 13)
H Li B C N O F Ne Al Si P S Cl
40. Basis Set "DZP (Dunning)" (number of atoms 13)
H Li B C N O F Ne Al Si P S Cl
41. Basis Set "DZP + Diffuse (Dunning)" (number of atoms 7)
H B C N O F Ne

42. Basis Set "cc-pVDZ" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
43. Basis Set "cc-pVTZ" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
44. Basis Set "cc-pVQZ" (number of atoms 22)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se
45. Basis Set "cc-pV5Z" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
46. Basis Set "cc-pV6Z" (number of atoms 14)
H He B C N O F Ne Al Si P S Cl Ar
47. Basis Set "pV6Z" (number of atoms 4)
H C N O
48. Basis Set "cc-pVDZ(seg-opt)" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
49. Basis Set "cc-pVTZ(seg-opt)" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
50. Basis Set "cc-pVQZ(seg-opt)" (number of atoms 4)
H He C O
51. Basis Set "cc-pCVDZ" (number of atoms 12)
Li Be B C N O F Ne Na Mg Si Cl
52. Basis Set "cc-pCVTZ" (number of atoms 13)
Li Be B C N O F Ne Na Mg Si P Cl
53. Basis Set "cc-pCVQZ" (number of atoms 13)
Li Be B C N O F Ne Na Mg Si P Cl
54. Basis Set "cc-pCV5Z" (number of atoms 8)
B C N O F Ne Na Mg
55. Basis Set "cc-pCV6Z" (number of atoms 1)
O
56. Basis Set "aug-cc-pVDZ" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
57. Basis Set "aug-cc-pVTZ" (number of atoms 23)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br
58. Basis Set "aug-cc-pVQZ" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
59. Basis Set "aug-cc-pV5Z" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
60. Basis Set "aug-cc-pV6Z" (number of atoms 14)
H He B C N O F Ne Al Si P S Cl Ar

61. Basis Set "aug-cc-pCVDZ" (number of atoms 7)
B C N O F Na Mg
62. Basis Set "aug-cc-pCVTZ" (number of atoms 8)
B C N O F Ne Na Mg
63. Basis Set "aug-cc-pCVQZ" (number of atoms 6)
B C N O F Ne
64. Basis Set "aug-cc-pCV5Z" (number of atoms 5)
B C N O F
65. Basis Set "d-aug-cc-pVDZ" (number of atoms 8)
H He B C N O F Ne
66. Basis Set "d-aug-cc-pVTZ" (number of atoms 8)
H He B C N O F Ne
67. Basis Set "d-aug-cc-pVQZ" (number of atoms 8)
H He B C N O F Ne
68. Basis Set "d-aug-cc-pV5Z" (number of atoms 8)
H He B C N O F Ne
69. Basis Set "d-aug-cc-pV6Z" (number of atoms 5)
H B C N O
70. Basis Set "cc-pV(D+d)Z" (number of atoms 6)
Al Si P S Cl Ar
71. Basis Set "cc-pV(T+d)Z" (number of atoms 6)
Al Si P S Cl Ar
72. Basis Set "cc-pV(Q+d)Z" (number of atoms 6)
Al Si P S Cl Ar
73. Basis Set "cc-pV(5+d)Z" (number of atoms 6)
Al Si P S Cl Ar
74. Basis Set "cc-pV(6+d)Z" (number of atoms 6)
Al Si P S Cl Ar
75. Basis Set "aug-cc-pV(D+d)Z" (number of atoms 6)
Al Si P S Cl Ar
76. Basis Set "aug-cc-pV(T+d)Z" (number of atoms 6)
Al Si P S Cl Ar
77. Basis Set "aug-cc-pV(Q+d)Z" (number of atoms 6)
Al Si P S Cl Ar
78. Basis Set "aug-cc-pV(5+d)Z" (number of atoms 6)
Al Si P S Cl Ar
79. Basis Set "aug-cc-pV(6+d)Z" (number of atoms 6)
Al Si P S Cl Ar

80. Basis Set "Feller Misc. CVDZ" (number of atoms 1)
K
81. Basis Set "Feller Misc. CVTZ" (number of atoms 1)
K
82. Basis Set "Feller Misc. CVQZ" (number of atoms 1)
K
83. Basis Set "GAMESS VTZ" (number of atoms 16)
H Be B C N O F Ne Na Mg Al Si P S Cl Ar
84. Basis Set "GAMESS PVTZ" (number of atoms 8)
H Be B C N O F Ne
85. Basis Set "Partridge Uncontr. 1" (number of atoms 36)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb Sr
86. Basis Set "Partridge Uncontr. 2" (number of atoms 34)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr
87. Basis Set "Partridge Uncontr. 3" (number of atoms 28)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn
88. Basis Set "Ahlrichs VDZ" (number of atoms 36)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr
89. Basis Set "Ahlrichs pVDZ" (number of atoms 36)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr
90. Basis Set "Ahlrichs VTZ" (number of atoms 36)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr
91. Basis Set "Ahlrichs TZV" (number of atoms 34)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr
92. Basis Set "Binning/Curtiss SV" (number of atoms 6)
Ga Ge As Se Br Kr
93. Basis Set "Binning/Curtiss VTZ" (number of atoms 6)
Ga Ge As Se Br Kr
94. Basis Set "Binning/Curtiss SVP" (number of atoms 6)
Ga Ge As Se Br Kr
95. Basis Set "Binning/Curtiss VTZP" (number of atoms 6)
Ga Ge As Se Br Kr
96. Basis Set "McLean/Chandler VTZ" (number of atoms 8)
Na Mg Al Si P S Cl Ar
97. Basis Set "SV + Rydberg (Dunning-Hay)" (number of atoms 9)
H Li Be B C N O F Ne
98. Basis Set "SVP + Rydberg (Dunning-Hay)" (number of atoms 9)
H Li Be B C N O F Ne

