

## LUND UNIVERSITY

A first-principles approach to protein-ligand interaction

Söderhjelm, Pär

2009

Link to publication

Citation for published version (APA):

Söderhjelm, P. (2009). A first-principles approach to protein–ligand interaction. [Doctoral Thesis (compilation), Computational Chemistry]. Department of Theoretical Chemistry, Lund University.

Total number of authors:

#### General rights

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights. • Users may download and print one copy of any publication from the public portal for the purpose of private study

or research.

You may not further distribute the material or use it for any profit-making activity or commercial gain
 You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: https://creativecommons.org/licenses/

#### Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

**PO Box 117** 221 00 Lund +46 46-222 00 00

## A first-principles approach to protein–ligand interaction

Thesis submitted for the degree of Doctor of Philosophy by

#### Pär Söderhjelm

Department of Theoretical Chemistry



UNIVERSITET

AKADEMISK AVHANDLING för avläggande av filosofie doktorsexamen vid naturvetenskapliga fakulteten, Lunds Universitet, kommer att offentligen försvaras i

Hörsal B, Kemicentrum, fredagen den 6 Februari 2009, kl 10.15

Fakultetens opponent är Dr. Jean-Philip Piquemal Laboratoire de Chimie Théorique, Université Pierre et Marie Curie, Paris, Frankrike

#### © Pär Söderhjelm, 2009

Lund University, Sweden Doctoral dissertation in Theoretical Chemistry

ISBN 978-91-628-7665-4

Printed by Media Tryck, Lund

## Contents

P	refac	е		$\mathbf{v}$	
P	opulä	irvetei	nskaplig sammanfattning	vii	
$\mathbf{L}\mathbf{i}$	ist of	paper	°S	xi	
A	cknov	wledge	ements	xiii	
1	Che	emistr	y from a theoretical perspective	1	
	1.1	The e	lectronic problem: Quantum chemistry	. 2	
		1.1.1	The basis set problem	. 4	
		1.1.2	The electron-correlation problem	. 5	
	1.2	The n	uclear problem: Statistical mechanics	. 7	
		1.2.1	Physical description	. 7	
		1.2.2	Computation	. 9	
<b>2</b>	Inte	ermole	ecular interactions	11	
	2.1	The s	upermolecular approach	. 11	
		2.1.1	Electron correlation	. 11	
		2.1.2	Basis set	. 13	
	2.2	Pertu	rbative approach	. 14	
		2.2.1	Qualitative picture	. 15	
		2.2.2	Quantum-chemical perturbation theory	. 16	
		2.2.3	Molecular mechanics	. 19	
	2.3	Molec	cular mechanics based on quantum chemistry	. 20	
		2.3.1	Electrostatic interaction energy	. 21	
		2.3.2	Polarization	. 22	
		2.3.3	Dispersion	. 29	
		2.3.4	Repulsion	. 30	
		2.3.5	Charge transfer	. 34	

3	Pro	tein–ligand affinities	35
	3.1	Theory	36
	3.2	Path-based methods	37
		3.2.1 Relative affinities	37
		3.2.2 Absolute affinities	37
		3.2.3 Potential of mean force	38
	3.3	End-point methods	39
		3.3.1 Perturbations from single simulations	39
		3.3.2 Linear interaction energy (LIE)	40
		3.3.3 MM/PBSA	40
	3.4	Solvation	42
		3.4.1 Polar solvation	43
		3.4.2 Non-polar solvation	44
	3.5	Potential energy	44
<b>4</b>	Sun	nmary of the papers	<b>47</b>
	4.1	Paper I: Exchange-repulsion energy	47
	4.2	Paper II: Multipoles and polarizabilities	48
	4.3	Papers III and IV: Polarization models	49
		4.3.1 Densities	49
		4.3.2 Energies	51
	4.4	Paper V, VI, and VII: The PMISP method	52
		4.4.1 Methods	53
		4.4.2 Results	53
	4.5	Conclusions and outlook	55
	Bib	liography	57

## Preface

When the subject of force fields for protein simulations was reviewed some years ago by Ponder and Case [1], the authors noted that "without further research into the accuracy of force-field potentials, future macromolecular modeling may well be limited more by validity of the energy functions, particularly electrostatic terms, than by technical ability to perform the computations. For many calculations related to ligand binding, drug design, and protein structure prediction, accuracy of the underlying potential functions is critical."

The current work is my contribution to this problem. The thesis describes some basic attempts to get a grip of how accurate force fields are, how accurate they can become, and where to put in the effort to improve them, all in the context of using them for predicting interaction energies in biological systems.

The physical laws underlying chemistry have been known for about 80 years [2]. What prevents us from calculating e.g. protein–ligand affinities directly from these laws, i.e. from *first principles*, is essentially that they lead to equations that cannot be solved exactly. Therefore, this thesis explores an approach, in which the intermolecular potentials are purely based on quantum chemistry, but approximated in several ways. In principle, the accuracy of each approximation can be thoroughly tested, so that the quantum-chemical value can be approached. In contrast, most approaches to binding affinities are empirical, using experimental data on at least some level. A drawback of such approaches are that they are not systematically improvable.

The structure of the thesis is as follows. The first chapter introduces the constitution of matter as it appears to a computational chemist like me. The second chapter deals with calculating interaction energies, focusing on how to use quantum chemistry to derive simpler methods that can be used for larger systems. The third chapter addresses the protein–ligand binding problem, which in addition to interaction energies encompasses several difficult problems, such as sampling and solvation. The fourth chapter gives a summary of the papers included in the thesis and the conclusions I have drawn from them. The papers are attached at the end of the thesis.

The purpose of the first three chapters is to give an overview of the theory and computational approaches that underlie my own work. My goal has been to present a non-mathematical description, giving a personal (and far from exhaustive) selection of references that can provide details to the interested reader. There are indeed many formulas, but that does not mean you have to understand them to read the thesis. They are there simply because I like formulas. Knowing that there is a formula and seeing what quantities are in it is a great comfort to me, even if I do not have a clue how to evaluate it.

Lund, December 2008

Pär Söderhjelm

## Populärvetenskaplig sammanfattning

Många biologiska processer på molekylär nivå börjar med att två molekyler binder till varandra på ett speciellt sätt. Exempelvis måste ett enzym, d.v.s. ett protein som är designat för att utföra en viss kemisk reaktion, binda till just den molekylen som ska reagera. Ett annat exempel är vårt luktsinne, där doftämnet måste binda till en receptor, som också är ett protein, för att utlösa en nervsignal till hjärnan. Den mindre molekylen som binder till proteinet kallas ofta *ligand*.

Denna bindingsprocess brukar kallas molekylär igenkänning, men har naturligtvis ingenting med en intelligent process att göra. Molekylerna följer en bana som bestäms av fysikens lagar där de flyter runt i en trög soppa av vattenmolekyler, joner och andra små och stora molekyler. På sin väg innan de träffar "den rätte" hinner de stöta ihop med tusentals andra tänkbara partners, men om bindningsstyrkan är för liten fastnar de inte alls eller bara kanske en kort stund. Om bindningsstyrkan däremot är stor sitter molekylerna ihop länge. Den här avhandlingen handlar om att teoretiskt beräkna bindningsstyrkan mellan ett protein och en annan molekyl.

sådana beräkningar är bl.a. av stort intresse för läkemedelsindustrin. Om man kan hitta en molekyl som binder starkare till ett enzym än den naturliga molekyl som enzymet egentligen ska binda till, så kommer den att blockera enzymet så att den naurliga molekylen inte får plats och ingen reaktion kan ske, vilket ibland kan vara ett sätt att behandla en sjukdom. Förutsättningen för att en sådan molekyl ska fungera som läkemedel är förstås också att man lyckas få in den i cellen och att den inte binder till andra proteiner och på så vis ställer till oreda.

Trots att proteiner är jättestora för att vara molekyler är de små med våra mått mätt, bara några nanometer (milliondels millimeter). Vi kan därför inte ens med mikroskop se hur det ser ut när en ligand binder till ett protein. Med datorns hjälp kan vi däremot göra oss en konstgjord bild av det hela (Fig. 1). Datorn är också det redskap vi behöver för att uppskatta bindningsstyrkan.

Som en första approximation brukar man kunna anta att för att liganden ska binda starkt, så ska den "passa in" i proteinet, som en nyckel i ett lås.



Figure 1: En ögonblicksbild av ett protein (avidin) som bundit in en ligand (biotin), färgad lila på bilden. Proteinets atomer har förminskats för att man ska se liganden som egentliger ligger helt inbäddad. Varje färg motsvarar en sorts atom: vit=väte, röd=syre, ljusblå=kol, blå=kväve och gul=svavel. Allt tomrum är i verkligheten fyllt av vattenmolekyler.

Ligandens yta ska alltså vara som en avgjutning av proteinets yta, så att de kan ligga intill varandra utan att det blir några hål. Detta beror på att de flesta sorters molekyler attraherar varandra när man kommer tillräckligt nära. Det är t.ex. sådana krafter som gör att en ödla kan gå uppför en lodrät vägg.

För att få någorlunda bra resultat krävs mer noggranna beräkningsmodeller som också tar hänsyn till den kemiska strukturen på både proteinet och liganden. Det räcker alltså inte att ytorna ligger intill varandra, det måste också vara ytor som passar ihop, t.ex. en positivt och en negativt laddad yta. Detta kan beskrivas med hjälp av ett *kraftfält*.

Ett kraftfält kan ses som en topografisk karta över t.ex. ett bergslandskap. Det talar om hur fördelaktigt det är för varje par av atomer att befinna sig nära varandra, eller rättare sagt hur *energin* varierar beroende på avståndet mellan dem. Enligt en grundläggande fysikalisk princip strävar systemet efter att ha så låg energi som möjligt, d.v.s. komma så långt ner i dalgångarna som möjligt. Atomerna är dock aldrig stilla utan kan snarare betraktas som ett gäng dagisbarn som när de springer ner för backarna får sådan fart att de fortsätter uppför nästa backe med oförtröttlig energi. Denna *dynamik* måste man normalt också ta hänsyn till i sina beräkningar. Endast om man kyler ner alltihop till absoluta nollpunkten  $(-273^{\circ}C)$  börjar atomerna lugna ner sig.

Den rätta kartan, den som motsvarar verkligheten, är tyvärr omöjlig att finna eftersom vi inte kan se hur atomerna rör sig. Normalt nöjer man sig med en karta som ger hyfsade resultat i genomsnitt. Detta är bekvämt, för när resultatet blir fel kan man alltid skylla på att kartan var dålig just där.

Det finns dock en annan utgångspunkt, och det är det som denna avhandling handlar om. Om vi tränger under skinnet på atomerna så ser vi att de egentligen består av en positivt laddad *kärna* samt ett antal negativa *elektroner* som snurrar runt kärnan. Vidare är elektronerna inte tvungna att stanna runt "sin" kärna utan snurrar fritt runt alla kärnor i molekylen och till och med mellan molekylerna. Lyckligtvis finns det en fysikalisk teori, *kvantmekaniken*, som beskriver hur elektronerna fördelar sig och hur molekylerna därmed växelverkar med varandra. Vi kan alltså räkna ut hur kartan ser ut!

Tyvärr är kvantmekanik praktiskt tillämpbar endast på små molekyler. Att räkna ut energin noggrant i en punkt på vår protein–ligand-karta skulle ta flera miljarder år med dagens datorer. Och en miljondels miljondels sekund senare har alla dagisbarnen flyttat sig och man får räkna om alltihop igen.

En lösning på detta är att i huvudsak använda den förenklade bilden med atomer, men att utnyttja den komplicerade bilden med elektroner för att förbättra kartan lite i taget. Olika varianter på detta tema behandlas i de olika artiklarna i denna avhandling. Exempelvis handlar den första artikeln om hur repulsionen mellan molekylerna kan uppskattas utifrån hur mycket elektronerna tillhörande vardera molekylen överlappar med varandra, medan den andra handlar om hur noggrant man kan uppskatta den elektriska växelverkan mellan t.ex. positiva och negativa molekylytor med hjälp av noggranna räkningar på vardera ytan. De sista artiklarna handlar om hur man kan använda den komplicerade bilden för att räkna ut växelverkan mellan liganden och små bitar av proteinet, en i taget, medan effekter som beror på hela proteinet beräknas med den enkla bilden.

Resultatet av dessa ansträngningar skulle kunna bli ett kraftfält för protein– ligand-interaktionsenergier som är helt byggt på fysikaliska principer (*first principles*), därav titeln på denna avhandling. Tyvärr är vi inte där ännu, men min förhoppning är att dessa små steg kan hjälpa till att så småningom uppnå detta mål.

## List of papers

- I. Comparison of overlap-based models for approximating the exchangerepulsion energy
  P. Söderhjelm, G. Karlström, and U. Ryde, *Journal of Chemical Physics*, **104**, 244101 (2006)
- II. Accuracy of distributed multipoles and polarizabilities: Comparison between the LoProp and MpProp models
  P. Söderhjelm, J. W. Krogh, G. Karlström, U. Ryde, and R. Lindh, Journal of Computational Chemistry, 28, 1083 (2007)
- III. Accuracy of typical approximations in classical models of intermolecular polarization
  P. Söderhjelm, A. Öhrn, U. Ryde, and G. Karlström, Journal of Chemical Physics, 128, 014102 (2008)
- IV. On the coupling of intermolecular polarization and repulsion through pseudo-potentials
  P. Söderhjelm and A. Öhrn, *Chemical Physics Letters*, DOI:10.1016/j.cplett.2008.11.074 (2008)
- V. How accurate may a force field become? A polarizable multipole model combined with fragment-wise quantum-mechanical calculations
  P. Söderhjelm and U. Ryde, Journal of Physical Chemistry A, in press
- VI. Calculation of protein–ligand interaction energies by a fragmentation approach combining high-level quantum chemistry with classical many-body effects
  P. Söderhjelm, F. Aquilante, and U. Ryde, Submitted to Journal of Physical Chemistry B
- VII. Ligand affinities estimated by quantum chemical calculations P. Söderhjelm, J. Kongsted, and U. Ryde, Manuscript

Other papers, not included in the thesis.

- Ligand affinities predicted with the MM/PBSA method: dependence on the simulation method and the force field
   A. Weis, K. Katebzadeh, P. Söderhjelm, I. Nilsson, and U. Ryde, Journal of Medicinal Chemistry, 49, 6596 (2006)
- Conformational dependence of charges in protein simulations P. Söderhjelm and U. Ryde, Journal of Computational Chemistry, DOI: 10.1002/jcc.21097
- Proton transfer at metal sites in proteins studied by quantum mechanical free-energy perturbations
   M. Kaukonen, P. Söderhjelm, J. Heimdal, and U. Ryde, Journal of Chemical Theory and Computation, 4, 985 (2008)
- A QM/MM-PBSA method to estimate free energies for reactions in proteins

M. Kaukonen, P. Söderhjelm, J. Heimdal, and U. Ryde, Journal of Physical Chemistry B, **112**, 12537 (2008)

- On the performance of quantum chemical methods to predict solvatochromic effects. The case of acrolein in aqueous solution K. Aidas, A. Mogelhoj, E. Nilsson, M. S. Johnson, K. V. Mikkelsen, O. Christiansen, P. Söderhjelm, and J. Kongsted, Journal of Chemical Physics, **128**, 194503 (2008)
- Protein influence on electronic spectra modelled by multipoles and polarisabilities

P. Söderhjelm, C. Husberg, A. Strambi, M. Olivucci, and U. Ryde, Journal of Chemical Theory and Computation, in press

- The ozone ring closure as a test for multi-state multi-configurational second order perturbation theory (MS-CASPT2)
   L. De Vico, L. Pegado, J. Heimdal, P. Söderhjelm, and B. O. Roos, *Chemical Physics Letters*, 461, 136 (2008)
- Combined computational and crystallographic study of the oxidised states of [NiFe] hydrogenase
   P. Söderhjelm and U. Ryde,
   Journal of Molecular Structure: THEOCHEM, 770, 199 (2006)
- How accurate are really continuum solvation models for druglike molecules?

J. Kongsted, P. Söderhjelm, and U. Ryde, Submitted to *Journal of Chemical Theory and Computation* 

## Acknowledgements

Först vill jag tacka **Ulf** för ditt varma ledarskap. Din dörr har alltid varit öppen, både bildligt och bokstavligt, och inget har varit för smått eller stort för att ta upp med dig. När jag drivit iväg åt mitt håll har du inte hållit mig tillbaka, utan tvärtom själv följt efter för att kunna vara det stöd som jag själv inte alltid har förstått att jag behövt. Du representerar även en nästan ofattbar effektivitet, som ändå lämnar plats för eftertanke när det behövs. Jag tror också att ditt pragmatiska synsätt på den akademiska världen och på "livspusslet" samt alla konkreta tips kommer att vara till stor nytta för mig.

Ett stort tack också till **Gunnar** för dina många goda idéer och för att du alltid haft tid att svara på mina dumma frågor. Jag har väl varit lite envis och gjort saker på mitt sätt, men å andra sidan, det är du också...

Anders för alla givande samtal. Jag tror och hoppas att lite av din "precision" i arbetet har smittat av sig, även om ditt ordningssinne aldrig gjorde det. Vi lärde oss en del den hårda vägen, men jag gillade verkligen att arbeta med dig och hoppas vi någon gång kan återta samarbetet.

Biogruppen: Jacob för gott samarbete, synd att vi inte hade mer överlapp. Markus för att du tog tag i hydrogenas-tråden och verkligen fick något gjort. Jimmy för samarbeten och allt möjligt annat, inte minst alla goda skratt. Lubos för att du tog hand om mig när jag kom till Lund och delade med dig av din syn på vetenskap – du spred verkligen värme. Kasper för din härliga humor och dina goda råd. Kristina för att du ofta lättade upp stämningen. Patrik för din sköna stil och dina försök att lära mig dansa. Yawen och Magnus för trevlig samvaro. Thomas för all hjälp – önskar bara att jag hade lyssnat mer på dig. Aaron och Kambiz för gott samarbete. Torben för god hjälp i början. Samuel i slutet.

**Francesco** för hjälp med Molcas och för allt roligt genom åren, för sällskap sena jobbkvällar och tennismatcher som det gärna kunde blivit fler av. **Luca** för att du mot alla odds fick igenom ozonarbetet och naturligtvis för alla trevliga fester, äventyr på Sicilien och härliga skratt. **Asbjørn** för din avslappnade attityd. Det blev inte så mycket konkret samarbete, men mycket småsaker har klarnat med din hjälp och Norge var kul. **Daniel** för att du hjälpte mig igång med NEMO.

Per-Åke och Björn för alla utomordentliga kurser, både hemma och i Torre

Normana. Roland för Molcas-insikter och för att du vågade släppa mig lös i koden. Valera och Per-Olof för all möjlig hjälp med Molcas. Jesper för gott samarbete. Bo för att du säger vad du tycker, Mikael och Martin för Machjälp, och hela statmek-folket för många intressanta seminarier. Luis för fint samarbete. Roland Kjellander & co. för en inspirerande kurs.

**Matlaget** för alla härliga middagar och för att ni lät mig utveckla min experimentella sida. Och för diskussioner om allt från vetenskapliga principer till get-relaterade frågor. **OMM Journal club** för att ni vidgade mina horisonter.

All other people that have been around and made this department to such a nice place, especially Juraj, Ajitha, Sasha, Thomas B., Takashi, Claudio, Giovanni, Tomas, Dorota, and Sergey.

Alla på **BPC** för givande kafferaster. Särskilt **Houman**, **Stina**, **Olga**, **Ingemar**, **Erik**, **Mikael** och **Carl** för allehanda trevligt, **Wei-Feng** och **Torbjörn** för att ni försökte lära mig spela badminton och **Sara** för att du drog med mig till orienteringen.

Ett särskilt tack till **Eva** och **Bodil** för all hjälp med pappersarbetet, i synnerhet allt krångel i samband med pappaledigheten. Ni är problemlösare av sällan skådat slag och har också stor del i den mysiga stämningen vi haft på avdelningarna.

Jag vill också tacka er på AstraZeneca, särskilt **Ingemar** för din helhetssyn och ärliga intresse, **Lars** för att du fått mig att börja förstå solvation, **Ola** för goda idéer och användbara litteraturtips och **Anders** för din expertkunskap. Ni har alla bidragit till detta arbete och också fått mig att inse att vetenskapliga principer också gäller på ett stort företag. Jag hoppas att ni får ut något av mitt arbete, trots att det inte blev så läkemedelsinriktat som vi alla avsåg från början.

Henning för att du gav mig något helt annat än avhandlingen att tänka på, och Lotta och Staffan för allt barnvaktande som lät mig tänka på avhandlingen igen. Överhuvudtaget min familj och svärfamilj för allt stöd jag fått genom dessa år. Och alla vänner som genom er nyfikenhet ständigt fått mig att reflektera över vad jag gör. Slutligen, tack Emma att du stått ut med att mitt humör beror på svaren på beräkningarna och kötiden på datorklustren. Tack för att du får mig att behålla helhetssynen på livet genom att få alla stunder att bli betydelsefulla.

## Chapter 1

# Chemistry from a theoretical perspective

From a chemist's point of view, matter consists of atomic nuclei and electrons. As these particles are quite small, they obey the laws of quantum mechanics. Thus, instead of describing a system by the coordinates and velocity of each particle (as in classical mechanics), we must use a *wave function*, i.e. a function of the positions of all particles

$$\Psi = \Psi(\boldsymbol{r}_1, \boldsymbol{r}_2, \dots \boldsymbol{r}_n, \boldsymbol{R}_1, \boldsymbol{R}_2, \dots \boldsymbol{R}_N), \qquad (1.1)$$

where  $r_i$  and  $R_i$  are electronic and nuclear positions, respectively, and we have assumed that we are dealing with an isolated, unperturbed system so that the wave function is time-independent (apart from an ignored phase factor).

The connection to the classical picture is provided through the Born interpretation, stating that the probability of finding each particle *i* in a volume element  $d\mathbf{V}_i$  at position  $\mathbf{r}_i$  is equal to  $|\Psi|^2 d\mathbf{V}_1 d\mathbf{V}_2 \dots d\mathbf{V}_{n+N}$ .

One of the postulates of quantum mechanics states that the wave function must satisfy the Schrödinger equation:

$$\widehat{H}\Psi = E\Psi,\tag{1.2}$$

where  $\widehat{H}$  is the Hamiltonian operator describing the system and the eigenvalue E is the particular energy corresponding to  $\Psi$ . For our system of nuclei and electrons, the Hamiltonian is given by

$$\hat{H} = -\frac{1}{2} \sum_{i=1}^{n} \nabla_{i}^{2} - \frac{1}{2} \sum_{A=1}^{N} \frac{1}{M_{A}} \nabla_{A}^{2} - \sum_{i=1}^{n} \sum_{A=1}^{N} \frac{Z_{A}}{|\mathbf{r}_{i} - \mathbf{R}_{A}|} + \sum_{i} \sum_{j>i} \frac{1}{|\mathbf{r}_{i} - \mathbf{r}_{j}|} + \sum_{A} \sum_{B>A} \frac{Z_{A} Z_{B}}{|\mathbf{R}_{A} - \mathbf{R}_{B}|},$$
(1.3)

where i, j count over electrons, A, B count over nuclei, and  $M_A$  and  $Z_A$  are nuclear masses and charges, respectively (atomic units are used in this and the following equations).

As the nuclei are much heavier than the electrons ( $M_A = 1854$  for hydrogen), it is usually possible to assume that the electrons move independently of the nuclear motion, and thus the total wave function can be separated into a product of nuclear and electronic wave functions,

$$\Psi = \Psi_{nuc}(\mathbf{R}_1, \mathbf{R}_2, ... \mathbf{R}_N) \Psi_{ele}(\mathbf{r}_1, \mathbf{r}_2, ... \mathbf{r}_n; \{\mathbf{R}_A\}),$$
(1.4)

where the nuclear coordinates have been included as a collective argument to  $\Psi_{ele}$  as a reminder that the electronic Hamiltonian determining  $\Psi_{ele}$  depends on the nuclear positions through their electric potential. This separation, known as the *Born–Oppenheimer approximation*, turns out to be extremely practical. The nuclear and electronic problems can now be treated separately, with the Hamiltonian for the nuclear problem reducing to

$$\widehat{H}_{nuc} = -\frac{1}{2} \sum_{A=1}^{N} \frac{1}{M_A} \nabla_A^2 + E_{pot}(\mathbf{R}_1, \mathbf{R}_2, ..., \mathbf{R}_N), \qquad (1.5)$$

where  $E_{pot}$  is the sum of the nuclear repulsion and the eigenvalue of the electronic Hamiltonian formed with these particular nuclear coordinates. In the following two sections, the electronic and nuclear problems are treated separately. The nuclear problem will lead us into statistical mechanics.

#### 1.1 The electronic problem: Quantum chemistry

The electronic Hamiltonian is given by

$$\widehat{H}_{ele} = -\frac{1}{2} \sum_{i=1}^{n} \nabla_i^2 - \sum_i \sum_{A}^{N} \frac{Z_A}{|\mathbf{r}_i - \mathbf{R}_A|} + \sum_{i=1}^{n} \sum_{j>i}^{n} \frac{1}{|\mathbf{r}_i - \mathbf{r}_j|}$$
(1.6)

Unfortunately, it is impossible to find exact solutions to the Schrödinger equation with this Hamiltonian for anything more complicated than the  $H_2^+$  ion (n = 1, N = 2), so we are left with doing approximations. A very useful theorem in this context is the *variational principle*, which states that for any approximate wave function, the expectation value of  $\hat{H}_{ele}$  will be larger than that obtained with the correct wave function. This reduces the problem to devising a set of trial wave functions and using the one with lowest energy as the best approximation to the real wave function.

The most common starting point for such procedure is to assume that each electron has its own one-electron wave function, called an *orbital*. Because the electrons are indistinguishable fermions, the antisymmetry principle must be obeyed, and thus a direct product of such orbitals is not a valid wave function. Instead, the simplest allowed wave function is an anti-symmetrized product, also called a *Slater determinant*:

$$\Psi_{ele}(\boldsymbol{r}_1, \boldsymbol{r}_2, ... \boldsymbol{r}_n) \propto \mathcal{A}\left[\psi_1(\boldsymbol{r}_1)\psi_1(\boldsymbol{r}_2)...\psi_n(\boldsymbol{r}_n)\right]$$
(1.7)

where the antisymmetrizer  $\mathcal{A}$  is defined by

$$\mathcal{A} = \frac{1}{n!} \sum_{P \in S_n} (-1)^{\pi} \widehat{P}$$
(1.8)

where the sum goes over all n! permutations of electron labels 1...n and the operator  $\hat{P}$  changes the labels accordingly, with  $\pi$  being the number of transpositions in the permutation. For example,

$$\mathcal{A}[\psi_1(\mathbf{r}_1)\psi_2(\mathbf{r}_2)] = \frac{1}{2} \left( \psi_1(\mathbf{r}_1)\psi_2(\mathbf{r}_2) - \psi_1(\mathbf{r}_2)\psi_2(\mathbf{r}_1) \right)$$
(1.9)

Inserting Eq. 1.7 into the Schrödinger equation, applying the variational principle, and integrating over spin (assuming an even number of electrons) gives the closed-shell Hartree–Fock equations

$$-\frac{1}{2}\nabla^{2}\psi_{j}(\boldsymbol{r}) - \sum_{A}^{N} \frac{Z_{A}\psi_{j}(\boldsymbol{r})}{|\boldsymbol{r} - \boldsymbol{R}_{A}|} + \sum_{i=1}^{n/2} \left[ 2\int \frac{|\psi_{i}(\boldsymbol{r}')|^{2}\psi_{j}(\boldsymbol{r})}{|\boldsymbol{r}' - \boldsymbol{r}|} d\boldsymbol{r}' - \int \frac{\psi_{i}^{*}(\boldsymbol{r}')\psi_{j}(\boldsymbol{r}')\psi_{i}(\boldsymbol{r})}{|\boldsymbol{r}' - \boldsymbol{r}|} d\boldsymbol{r}' \right] = \epsilon_{j}\psi_{j}(\boldsymbol{r})$$

$$(1.10)$$

which is an eigenvalue equation that can be solved to obtain a (in principle infinite) set of  $\{\epsilon_j, \psi_j\}$  pairs, where the  $\psi_j$  corresponding to the n/2 smallest eigenvalues are the *occupied orbitals* that builds up the wave function and the rest are called *virtual orbitals*. Note that the occupied orbitals occur in the equations, so the equations must be solved iteratively.

To make the physical interpretation of Hartree–Fock theory more apparent, Eq. 1.10 can be written in terms of two one-electron operators

$$\hat{h}\psi_j + \hat{v}^{HF}\psi_j = \epsilon_j\psi_j \tag{1.11}$$

where  $\hat{h}$  takes care of the two first terms of Eq. 1.10, i.e. is the sum of kinetic energy operator and nuclear attraction operator, and  $\hat{v}^{HF}$  is an effective operator representing the mean field of the other electrons. It consists of the remaining two terms of Eq. 1.10: an intuitive *Coulomb term*, as well as an *exchange term* that arises from the anti-symmetry principle. Note that the inclusion of the electron's interaction with itself is only apparent because it cancels between the Coulomb and exchange terms.

After finding self-consistent orbitals, the Hartree–Fock energy is given by

$$E_{HF} = 2\sum_{i=1}^{n/2} \left\langle \psi_i | \hat{h} | \psi_i \right\rangle + \sum_{i=1}^{n/2} \sum_{j=1}^{n/2} \left( 2 \left\langle ij | ij \right\rangle - \left\langle ij | ji \right\rangle \right),$$
(1.12)

where we have introduced the standard Dirac notation for integrals

$$\left\langle \psi_a | \hat{h} | \psi_b \right\rangle = \int \psi_a^*(\boldsymbol{r}) \hat{h} \psi_b(\boldsymbol{r}) d\boldsymbol{r}$$
 (1.13)

and used the physicist's notation for two-electron repulsion integrals:

$$\langle ab|cd\rangle = \int \int \frac{\psi_a^*(\boldsymbol{r}_1)\psi_b^*(\boldsymbol{r}_2)\psi_c(\boldsymbol{r}_1)\psi_d(\boldsymbol{r}_2)}{|\boldsymbol{r}_1 - \boldsymbol{r}_2|} d\boldsymbol{r}_1 d\boldsymbol{r}_2$$
(1.14)

The electron density is simply given by

$$\rho(\mathbf{r}) = 2\sum_{i}^{n/2} |\psi_i(\mathbf{r})|^2$$
(1.15)

The importance of the Hartree–Fock method cannot be overestimated, because it provides the very useful picture of molecular orbitals and usually gives good qualitative results for many molecules and their interactions. Moreover, it is the basis for more advanced treatments, *post-Hartree–Fock* methods.

#### 1.1.1 The basis set problem

Exact solutions to Eq. 1.10 can only be found for atoms, so in practice one solves Eq. 1.10 by expanding the orbitals in a set of basis functions, i.e.

$$\psi_{i} = \sum_{\mu=1}^{K} C_{\mu i} \phi_{\mu} \tag{1.16}$$

Insertion of this expansion into the Hartree–Fock equations gives a matrix eigenvalue problem to solve in each iteration [3]. This is very suitable for computation. As expected, in this finite basis, only K orbitals are obtained.

Clearly, the number of basis functions and their specific form determines the quality of the obtained orbitals and thereby the results. Most calculations utilize atomic orbital basis sets, which are inspired by the eigenfunctions of the hydrogen atom. Although basis functions of Slater type  $(\exp[-ar])$  describe the orbitals better, most quantum-chemical software use linear contractions of Gaussian functions  $(\exp[-ar^2])$  instead, because the integrals become significantly easier to compute.

It is important that the basis set describes the valence orbitals well (preferably by at least three basis functions each, *triple valence*) and that it includes functions with higher angular momentum (*polarization functions*). For properties and interaction energies, which are the main interest in this thesis, it is also important that the outer region (*tail*) of the electron density is well described. This is normally done by including some functions with much slower decay (*diffuse functions*), giving an *augmented* basis set.

#### 1.1.2 The electron-correlation problem

In Hartree–Fock theory, the electron repulsion is only treated in an average manner, so the instantaneous *electron correlation* is missing. For example, when a hydrogen molecule is stretched out, the two electrons (which by antisymmetry are forbidden to simply "choose" one nucleus each) cannot lower their energy by avoiding to spend time at the same nucleus, because in the mean-field approach the repulsion is the same whichever nucleus they choose to be near. There are several methods to include electron correlation in quantum-chemical calculations, but only the two employed in this thesis will be described.

#### Perturbation theory

Perturbation theory is a general approximation method in quantum mechanics. In its simplest form, it involves partitioning the Hamiltonian into an "easy" part  $(\hat{H}_0)$  that we already know the solution to, and a "tricky" part (the perturbation  $\hat{V}$ ) that we want to approximate by exploiting the solutions to  $\hat{H}_0$ . The advantage of this approach is that if the perturbation is turned on gradually, i.e.

$$\widehat{H} = \widehat{H}^{(0)} + \lambda \widehat{V}, \qquad (1.17)$$

where  $\lambda$  goes from 0 to 1, the ground-state energy will vary as

$$E_0 = E_0^{(0)} + \lambda E_0^{(1)} + \lambda^2 E_0^{(2)} \dots$$
(1.18)

(i.e. a normal Taylor expansion) and if the perturbation is small enough, each  $E^{(n)}$  will be smaller in magnitude than the preceding one. By inserting Eq. 1.17 into the Schrödinger equation and collecting terms of the same *order* (i.e. with the same  $\lambda$ -dependence), we can obtain expressions for the various energy corrections [4]. The results for the first- and second-order corrections are

$$E_0^{(1)} = \left\langle \Psi_0 | \widehat{V} | \Psi_0 \right\rangle$$
$$E_0^{(2)} = \sum_{n \neq 0} \frac{\left\langle \Psi_0 | \widehat{V} | \Psi_n \right\rangle \left\langle \Psi_n | \widehat{V} | \Psi_0 \right\rangle}{E_0 - E_n},$$
(1.19)

where  $\Psi_0, \Psi_1, \Psi_2, ...$  are all eigenfunctions to  $\hat{H}^{(0)}$  and  $E_0, E_1, E_2, ...$  are the corresponding eigenvalues (energies).