99. Basis Set "SVP + Diffuse + Rydberg" (number of atoms 9)
H Li Be B C N O F Ne
100. Basis Set "DZ + Rydberg (Dunning)" (number of atoms 12)
H B C N O F Ne Al Si P S Cl
101. Basis Set "DZP + Rydberg (Dunning)" (number of atoms 12)
H B C N O F Ne Al Si P S Cl
102. Basis Set "DZ + Double Rydberg (Dunning-Hay)" (number of atoms 12)
H B C N O F Ne Al Si P S Cl
103. Basis Set "SV + Double Rydberg (Dunning-Hay)" (number of atoms 9)
H Li Be B C N O F Ne
104. Basis Set "wachters+f" (number of atoms 9)
Sc Ti V Cr Mn Fe Co Ni Cu
105. Basis Set "Bauschlicher ANO" (number of atoms 9)
Sc Ti V Cr Mn Fe Co Ni Cu
106. Basis Set "NASA Ames ANO" (number of atoms 12)
H B C N O F Ne Al P Ti Fe Ni
107. Basis Set "NASA Ames cc-pVTZ" (number of atoms 1)
Ti
108. Basis Set "NASA Ames cc-pVQZ" (number of atoms 1)
Ti
109. Basis Set "NASA Ames cc-pV5Z" (number of atoms 1)
Ti
110. Basis Set "NASA Ames cc-pCVTZ" (number of atoms 1)
Ti
111. Basis Set "NASA Ames cc-pCVQZ" (number of atoms 1)
Ti
112. Basis Set "NASA Ames cc-pCV5Z" (number of atoms 1)
Ti
113. Basis Set "Sadlej pVTZ" (number of atoms 19)
H Li Be C N O F Na Mg Si P S Cl K Ca Br Rb Sr I
114. Basis Set "WTBS" (number of atoms 84)
He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba La Ce Pr Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Lu
Hf Ta W Re Os Ir Pt Au Hg Tl Pb Bi Po At Rn

Resolution of Identity (RI) fitting basis sets:

1. Basis Set "cc-pVDZ-fit2-1" (number of atoms 10)
H He Li Be B C N O F Ne
2. Basis Set "cc-pVTZ-fit2-1" (number of atoms 10)
H He Li Be B C N O F Ne

Density functional specific basis sets:

1. Basis Set "DZVP (DFT Orbital)" (number of atoms 54)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe
2. Basis Set "DZVP2 (DFT Orbital)" (number of atoms 25)
H He Li Be B C N O F Al Si P S Cl Ar Sc Ti V Cr Mn Fe Co Ni Cu Zn
3. Basis Set "TZVP (DFT Orbital)" (number of atoms 11)
H C N O F Al Si P S Cl Ar

Density functional Coulomb and Exchange fitting basis sets:

1. Basis Set "DGauss A1 DFT Coulomb Fitting" (number of atoms 54)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe
2. Basis Set "DGauss A1 DFT Exchange Fitting" (number of atoms 54)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe
3. Basis Set "DGauss A2 DFT Coulomb Fitting" (number of atoms 28)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Sc Ti V Cr Mn Fe Co Ni Cu Zn
4. Basis Set "DGauss A2 DFT Exchange Fitting" (number of atoms 64)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Sc Ti V Cr Mn Fe Co Ni Cu Zn
5. Basis Set "Ahlrichs Coulomb Fitting" (number of atoms 70)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba Hf Ta W Re Os Ir Pt Au Hg Tl Pb Bi Po At
6. Basis Set "DeMon Coulomb Fitting" (number of atoms 54)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb
Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe

Effective core potentials and their respective basis sets:

1. Basis Set "Hay-Wadt MB (n+1) ECP" (number of atoms 32)
K Ca Sc Ti V Cr Mn Fe Co Ni Cu Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cs Ba La Ta W Re Os Ir Pt Au
2. ECP "Hay-Wadt MB (n+1) ECP" (number of atoms 32)
K Ca Sc Ti V Cr Mn Fe Co Ni Cu Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cs Ba La Ta W Re Os Ir Pt Au
3. Basis Set "Hay-Wadt VDZ (n+1) ECP" (number of atoms 32)
K Ca Sc Ti V Cr Mn Fe Co Ni Cu Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cs Ba La Ta W Re Os Ir Pt Au
4. ECP "Hay-Wadt VDZ (n+1) ECP" (number of atoms 32)
K Ca Sc Ti V Cr Mn Fe Co Ni Cu Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cs Ba La Ta W Re Os Ir Pt Au
5. Basis Set "LANL2DZ ECP" (number of atoms 67)
H Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb Sr
Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba La Hf Ta W Re Os Ir Pt Au U Np Pu