In the perturbational treatment of electron correlation, usually called  $M \emptyset ller-Plesset$  perturbation theory, the perturbation represents the difference between the real electron repulsion and the mean-field electron repulsion, i.e.

$$\widehat{V} = \sum_{i=1}^{n} \sum_{j>i}^{n} \frac{1}{|\boldsymbol{r}_i - \boldsymbol{r}_j|} - \sum_{i=1}^{n} \widehat{v}_i^{HF}, \qquad (1.20)$$

where  $\hat{v}_i^{HF}$  acts on the *i*th electron. At first order, one simply removes the double-counting of electron interactions in the second term of Eq. 1.20 and thus recovers the Hartree–Fock energy. At second order, one obtains

$$E_0^{(2)} = \sum_{i,j} \sum_{a,b} \frac{2|\langle ij|ab\rangle|^2 - \langle ij|ab\rangle\langle ab|ji\rangle}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b},$$
(1.21)

where i, j count over occupied orbitals, a, b count over virtual orbitals, and  $\epsilon_i$  is the orbital energy of orbital  $\psi_i$ . The sum of the uncorrelated (Hartree–Fock) energy and the approximate correlation energy in Eq. 1.21 is usually called the MP2 energy, and will be used as a reference level in most of this thesis.

Although MP2 recovers most of the correlation energy, the remaining difference is still rather large, as can be expected because the orbitals are still optimized for the mean-field situation. However, it turns out that for relative energies, which are the only important energies in chemistry, MP2 is usually a good approximation. For more accurate results, the perturbation series can be continued (MP3, MP4, etc.), but in general it is better to include a variational optimization of the correlated wave function and only use perturbation theory as a small correction, as in the *coupled cluster singles and doubles with perturbative triples* (CCSD(T)) method.

#### Density functional theory

A more empirical approach to quantum chemistry starts from the Hohenberg– Kohn theorem [5], which states that the ground-state energy (and all groundstate electronic properties) is uniquely determined by the electron density  $\rho(\mathbf{r})$ , so that, for practical purposes, the much more complicated many-electron wave function  $\Psi_{ele}(\mathbf{r}_1, \mathbf{r}_2, ..., \mathbf{r}_n)$  is not needed. Exactly how the energy depends on  $\rho(\mathbf{r})$ , i.e. the *density functional*  $E[\rho(\mathbf{r})]$  is unfortunately unknown.

To exploit the good treatment of e.g. the kinetic energy in wave-function theory, most practical implementations of density functional theory still follows a Hartree–Fock like approach called the Kohn–Sham method [6]. It simply replaces the one-electron operator  $\hat{v}^{HF}$  by an empirical operator  $\hat{v}^{KS}$  that depends only on the electron density, i.e. not on the individual orbitals. As the Coulomb part of  $\hat{v}^{HF}$  already fulfils this requirement, only the exchange part is usually changed, so that

$$\hat{v}^{KS} = \int \frac{\rho(\mathbf{r}')}{|\mathbf{r}' - \mathbf{r}|} d\mathbf{r}' + \hat{v}^{XC}[\rho(\mathbf{r})], \qquad (1.22)$$

where  $v^{\hat{X}C}$  is usually called the exchange–correlation functional. As the name suggests, it contains not only an approximate exchange contribution but also an approximate correlation contribution. In the simplest case, the *local density approximation*, it is a simple function of the density and relates to the exchange–correlation energy of the uniform electron gas. However, the most commonly used functionals depend also on the gradient of the density. A large variety of exchange–correlation functionals have been developed, and the explicit parameterization of these may build on comparisons with either high-level wave-function methods or experimental information. In this thesis, density functional theory is used only to generate an electron density that includes electron correlation, and thus most available functionals will give similar results.

#### 1.2 The nuclear problem: Statistical mechanics

#### 1.2.1 Physical description

In the previous section, we saw how to (approximately) compute the potential energy surface that determines the motion of the atomic nuclei. A curious consequence of the mathematical properties of the Schrödinger equation combined with the anti-symmetry principle for electrons is the formation of aggregates of several nuclei, which we normally call *molecules*. Molecules are characterized by *covalent bonds* between the nuclei. These bonds are strong, i.e. the force constants for the energy wells are large. In contrast, the *non-covalent bonds*, which form the intermolecular interactions of special interest in this thesis, are usually weaker with more shallow and diffuse energy wells.

As seen from Eqs. 1.4 and 1.5, the nuclear motion should be described by a wave function  $\Psi_{nuc}(\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N)$ . Such treatment is necessary for e.g. determining the vibrational energy levels of a molecule. For a small molecule, solving the nuclear Schrödinger equation gives the result that the nuclei oscillate around an equilibrium geometry, which is the global minimum of the potential energy surface and can be found by geometry-optimization techniques [7] normally integrated in the quantum-chemical softwares. Except for this *zero-point vibration*, most systems actually behave rather classically, i.e. the nuclei moves on the potential energy surface like a couple of balls on a curved (multi-dimensional) surface. Whenever nuclear quantum effects (e.g. tunneling) are important, they can normally be treated as corrections to this picture.

When we consider a macroscopic system, the notion of eigenstates  $\Psi_{nuc}$  becomes completely meaningless, because the energy levels are so densely spaced that the system changes state forth and back in a rather chaotic manner. The behavior of such systems is therefore governed by the laws of *statistical mechanics*. A basic postulate is that all microscopic states with the same energy are equally probable. For an isolated system, only states with a given energy are allowed (by the law of energy conservation) so they all occur with same probability. On the other hand, for a system that can transfer heat with the surroundings (at constant temperature), the probability depends on the energy. This is not because states with lower energy are "better" but simply because if the energy of the system is smaller, that implies that the energy of the surroundings is larger (again by the law of energy conservation) and thus a greater number of microstates of the surroundings are available. The probability for the system to be in one particular state i is therefore given by the Boltzmann distribution:

$$p_i = \frac{\exp\left(-\frac{E_i}{k_B T}\right)}{\sum_j \exp\left(-\frac{E_j}{k_B T}\right)},\tag{1.23}$$

where  $E_i$  is the energy of state *i*,  $k_B$  is the Boltzmann constant, and *T* is the temperature. The denominator of Eq. 1.23 is called the *partition function*, *Q*. Thus, the average energy of the system is a competition between the number of states that have this energy (favoring high energies) and the Boltzmann factor (favoring low energies).

A process for which Q increases is known as a *spontaneous* process. For historical reasons, it is customary to convert Q to a *free energy* (Helmholtz in case of constant volume or Gibbs in case of constant pressure; for condensed systems they are roughly the same so we will not make a distinction) by the formula

$$G = -k_B T \ln Q \tag{1.24}$$

As seen from Eq. 1.23, the value of Q can be increased in two principal ways: Either the number of states at each energy can be increased or the energy levels themselves can be decreased. The first option roughly corresponds to an increase in entropy whereas the second option corresponds to a decrease in enthalpy, i.e. the two driving forces for spontaneous transitions that we are used to from classical thermodynamics.

If the system is treated fully classically, the sum over states Q can be replaced by an integral and further partitioned into one part that depends only on the momenta of the nuclei and another that depends only on the coordinates. Moreover, the momentum part can be integrated out to obtain the following classical probability density:

$$P(\boldsymbol{R}_1,...\boldsymbol{R}_N) = \frac{\exp\left(-\frac{E_{pot}(\boldsymbol{R}_1,...\boldsymbol{R}_N)}{k_BT}\right)}{\int_{-\infty}^{\infty}...\int_{-\infty}^{\infty}\exp\left(-\frac{E_{pot}(\boldsymbol{R}'_1,...\boldsymbol{R}'_N)}{k_BT}\right)d\boldsymbol{R}'_1...d\boldsymbol{R}'_N} = \frac{\exp\left(-\frac{E_{pot}}{k_BT}\right)}{Z},$$
(1.25)

where Z is sometimes called the configuration integral and is related to Q by a multiplicative constant (depending on e.g. the nuclear masses) and thus gives identical free energy differences.

By mere statistics, the probability distribution over the energy is more peaked the more degrees of freedom there are in the system (except for quantum effects such as electronic states). Obviously, for a macroscopic system at room temperature, the global energy minimum is of no interest, because there will be only one or a few such microstates. All occurring microstates will have almost the same energy, given by a sum of the kinetic energy  $3Nk_BT/2$  and a systemdependent potential energy. On the other hand, for a microscopic system, e.g. a single protein molecule (of the size 10–100 Å), there is a finite number of available conformations and the energy fluctuates.

#### 1.2.2 Computation

The microstates with highest probability contribute most to the observed properties of the system. Thus, for assessing such properties theoretically, we need a method to generate a representative collection of microstates, usually called an *ensemble*. The best statistics is obtained if the probability for a microstate to occur in the ensemble is given by Eq. 1.23, because then a macroscopic property can be computed as a simple average over the members of the ensemble. The two most common methods for generating ensembles are *molecular dynamics* (MD) and *Metropolis Monte Carlo* (MC) simulations. In the MD approach, the system is propagated in time by stepwise integrating Newton's equations of motion (a small modification is needed to allow for heat exchange with the surroundings). In the MC approach, a random change is done at each step and the change is accepted or rejected by an energy criterion designed to guarantee a Boltzmann-weighted ensemble. From the *ergodic hypothesis* follows that the (infinite) ensembles generated by these methods are equivalent.

Statistical properties, such as free energies or entropies, are not possible to express as such averages [8]. However, relative free energies are in principle obtainable from a single simulation. One just counts the number of ensemble members that can be classified as being of type A and type B, respectively, and then compute the free energy difference as

$$\Delta G(A \to B) = -k_B T \ln \frac{N_B}{N_A} \tag{1.26}$$

Unfortunately, this procedure is seldom applicable in practice because it may require too much sampling before the ratio converges. Usually, one is interested in two macroscopic states that are separated by an energy barrier, and if the barrier is significantly higher than  $k_BT$ , the simulation will mainly stay on the same side of the barrier. Several solutions to this problem have been suggested, e.g. artificially suppressing the barrier (umbrella sampling) [8] or running the simulation at several temperatures (replica exchange) [9].

If the two states of interest have different Hamiltonians, e.g. being chemically different, the simple counting approach will not work. A possible solution is *thermodynamic integration* (TI), in which the Hamiltonian is written as a continuous function of a coupling parameter  $\lambda$  taking values from 0 to 1:

$$H(\lambda) = (1 - \lambda)^m H_A + \lambda^m H_B, \qquad (1.27)$$

where one normally uses m = 1. The exact free energy difference is then given by

$$\Delta G(A \to B) = \int_0^1 \left\langle \frac{\partial H(\lambda)}{\partial \lambda} \right\rangle_{\lambda} d\lambda \tag{1.28}$$

where  $\langle ... \rangle_{\lambda}$  denotes an average over the ensemble generated with  $H(\lambda)$ .

A related method is *free energy perturbation* (FEP) [10], in which the same difference is computed by

$$\Delta G(A \to B) = -k_B T \ln \left\langle \exp\left(\frac{E_B - E_A}{k_B T}\right) \right\rangle_A.$$
 (1.29)

Although this technique does not directly solve the problem with overcoming barriers, it is often possible to divide the perturbation into several small parts and evaluate each part separately, making it similar to the thermodynamic integration approach. In fact, if one takes the average over the two "perturbation directions" and Taylor-expands the exponential function, one obtains [11]

$$\Delta G(A \to B) \approx \frac{1}{2} \left( \left\langle E_B - E_A \right\rangle_A + \left\langle E_B - E_A \right\rangle_B \right), \tag{1.30}$$

which is in fact equivalent to using the one-interval trapezoid rule when evaluating the integral in Eq. 1.28. This limiting approximation common to both TI and FEP is sometimes called the *linear response approximation* (see section 3.3).

## Chapter 2

## Intermolecular interactions

In this chapter, we are interested in the potential energy as a function of the *intermolecular* degrees of freedom, whereas we assume that the *intramolecular* degrees of freedom, i.e. bond lengths, bond angles, and dihedral angles, are fixed. For simplicity, we consider a dimer of two molecules, which we denote A and B.

The A-B dimer can be treated as any other system of nuclei and electrons, i.e. as a *supermolecule*. As a result, the two monomers are no longer distinguishable, as illustrated in the right part of Fig. 2.1. On the other hand, the weak character of the intermolecular interactions sometimes allows us to focus on the actual interaction. The latter, *perturbative* approach, is illustrated in the left part of Fig. 2.1. We will see that it opens up immense possibilities for approximations.

#### 2.1 The supermolecular approach

In Section 1.1, we saw how to compute the energy of an arbitrary configuration of nuclei with a number of electrons around them. Although it must be done in an approximate way, the accuracy is systematically improvable. The most straight-forward way to calculate the interaction energy is therefore to calculate the energy of the AB dimer and subtract the A and B monomer energies. This gives the *supermolecular* interaction energy.

#### 2.1.1 Electron correlation

The supermolecular approach is applicable to any (size-consistent) level of theory, but of course the quality of the results depends on the applied theory. For interaction energies, electron correlation is usually important, giving rise to e.g. the dispersion attraction (see Section 2.2.1). Thus, to obtain quantitative results, Hartree–Fock theory is inadequate: at least second-order perturbation



Figure 2.1: Illustration of the supermolecular and perturbative approaches to interaction energies

theory (MP2) must be applied. In fact, the interaction energy converges slowly with respect to the included level of electron correlation and current benchmark studies typically involve CCSD(T) theory. However, this limits the applicability to clusters with 20–30 atoms, so much effort has been spent on determining the accuracy of MP2 and other methods in relation to the CCSD(T) method.

Such studies have shown that MP2 (at the complete basis set limit) often overestimates the correlation contribution to the interaction energy [12]. The *spin-component-scaled MP2* (SCS-MP2) method [13] is a partly empirical method to improve the MP2 method with no extra cost. It is based on a separate scaling of the correlation-energy contributions from electron pairs with parallel and antiparallel spin, respectively. Although SCS-MP2 seems to improve the interaction energy of complexes involving aromatic stacking, the description of hydrogen bonding becomes worse than with standard MP2 [14].

A computationally very attractive solution would be to use density functional theory (DFT). However, using exchange–correlation functionals that involve only the local density and its gradient, it is impossible to model the longrange electron correlation responsible for the dispersion. In many cases, e.g. hydrogen bonding, this deficiency is partly cancelled by the overestimation of charge transfer and other non-additive effects [15], so that the full interaction potential is reasonable. However, for e.g. aromatic interactions, the result is poor. Various solutions to this problem have been proposed [16], including methods based on perturbation theory, new exchange–correlation functionals, and empirical methods treating dispersion as in force fields (see Section 2.3.3). Although promising results have been obtained, supermolecular DFT energies will not be further discussed in this thesis.

There are in fact completely different approaches to electronic structure calculation that intrinsically include electron correlation. In particular, the Diffusion Monte Carlo (DMC) method has been applied to the calculation of interaction energies with high accuracy [17]. This method treats the time-dependent Schrödinger equation as a diffusion equation and propagates the wave function by Monte Carlo integration methods towards the exact solution. Interestingly, the only systematic error in DMC, the fixed-node approximation, effectively cancels out when computing interaction energies by the supermolecular approach.

#### 2.1.2 Basis set

Unfortunately, supermolecular interaction energies computed at a post-Hartree– Fock level also converge very slowly with basis set. A particular difficulty is the *basis set superposition error* (BSSE), which originates from the fact that the basis set of the dimer can describe the monomer charge density better than the monomer basis set, simply because it is larger. By the variational principle, the BSSE contribution to the interaction energy is always negative. Although the contribution vanishes at the complete basis limit, it is substantial for all normal basis sets [18].

The most common way to address the BSSE is the *counterpoise proce*dure [19], in which all energies are calculated in the same basis set, the dimer basis set:

$$E_{AB}^{sup} = E_{A+B} - E_{A+(B)} - E_{B+(A)}$$
(2.1)

where  $E_{X+(Y)}$  denotes the energy of monomer X when including the basis functions of monomer Y (ghost orbitals) without the nuclear charges. There has been significant debate whether the counterpoise procedure is the best way to eliminate the BSSE [18]. A common argument against it is that the electrons are "too free" in the monomer calculations, because they can occupy also the space that, in the dimer, is occupied by the other monomer and therefore forbidden by the antisymmetry principle. For this reason, it has been suggested to use only the virtual orbitals of the other monomer as ghost orbitals. However, it has been shown that the counterpoise-corrected result is indeed a pure interaction energy [18]; the restriction of the available space is a real (repulsive) effect that should be included.

A remaining concern is the *higher-order BSSE* [20], which refers to the modified properties of the monomers in the dimer calculation due to the (asymmetrically) extended basis set. However, when computing interaction energies, one can in fact benefit from this effect (for example, the polarizabilities become better), provided that the static density is sufficiently well described in the monomer basis set [20]. On the other hand, when computing the effect of interactions on various molecular properties, one should correct for the higher-error BSSE by also computing the monomer properties in the dimer basis set [21] (see paper III).

Even when the counterpoise procedure is applied, very large basis sets are required before interaction energies are converged. For example, with the augcc-pVTZ basis set, an underestimation of the correlation energy contribution by 5–10% is typical [22]. A solution to this problem is *complete basis set* (CBS) extrapolation, in which the complete basis set value is estimated from a series of calculations with affordable basis sets. This has been shown to significantly improve the results [23]. Several extrapolation schemes have been devised, but the most common one fits the correlation energies  $E_n$  obtained at two or more different basis set levels (e.g. aug-cc-pVTZ, n = 3; and aug-cc-pVQZ, n = 4) to the simple expression [24]:

$$E_n = E_{CBS} + An^{-3} \tag{2.2}$$

where  $E_{CBS}$  and A are fitting parameters,  $E_{CBS}$  being the sought correlation energy at the CBS limit.

The basis-set convergence can be somewhat improved by supplementing the conventional dimer-centered basis set with functions located between the monomers [25]. A more rigorous way, directly addressing the correlation energy, is to include terms into the wave-function that depend explicitly on the inter-electronic distances, as in the MP2-R12 method [26]. Although such treatments formally leads to a large number of 3- and 4-electron integrals, the computation of these can be avoided by resolution-of-identity approaches. It has been shown that interaction energies obtained with the MP2-R12 method with an augmented double-zeta basis set are already closer to the CBS limit than the conventional MP2/aug-cc-pV5Z energies [22]. Similar methods have been devised at the CCSD(T) level [27, 28].

A more pragmatic approach, which is widely used, is to rely on error cancellation between the inadequate treatment of electron correlation and the limited basis set. For example, it has been suggested to use a smaller basis set (e.g. cc-pVTZ) together with MP2 to compensate for the overestimated correlation energy [29], but the situation is complicated by the fact that for hydrogenbonded complexes, the inclusion of diffuse functions give better results [12]. Extrapolations to both the CBS and to the CCSD(T) level (with a modest basis set) therefore seems to be a more reliable approach [16].

In conclusion, the counterpoise-corrected supermolecular approach is well established as a reliable method for calculating interaction energies. As long as the system of interest is small enough to allow for a large basis set and a good treatment of electron correlation, it is the preferred method for most types of interactions. Moreover, for evaluating the accuracy of simpler methods, the situation is even better. Within a given theory and basis set, the counterpoisecorrected supermolecular result can often be regarded as exact. It will therefore be used as a reference for several of the methods discussed in this thesis.

#### 2.2 Perturbative approach

A completely different picture is obtained if one directly calculates the interactions between molecules. In most such descriptions, one obtains the total



Figure 2.2: Decomposed interaction energy for the hydrogen bond between water and the oxygen atom of propionamide as a function of the intermolecular separation.

interaction energy as a sum of terms, each with a distinct physical meaning. This can be useful for interpreting which physical effects are dominant in the studied complex and, most importantly, it can be used to derive simplified, computationally cheaper methods. Nevertheless, it should be remembered that only the total interaction energy is an observable, so any decomposition is ambiguous.

#### 2.2.1 Qualitative picture

There are four main contributions to the total interaction energy [30]: electrostatic energy, induction energy, dispersion energy, and exchange-repulsion energy, i.e.

$$E_{tot} = E_{ele} + E_{ind} + E_{disp} + E_{rep} \tag{2.3}$$

The typical magnitude and distance-dependence of each of these terms is illustrated in Fig. 2.2.

The electrostatic energy is easily understandable in terms of Coulomb's law: a negatively charged part of one monomer attracts a positively charged part of the other, whereas two like charges repel each other. In cases where molecules (or interacting functional groups) are charged or polar, the electrostatic energy tends to dominate the interaction and thus for the most probable configurations, the electrostatic energy is attractive.

The induction energy (or polarization energy) is the energy change (always negative) obtained by polarizing each monomer wave function in response to the electric field from the other monomer. The induction is usually the dominant attractive contributions if one molecule is polar (or charged) and the other non-polar, but its important role in e.g. hydrogen bonding has also been demonstrated [31]. It is the major contributor to many-body (non-additive) effects, i.e. the fact that the interaction energy of a cluster does not equal the sum of

pairwise interaction energies. The role of induction in this context depends on that each monomer responds to the *total* electric field from the other monomers. If, for example, the fields from two neighboring monomers cancel each other, the total induction energy vanishes, whereas the attractive induction contribution to each of the two pairwise interaction energies may be substantial.

The dispersion is usually the dominant attractive contribution if both molecules are non-polar. It arises from the coupling of instantaneous fluctuations in the monomer charge distributions [4]. If an instantaneous dipole arises in one monomer, it will induce a dipole in the other monomer giving an instantaneous "energy kick". Clearly, this is an electron-correlation effect: if the fluctuations on one monomer are averaged they vanish and can not give any interaction.

The exchange-repulsion term has no classical counterpart. Nevertheless, it is extremely important for the constitution of matter, as it balances the attractive forces and prevents the electron clouds of closed-shell molecules to overlap, effectively causing the *shape* of the molecules. The reason for this behavior of molecules is the Pauli principle, stating that two electrons cannot occupy the same quantum state. This gives rise to an effective repulsive force between electrons of the same spin. The Pauli principle, in turn, is a direct consequence of the anti-symmetry principle for fermions.

#### 2.2.2 Quantum-chemical perturbation theory

The formal description of interaction energies is most conveniently obtained by perturbation theory. The interaction part of the total Hamiltonian is given by

$$\hat{V} = \hat{H} - \hat{H}_A - \hat{H}_B \tag{2.4}$$

or, explicitly

$$\hat{V} = \sum_{I \in A} \sum_{J \in B} \frac{Z_I Z_J}{r_{IJ}} - \sum_{I \in A} \sum_{j \in B} \frac{Z_I}{r_{Ij}} - \sum_{i \in A} \sum_{J \in B} \frac{Z_J}{r_{iJ}} - \sum_{i \in A} \sum_{j \in B} \frac{1}{r_{ij}}$$
(2.5)

where I, J denote nuclei, i, j denote electrons, Z is the nuclear charge, and r the interparticle distance.

#### Polarization approximation

If one applies regular Rayleigh–Schrödinger (RS) perturbation theory (Eq. 1.19), using  $\hat{H}_0 = \hat{H}_A + \hat{H}_B$  and  $\Psi_0 = \psi_0^A \psi_0^B$ , i.e. the direct product of unperturbed monomer wave functions, one obtains at first order

$$E^{(1)} = \left\langle \psi_0^A \psi_0^B | \hat{V} | \psi_0^A \psi_0^B \right\rangle = \sum_{I \in A} \sum_{J \in B} \frac{Z_I Z_J}{r_{IJ}} - \sum_{I \in A} \int \frac{Z_I \rho_B(\mathbf{r})}{|\mathbf{r} - \mathbf{r}_I|} d\mathbf{r} - \sum_{J \in B} \int \frac{Z_J \rho_A(\mathbf{r})}{|\mathbf{r} - \mathbf{r}_J|} d\mathbf{r} + \int \frac{\rho_A(\mathbf{r}_1) \rho_B(\mathbf{r}_2)}{|\mathbf{r}_1 - \mathbf{r}_2|} d\mathbf{r}_1 d\mathbf{r}_2$$
(2.6)

which is easily recognized as the classical electrostatic interaction energy. At second order, one obtains

$$E^{(2)} = -\sum_{a>0} \frac{\left|\left\langle\psi_{0}^{A}\psi_{0}^{B}|\hat{V}|\psi_{a}^{A}\psi_{0}^{B}\right\rangle|^{2}}{E_{a}^{A} - E_{0}^{A}} - \sum_{b>0} \frac{\left|\left\langle\psi_{0}^{A}\psi_{0}^{B}|\hat{V}|\psi_{0}^{A}\psi_{b}^{B}\right\rangle|^{2}}{E_{b}^{B} - E_{0}^{B}} - \sum_{a>0} \sum_{b>0} \frac{\left|\left\langle\psi_{0}^{A}\psi_{0}^{B}|\hat{V}|\psi_{a}^{A}\psi_{b}^{B}\right\rangle|^{2}}{E_{a}^{A} - E_{0}^{A} + E_{b}^{B} - E_{0}^{B}},$$

$$(2.7)$$

where  $\psi_a^A$  and  $\psi_b^B$  are the wave functions of the excited states of monomers A and B, respectively, and  $E_a^A$  and  $E_b^B$  are the corresponding eigenvalues (energies). The three terms arise from a convenient classification of the excitations of the dimer into three classes. The first term, involving excitations of monomer A, can be recognized as the polarization of A by the static charge distribution of B and the second term analogously as the polarization of B by A. The third term, involving simultaneous excitations of monomer A and B, is the dispersion energy. At higher orders, one obtains e.g. terms corresponding to the coupling of polarization in monomers A and B.

The RS treatment neglects exchange effects, i.e. the Pauli principle is not imposed between the molecules. This approximation is usually called the *polarization approximation* and its consequences for force field development will be discussed in paper III. Obviously, the polarization approximation is only valid in the long-range limit of intermolecular interactions.

#### Symmetry-adapted perturbation theory

Several perturbative treatments that include exchange effects have been devised [32, 33]. Observing that the direct product  $\Psi_0$  is normally very different from the "true" dimer wave function [33], it is clear that  $\hat{V}$  cannot be considered as a small perturbation. A more reasonable choice of ground-state function is the antisymmetrized product  $\mathcal{A}\Psi_0$ , where  $\mathcal{A}$  is the antisymmetrizer, defined in Eq. 1.8.

However, as  $\mathcal{A}\Psi_0$  is not an eigenfunction of  $\hat{H}_0$ , one must either modify the partitioning of the total Hamiltonian so that  $\mathcal{A}\Psi_0$  becomes an eigenfunction of the new  $\hat{H}_0$ , or modify the actual RS perturbation scheme. The first option is used in *intermolecular perturbation theory* (IMPT) [34], whereas the second option leads to the more used symmetry-adapted perturbation theories (SAPT).

Several SAPT theories have been developed [33,35]. The simplest and most used variant is the *symmetrized Rayleigh-Schrödinger* (SRS) theory. It uses *weak symmetry forcing*, which means that the antisymmetrizer is only used in the energy expressions. Thus, the perturbed wave functions obtained with SRS will be identical to these obtained within the polarization approximation. The first- and second-order contributions to the energy are given by

$$E_{SAPT}^{(1)} = N_0 \left\langle \Psi_0 | \mathcal{A} \hat{V} | \Psi_0 \right\rangle$$

$$E_{SAPT}^{(2)} = -\sum_{m>0} \frac{N_0 \left\langle \Psi_0 | \mathcal{A} \hat{V} | \Psi_m \right\rangle \left\langle \Psi_m | \hat{V} | \Psi_0 \right\rangle}{E_m - E_0} - E_1^{SAPT}$$
(2.8)

where  $N_0 = \langle \Psi_0 | \mathcal{A} \Psi_0 \rangle$ . By a decomposition of the antisymmetrizer into interand intramonomer parts, the energy at each order can be divided into one term that is identical to the energy from polarization theory and one term that represents the *exchange* contribution. Moreover, a decomposition of  $E_2^{SAPT}$ similar to that done in Eq. 2.7 can be performed. Thus, the SAPT interaction energy can be written as

$$E_{SAPT} = E_{ele}^{(1)} + E_{exch}^{(1)} + E_{ind}^{(2)} + E_{disp}^{(2)} + E_{exch-ind}^{(2)} + E_{exch-disp}^{(2)} + E^{higher}.$$
 (2.9)

This expression clearly resembles Eq. 2.3, although the terms do not need to be equally defined, as will be discussed in Section 2.3.

In principle, the monomer wave functions inserted into perturbation theory should be the exact monomer wave functions. In practice, however, they are always approximate. The most common type of monomer wave functions are Hartree–Fock wave functions. In this case,  $E_{SAPT}^{(1)}$  becomes almost identical to the *Heitler–London* interaction energy, which is the energy obtained in a supermolecular HF calculation if the unperturbed monomer orbitals are used without any subsequent iterations. In principle, the equality requires a complete basis, but in fact it holds for any basis set if the dimer-centered basis set is used when computing the monomer wave functions [36]. Such procedure is often used in SAPT, although more effective basis sets have been proposed [37].

Contributions to the interaction energy from the intramonomer electron correlation can be calculated by a double perturbation approach [33], with one perturbation being of the type in Eq. 1.20 and the other as in Eq. 2.5. SAPT can be used as a stand-alone method to compute interaction energies, but is often used together with a supermolecular calculation at a lower level of theory, because the latter includes all higher-order terms, some of which are cumbersome to compute in SAPT.

#### Supermolecular decomposition schemes

Interaction energies can also be decomposed within the supermolecular approach. The advantage of this is that all higher-order terms are included. Numerous decomposition schemes have been developed over the years, but in this thesis only the Kitaura–Morokuma (KM) scheme [38] and the *restricted virtual space* (RVS) scheme [39, 40] will be considered. Both give the Heitler–London energy, i.e. the sum of electrostatic and exchange-repulsion terms, but they

differ in the decomposition of the remaining term, sometimes called the *defor*mation energy. The main difference between the KM and RVS schemes [41] is that the polarization energy in the KMD scheme is computed without account of the Pauli effects, and thus the model breaks down as the intermolecular distance becomes short or as the basis set becomes more complete [42]. The RVS method includes Pauli effects by letting each monomer be polarized under the anti-symmetry constraints imposed by the frozen molecular orbitals of the other monomer. Thus, the polarization term is more physical and has better convergence properties. However, the use of frozen orbitals prohibits the selfconsistent treatment of the polarization of both molecules (see paper III for a validation). In this respect, the RVS model resembles the SAPT methods in that higher-order polarization terms are ignored.

#### 2.2.3 Molecular mechanics

A molecular mechanics force field is a simplified description of molecules, where one has eliminated the electronic degrees of freedom and only consider the interactions and movements of atomic nuclei. If the force field allows the molecules to be flexible, the energy expression contains bonded terms in addition to the non-bonded terms that are always present. A typical functional form [11] of the bonded terms is

$$E_{bonded} = \sum_{i}^{bonds} \frac{k_i}{2} (l_i - l_{i,0})^2 + \sum_{i}^{angles} \frac{k_i}{2} (\theta_i - \theta_{i,0})^2 + \sum_{i}^{torsions} \frac{V_i}{2} (1 + \cos(n_i \omega_i - \gamma_i)),$$
(2.10)

where  $l_i$  and  $\theta_i$  are bond lengths and angles, respectively,  $l_{i,0}$  and  $theta_{i,0}$  are the corresponding reference values,  $k_i$  is the force constant,  $\omega_i$  is the torsional angle, and  $V_i$ ,  $n_i$ , and  $\gamma_i$  are constants.

The non-bonded contribution normally follows the decomposition in Eq. 2.3 (a polarizable force field), or omits the induction term (a non-polarizable force field). We will frequently refer to a standard molecular-mechanics force field, by which we imply the following functional form for the non-bonded terms:

$$E_{non-bonded} = \sum_{i}^{atoms} \sum_{j>i}^{atoms} \left[ \frac{q_i q_j}{r_{ij}} - \frac{A_{ij}}{r_{ij}^6} + \frac{B_{ij}}{r_{ij}^{12}} \right],$$
 (2.11)

where  $r_{ij}$  is the distance between atoms *i* and *j*,  $q_i$  is the partial charge of atom *i*, and  $A_{ij}$  and  $B_{ij}$  are fitted parameters. The first term is the electrostatic term, whereas the other two together constitutes a Lennard-Jones interaction, roughly corresponding to the dispersion and repulsion terms. This is the most common functional form used in force fields for biomolecular systems, e.g. Amber, CHARMM, OPLS, and GROMOS.

When going beyond standard force fields, a multitude of variants exist. The specific functional forms that can be used for each non-bonded term, as well as

how to extract the needed parameters from quantum-chemical calculations, will be discussed in the next section.

A force field can be based on quantum chemistry (see section 2.3), experimental results, or (most commonly) a combination of both. In this thesis, we will focus on methods purely based on quantum chemistry. Such approach has an important philosophical value: If we reproduce a wide range of experimental quantities with a purely theoretical model, we can deduce that the model is reasonably correct and base our physical theories upon the model. The same is not true if we have included experimental knowledge into the force field: the model then simply becomes a predictive tool, regardless of how much physics it contains.

On the other hand, the inclusion of experimental data when constructing the force fields has one clear advantage: If one accepts that the underlying physics is too complex to be studied in detail, one can hope to capture some of the missing effects by including experimental data. This applies in particular to the approximations done in quantum chemistry (i.e. insufficient basis set and treatment of electron correlation), but with similar reasoning one can also try to include other effects (e.g. from solvent or entropy) so that the potential energy surface becomes an effective potential (potential of mean force).

Most molecular-mechanics studies are performed using non-polarizable force fields, partly because of the reduced computational cost and partly because of the vast experience of non-polarizable force fields collected through various applications. The reason for the success of non-polarizable force fields is not that the induction energy is negligible in all these applications, but that only the total interaction energy matters. To reproduce experimental data in the condensed phase (e.g. in water solution, where polarization is known to be important), the charges must typically be enhanced to simulate an average polarization [43]. Clearly, this limits the transferability of the force fields to other types of systems.

A polarizable force field, on the other hand, builds more solidly on physical principles and each term can in principle be related to the corresponding term in a quantum-chemical perturbation treatment. Parameters of polarizable force fields therefore have the prospect to be much more transferable [44]. However, the molecular-mechanics picture is still extremely simplified and only a few studies have explicitly demonstrated the greater transferability [45].

#### 2.3 Molecular mechanics based on quantum chemistry

We are now ready to tackle one of the main themes of this thesis, namely how to extract molecular mechanics parameters from quantum chemistry, both from monomer calculations, supermolecular calculations, and perturbation theory. We will frequently refer to the *Sum of Interactions Between Fragments Ab initio computed* (SIBFA) method [46], the *effective fragment potential* (EFP) method [47], and the NEMO method [48, 49], three prominent examples of molecular mechanics force fields developed along these lines. Other examples include the PFF [50] and QMPFF [51] methods. Out of the force fields that can be directly applied for protein simulations, Amoeba [52] is probably the most theoretically founded, whereas the Amber [53] and CHARMM [54] polarizable variants are more empirical. More detailed accounts of polarizable force fields can be found in recent reviews [55, 56].