6. ECP "LANL2DZ ECP" (number of atoms 58)
Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba La Hf Ta W Re Os Ir Pt Au U Np Pu
7. Basis Set "LANL2DZdp ECP" (number of atoms 19)
H C N O F Si P S Cl Ge As Se Br Sn Sb Te I Pb Bi
8. ECP "LANL2DZdp ECP" (number of atoms 13)
Si P S Cl Ge As Se Sn Sb Te I Pb Bi
9. Basis Set "SBKJC VDZ ECP" (number of atoms 73)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba La Ce Hf Ta W Re Os Ir Pt Au Hg Tl Pb Bi Po At Rn
10. ECP "SBKJC VDZ ECP" (number of atoms 71)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba La Ce Hf Ta W Re Os Ir Pt Au Hg Tl Pb Bi Po At Rn
11. Basis Set "CRENBL ECP" (number of atoms 116)
H Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba La Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Lu Hf Ta W Re Os Ir Pt Au Hg Pb Bi Po At Rn Fr Ra Ac Th Pa U Np Pu Am Cm Bk Cf Es Fm Md No Lr Rf Db Sg Bh Hs Mt Un Uu Ub Ut Uq Up Uh Us Uo
12. ECP "CRENBL ECP" (number of atoms 115)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd In Sn Sb Te I Xe Cs Ba La Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Lu Hf Ta W Re Os Ir Pt Au Hg Pb Bi Po At Rn Fr Ra Ac Th Pa U Np Pu Am Cm Bk Cf Es Fm Md No Lr Rf Db Sg Bh Hs Mt Un Uu Ub Ut Uq Up Uh Us Uo
13. Basis Set "CRENBS ECP" (number of atoms 50)
Sc Ti V Cr Mn Fe Co Ni Cu Zn Y Zr Nb Mo Tc Ru Rh Pd Ag Cd La Hf Ta W Re Os Ir Pt Au Hg Pb Bi Po At Rn Rf Db Sg Bh Hs Mt Un Uu Ub Ut Uq Up Uh Us Uo
14. ECP "CRENBS ECP" (number of atoms 50)
Sc Ti V Cr Mn Fe Co Ni Cu Zn Y Zr Nb Mo Tc Ru Rh Pd Ag Cd La Hf Ta W Re Os Ir Pt Au Hg Pb Bi Po At Rn Rf Db Sg Bh Hs Mt Un Uu Ub Ut Uq Up Uh Us Uo
15. Basis Set "Stuttgart RLC ECP" (number of atoms 57)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Zn Ga Ge As Se Br Kr Rb Sr In Sn Sb Te I Xe Cs Ba Hg Tl Pb Bi Po At Rn Ac Th Pa U Np Pu Am Cm Bk Cf Es Fm Md No Lr
16. ECP "Stuttgart RLC ECP" (number of atoms 57)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Zn Ga Ge As Se Br Kr Rb Sr In Sn Sb Te I Xe Cs Ba Hg Tl Pb Bi Po At Rn Ac Th Pa U Np Pu Am Cm Bk Cf Es Fm Md No Lr
17. Basis Set "Stuttgart RSC 1997 ECP" (number of atoms 64)
K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd Cs Ba Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Hf Ta W Re Os Ir Pt Au Hg Ac Th Pa U Np Pu Am Cm Bk Cf Es Fm Md No Lr Db
18. ECP "Stuttgart RSC 1997 ECP" (number of atoms 64)
K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Rb Sr Y Zr Nb Mo Tc Ru Rh Pd Ag Cd Cs Ba Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Hf Ta W Re Os Ir Pt Au Hg Ac Th Pa U Np Pu Am Cm Bk Cf Es Fm Md No Lr Db

19. Basis Set "Stuttgart RSC Segmented/ECP" (number of atoms 15)
La Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Lu
20. ECP "Stuttgart RSC Segmented/ECP" (number of atoms 15)
La Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Lu
21. Basis Set "Stuttgart RSC ANO/ECP" (number of atoms 15)
La Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Lu
22. ECP "Stuttgart RSC ANO/ECP" (number of atoms 15)
La Ce Pr Nd Pm Sm Eu Gd Tb Dy Ho Er Tm Yb Lu
23. Basis Set "SDB-cc-pVTZ" (number of atoms 12)
Ga Ge As Se Br Kr In Sn Sb Te I Xe
24. ECP "SDB-cc-pVTZ" (number of atoms 12)
Ga Ge As Se Br Kr In Sn Sb Te I Xe
25. Basis Set "SDB-cc-pVQZ" (number of atoms 12)
Ga Ge As Se Br Kr In Sn Sb Te I Xe
26. ECP "SDB-cc-pVQZ" (number of atoms 12)
Ga Ge As Se Br Kr In Sn Sb Te I Xe
27. Basis Set "SDB-aug-cc-pVTZ" (number of atoms 10)
Ga Ge As Se Br In Sn Sb Te I
28. ECP "SDB-aug-cc-pVTZ" (number of atoms 30)
Ga Ge As Se Br In Sn Sb Te I Ga Ge As Se Br In Sn Sb Te I
29. Basis Set "SDB-aug-cc-pVQZ" (number of atoms 10)
Ga Ge As Se Br In Sn Sb Te I
30. ECP "SDB-aug-cc-pVQZ" (number of atoms 10)
Ga Ge As Se Br In Sn Sb Te I

Douglas-Kroll (DK) all-electron basis sets:

1. Basis Set "cc-pVDZ DK" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
2. Basis Set "cc-pVTZ DK" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
3. Basis Set "cc-pVQZ DK" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
4. Basis Set "cc-pV5Z DK" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr

Dyall's Modified Dirac (DmD) all-electron basis sets:

1. Basis Set "cc-pvdz fi sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr

2. Basis Set "cc-pvdz fi sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
3. Basis Set "cc-pvdz fi sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
4. Basis Set "cc-pvdz pt sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
5. Basis Set "cc-pvdz pt sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
6. Basis Set "cc-pvdz pt sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
7. Basis Set "cc-pvtz fi sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
8. Basis Set "cc-pvtz fi sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
9. Basis Set "cc-pvtz fi sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
10. Basis Set "cc-pvtz pt sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
11. Basis Set "cc-pvtz pt sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
12. Basis Set "cc-pvtz pt sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
13. Basis Set "cc-pvqz fi sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
14. Basis Set "cc-pvqz fi sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
15. Basis Set "cc-pvqz fi sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
16. Basis Set "cc-pvqz pt sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
17. Basis Set "cc-pvqz pt sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
18. Basis Set "cc-pvqz pt sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
19. Basis Set "cc-pv5z fi sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
20. Basis Set "cc-pv5z fi sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr

21. Basis Set "cc-pv5z fi sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
22. Basis Set "cc-pv5z pt sf fw" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
23. Basis Set "cc-pv5z pt sf lc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
24. Basis Set "cc-pv5z pt sf sc" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr

Polarization functions:

1. Basis Set "STO-3G* Polarization" (number of atoms 8)
Na Mg Al Si P S Cl Ar
2. Basis Set "3-21G* Polarization" (number of atoms 8)
Na Mg Al Si P S Cl Ar
3. Basis Set "6-31G* Polarization" (number of atoms 27)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Co Ni Cu Zn
4. Basis Set "6-31G** Polarization" (number of atoms 29)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Co Ni Cu Zn
5. Basis Set "6-311G* Polarization" (number of atoms 23)
Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr I
6. Basis Set "6-311G** Polarization" (number of atoms 25)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr I
7. Basis Set "Pople (2d,2p) Polarization" (number of atoms 26)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Ga Ge As Se Br Kr
8. Basis Set "Pople (2df,2pd) Polarization" (number of atoms 12)
H He Li Be B C N O F Ne K Ca
9. Basis Set "Pople (3df,3pd) Polarization" (number of atoms 18)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar
10. Basis Set "HONDO7 Polarization" (number of atoms 17)
H He Be B C N O F Ne Na Mg Al Si P S Cl Ar
11. Basis Set "Huzinaga Polarization" (number of atoms 22)
K Ca Ga Ge As Se Br Kr Rb Sr In Sn Sb Te I Xe Tl Pb Bi Po At Rn
12. Basis Set "DHMS Polarization" (number of atoms 15)
H He Li Be B C N O F Ne Al Si P S Cl
13. Basis Set "Binning-Curtiss (1d) Polarization" (number of atoms 6)
Ga Ge As Se Br Kr
14. Basis Set "Binning-Curtiss (df) Polarization" (number of atoms 6)
Ga Ge As Se Br Kr

15. Basis Set "Ahlrichs Polarization" (number of atoms 36)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca Sc Ti V Cr Mn Fe Co Ni Cu Zn Ga Ge As Se Br Kr
16. Basis Set "Glendening Polarization" (number of atoms 3)
K Rb Cs
17. Basis Set "Blaudeau Polarization" (number of atoms 2)
K Ca
18. Basis Set "LANL2DZdp ECP Polarization" (number of atoms 19)
H C N O F Si P S Cl Ge As Se Br Sn Sb Te I Pb Bi
19. Basis Set "Dunning-Hay Rydberg" (number of atoms 12)
B C N O F Ne Al Si P S Cl Ar
20. Basis Set "Dunning-Hay Double Rydberg" (number of atoms 12)
B C N O F Ne Al Si P S Cl Ar

Diffuse functions:

1. Basis Set "Pople-style Diffuse" (number of atoms 19)
H Li Be B C N O F Ne Na Mg Al Si P S Cl Ar K Ca
2. Basis Set "Dunning-Hay Diffuse" (number of atoms 9)
H Li Be B C N O F Ne
3. Basis Set "aug-cc-pVDZ Diffuse" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
4. Basis Set "aug-cc-pVTZ Diffuse" (number of atoms 23)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br
5. Basis Set "aug-cc-pVQZ Diffuse" (number of atoms 24)
H He Li Be B C N O F Ne Na Mg Al Si P S Cl Ar Ga Ge As Se Br Kr
6. Basis Set "aug-cc-pV5Z Diffuse" (number of atoms 20)
H He B C N O F Ne Al Si P S Cl Ar Ga Ge As Se Br Kr
7. Basis Set "aug-cc-pV6Z Diffuse" (number of atoms 14)
H He B C N O F Ne Al Si P S Cl Ar
8. Basis Set "d-aug-cc-pVDZ Diffuse" (number of atoms 8)
H He B C N O F Ne
9. Basis Set "d-aug-cc-pVTZ Diffuse" (number of atoms 8)
H He B C N O F Ne
10. Basis Set "d-aug-cc-pVQZ Diffuse" (number of atoms 8)
H He B C N O F Ne
11. Basis Set "d-aug-cc-pV5Z Diffuse" (number of atoms 8)
H He B C N O F Ne
12. Basis Set "d-aug-cc-pV6Z Diffuse" (number of atoms 5)
H B C N O

13. Basis Set "aug-cc-pV(D+d)Z Diffuse" (number of atoms 6)
Al Si P S Cl Ar
14. Basis Set "aug-cc-pV(T+d)Z Diffuse" (number of atoms 6)
Al Si P S Cl Ar
15. Basis Set "aug-cc-pV(Q+d)Z Diffuse" (number of atoms 6)
Al Si P S Cl Ar
16. Basis Set "aug-cc-pV(5+d)Z Diffuse" (number of atoms 6)
Al Si P S Cl Ar
17. Basis Set "aug-cc-pV(6+d)Z Diffuse" (number of atoms 6)
Al Si P S Cl Ar
18. Basis Set "SDB-aug-cc-pVTZ Diffuse" (number of atoms 10)
Ga Ge As Se Br In Sn Sb Te I
19. Basis Set "SDB-aug-cc-pVQZ Diffuse" (number of atoms 10)
Ga Ge As Se Br In Sn Sb Te I

Core-valence functions:

1. Basis Set "Core/val. Functions (cc-pCVDZ)" (number of atoms 12)
Li Be B C N O F Ne Na Mg Si Cl
2. Basis Set "Core/val. Functions (cc-pCVTZ)" (number of atoms 13)
Li Be B C N O F Ne Na Mg Si P Cl
3. Basis Set "Core/val. Functions (cc-pCVQZ)" (number of atoms 13)
Li Be B C N O F Ne Na Mg Si P Cl
4. Basis Set "Core/val. Functions (cc-pCV5Z)" (number of atoms 8)
B C N O F Ne Na Mg
5. Basis Set "Core/val. Functions (cc-pCV6Z)" (number of atoms 1)
O

Appendix B

Sample input files

B.1 Water SCF calculation and geometry optimization in a 6-31g basis

The input file in section 2 performs a geometry optimization in a single task. A single point SCF energy calculation is performed and then restarted to perform the optimization (both could of course be performed in a single task).