#### 2.3.1 Electrostatic interaction energy

A natural goal of an electrostatic model is to reproduce the exact electrostatic energy, Eq. 2.6, accurately, but with less computational effort. This can be done by replacing the continuous monomer charge density,  $\rho(r)$ , with a set of discrete charges and possibly higher multipoles. The interaction energy between two such *multicenter-multipole expansions* is given by [57]:

$$E_{ele} = \sum_{i}^{A} \sum_{j}^{B} \left[ \frac{q_i q_j}{|\boldsymbol{r}_{ij}|} + (q_i \boldsymbol{\mu}_j - q_j \boldsymbol{\mu}_i) \cdot \boldsymbol{\nabla} \left( \frac{1}{\boldsymbol{r}_{ij}} \right) + (q_i \boldsymbol{\Theta}_j - \boldsymbol{\mu}_i \boldsymbol{\mu}_j + q_j \boldsymbol{\Theta}_i) \cdot \boldsymbol{\nabla} \boldsymbol{\nabla} \left( \frac{1}{\boldsymbol{r}_{ij}} \right) \dots \right],$$
(2.12)

where  $\mathbf{r}_{ij}$  is the distance vector from center *i* to center *j*, and  $q_i$ ,  $\boldsymbol{\mu}_i$ , and  $\boldsymbol{\Theta}_i$  are the charge, dipole, and quadrupole in center *i*, respectively. In standard force fields, only the first term is used.

#### **Determining multipoles**

There are two main approaches for extracting multipoles form a quantumchemical calculation, viz. to fit the multipoles to reproduce the electrostatic potential (ESP) around the molecule or to directly analyze the charge density.

In the former case, one typically selects a large number of points around the molecule and minimizes the mean squared deviation between the quantumchemical ESP and the ESP generated from atomic charges. Unfortunately, the charges depend significantly on the way the points are selected [58] and especially charges of buried atoms are ill-defined. A way to reduce these problems is to fit the charges locally to the ESP from density-derived multipoles [59]. In force fields for flexible molecules, the ESP charges are typically restrained [60] or averaged over several conformations [61,62] to reduce their conformational dependence. The ESP method has also been used for higher multipoles [63].

Methods that analyze the charge density include the Mulliken [64] and the Löwdin [65] population analyses, the distributed multipole analysis (DMA) [66] and similar approaches [67, 68], and the natural atomic orbitals (NAO) analysis [69]. These methods avoid the arbitrary selection of points, but instead they are quite sensitive to the basis set used in the quantum-chemical calculation. Several solutions to this problem have been proposed, e.g. the atoms
in molecules (AIM) scheme [70, 71] and other topological partitioning methods [72–74], as well as improved population analysis methods [75, 76]. The convergence of distributed multipole expansions has been studied for both electrostatic energies [77] and the ESP (see paper II). It should be noted that multipoles obtained from the density are usually significantly worse than ESPderived multipoles (of the same order) at reproducing electrostatic energies [78]; therefore an "extra" multipole order must typically be used.

#### **Beyond multipoles**

Regardless of how multipoles are obtained, they cannot reproduce the electrostatic energy when the overlap between the charge densities of the monomers is significant. Unfortunately, this is the case for most interactions of interest, e.g. hydrogen bonds in their energy minimum. The difference between the true electrostatic energy and the multipolar interaction energy is usually called *charge penetration energy*, and is attractive for most interactions (because the repulsion between the electron clouds is the first contribution to be damped when the monomers start to overlap). In most force fields, this term is absorbed into the repulsion energy, as it has a similar dependence on the overlap.

Electrostatic models that explicitly includes the charge penetration have appeared. Noting that (for Gaussian basis sets) the exact electron density is a sum of Gaussians distributed throughout the molecule, one solution is to reduce the number of Gaussians by a fitting procedure [79–81]. A similar approach is the Gaussian Electrostatic Model (GEM) [82–85], which, inspired by the density fitting method [86], first employed an analytical minimization of the Coulomb metric for the density deviation, but later changed to a numerical (grid-based) fitting [85]. GEM has been combined with the SIBFA force field [46]. Other groups have used pure numerical integration [87,88] or combinations of numerical integration and multipole integration [89,90]. Within the EFP approach, a reduction of the Coulomb integrals to orbital overlap integrals (which are much easier to compute) has been done for s-functions [91], as well as the more pragmatic approach to include damping functions for the multipoles [92]. A similar method [93] has also been used in SIBFA. The drawback with all these methods, except the last two, is the increased computational time taken for evaluating the electrostatic energy, compared to a pure multipole approach.

#### 2.3.2 Polarization

Polarization models [44] are a central theme in this thesis. The polarization problem is significantly more complex than the electrostatic interaction. The reason that so much effort has been spent on improving the electrostatic interaction is not its complexity, but the fact that it is usually larger in magnitude and more long-ranged and therefore more important to model accurately (another reason is of course the dominance of non-polarizable force fields). One difficulty with polarization is the lack of unambiguous reference quantities. A naive idea would be to calculate the exact electric field from each monomer, apply that in the quantum-chemical calculation of the other monomer, and iterate until self-consistency. This is the *exact electric field polarization* (EPOL) method defined in paper III, and it is (for Hartree–Fock theory) identical to the polarization energy in the Kitaura–Morokuma decomposition [38]. The problem with this approach is the same as for the RS perturbation theory, i.e. that the obtained wave function does not obey the anti-symmetry principle. The consequences of this problem will be thoroughly discussed in paper III.

Another difficulty is that one normally wants to describe the polarization as a classical physical process, whereas it is really a quantum-mechanical process (as can be seen from Eq. 2.7) describing the response of the wave function to an applied electric field. The classical picture arises from the *linear response approximation*, stating that when a molecular system is subjected to an electric field  $\mathbf{F}$ , the induced dipole moment is given by

$$\boldsymbol{\mu}_{ind} = \boldsymbol{\alpha} \cdot \boldsymbol{F} \qquad (\text{e.g. } \boldsymbol{\mu}_x = \alpha_{xx}F_x + \alpha_{xy}F_y + \alpha_{xz}F_z) \qquad (2.13)$$

where  $\alpha$  is called the *polarizability tensor* of the system. The advantage of this description is that it allows us to ignore the quantum-chemical details of the charge redistribution. Certainly, the value of  $\alpha$  must be computed quantum-chemically (typically by applying weak electric fields), but this is only done once. For developing and testing polarization models, it is useful to go back to the quantum-chemical description (see paper III). The polarizability tensor  $\alpha$  is often replaced by a scalar quantity (the *isotropic polarizability*), defined by  $(\alpha_{xx} + \alpha_{yy} + \alpha_{zz})/3$ . For an atom, this substitution is of course exact.

For two interacting molecules, linear response can usually still be assumed, but the complication is twofold. First, the electric field is not homogeneous, i.e. it varies in different parts of the molecule. Second, the induced dipole moment is not a very useful quantity, because it does not specify where in the molecule the largest response occurs. A solution to these problems is the *distributed point polarizability model*, which is the most common model in polarizable force fields, used in both SIBFA, EFP, and NEMO. In this model, each molecule contains a set of polarizabilities. At each center i, i.e. at the position of the polarizability  $\alpha_i$ , the induced dipole  $\mu_i^{ind}$  is given by:

$$\boldsymbol{\mu}_{i}^{ind} = \boldsymbol{\alpha}_{i} \cdot \boldsymbol{F}_{i} = \boldsymbol{\alpha}_{i} \cdot \left[ \boldsymbol{F}_{i}^{stat} - \sum_{j \neq i} \boldsymbol{\mu}_{j}^{ind} \boldsymbol{\nabla} \boldsymbol{\nabla} \left( \frac{1}{\boldsymbol{r}_{ij}} \right) \right]$$
(2.14)

where the electric field  $E_i$  in that position has been divided into contributions from the static charge distribution and from other induced dipoles. Writing up Eq. 2.14 for each center *i* defines a linear system of equations, which can be solved through matrix inversion or by iteration. As the static term tends to dominate Eq. 2.14, the iterative approach usually converges in ~10 iterations. The induction energy is given by

$$E^{ind} = -\frac{1}{2} \sum_{i=1}^{N} \boldsymbol{\mu}_i \cdot \boldsymbol{E}_i^{stat}$$
(2.15)

where the factor 1/2 comes from the fact that the (positive) work of polarization cancels half of the energy of the dipoles in the field. Note that the energy contribution from a pair of induced dipoles is completely cancelled by the work of polarization, so Eq. 2.15 contains only the static field (but the coupling is anyway included through the values of  $\mu_i$ ).

The static field  $E_i^{stat}$  may in principle come from the exact charge distribution (as in paper III), but is normally computed from the multipole expansion:

$$\boldsymbol{E}_{i}^{stat} = -\sum_{j \neq i} \left[ q_{j} \boldsymbol{\nabla} \left( \frac{1}{\boldsymbol{r}_{ij}} \right) + \boldsymbol{\mu}_{j} \cdot \boldsymbol{\nabla} \boldsymbol{\nabla} \left( \frac{1}{\boldsymbol{r}_{ij}} \right) + \boldsymbol{\Theta}_{j} \cdot \boldsymbol{\nabla} \boldsymbol{\nabla} \boldsymbol{\nabla} \left( \frac{1}{\boldsymbol{r}_{ij}} \right) \dots \right] \quad (2.16)$$

where, in the simplest versions of polarizable force fields, only the first term is used.

In Eqs. 2.14 and 2.16, the contribution to the field from all surrounding centers are treated equally. In practice, one often omits or damps contributions from certain (typically close-lying) induced or static multipoles, but the reasons for doing so vary in the two cases. For the coupling between polarizabilities in Eq. 2.14, the induced dipoles may become infinite at small distances (the "polarization catastrophe"), and the polarization can become unphysical even before that happens. Although it is possible to reproduce the molecular polarizability by a set of fully coupled atomic polarizabilities, significantly smaller values must be used than with an uncoupled set [94]. An alternative is to damp the dipole interaction tensor for overlapping charge distributions. The most used such damping is the one by Thole [95], sometimes reparameterized [96]. Other options include omitting contributions from atoms that are directly bonded [97] or separated by less than three bonds [53], for example. In most methods for deriving polarizabilities that are based directly on quantum-chemical calculations (see below), the coupling is ignored between polarizabilities in the same molecule, because it is already implicitly included through the quantum-chemical calculations. The explicit and implicit coupling schemes have been compared [98] and will also be discussed in Paper V. To reduce the computational cost, it has been suggested to omit all polarizability coupling (i.e. not only the intramolecular part) and instead slightly enhance the charges [99] or to include the coupling to first order [100]. These approximations assume that the effect of the coupling is a minor part of the total effect of polarization.

On the other hand, for the response to the static field, which we will refer to as *static polarization*, the choice of which interactions to include in Eq. 2.16 (and possible damping) is more a question of definition. Indeed, for a rigid molecule, it is simple to change between on one hand a representation where the static field from multipoles in the same molecule are included in Eq. 2.16, and on the other hand a representation where all such contributions are excluded and instead modeled by permanent dipoles, i.e. a "pre-polarized" molecule. Although these representations give equal results, the latter can be preferable for qualitative conclusions, as will be discussed in paper V. For a flexible molecule, the static polarization can be used as a way to model how the charge distribution (or the multipole representation of it) varies with changes in the molecular geometry [52, 101–104].

#### Determining distributed point polarizabilities

The determination of distributed point polarizabilities from quantum-chemical calculations is more cumbersome than for multipole moments [44], mainly because of two issues.

First, the response is a more complicated property. The electrostatic properties are simply determined by the charge density  $\rho(\mathbf{r})$ ; and distributed multipole moments are simply expectation values over local partitions of  $\rho(\mathbf{r})$ . However, for fully describing the polarization of a system, we need to know the change in charge density for an arbitrary external potential. Assuming linear response, this can be expressed as

$$\Delta \rho(\mathbf{r}) = \int \alpha(\mathbf{r}, \mathbf{r}') V(\mathbf{r}') d\mathbf{r}', \qquad (2.17)$$

where  $V(\mathbf{r'})$  is the external potential and  $\alpha(\mathbf{r}, \mathbf{r'})$  is the charge density susceptibility function, describing the change in density in one point caused by a certain electric potential in another point. A direct partitioning and "multipole expansion" of the charge density susceptibility function (using e.g. coupled perturbed Hartree–Fock theory) therefore leads to both local and non-local point polarizabilities [105], the latter giving the induced dipole at one center caused by the electric field at another center. Such description is awkward in practice, because the number of polarizabilities varies as  $N^2$  with the number of centers N). However, if the partitioning is done properly, the non-local polarizabilities can be eliminated [106]. A simpler approach is to use homogeneous electric fields for the perturbation and assume from the outset that the charge density susceptibility function is local [76, 107].

Second, a direct analysis of  $\Delta \rho(\mathbf{r})$  will give monopoles as well as dipoles (higher-order multipoles are normally discarded). This represents the "charge flow" due to the polarization, which is a realistic process but cannot be modeled in a pure point-polarizability model. Therefore, they are normally converted to point polarizabilities. As this process is of course approximate, it is advantageous for a partitioning scheme to give small monopole polarizabilities [108]. In another context, monopole polarizabilities have been used to characterize chemical bonds [109].

The partitioning of the response can mainly be done by the same two approaches as the density partitioning for multipoles. In the grid-based approach,

one typically runs a large set of calculations each having a perturbing point charge at one grid point. The induced electric potential at that particular grid point [110] or at all grid points simultaneously [111, 112] is used to fit a set of polarizabilities up to an arbitrary multipole order. The main advantage of this procedure is that the set of dipole polarizabilities can be optimized to describe effects related to higher-order polarizabilities as well. A drawback is that it is unable to assign polarizabilities reliably to "buried" atomic sites.

It was early recognized that basis-space partitioning of the response can lead to large charge flow and unphysical polarizabilities, especially for basis sets containing diffuse functions [113]. The reason is the same that causes the basis-set dependence of distributed multipoles, i.e. that when adding together basis sets for each atom, the resulting basis set may be "overcomplete" (illconditioned). Thus, the application of even a very weak electric field can cause significant changes in the MO coefficients without the total density being affected very much. As for multipoles, one solution to this problem is to use a real-space (topological) partitioning [114,115]. More computationally appealing methods include a density fitting approach [108], a localized molecular orbitals approach [107], and the LoProp approach [76] (see papers II and III), in which a more physical partitioning reduces the numerical problems and the charge flow terms are minimized by a Lagrange multiplier approach. A quite different approach is to express the polarizability in terms of force operators that can be decomposed into nuclear contributions without reference to any basis set [116].

The numerical problems can also be avoided by not performing the perturbed calculations explicitly, as in the uncoupled Hartree–Fock approximation [117], which is also much faster. Although the molecular polarizability is not reproduced exactly in such approaches, they may still be useful in practice, because higher-order molecular polarizabilities (which crucially depend on the distribution method) may be important as well. Moreover, a simple scaling is often possible (see paper II).

All the methods considered above are "method-consistent", i.e. the polarizabilities obtained at a certain quantum-chemical level (and a certain basis set) give interaction energies comparable to supermolecular results at the same level. However, it may sometimes be possible to cure deficiencies in the interaction model by employing polarizabilities computed at another level. For example, it has been suggested to avoid diffuse functions when computing polarizabilities, in order to capture some of the Pauli effects present in the condensed state [50, 118].

#### Beyond point polarizabilities

As has already been mentioned, non-local (two-site) polarizabilities, monopole (charge flow) polarizabilities, and higher-order polarizabilities can be derived by similar methods. In fact, dipole–quadrupole polarizabilities have been found to improve the description of the response [119]. Explicit inclusion of charge flow

does not seem to give a significant effect for most systems, as will be demonstrated in paper III. An illustrative example is given in Figure 2.3, which shows the response of a formamide molecule with two different polarization models when interacting strongly. Although the response looks quite different for the model including charge flow, the induced electric potential around the molecule is similar (see paper III). On the other hand, monopole polarizabilities are expected to be advantageous for describing conjugated  $\pi$ -systems, for example.



Figure 2.3: Response of a monomer of the formamide dimer from paper III, using two different polarization models: the polarizability model (left) with only dipole polarizabilities, and the LPOL model (right) with charge flow. Numbers denote induced charges in thousandths of the elementary charge. Arrows denote induced dipoles (1 cm  $\approx 0.084$  a.u.)

An alternative to the point-polarizability model is the *shell model* (also called the *Drude oscillator model*), in which each polarizability is modeled as a core charge +q and a shell charge -q connected by a harmonic spring with force constant k. These parameters can be derived by the same methods as the polarizabilities (see above), because of the direct relation:

$$\alpha = \frac{q^2}{k} \tag{2.18}$$

where  $\alpha$  is the isotropic polarizability. The reason for using Shell models is often to avoid the interaction tensor in Eq. 2.14 or to avoid the iterative procedure by treating the shell separation as a dynamical variable in a simulation. However, by letting the interaction site for the repulsion coincide with the shell charge, one can also model the coupling between induction and repulsion in an approximate but physically realistic way [44].

In contrast, pure fluctuating charge models are normally not based on polarizabilities. A common way to determine the charges is *electronegativity equalization*, in which an energy expression of the type

$$E(q_1, q_2, ...q_N) = \sum_{i}^{N} \left( \chi_i q_i + \frac{1}{2} J_i q_i^2 \right) + \sum_{i < j}^{N} \frac{f(r_{ij}) q_i q_j}{r_{ij}}, \qquad (2.19)$$

where N is the number of atoms,  $\chi_i$  and  $J_i$  are the standard electronegativity and hardness, respectively, and f(r) is a short-range damping function, is minimized under the constraint that the sum of the charges  $\{q_i\}$  is conserved. The advantage of fluctuating charge models is primarily the computational efficiency and simplicity, but they can also be coupled to the repulsion [44]. On the other hand, they obviously have problems with describing the response of symmetric systems (e.g. benzene), and are therefore sometimes combined with polarizabilities. A comparison of such mixed models with point-polarizability models [120] showed no particular advantage of including fluctuating charges.

#### 2.3.3 Dispersion

A multipole expansion of the interaction operator in the third term of Eq. 2.7 leads to the following asymptotic series for the dispersion energy between two atoms A and B:

$$E_{disp} = -\frac{C_6}{R^6} - \frac{C_8}{R^8} - \frac{C_{10}}{R^{10}}\dots$$
(2.20)

where R is the interatomic distance and  $C_n$  are known as dispersion coefficients. In the well-known London formula, the series is truncated after the first term and the dispersion coefficient is estimated [4] by

$$C_6 = \frac{3}{2} \alpha_A \alpha_B \left( \frac{I_A I_B}{I_A + I_B} \right) \tag{2.21}$$

where  $\alpha_X$  and  $I_X$  are the polarizability and ionization potential of monomer X, respectively. This approach can be directly extended to molecules by using a sum over atom-atom pairs and either treating the dispersion coefficients as adjustable parameters (as in standard force fields) or calculating them by a London-like formula [49].

To account for the effects of overlapping charge distributions, each term in Eq. 2.20 can be damped, e.g. by the much used Tang–Toennies model [121] or by models fitted to symmetry-adapted perturbation theory [122].

#### Determining dispersion coefficients

A direct calculation of the distributed dispersion coefficients from quantum chemistry usually proceeds via the frequent-dependent polarizabilities  $\alpha(i\nu)$ . In EFP, the dispersion is described using a sum over LMO (local molecular orbital) pairs involving anisotropic polarizabilities [123]:

$$E_{disp} = \sum_{j \in A} \sum_{k \in B} \sum_{\alpha\beta\gamma\delta} T_{\alpha\beta}^{jk} T_{\gamma\delta}^{jk} \int_0^\infty \alpha_{\alpha\gamma}^j(i\nu) \alpha_{\beta\delta}^k(i\nu) d\nu, \qquad (2.22)$$

where  $T_{xx}^{jk}, T_{xy}^{jk}, \dots$  are elements of the dipole interaction tensor for the distance between LMO centers j and k, and  $\alpha_{xx}^{j}, \alpha_{xy}^{j}, \dots$  are elements of the anisotropic frequency-dependent polarizability tensor located at center j. For practical reasons, one normally uses isotropic polarizabilities in the EFP dispersion, so that Eq. 2.22 reduces to a simple sum of  $R_{jk}^{-6}$  terms. Similar expressions have been used by others and augmented with  $R^{-8}$  and  $R^{-10}$  terms arising from the corresponding treatment of higher-order polarizabilities [124]. Analogously, the charge-flow polarizabilities give rise to lower-order terms (down to  $R^{-2}$ ) that cancel at long range, but including such terms may give numerical errors [124]. A non-expanded treatment of the dispersion, including higher-order terms as well as overlap effects, can be formulated in terms of frequency-dependent density susceptibilities [125], but to a significantly higher computational cost. The calculation of frequency-dependent polarizabilities (or susceptibilities) requires at least time-dependent Hartree–Fock (TDHF) or density functional theory (TDDFT). Distributing the polarizabilities over the molecule can be performed by the same type of methods as for static polarizabilities (see Section 2.3.2). A computationally simpler approach to distributed dispersion coefficients exploits a connection to the dipole moment of the exchange hole [126], which can also be localized using standard methods [127, 128].

#### 2.3.4 Repulsion

Similar to the case of polarization, there is no unambiguous definition of repulsion. One possible definition is the Hartree–Fock exchange–repulsion energy. This approach is taken by the EFP and SIBFA force fields and allows for a systematic improvement of each term based on physics. Another common approach, used in NEMO, is to use the repulsion as a rest term [49], defined e.g. as

$$E_{rep} = E_{tot} - E_{ele} - E_{ind} - E_{disp}.$$
(2.23)

It is then fitted to supermolecular energies for a set of interacting dimers [129]. With this definition, the repulsion will also contain overlap corrections to the induction, dispersion, and electrostatic energies, as well as charge transfer (see below). The advantages of this approach are that simpler expressions can be used for all terms and that errors inherent in each term can be effectively corrected for, but on the other hand the transferability is not guaranteed.

The Hartree–Fock exchange–repulsion energy can be obtained by subtracting the electrostatic energy from the Heitler–London interaction energy (or equivalently, by the Kitaura–Morokuma decomposition [38] scheme). It is often divided into an exchange contribution and a repulsion contribution [130]. The former is always attractive and is simply the sum of exchange contributions (second term in the i, j-sum of Eq. 1.12) from orbitals on different monomers. The latter is repulsive and always larger in absolute value, thus causing a net repulsion. Its origin can be qualitatively understood in terms of a depletion of electron density in the overlap region due to the antisymmetrization [131].

The common modeling of the repulsion by a term proportional to  $R^{-12}$  is mainly due to computational convenience; an  $R^{-9}$  or  $R^{-10}$  dependence has also been used [11]. A more realistic repulsion potential decays exponentially, as in the *Born–Mayer* model:

$$E_{rep} = \sum_{i \in A} \sum_{j \in B} \kappa_{ij} \exp\left(-\alpha_{ij} R_{ij}\right), \qquad (2.24)$$

where  $\alpha_{ij}$  and  $\kappa_{ij}$  are atom-atom parameters and  $R_{ij}$  is the interatomic distance. If needed, the potential can be modified to include the anisotropy of each atom-atom term [132], e.g. by

$$E_{rep} = \sum_{i \in A} \sum_{j \in B} \exp\left(-\alpha_{ij}(\omega_{ij})\left[R_{ij} - \rho_{ij}(\omega_{ij})\right]\right), \qquad (2.25)$$

where  $\omega_{ij}$  is the relative orientation of the local axes of atoms *i* and *j*, and  $\alpha_{ij}$  and  $\rho_{ij}$  are functions of  $\omega_{ij}$ . Thus, the main problem lies in determining the parameters.

#### Determining repulsive parameters

Asymptotically, the exchange repulsion is proportional to  $S^2$ , where S is the overlap integral of the monomer wave functions [133]. For Hartree–Fock monomer wave functions,  $S^2$  is given by

$$S^{2} = \sum_{i}^{A} \sum_{j}^{B} \langle \psi_{i} | \psi_{j} \rangle^{2}$$
(2.26)

where  $\psi_i$  and  $\psi_j$  are occupied molecular orbitals (MOs) of the isolated monomers A and B.

For further analysis, it is practical to start from the perturbative definition of exchange repulsion (i.e. the  $E_{exch}^{(1)}$  term of Eq. 2.9), which differs from the Hartree–Fock exchange repulsion only at order  $S^4$  [134]. An exact expression for  $E_{exch}^{(1)}$  has been given [134], but more computationally feasible formulas are usually based on an (infinite) expansion in terms of overlap integrals. For a dimer, only even powers of S occurs in the expansion and expressions for the  $S^2$  and  $S^4$  orders have been given. By approximating the  $S^2$  expression in various ways (e.g. the spherical Gaussian overlap approximation [135]), a parameter-free exchange-repulsion model [136] has been included in the EFP method. More approximate models are based directly on the squared orbital overlap of Eq. 2.26 [137, 138], sometimes modified by introducing an explicit distance dependence [139, 140].

Because the exact computation of  $S^2$  requires the wave functions of the monomers, it is rather impractical to use in a force field. Therefore, many models are instead related to the density overlap:

$$\Omega = \int \rho_A(\mathbf{r}) \rho_B(\mathbf{r}) d\mathbf{r}$$
(2.27)

where  $\rho_A$  and  $\rho_B$  are the electron densities of the isolated monomers. Approximate proportionality between the exchange repulsion and  $\Omega$  was found for both atomic [141] and molecular dimers [87,142]. The comparison of repulsion models based on orbital overlap ( $S^2$ ) and density overlap ( $\Omega$ ) is the subject of paper I.

The gain of relating the repulsion to the density overlap instead of the orbital overlap might not be immediately apparent. However, the density overlap is significantly easier to approximate. As in the charge-penetration case, the number of Gaussians used to represent the molecular density can be significantly reduced [80]. The density overlap model can also be used together with density-fitting techniques [83]. Moreover, the integral in Eq. 2.27 is amenable to numerical integration [87]. If one divides the monomer density into atomic contributions, each approximated by a spherical exponentially decaying function, the overlap model provides a way to directly calculate the repulsion parameters in simple potentials, e.g. Eq. 2.24 (see paper I). A related approach is used in NEMO to derive atom-wise parameters from the trace of the distributed quadrupole moments [129].

Another way of determining parameters is to place a probe atom at many different positions around the molecule and fit atomic parameters (possibly anisotropic) to the resulting exchange-repulsion energy [143]. This method may capture some system-specific effects while still only performing monomer calculations. As for the decomposed density overlap approach, accurate combination rules are required to obtain the atom-atom properties of Eqs. 2.24 or 2.25.

#### Beyond additive repulsion

Perturbation theory has shown that the true exchange repulsion, in contrast to the models above, is non-additive [144], and moreover that it couples with the induction. This is not surprising, considering that the repulsion, like the polarization, influences the charge density. As an example, the charge redistribution in the water dimer caused by the mere antisymmetrization of the dimer wave function is shown in the left part of Fig. 2.4, with the lighter surface representing a decreased and the darker an increased electron density. For comparison, the fully polarized redistribution is shown in the right part of Fig. 2.4. As can be seen, the antisymmetrization gives a significant contribution.



Figure 2.4: Isodensity surfaces (light for negative and dark for positive) for the change in density upon formation of the water dimer. The left part shows the antisymmetrized but unpolarized wave function whereas the right part shows the fully polarized wave function.

For these reasons, one would ideally like to model the repulsion on a similar level as the polarization, that is, one would need a way to "polarize" the molecule by repulsion. However, because the repulsion is a non-classical effect, there is nothing corresponding to point-polarizabilities. One solution, which was recently applied in a non-additive force field [145], is to model the density depletion as an *exchange-quadrupole*. Another solution is to directly compute the coupling from orbital overlap expressions [146].

A possible approach to develop this type of models is to start from a description in which one of the monomers is treated by quantum chemistry and the other monomer (B) by a simpler model. This description is also directly useful for treating the influence from the environment on a quantum-chemical calculation [147–150], e.g. for calculating solvent effects on spectra- [151]. The approach is analogous to conventional QM/MM methods [152], in which electrostatic (but usually not repulsive) effects are captured by embedding the system of interest in point charges.

The repulsive effect can be modeled as a pseudopotential, which is added to the Hamiltonian of the quantum-chemical system. Pseudopotentials were first introduced in studies of atoms as a way to remove the explicit treatment of the core electrons. A common choice (for monomer A) is

$$\widehat{V}_A = \sum_i d_i \left| \psi_i^B \right\rangle \left\langle \psi_i^B \right| \tag{2.28}$$

where the sum runs over the occupied orbitals  $|\psi_i^B\rangle$  of monomer B and  $d_i$  is a parameter which often depends on the orbital energy. For determining the parameters  $d_i$ , theoretical analyses exist for the pure repulsion in simple systems, but normally a fitting or qualitative arguments are applied. In this way, one can also account for the attractive exchange (which is costly to add separately), or even the effective repulsion in Eq. 2.23. A more advanced pseudopotential has been derived from the variational derivative of the EFP repulsion [146].

Pseudopotentials based on the density overlap have also been developed [153,

154]. The effect can also be modeled by standard density functionals [155].

#### 2.3.5 Charge transfer

The charge-transfer term is usually described as the energy lowering due to excitation of electrons from the occupied orbitals of one monomer into the virtual orbitals of the other [156], in contrast to the induction energy, which involves excitations to the virtual orbitals of the same molecule. Obviously, as the basis set is enlarged and approaches completeness, the two virtual spaces will practically coincide and the charge transfer will either vanish [39,40] or double-count the induction energy [38] depending on the exact definition employed. Thus, the charge-transfer energy can be seen as an ill-defined part of the induction energy [157], or simply be defined as the part of the interaction energy that cannot be described when each monomer uses its own basis set. A similar definition has been used to compute the charge-transfer within intermolecular perturbation theory [157], as the standard definitions suffer from contamination from BSSE.

It has been shown that the charge-transfer energy has an exponential decay with intermolecular distance (like the exchange repulsion) and that it is significant for normal hydrogen bonds using typical basis sets [157]. Although one can argue that it is a mathematical artifact and could be lumped together with other garbage terms, such treatment may interfere with the goal of transferability. Therefore, explicit estimates of the charge-transfer energies based on monomer properties are included in both the SIBFA and EFP methods, although the expressions used are remarkably different [140, 156]. Much simpler approaches have also been devised, e.g. using a Born–Meyer term with negative sign [158].

## Chapter 3

# Protein–ligand affinities

Mutual molecular recognition is the starting point for almost all processes in biological systems [159]. An enzyme must bind the correct substrate for a reaction to take place, a receptor must bind the correct compound to trigger a signal, etc. From a computational point of view, this is a very difficult problem, because the protein is large, both molecules are flexible, and water (which is in itself a difficult liquid to model) plays an important role. At the same time, however, it is an intriguing problem, because it is one of the rare quantities that are both well-defined at a microscopic level and at the same time possible to measure accurately experimentally.

Not surprisingly, this problem is of great interest to the pharmaceutical industry, because many drugs are competitors with natural compounds in binding to proteins. The prediction of the binding free energy of a ligand to a protein has been described as the Holy Grail of rational drug design [159]. For the calculation to be really useful in drug development, an accuracy of  $\sim 4$  kJ/mol is required. This can appear as an impossible wish, considering the complexity of the process, but in fact, several studies have reported such accuracies, although methods that achieve this for an arbitrary protein–ligand complex are still far away.

Depending on the development phase, one might have different starting points for the problem. Some methods, e.g. QSAR and pharmacophore modeling, can be applied even if the structure of the protein has not been determined. Predicting the binding pose of the ligand to the protein, *docking*, requires that the structure is known. The subsequent prediction of the binding free energy is called *scoring*. Many types of *scoring functions* have been developed [159], including *empirical*, based on a weighted sum of simple structural properties, *knowledge-based*, based on a statistical analysis of known protein–ligand structures, and *physical*, based on a potential-energy surface. Docking and scoring is normally an interrelated process, but specific tests of scoring functions can e.g. be performed by using experimentally determined binding poses. In the following, it is assumed that the binding pose is known. Moreover, to comply with the main thread of the thesis, only physical methods will be discussed. A comparison of the accuracy of scoring functions [160] showed that a physical approach is not guaranteed to give the best results. Only if the model includes all important effects in a balanced way and treated with as few approximations as possible, such distinction can be expected.

### 3.1 Theory

For the dissociation of a protein-ligand (PL) complex into the unbound protein (P) and ligand (L) in aqueous solution, the equilibrium constant  $K_d$  is given by

$$K_d = \frac{[P][L]}{[PL]} \tag{3.1}$$

where [...] denotes the concentration. Experimentally,  $K_d$  is the concentration of ligand necessary to make a half of the receptor sites occupied. The standard free energy of binding (per molecule) is then defined as

$$\Delta G_{PL}^{\circ} = -k_B T \ln(K_d/C^{\circ}) \tag{3.2}$$

where  $C^{\circ}$  is the standard concentration (1 M). The work associated with the volume change of the process is usually negligible, so no differentiation between Gibbs and Helmholtz free energies is done.

Based on the classical probability distribution defined in Eq. 1.25, the free energy is given [161] by

$$\Delta G_{PL}^{\circ} = -k_B T \ln \left( \frac{C^{\circ}}{8\pi^2} \frac{Z_{PL+S} Z_S}{Z_{P+S} Z_{L+S}} \right).$$
(3.3)

Here, each configuration integral is given by

$$Z_X = \int \exp\left[-E(\mathbf{r}_X)/k_B T\right] d\mathbf{r}_X \tag{3.4}$$

where  $\mathbf{r}_X$  are the coordinates of the system X, for example all protein and solvent coordinates in the case X = P + S. It is possible to integrate out the explicit solvent coordinates by using the solvation free energy  $W(\mathbf{r}_X)$  for a given solute configuration. Thus, Eq. 3.3 is equivalent to

$$\Delta G_{PL}^{\circ} = -k_B T \ln \left( \frac{C^{\circ}}{8\pi^2} \frac{Z_{PL}}{Z_P Z_L} \right).$$
(3.5)

where

$$Z_X = \int \exp\left(-\left[E(\mathbf{r}_X) + W(\mathbf{r}_X)\right]/k_B T\right) d\mathbf{r}_X.$$
 (3.6)

This is the starting point for implicit-solvent approaches.

For reasonably small complexes, it is possible to directly approximate the configuration integral in Eq. 3.6 by a sum of integrals of the predominant states [162] found by a conformational search. Most other methods, however, rely on simulation techniques to find suitable ensembles from which the free energies can be extracted [163]. Usually, one distinguishes between path-based methods, which include sampling of various (usually unphysical) intermediate states, and end-point methods, which only involve simulations of the free and bound states.