B.1.1 Job 1. Single point SCF energy

```
start h2o
title "Water in 6-31g basis set"

geometry units au
  O      0.00000000    0.00000000    0.00000000
  H      0.00000000    1.43042809   -1.10715266
  H      0.00000000   -1.43042809   -1.10715266
end
basis
  H library 6-31g
  O library 6-31g
end
task scf
```

The final energy should be -75.983998.

B.1.2 Job 2. Restarting and perform a geometry optimization

```
restart h2o
title "Water geometry optimization"

task scf optimize
```

There is no need to specify anything that has not changed from the previous input deck, though it will do no harm to repeat it.

B.2 Compute the polarizability of Ne using finite field

B.2.1 Job 1. Compute the atomic energy

```
start ne
title "Neon"
geometry; ne 0 0 0; end
basis spherical
  ne library aug-cc-pvdz
end
scf; thresh 1e-10; end
task scf
```

The final energy should be -128.496350.

B.2.2 Job 2. Compute the energy with applied field

An external field may be simulated with point charges. The charges here apply a field of magnitude 0.01 atomic units to the atom at the origin. Since the basis functions have not been reordered by the additional centers we can also restart from the previous vectors, which is the default for a restart job.

```
restart ne
title "Neon in electric field"
geometry units atomic
  bq1 0 0 100 charge 50
  ne 0 0 0
  bq2 0 0 -100 charge -50
end
task scf
```

The final energy should be -128.496441, which together with the previous field-free result yields an estimate for the polarizability of 1.83 atomic units. Note that by default NWChem does not include the interaction between the two point charges in the total energy (section 6).

B.3 SCF energy of H₂CO using ECPs for C and O

The following will compute the SCF energy for formaldehyde with ECPs on the Carbon and Oxygen centers.

```
title "formaldehyde ECP deck"

start ecphho

geometry units au
C      0.000000  0.000000 -1.025176
O      0.000000  0.000000  1.280289
H      0.000000  1.767475 -2.045628
H      0.000000 -1.767475 -2.045628
```

end

basis

```

C SP
  0.1675097360D+02 -0.7812840500D-01  0.3088908800D-01
  0.2888377460D+01 -0.3741108860D+00  0.2645728130D+00
  0.6904575040D+00  0.1229059640D+01  0.8225024920D+00
C SP
  0.1813976910D+00  0.1000000000D+01  0.1000000000D+01
C D
  0.8000000000D+00  0.1000000000D+01
C F
  0.1000000000D+01  0.1000000000D+01
O SP
  0.1842936330D+02 -0.1218775590D+00  0.5975796600D-01
  0.4047420810D+01 -0.1962142380D+00  0.3267825930D+00
  0.1093836980D+01  0.1156987900D+01  0.7484058930D+00
O SP
  0.2906290230D+00  0.1000000000D+01  0.1000000000D+01
O D
  0.8000000000D+00  0.1000000000D+01
O F
  0.1100000000D+01  0.1000000000D+01
H S
  0.1873113696D+02  0.3349460434D-01
  0.2825394365D+01  0.2347269535D+00
  0.6401216923D+00  0.8137573262D+00
H S  1 1.00
  0.1612777588D+00  0.1000000000D+01

```

end

ecp

```

C nelec 2
C ul
  1      80.0000000    -1.60000000
  1      30.0000000    -0.40000000
  2       0.5498205    -0.03990210
C s
  0       0.7374760     0.63810832
  0     135.2354832    11.00916230
  2       8.5605569    20.13797020
C p
  2      10.6863587    -3.24684280
  2      23.4979897     0.78505765
O nelec 2
O ul
  1      80.0000000    -1.60000000
  1      30.0000000    -0.40000000
  2       1.0953760    -0.06623814
O s
  0       0.9212952     0.39552179

```

```

      0      28.6481971      2.51654843
      2      9.3033500      17.04478500
    O p
      2      52.3427019      27.97790770
      2      30.7220233     -16.49630500
end

scf
  vectors input hcore
  maxiter 20
end

task scf

```

This should produce the following output:

```

      Final RHF  results
      -----
      Total SCF energy =    -22.507927218024
      One electron energy =   -71.508730162974
      Two electron energy =    31.201960019808
      Nuclear repulsion energy =   17.798842925142

```

B.4 MP2 optimization and CCSD(T) on nitrogen

The following performs an MP2 geometry optimization followed by a CCSD(T) energy evaluation at the converged geometry. A Dunning correlation-consistent triple-zeta basis is used. The default of Cartesian basis functions must be overridden using the keyword `spherical` on the `BASIS` directive. The `1s` core orbitals are frozen in both the MP2 and coupled-cluster calculations (note that these must separately specified). The final MP2 energy is -109.383276, and the CCSD(T) energy is -109.399662.

```

start n2

geometry
  symmetry d2h
  n 0 0 0.542
end

basis spherical
  n library cc-pvtz
end

mp2
  freeze core
end

task mp2 optimize

```

```
ccsd
  freeze core
end
```

```
task ccsd(t)
```


Appendix C

Examples of geometries using symmetry

Below are examples of the use of the SYMMETRY directive in the compound GEOMETRY directive (Section 6). The z axis is always the primary rotation axis. When in doubt about which axes and planes are used for the group elements, the keyword `print` may be added to the SYMMETRY directive to obtain this information.

C.1 C_s methanol

The σ_h plane is the xy plane.

```
geometry units angstroms
  C   0.11931097   -0.66334875   0.00000000
  H   1.20599017   -0.87824237   0.00000000
  H   -0.32267592  -1.15740001   0.89812652
  O   -0.01716588   0.78143468   0.00000000
  H   -1.04379735   0.88169812   0.00000000
symmetry cs
end
```

C.2 C_{2v} water

The z axis is the C_2 axis and the σ_v may be either the xz or the yz planes.

```
geometry units au
  O   0.00000000   0.00000000   0.00000000
  H   0.00000000   1.43042809  -1.10715266
symmetry group c2v
end
```

C.3 D_{2h} acetylene

Although acetylene has symmetry $D_{\infty h}$ the subgroup D_{2h} includes all operations that interchange equivalent atoms which is what determines how much speedup you gain from using symmetry in building a Fock matrix.

The C_2 axes are the x, y, and z axes. The σ planes are the xy, xz and yz planes. Generally, the unique atoms are placed to use the z as the primary rotational axis and use the xz or yz planes as the σ plane.

```
geometry units au
  symmetry group d2h
  C      0.000000000    0.000000000   -1.115108538
  H      0.000000000    0.000000000   -3.106737425
end
```

C.4 D_{2h} ethylene

The C_2 axes are the x, y, and z axes. The σ planes are the xy, xz and yz planes. Generally, the unique atoms are placed to use the z as the primary rotational axis and use the xz or yz planes as the σ plane.

```
geometry units angstroms
  C 0 0 0.659250
  H 0 0.916366 1.224352
  symmetry d2h
end
```

C.5 T_d methane

For ease of use, the primary C_3 axis should be the x=y=z axis. The 3 C_2 axes are the x, y, and z.

```
geometry units au
  c      0.0000000    0.0000000    0.0000000
  h      1.1828637    1.1828637    1.1828637
  symmetry group Td
end
```

C.6 I_h buckminsterfullerene

One of the C_5 axes is the z axis and the point of inversion is the origin.

```
geometry units angstroms # Bonds = 1.4445, 1.3945
  symmetry group Ih
  c      -1.2287651    0.0    3.3143121
end
```

C.7 S_4 porphyrin

The S_4 and C_2 rotation axis is the z axis. The reflection plane for the S_4 operation is the xy plane.

```

geometry units angstroms
symmetry group s4

fe          0.000  0.000  0.000
h           2.242  6.496 -3.320
h           1.542  4.304 -2.811
c           1.947  6.284 -2.433
c           1.568  4.987 -2.084
h           2.252  8.213 -1.695
c           1.993  7.278 -1.458
h           5.474 -1.041 -1.143
c           1.234  4.676 -0.765
h           7.738 -1.714 -0.606
c           0.857  3.276 -0.417
h           1.380 -4.889 -0.413
c           1.875  2.341 -0.234
h           3.629  3.659 -0.234
c           0.493 -2.964 -0.229
c           1.551 -3.933 -0.221
c           5.678 -1.273 -0.198
c           1.656  6.974 -0.144
c           3.261  2.696 -0.100
n           1.702  0.990 -0.035
end

```

C.8 D_{3h} iron penta-carbonyl

The C_3 axis is the z axis. The σ_h plane is the xy plane. One of the perpendicular C_2 axes is the x=y axis. One of the σ_v planes is the plane containing the x=y axis and the z axis. (The other axes and planes are generated by the C_3 operation.)

```

geometry units au
symmetry group d3h

fe          0.0          0.0          0.0

c           0.0          0.0          3.414358
o           0.0          0.0          5.591323

c           2.4417087    2.4417087    0.0
o           3.9810552    3.9810552    0.0
end

```

C.9 D_{3d} sodium crown ether

The C_3 axis is the z axis. The point of inversion is the origin. One of the perpendicular C_2 axes is the x=y axis. One of the σ_d planes is the plane containing the -x=y axis and the z axis.

Note that the oxygen atom is rotated in the x-y plane 15 degrees away from the y-axis so that it lies in a mirror plane. There is a total of six atoms generated from the unique oxygen, in contrast to twelve from each of the carbon and hydrogen atoms.

```

geometry units au
symmetry D3d

NA      .0000000000    .0000000000    .0000000000
O       1.3384771885    4.9952647969    .1544089284
H       6.7342048019   -0.6723850379    2.6581562148
C       6.7599180056   -0.4844977035    .6136583870
H       8.6497577017    0.0709194071    .0345361934

end

```

C.10 C_{3v} ammonia

The C_3 axis is the z axis. One of the σ_v planes is the plane containing the x=y axis and the z axis.

```

geometry units angstroms
N 0      0      -0.055
H 0.665  0.665  -0.481
symmetry c3v
end

```

C.11 D_{6h} benzene

The C_6 axis is the z axis. The point of inversion is the origin. One of the 6 perpendicular C_2' axes is the x=y axis. (-x=y works as a C_2'' axis.) The σ_h plane is the xy plane. The σ_d planes contain the C_2'' axis and the z axis. The σ_v planes contain the C_2' axis and the z axis.

```

geometry units au
C 1.855  1.855  0
H 3.289  3.289  0
symmetry D6h
end

```

C.12 C_{3h} BO_3H_3

The C_3 axis is the z axis. The σ_h plane is the xy plane.

```

geometry units au
  b 0 0 0
  o 2.27238285 1.19464491 0.00000000
  h 2.10895420 2.97347707 0.00000000
symmetry C3h
end

```

C.13 D_{5d} ferrocene

The C_5 axis is the z axis. The center of inversion is the origin. One of the perpendicular C_2 axes is the x axis. One of the σ_d planes is the yz plane.

```

geometry units angstroms
  symmetry d5d

  fe 0 0 0
  c 0 1.194 1.789
  h 0 2.256 1.789
end