## 3.2 Path-based methods

#### 3.2.1 Relative affinities

The most common approach to rigorously calculate the difference in binding free energy between two ligands is the *double decoupling* method [164]. It is based on the thermodynamic cycle shown in Fig. 3.1. Instead of directly calculating the binding affinities for the L and L' ligands ( $\Delta G_{PL}$  and  $\Delta G_{PL'}$ ), one calculates the free energy of transforming L to L' in the complex ( $\Delta G_1$ ) and in solution ( $\Delta G_2$ ) and uses the energy conservation to obtain

$$\Delta G_{PL'} - \Delta G_{PL} = \Delta G_1 - \Delta G_2 \tag{3.7}$$

The connection with the partition function expressions is trivial. Using Eq. 3.3 gives

$$\Delta G_{PL'}^{\circ} - \Delta G_{PL}^{\circ} = -k_B T \ln \frac{Z_{PL'+S}}{Z_{PL+S}} - \left(-k_B T \ln \left(\frac{Z_{L'+S}}{Z_{L+S}}\right) = \Delta G_1^{\circ} - \Delta G_2^{\circ} \quad (3.8)$$

Each of the free energies  $\Delta G_1$  and  $\Delta G_2$  is usually calculated by free energy perturbation (Eq. 1.29) or thermodynamic integration (Eq. 1.28), where one gradually transforms one ligand into the other by Eq. 1.27. The double decoupling approach is a well established technique for obtaining free energy differences between similar ligands. It has been found to give significantly better results than simple scoring functions [165]. Most applications have used molecular dynamics simulation [166], but Monte Carlo simulation have also been used [167].

#### 3.2.2 Absolute affinities

Calculating absolute binding affinities is a more challenging problem. It can be addressed by a similar type of thermodynamic cycle, which is shown in Figure 3.2 [168]. Here, the dashed line indicates that the ligand does not interact with the surroundings, i.e. it is in the gas phase. The upper right state was introduced later [161] as a convenient intermediate to account for the standard concentration while still allowing for a straight-forward computation; in this



Figure 3.1: Thermodynamic cycle used for calculating relative binding affinities.

state the ligand is not interacting but still constrained to occupy the binding pocket. The method is usually called the *double annihilation* method, although the reuse of the name *double decoupling* was suggested when the  $\Delta G_3$  term was added.

The quantities  $\Delta G_1$  and  $\Delta G_2$  are usually calculated by e.g. free energy perturbation, whereas  $\Delta G_3$  can be estimated analytically. The method (with or without  $\Delta G_3$ ) has been applied to many systems and seems to give good results when pushed to the convergence limit [169], although this requires substantial computational time. The importance of the particular form of the constraints has been noted [170, 171]. Another way of avoiding the problematic end point of the  $\Delta G_1$  path is to perturb the ligand into water [172]. Although most simulations of both relative and absolute binding affinities employ explicit water, implicit solvent methods or hybrid methods [171] can also be used.

#### 3.2.3 Potential of mean force

It can be difficult to perform standard decoupling methods if, for example, the ligand is highly charged or flexible [173]. An alternative approach is to transfer the ligand from solution to the binding site by pulling it along a predefined path while averaging over the rest of the system trough simulation. This can be done by umbrella sampling [174] or more sophisticated approaches, in which one sequentially introduces and removes translational, rotational, and conformational restraints [173]. Several other techniques for improving the convergence and efficiency of free energy calculations have also been proposed [175].



Figure 3.2: Thermodynamic cycle used for calculating absolute binding affinities.

## 3.3 End-point methods

#### 3.3.1 Perturbations from single simulations

The procedure of performing several simulations for each ligand using various  $\lambda$  values can be quite time-consuming in the double decoupling method. On the other hand, if the perturbations could be done in one step ( $\lambda = 0 \rightarrow \lambda = 1$ ), the relative free energies between all ligands could be computed from only one simulation (with  $\lambda = 0$ ). Typically, the largest problem with using one-step perturbations is that the space needed for the addition of a new group of atoms is already occupied by the solvent or the protein. Therefore, attempts have been made to use an artificial reference compound containing "soft" atoms, so that the ensemble generated with the reference compound is wide enough to have sufficient overlap with all the corresponding ensembles for the real ligands [176]. This approach has been quite successful, even for structurally diverse ligands [177]. Similar simplifications have also been reported [178].

Warshel et al. have also applied free-energy perturbation methods [179], but have also found that the *linear response approximation* (LRA; Eq. 1.30) is a good approximation for the electrostatic contribution to the free energy [180]. The neglect of the higher-order terms in Eq. 1.30 corresponds to the assumption that the free energy potentials with respect to an electric field are harmonic and have equal curvature [181].

#### 3.3.2 Linear interaction energy (LIE)

A more empirical approach to binding affinities based on LRA is the linear interaction energy method [181]:

$$\Delta G_{PL} = \alpha \left( \left\langle E_{int}^{vdw} \right\rangle_{PLS} - \left\langle E_{int}^{vdw} \right\rangle_{LS} \right) + \beta \left( \left\langle E_{int}^{ele} \right\rangle_{PLS} - \left\langle E_{int}^{ele} \right\rangle_{LS} \right), \quad (3.9)$$

where  $E_{int}^{vdw}$  and  $E_{int}^{ele}$  are the Van der Waals and electrostatic parts of the interaction energy between the ligand and the surroundings, the latter being only solvent in the LS ensembles but protein and solvent in the PLS ensemble. Based on the linear response approximation, the constant  $\beta$  was originally set to 1/2, but a direct FEP computation of the response for a wide range of molecules showed that  $\beta$  is system-dependent and deviates from 1/2 for nonionic systems [182]. Thus a standard LIE model was devised that uses  $\beta = 0.43$ , 0.37, and 0.33 for neutral ligands with zero, one, and two or more hydroxyl groups, respectively, and  $\alpha = 0.18$ . This model has provided good results [183] for relative (and in some cases even absolute) binding affinities and has been found rather independent of the force field [184]. Nevertheless, many studies treat both  $\alpha$  and  $\beta$  as adjustable parameters. The LIE method has also been used with implicit solvent models [185, 186].

#### 3.3.3 MM/PBSA

In the MM/PBSA approach [187], the binding affinity is estimated by starting from the thermodynamic cycle in Figure 3.3 [188] and then approximating the gas-phase binding free energy by the sum of the average interaction energy and the change in configurational entropy upon binding, i.e.

$$\Delta G_{PL} = \Delta G_{PL}^{gas} + \left( \Delta G_{PL}^{solv} - \Delta G_{P}^{solv} - \Delta G_{L}^{solv} \right) \\ \approx \left\langle E_{int}^{MM} \right\rangle_{PLS} - T \left\langle \Delta S_{config} \right\rangle_{PLS} + \left\langle \Delta \Delta G_{solv} \right\rangle_{PLS}$$
(3.10)

where  $\langle ... \rangle_{PLS}$  indicates that the averaging is done using snapshots (typically 20–100) from a simulation of the complex in solvent, usually explicit solvent. The interaction energy  $(E_{int}^{MM})$  is computed with the water molecules removed, whereas the change in solvation energy upon binding  $(\Delta\Delta G_{solv})$  is calculated using an implicit solvation model.  $\Delta S$  denotes the change in configurational entropy upon binding and T is the temperature.

For the interaction energy, a standard MM force field is normally used. A polarizable force field has also been tried [189], although without fully self-consistent treatment of the solvation. QM/MM interaction energies based on semiempirical or density functional theory have also been used [190, 191]. In paper VII, an attempt to use a higher-level QM method is presented.

For the solvation energy, the most common choice is the PBSA method, i.e. the Poisson–Boltzmann (PB) method for the electrostatic part combined with an estimate based on the surface area for the non-polar part. Several



Figure 3.3: Thermodynamic cycle used in MM/PBSA.

studies have been performed with the Generalized Born method replacing PB, although these methods can sometimes give quite different results [189, 192]. In paper VII, the polarizable continuum method (PCM) will be used, because it has advantages in combination with a polarizable force field.

Estimating the change in configurational entropy is a very difficult problem. In fact, for computing relative affinities for similar ligands, it is probably better to neglect this term [187], but for absolute affinities it is an essential term. It is normally computed by normal mode analysis or quasi-harmonic analysis. In the former case, one typically truncates and minimizes a small part of the complex and computes the vibrational, rotational, and translational entropy changes using this truncated model. For more stable results, a fixed buffer region can be included [193]. The quasi-harmonic analysis, on the other hand, is done for the trajectories from the simulation. A problem with this approach is limited sampling and that degrees of freedom with several local minima are treated as a single wide potential well. By using methods that treat the local minima separately, it has been found that the most important contribution to the entropy loss comes from the energy wells becoming narrower, whereas the reduced number of conformations available is less important [194]. This is in contrast to the assumption in many simple scoring functions that the entropy can be estimated by counting the number of rotatable bonds.

The statistical-mechanical basis of the MM-PBSA method departing from the exact Eq. 3.5 has been described [161,195]. The most approximate assumption is probably that the energetic landscape of each species has an energy and configurational volume that can be determined from a simulation, i.e.

$$Z_X = \int \exp\left(-\left[E(\mathbf{r}_X) + W(\mathbf{r}_X)\right]/k_B T\right) d\mathbf{r}_X. \approx z_X^{int} \exp\left(-\langle E + W \rangle_X/k_B T\right),$$
(3.11)

where  $z_X$  is the internal configuration integral (causing the  $T\Delta S$  term), and  $\langle ... \rangle_X$  denotes that the average is computed for the particular species. This means that the most rigorous way of computing the averages in Eq. 3.10 is by performing separate simulations of the complex, the protein and the ligand. However, as this approach usually gives severe convergence problems due to the intramolecular contributions from the protein [195], most applications have used only one simulation and thus have taken the protein and ligand configurations from the simulation of the complex. On the other hand, separate simulations were found to give significantly better results in at least one application [196]. It has also been found that the choice of force field and solvent treatment in the simulations have a significant effect on the results [189]. Interestingly, similar or even better results can sometimes be obtained using a single minimized structure [197].

## 3.4 Solvation

Water plays many important roles in the protein–ligand binding process. First, it acts as a dielectric continuum and screens the electrostatic interaction significantly. This is especially evident in the MM/PBSA method, where the vacuum electrostatic energy can be over a thousand kJ/mol, but is almost completely cancelled by the solvation contribution (see paper VII). Water can also play a more direct role by forming specific hydrogen bonds with the ligand or protein that must be broken when the complex is formed. This is a much more difficult problem to address. Unfortunately, the importance of solvation means that the treatment of solvation can introduce huge errors in calculations of binding affinities, regardless of whether an explicit or implicit solvent model is used.

There are mainly three problems with explicit water treatments. First, they rely completely on molecular mechanics solute–water and water–water potentials, which can be quite different depending on the force field [198]. Second, due to the strong hydrogen-bond network in water, the sampling time (and thus the computational load) required for adequately sampling the solvent degrees of freedom is substantial. Third, the treatment of long-range electrostatic effects is non-trivial. Most studies employ a periodic boundary condition, for which Ewald summation seems to give rather accurate results, but this model has been questioned [180]. The last problem can in principle be solved by combining the explicit model with a continuum treatment [199]. In spite of these problems, explicit models have been able to reproduce solvation energies of e.g. amino-acid models with remarkable accuracy (although differences of up to 6 kJ/mol depending on the force field were obtained) [200].

Implicit water models, on the other hand, suffer from the lack of specific interactions, e.g. hydrogen bonds. Moreover, they are usually not directly connected to a specific potential and are therefore dependent on other types of parameters, typically charges and cavity radii. The advantages are the fast evaluation and the possibility of easily fitting the models directly to experimental solvation energies. A quantitative comparison of explicit and implicit solvation reported good agreement in the context of binding affinities [201]. A different approach to solvation is methods based on integral equation theory, e.g. the 3D-RISM method [202]. In contrast to standard implicit models, they are based on a specific potential and can therefore incorporate some specific effects.

Most implicit solvent models are based on the decomposition

$$\Delta G_{solv} = \Delta G_{solv}^{polar} + \Delta G_{solv}^{non-polar} \tag{3.12}$$

where  $\Delta G_{solv}^{polar}$  is based on roughly the same physics in the various models, whereas  $\Delta G_{solv}^{non-polar}$  can involve quite different physical effects.

The physical motivation for Eq. 3.12 is that the solute–solvent potential energy is usually based on separate electrostatic and non-electrostatic terms [203]. Thus, the solvation can be described as a two-step process: first, the non-electrostatic terms of the potential are gradually turned on, creating a cavity in the solvent, and then the electrostatic term is gradually turned on. The first step gives  $\Delta G_{solv}^{non-polar}$  and the second step gives  $\Delta G_{solv}^{polar}$ . These quantities can be directly calculated by e.g. thermodynamic integration (Eq. 1.28), but the decomposition is ambiguous – for example, different (usually divergent) results are obtained if the steps are reversed [203]

#### 3.4.1 Polar solvation

The polar part of the solvation is typically based on modeling the solvent as a dielectric continuum, characterized by the macroscopic dielectric constant (i.e.  $\epsilon_r = 80$  for water). The simplest case is the solvation energy of a point charge q in a spherical cavity with radius R. This is usually called the Born energy:

$$\Delta G_{solv} = \frac{q^2}{2R} \left( \frac{1}{\epsilon_r} - 1 \right) \tag{3.13}$$

Any useful model must typically reproduce this energy. There are analogous models for higher multipoles, but they are typically not sufficiently accurate for modeling molecules, as these have both a complicated charge distribution and a non-spherical shape.

The basis for the continuum models is the Poisson equation:

$$\nabla \cdot [\epsilon(\mathbf{r}) \nabla \phi(\mathbf{r})] = -4\pi \rho(\mathbf{r}) \tag{3.14}$$

where  $\phi(\mathbf{r})$  is the electrostatic potential at a point  $\mathbf{r}$ ,  $\rho(\mathbf{r})$  is the charge density of the solute, and  $\epsilon(\mathbf{r})$  is the position-dependent dielectric constant.

For point-charge models of solutes, Eq. 3.14 is usually solved numerically on a grid using a finite-difference algorithm. The effect of a non-zero ionic strength can easily be included in such model, and it is therefore often referred to as the Poisson–Boltzmann (PB) method. The Generalized Born (GB) methods [204] can be seen as an approximation to the PB method, or alternatively as a direct extension of Coulomb's law in solution.

For more complicated charge densities, such as quantum-chemical models, it is more common to use apparent-surface-charge methods, such as the *polarizable continuum model* (PCM) [205]. The PCM model has recently been extended to larger systems treated by fragmentation approaches or molecular mechanics [206].

A different model, the *Langevin-dipoles model*, has been applied extensively in protein simulations [207]. It is built on a discrete (semi-explicit) representation of the solvent as a grid of rotatable point dipoles, but with the interactions modeled effectively by using the Langevin equation for a rotationally averaged dipole in an electric field. It has been used to demonstrate the connection between explicit and continuum models [208]. A continuum model based on the Langevin equation has also been developed [209], effectively capturing the non-linearity of the response without introducing the numerically cumbersome discreteness.

#### 3.4.2 Non-polar solvation

Widely different expressions for the non-polar solvation energy are used [210]. For example, the non-polar part of the PCM model is a sum of cavitation, dispersion, and repulsion terms, whereas a common approximation is to assume a linear relationship between the non-polar solvation energy and the solvent-accessible surface area. In paper VII, it is demonstrated that such models give significantly different results for binding affinities.

## 3.5 Potential energy

Having examined various methods for treating the statistical-mechanical problem, it is clear that the choice of potential energy surface is an essential part. In chapter 2 we saw how the intermolecular potential energy surface can be computed by various degrees of approximation. In the final part of this thesis, some examples of how these methods can be applied to the specific problem of protein–ligand interaction energies are given.

It is still an open question to what extent the potential energies (i.e. the force field) limits the accuracy of binding affinity predictions. Several rigorous binding affinity studies have mentioned inaccuracies in the force field as the major source of error [165, 167]. On the other hand, the necessity of including polarization has been questioned [55] for the reason that the sampling problem will anyway limit the accuracy. However, it is rather obvious that with the fast

advance in computational power, the potential-energy problem will sooner or later be dominant [1].

Most studies of protein–ligand interactions with a physical approach employ standard non-polarizable MM force fields [163, 211]. Among the polarizable force fields with a careful treatment of each physical term, several have been applied to protein–ligand interactions, including the SIBFA [46, 140, 212, 213], AMOEBA [52, 214], and PFF [50, 215] models. Although several of these studies reported a large effect of polarization on the binding affinities, it should be noted that this effect is not always attractive when the solvent is taken into account [214]. Because the quantum-chemical calculations used to derive the parameters cannot be run for the full protein, the assembly of properties from smaller calculations (typically amino acids) is an important part of such calculations [216–218] (see paper V).

A different approach to determine the potential energy is to directly exploit quantum-chemical calculations, as these are intrinsically transferable, and also systematically improvable [219]. Protein–ligand systems are in general too large for a full quantum-chemical treatment. Attempts have been made to use standard QM/MM methods [152] that treat only the ligand [190, 220] or preferably also the closest residues [221] by quantum chemistry and the rest of the protein by MM. However, the size of a quantum-chemical region containing the closest residues is typically 300-800 atoms, whereas most methods that treats dispersion interactions reasonably well (at least MP2 with a sufficiently large basis set) are limited to  $\sim 100$  atoms. A more fruitful strategy is to decompose the complex into smaller subsystems, which are treated more or less independently. Several such fragmentation methods have been applied to protein-ligand interactions, including the divide-and-conquer method [222, 223], the fragment molecular orbital (FMO) method [224–227], the molecular fractionation with conjugate caps (MFCC) method [228–230] and related approaches [231, 232]. However, only Ref. [232] has employed a reasonably accurate quantum-chemical method.

A combination of the polarizable molecular mechanics and fragmentation approaches will be presented in paper V, used to calculate a protein–ligand interaction energy in paper VI, and combined with the MM/PBSA method in paper VII, providing the first protein–ligand free-energy calculations using quantumchemical interaction energies at a reasonably high level of theory. However, an improvement of the treatment of solvation is probably necessary before an improvement in relation to experiment can be expected.