```

C.14 C_{4v} SF_5Cl

The C_4 axis is the z axis. The σ_v planes are the yz and the xz planes. The σ_d planes are: 1) the plane containing the x=y axis and the z axis and 2) the plane containing the -x=y axis and the z axis.

```

geometry units au
  S 0.00000000 0.00000000 -0.14917600
  Cl 0.00000000 0.00000000 4.03279700
  F 3.13694200 0.00000000 -0.15321800
  F 0.00000000 0.00000000 -3.27074500

  symmetry C4v
end

```

C.15 C_{2h} trans-dichloroethylene

The C_2 axis is the z axis. The origin is the inversion center. The σ_h plane is the xy plane.

```

geometry units angstroms
  C 0.65051239 -0.08305064 0
  Cl 1.75249381 1.30491767 0
  H 1.14820954 -1.04789741 0
  symmetry C2h
end

```

C.16 D_{2d} CH_2CCH_2

The C_2 axis is the z axis (z is also the S_4 axis). The x and y axes are the perpendicular C_2' s. The σ_d planes are: 1) the plane containing the x=y axis and the z axis and 2) the plane containing the -x=y axis and the z axis.

```

geometry units angstroms
symmetry d2d
c 0 0 0
c 0 0 1.300
h 0.656 0.656 1.857
end

```

C.17 D_{5h} cyclopentadiene anion

The C_5 axis is the z axis (z is also the S_5 axis). The y axis is one of the perpendicular C_2 axes. The σ_h plane is the xy plane and one of the σ_d planes is the yz plane.

```

charge -1
geometry units angstroms
symmetry d5h
c 0 1.1853 0
h 0 2.2654 0
end

```

C.18 D_{4h} gold tetrachloride

The C_4 axis is the z axis (z is also the S_4 axis). The C_2' axes are the x and y axes and the C_2'' axes are the x=y axis and the x=-y axis. The inversion center is the origin. The σ_h plane is the xy plane. The σ_v planes are the xz and yz planes and the σ_d planes are 1) the plane containing the x=-y axis and the z axis and 2) the plane containing the x=y axis and the z axis.

```

geometry units au
Au 0 0 0
Cl 0 4.033 0
symmetry D4h
end

```

Appendix D

Running NWChem

The command required to invoke NWChem is machine dependent, whereas most of the NWChem input is machine independent¹.

D.1 Sequential execution

To run NWChem sequentially on nearly all UNIX-based platforms simply use the command `nwchem` and provide the name of the input file as an argument (See section 2.1 for more information). This does assume that either `nwchem` is in your path or you have set an alias of `nwchem` to point to the appropriate executable.

Output is to standard output, standard error and Fortran unit 6 (usually the same as standard output). Files are created by default in the current directory, though this may be overridden in the input (section 5.2).

Generally, one will run a job with the following command:

```
nwchem input.nw >& input.out &
```

D.2 Parallel execution on UNIX-based parallel machines including workstation clusters using TCGMSG

These platforms require the use of the TCGMSG² `parallel` command and thus also require the definition of a process-group (or `procgrou`) file. The process-group file describes how many processes to start, what program to run, which machines to use, which directories to work in, and under which userid to run the processes. By convention the process-group file has a `.p` suffix.

The process-group file is read to end-of-file. The character # (hash or pound sign) is used to indicate a comment which continues to the next new-line character. Each line describes a cluster of processes and consists of the following whitespace separated fields:

```
userid hostname nslave executable workdir
```

¹Machine dependence within the input arises from file names, machine specific resources, and differing services provided by the operating system.

²Where required TCGMSG is automatically built with NWChem.

- `userid` – The user-name on the machine that will be executing the process.
- `hostname` – The hostname of the machine to execute this process. If it is the same machine on which `parallel` was invoked the name must match the value returned by the command `hostname`. If a remote machine it must allow remote execution from this machine (see man pages for `rlogin`, `rsh`).
- `nslave` – The total number of copies of this process to be executing on the specified machine. Only “clusters” of identical processes specified in this fashion can use shared memory to communicate. If no shared memory is supported on machine `<hostname>` then only the value one (1) is valid.
- `executable` – Full path name on the host `<hostname>` of the image to execute. If `<hostname>` is the local machine then a local path will suffice.
- `workdir` – Full path name on the host `<hostname>` of the directory to work in. Processes execute a `chdir()` to this directory before returning from `pbegin()`. If specified as a “.” then remote processes will use the login directory on that machine and local processes (relative to where `parallel` was invoked) will use the current directory of `parallel`.

For example, if your file `"nwchem.p"` contained the following

```
d3g681 pc 4 /msrc/apps/bin/nwchem /scr22/rjh
```

then 4 processes running NWChem would be started on the machine `pc` running as user `d3g681` in directory `"/scr22/rjh"`. To actually run this simply type:

```
parallel nwchem big_molecule.nw
```

N.B. : The first process specified (process zero) is the only process that

- opens and reads the input file, and
- opens and reads/updates the database.

Thus, if your file systems are physically distributed (e.g., most workstation clusters) you must ensure that process zero can correctly resolve the paths for the input and database files.

N.B. In releases of NWChem prior to 3.3 additional processes had to be created on workstation clusters to support remote access to shared memory. This is no longer the case. The `TCGMSG` process group file now just needs to refer to processes running NWChem.

D.3 Parallel execution on UNIX-based parallel machines including workstation clusters using MPI

To run with MPI, `parallel` should not be used. The way we usually run `nwchem` under MPI are the following

- using `mpirun`:

```
mpirun -np 8 $NWCHEM_TOP/bin/$NWCHEM_TARGET/nwchem input.nw
```

- If you have all nodes connected via shared memory and you have installed the `ch_shmem` version of MPICH, you can do

```
$NWCHEM_TOP/bin/$NWCHEM_TARGET/nwchem -np 8 h2o.nw
```

D.4 Parallel execution on MPPs

All of these machines require use of different commands in order to gain exclusive access to computational resources.

D.5 IBM SP

If using POE (IBM's Parallel Operating Environment) interactively, simply create the list of nodes to use in the file "host.list" in the current directory and invoke NWChem with

```
nwchem <input_file> -procs <n>
```

where *n* is the number of processes to use. Process 0 will run on the first node in "host.list" and must have access to the input and other necessary files. Very significant performance gains may be had by setting the following environment variables before running NWChem (or setting them using POE command line options).

- `setenv MP_EUILIB us` — dedicated user space communication over the switch (the default is IP over the switch which is much slower).
- `setenv MP_CSS_INTERRUPT yes` — enable interrupts when a message arrives (the default is to poll which significantly slows down global array accesses).

In addition, if the IBM is running PSSP version 3.1, or later

- `setenv MP_MSG_API lapi`, or
- `setenv MP_MSG_API mpi,lapi` (if using both GA and MPI)

For batch execution, we recommend use of the `llnw` command which is installed in `/usr/local/bin` on the EMSL/PNNL IBM SP. If you are not running on that system, the `llnw` script may be found in the NWChem distribution directory `contrib/loadleveler`. Interactive help may be obtained with the command `llnw -help`. Otherwise, the very simplest job to run NWChem in batch using Load Leveller is something like this

```
#!/bin/csh -x
# @ job_type          =    parallel
# @ class             =    small
# @ network.lapi     =    css0,not_shared,US
# @ input             =    /dev/null
# @ output            =    <OUTPUT_FILE_NAME>
# @ error             =    <ERROUT_FILE_NAME>
# @ environment       =    COPY_ALL; MP_PULSE=0; MP_SINGLE_THREAD=yes; MP_WAIT_MODE=yield;
# @ min_processors    =    7
# @ max_processors    =    7
# @ cpu_limit         =    1:00:00
# @ wall_clock_limit =    1:00:00
# @ queue
#

cd /scratch

nwchem <INPUT_FILE_NAME>
```

Substitute <OUTPUT_FILE_NAME>, <ERROUT_FILE_NAME> and <INPUT_FILE_NAME> with the *full* path of the appropriate files. Also, if you are using an SP with more than one processor per node, you will need to substitute

```
# @ network.lapi      = css0,shared,US
# @ node              = NNODE
# @ tasks_per_node    = NTASK
```

for the lines

```
# @ network.lapi      = css0,not_shared,US
# @ min_processors    = 7
# @ max_processors    = 7
```

where NNODE is the number of physical nodes to be used and NTASK is the number of tasks per node.

These files and the NWChem executable must be in a file system accessible to all processes. Put the above into a file (e.g., "test.job") and submit it with the command

```
llsubmit test.job
```

It will run a 7 processor, 1 hour job in the queue small. It should be apparent how to change these values.

Note that on many IBM SPs, including that at EMSL, the local scratch disks are wiped clean at the beginning of each job and therefore persistent files should be stored elsewhere. PIOFS is recommended for files larger than a few MB.

D.6 Cray T3E

```
mpprun -n <npes> $NWCHEM_TOP/bin/$NWCHEM_TARGET/nwchem <input_file>
```

where *npes* is the number of processors and *input_file* is the name of your input file.

D.7 Alpha systems with Quadrics switch

```
prun -n <npes> $NWCHEM_TOP/bin/$NWCHEM_TARGET/nwchem <input_file>
```

where *npes* is the number of processors and *input_file* is the name of your input file.

D.8 Windows 98 and NT

```
$NWCHEM_TOP/bin/win32/nw32 <input_file>
```

where *input_file* is the name of your input file. If you use WMPI, you must have a file named *nw32.pg* in the *\$NWCHEM_TOP/bin/win32* directory; the file must only contain the following single line

```
local 0
```

D.9 Tested Platforms and O/S versions

- IBM SP with Power 3 and Power 4 nodes, AIX 5.1 and PSSP 3.4; IBM RS6000 workstation, AIX 5.1. Xlf 8.1.0.0 and 8.1.0.1 are known to produce bad code.
- SGI R12000 IRIX 6.5
- SUN workstations with Solaris 2.6 and 2.8. Fujitsu SPARC systems (thanks to Herbert Früchtel) with Parallelnavi compilers.
- HP DEC alpha workstation , Tru64 V5.1, Compaq Fortran V5.3, V5.4.2, V5.5.1
- Linux with Intel x86 cpus. NWChem Release 4.5 has been tested on RedHat 6.x and 7.x, Mandrake 7.x. We have tested NWChem on Linux for the Power PC Macintosh with Yellow Dog 2.4. These all use the GCC compiler at different levels. The Intel Fortran Compiler version 7 is supported. The Portland Group Compiler has been tested in a less robust manner. Automatic generation of SSE2 optimized code is available when the Intel compiler is used (ifc vs g77 performances gain of 40% in some benchmarks) A somewhat Athlon optimized code can be generated under the GNU or Intel compilers by typing `make _CPU=k7`. GCC3 specific options can be turned on by typing `make GCC31=y`
- HP 9000/800 workstations with HPUX B.11.00. f90 must be used for compilation.
- Intel x86 with Windows 2000 has been tested with Compaq Visual Fortran 6.0 and 6.1 with WMPI 1.3 or NT-Mpich. NT-MPICH is available from <http://www-unix.mcs.anl.gov/~ashton/mpich.nt/>
- Intel IA64 under Linux (with Intel compilers version 7 and later) and under HPUX.
- Fujitsu VPP computers.