## Chapter 4

## Summary of the papers

## 4.1 Paper I: Exchange-repulsion energy

In paper I, we compare the ability of various repulsion models to reproduce the Hartree–Fock exchange-repulsion energy for a large set of molecular dimers. In particular, models based on orbital overlap, Eq. 2.26, and density overlap, Eq. 2.27, respectively, are considered.

The models contain 1 or 2 parameters, with the exception of the density overlap model with special scaling of certain element pairs (8 parameters) and the parameter-free EFP method. For each model, the parameters are fitted to the full set of dimers so that the average error is minimized.

The average error for each model is given in Table 4.1. The models assuming direct proportionality to the orbital overlap, energy-weighted orbital overlap, and density overlap, have similar errors. Introducing another parameter to model the deviation from proportionality at shorter distances reduces the error in both the orbital-overlap and density-overlap cases. However, the corrections have different signs, because the exchange repulsion is steeper than the orbital overlap but less steep than the density overlap. The similar results for the SIBFA-type of correction shows that the exact functional form is not critical.

For the density overlap models, larger errors are obtained for the ethanethiol– water dimers than for the other dimers, whereas no such effect is seen for the orbital overlap models. By comparing the contributions from each pair of molecular orbitals to the orbital and density overlap, respectively, we show that around the sulfur atom, destructive interference plays an important role, and this cannot be modeled by density overlap models. Therefore, it seems to be necessary to introduce element-dependent parameters if density-overlap models should become as accurate as the orbital overlap models, . The EFP method performs worse than the other methods, but considering that no fitting is needed, it is a useful method for unknown systems.

Table 4.1: Average error in  $\rm kJ/mol$  for each considered model. The number of parameters in each model is also given.

Method	# param.	Error
Orbital overlap:		
Proportional	1	1.21
Proportional (energy-weighted)	1	1.10
Distance-dependent	2	0.73
SIBFA-type	2	0.74
EFP	0	1.68
Density overlap:		
Proportional	1	1.15
Distance-dependent	2	0.97
Distance- and element-dependent	8	0.60

## 4.2 Paper II: Multipoles and polarizabilities

In paper II, the accuracy of distributed multipoles and polarizabilities are evaluated by comparing the static and induced electric potentials with those calculated from a quantum-chemical (DFT) calculation, for a set of 20 molecules. The comparison is done in 10,000 points chosen randomly in a suitable elementdependent distance range. For the evaluation of the induced potential, homogeneous fields are applied in three orthogonal directions.

In particular, the multipoles and polarizabilities obtained by two different approaches are compared: the MpProp method, in which the multipoles are derived through a Mulliken-like approach [67] and the polarizabilities using the uncoupled Hartree–Fock method [117] (but scaled to give the correct molecular polarizability), and the LoProp method [76], in which both the multipoles and polarizabilities are obtained using an orthogonal localized basis set and finite-field perturbations are used for the polarizabilities. All calculations are performed with three different basis sets.

The mean absolute error in the static and induced potentials for each molecule is shown in Fig. 4.1. For the basis sets without diffuse functions, the MpProp method gives better multipole moments, regardless of the multipole level. The reason for this is probably the better placement of the bond centers in the Mp-Prop method. On the other hand, with the basis set including diffuse functions (aug-cc-pVTZ), the LoProp method gives better results for  $L \leq 3$ , which are the most interesting levels in practice. In particular, the MpProp method fails for large aromatic molecules. The LoProp method also gives better polarizabilities for all basis sets, making it the preferred method for constructing intermolecular potentials. In fact, all results become worse with the larger basis set, probably because the neglect of charge penetration becomes a more severe approximation as the density becomes more diffuse.

The dependence of the error on the distance from the molecular surface is



Figure 4.1: Errors in the static potential for various multipole levels L and in the induced potential (using a field strength of 0.01 a.u.). Each point corresponds to one molecule. A probe charge of 1 a.u. is assumed.

also tested. It is found that, for typical interaction distances, there is no gain in using a higher multipole level than octupoles. In fact, the quadrupole level might be a more balanced choice when simultaneously considering the performance of the polarizabilities.

## 4.3 Papers III and IV: Polarization models

In **paper III**, the performance of the point-polarizability model is tested for a large set of molecular dimers. To separately evaluate the various approximations in the model, we use a hierarchy of models including various effects, as shown in Table 4.2. In **paper IV**, we introduce an additional model (ppEPOL) that is tested partly for the same systems.

#### 4.3.1 Densities

Because the induction energy is not an observable, we concentrate in paper III on the change in charge density due to the interaction. This is in itself an interesting quantity, because it is responsible for the major part of the manybody effects – one of the reasons for using polarizable force fields at all. Various ways to analyze this *induced density* are used, with the most quantitative being the point-wise comparison of the induced electric potential used in paper II.

	Non-expanded	Intramol.	Local field-	Pauli	Charge
Model	induced density	coupling	inhomogeneity	effects	transfer
Supermolecular	yes	$\mathbf{yes}$	yes	yes	$\mathbf{yes}$
RVS	yes	$\mathbf{yes}$	$\mathbf{yes}$	$\mathbf{yes}$	no
EPOL	yes	$\mathbf{yes}$	$\mathbf{yes}$	no	no
LPOL	yes	$\mathbf{yes}$	no	no	no
Polarizabilities	no	no	no	no	no
ppEPOL	yes	yes	yes	yes	no

Table 4.2: Summary of the considered models.



Figure 4.2: Relative mean error in the induced electric potential around the water dimer as a function of intermolecular separation.

The result for the water dimer is given in Fig. 4.2, which shows the relative error compared to the supermolecular calculation for various models as a function of the intermolecular separation. The most obvious feature is that the EPOL model, which uses the exact electric field as a perturbation in the quantumchemical calculation, gives the largest error for all interesting distances. A comparison with the RVS model shows that this is almost exclusively due to the neglect of Pauli effects (cf. Section 2.2.2). Interestingly, the LPOL model, which also neglects Pauli effects, gives much better results. Because the only difference between the EPOL and LPOL is that the latter uses a locally homogeneous (linearized) field, the results indicate that this approximation partly cancels the neglect of Pauli effects.

The difference between the results for the LPOL model (using quantumchemical calculations) and the simple polarizability model is much smaller. This indicates that the most important approximations in the polarizability model are not related to the actual response, but to the simplified potential that is applied (linearized and neglecting the perturbing molecule). The curves with the LPOL response expanded up to dipoles and quadrupoles, respectively, show that the remaining discrepancy comes entirely from the use of only dipoles in the polarizability model; in fact, dipole–quadrupole polarizabilities would be sufficient to reproduce the LPOL result. Finally, the ppEPOL model, which includes Pauli effects in an approximate way, gives the best results, indicating that better results can be obtained without relying on the error cancellation.

A statistical analysis of the large data set shows that the water dimer is rather typical for the considered systems. In general, the LPOL and polarizability models give lower errors (in average 9 and 10 %, respectively) than the EPOL model (16 %). However, the error cancellation between Pauli effects and local inhomogeneity effects is not always effective. For example, in the formamide dimer, the error for the polarizability model is ~90 % at the energy minimum, i.e. only slightly better than a non-polarizable model. The large data set also confirms that the LPOL and polarizability models give similar results. In particular, the neglect of intramolecular coupling of the polarizabilities seems to be a minor effect for the small molecules considered in this study.

### 4.3.2 Energies

With the introduction of the ppEPOL model in paper IV, we can also address the total interaction energy. In the ppEPOL model, Pauli effects are described by a pseudopotential (see Section 2.3.4), which models both the exchange-repulsion energy itself and its effect on the polarization in a consistent manner. As in the EPOL model, the exact electric field is used. For reproducing results at an electron-correlated level, the model must be augmented by a dispersion term (in paper IV taken as the exact expression from second-order perturbation theory), whereas the intramonomer correlation contributions are self-consistently included by using correlated monomer electron densities (*MP2 response densities*).

As established in paper I, a model that is proportional to the squared overlap  $(S^2)$  cannot reproduce the exchange repulsion over a large range of distances, and the same holds for standard pseudopotentials of the type in Eq. 2.28. Therefore, we introduce an  $S^4$ -dependent term in the pseudopotential. Although the resulting two-parameter model could presumably be fitted in the same way as in paper I, we avoid uncertainties in the parameters by fitting the model to each dimer separately, but still using only the zeroth-order Hartree–Fock exchange-repulsion energies. In practice, this means that we are mainly testing whether the polarization and its coupling to repulsion can be described by the ppEPOL method.

The results for the water dimer are shown in Fig. 4.3. The supermolecular curve is well reproduced, both at the HF and MP2 level. In particular, the ppEPOL model is significantly better than the uncoupled model that uses the EPOL model for the polarization and simply adds the exact Heitler–London energy to account for the exchange repulsion. This shows that the coupling



Figure 4.3: Intermolecular potential of the water dimer obtained with the ppEPOL model at the Hartree–Fock and MP2 levels and the corresponding supermolecular curves. The results without polarization (Heitler–London; HL) and with the EPOL result directly added to the HL energy (uncoupled) are also shown.

between polarization and repulsion is well described. The fact that this coupling in negative in sign for the minimum structures of the considered dimers suggests that it should rather be seen as a relaxation of the exchange-repulsion than as a damping of the polarization. However, at shorter distances the coupling becomes positive.

Encouragingly, the ppEPOL model gives good results also for the formamide dimer, which is a tough case for polarizability models. The water-chloride dimer appears slightly more difficult, but the agreement is still satisfactory. In all cases, a significant improvement of the induced density compared to the EPOL model is also obtained. For example, the error in the induced potential decreases from 128 to 6 % for the formamide dimer. Thus, the ppEPOLmodel appears as a useful model for developing and testing simpler polarization models including coupling between polarization and repulsion.

## 4.4 Paper V, VI, and VII: The PMISP method

In the last three papers, we explore a novel way of combining fragmented quantum-chemical and molecular-mechanical calculations to compute higly accurate interaction energies for large systems. **Paper V** describes the PMISP method, addresses several technical details, and compares the results with full supermolecular calculations, as well as with other fragment approaches. **Paper VI** presents the first calculation of a protein–ligand interaction energy performed at a reasonably high level of theory (MP2/aug-cc-pVTZ) and investigates the accuracy of various approximations. **Paper VII** combines the PMISP model with the polarizable continuum model (PCM) for treating solvation effects and



Figure 4.4: Example illustrating the MFCC procedure for cutting molecule A across the peptide bond and capping with  $-\text{COCH}_3$  and  $-\text{NHCH}_3$  groups. The result is two capped fragments  $A_1$  and  $A_2$  and a concap fragment  $A_3$ .

applies the resulting method in a MM/PBSA-like manner to a set of proteinligand complexes for which the experimental binding free energies are known.

#### 4.4.1 Methods

We try to approximate the quantum-chemical interaction energy (at a given level of theory) between a large molecule A and a small molecule B. In the *polarizable multipole interaction with supermolecular pairs* (PMISP) method, the interaction energy is decomposed into *classical* and *non-classical* parts, where the classical part is the sum of electrostatic and induction energies given by a polarizable multipole method. Furthermore, the large molecule A is divided into a set of *fragments*  $\{A_i\}$  by the *molecular fractionation with conjugate caps* (MFCC) procedure [228], as illustrated in Fig. 4.4.

A previously proposed method (which we denote pairwise additive, PA) approximates the total interaction energy by the sum of  $A_i$ -B interactions with the concap interactions subtracted [228]. However, this approach neglects the *non-additivity* (many-body effects) of the fragment interactions. Therefore, in the PMISP method, the MFCC procedure is applied only for the non-classical part, whereas the classical part is evaluated for the whole complex, i.e. including non-additivity. The *properties* (multipoles and polarizabilities), on the other hand, are assembled using the MFCC procedure.

When a full protein–ligand interaction energy is computed, the long-range non-classical contributions are estimated by a standard molecular mechanics force field, as described in paper VI. For the interaction energy in solvent (paper VII), it is assumed that only the classical part of the interaction energy is influenced by the solvent. A rescaling of the UAKS radii for the PCM method is performed by fitting to experimental solvation energies for a small data set.

#### 4.4.2 Results

In paper V, the PMISP model is tested against supermolecular calculations for a 216-atom model of avidin interacting with seven biotin analogues. Calculations are performed at the HF and MP2 levels using the 6-31G\* basis set. A

Level	Method	Error
HF	PMISP	5.8
$_{\mathrm{HF}}$	EMB-PMISP	5.0
$_{\mathrm{HF}}$	PA	19.4
$_{\mathrm{HF}}$	FMO2	6.4
$_{\mathrm{HF}}$	EE-PA	8.0
MP2	PMISP	6.9
MP2	PMISP-CE	3.3

Table 4.3: Mean absolute error (in kJ/mol) for the ligand set using several methods.

summary of the results is shown in Table 4.3. The error of the PMISP methods (5–7 kJ/mol) is small compared to the interaction energy (-256 kJ/mol in average). It is also significantly smaller than the error for the PA model, showing that many-body effects are important in this type of system. Moreover, a comparison with two other approaches to include many-body effects, the *fragment molecular orbital* method, FMO2 [224] and the *electrostatically embedded pairwise additive* method, EE-PA [233], shows that the PMISP method gives slightly smaller errors at a reduced computational cost. Interestingly, the PMISP error is not significantly larger at the MP2 level. This shows that the non-classical correlation contribution is almost pairwise additive. Therefore, the best results are obtained with the PMISP-CE method, which uses the PMISP procedure only for the correlation contribution.

A detailed investigation shows that the error introduced by the MFCC procedure is negligible, both for the assembly of properties and for the evaluation of the non-classical term. Instead, the three-body contributions to the coupling between the classical and non-classical terms are the main cause of the error. This can be addressed by an embedded model (EMB-PMISP), but the effect is overestimated because of the lack of Pauli effects. Introducing pseudopotentials would probably solve the problem.

Special attention in paper V is given to the problem of intramolecular polarization. It is shown that the neglect of intramolecular coupling of polarizabilities, which is a reasonable approximation in paper III, gives a significant error (in average  $\sim 8 \text{ kJ/mol}$ ) for the large molecule A. We also emphasize the importance of using a consistent definition of the electrostatic and induction interaction energy when comparing results of various force fields: the direct use of the electrostatic and polarization terms of the potential energy is not adequate.

The interaction energy for the full avidin-biotin complex, computed in paper VI, is given in Table 4.4 for PMISP/MM (with two different basis sets) and two Amber force fields. The basis-set dependence is significant, indicating that the 6-31G\* result is quite meaningless. Moreover, the two Amber force fields give widely different results, with the PMISP/MM estimate in between. The term-

Table 4.4: Interaction energies (in kJ/mol) of the avidin–biotin complex with the PMISP/MM method using two different basis sets and with two Amber force fields, divided into electrostatic  $(E^{ele})$ , induction  $(E^{ele})$ , and non-classical  $(E^{nc})$  energy.

	$MP2/6-31G^*$	MP2/aug-cc-pVTZ	Amber ff94	Amber ff02
$E^{ele}$	-1120	-1126	-1300	-1096
$E^{ind}$	-190	-286	0	-113
$E^{nc}$	50	-8	-143	-143
$E_{tot}$	-1261	-1419	-1443	-1353

wise comparison shows that the polarizable Amber ff02 force field significantly underestimates the induction energy, whereas this is partly cancelled by the non-classical term.

The approximation to take the long-range non-classical contributions from a standard force field is evaluated by varying the cutoff distance for the PMISP treatment. At 4 Å, this approximation gives an error of 16 kJ/mol, but this error can be reduced by including extra atoms for the strongest interacting groups. At 7 Å, the error is negligible. Various approximations to the classical term are also tested. It is found that MP2 properties can be replaced by B3LYP properties without loss of accuracy. Octupoles and quadrupoles can be neglected outside a distance of 2 and 5 Å from the ligand, respectively, and polarizability coupling outside of 15 Å. A significantly smaller basis set can also be used for the property calculations outside of 15 Å.

In the computation of binding free energies in paper VII, still for biotin analogues binding to avidin, we encounter several problems. In particular, the difference in non-polar solvation energies obtained with the PCM model deviates significantly from other estimates. Therefore, the solvent-accessible surface area (SASA) model is also tested. However, no matter which definition of non-polar solvation energy one uses, the results are disappointing when compared to experimental binding affinities, giving a mean absolute deviation of  $\sim 20 \text{ kJ/mol}$  for relative affinities. The reason might be that the solvation model is not sufficiently accurate to enable the improved interaction energies to make a difference in the total results.

## 4.5 Conclusions and outlook

The papers presented above can be said to represent three different approaches to a common goal: the calculation of protein–ligand energies from first principles. Papers I and II represent the traditional molecular-mechanical picture, in which properties are computed separately for each monomer by quantumchemical calculations, and then combined using simple expressions. The properties in paper I are molecular orbitals or electron densities, which are somewhat awkward to handle in a molecular-mechanics force field, but there are already such examples [47,83].

Papers III and especially paper IV develop a more sophisticated approach. Although it is shown that the polarizability model works well in most cases, this model is based on error cancellation, and thus there will always be systems for which it is inaccurate. Moreover, it cannot easily account for explicit coupling with the repulsion. A solution to this is found to be a model (ppEPOL) in which the monomers are treated quantum-chemically but the interactions are included by simple expressions. Although the main purpose of this method is as a reference for models using a simpler description of the monomers (i.e. of the previous type), we also believe that it has a potential on its own.

Finally, the last three papers represent an approach (PMISP) in which explicit quantum-chemical supermolecular calculations are performed, but only for small subsystems. The many-body effects are instead modeled by a pointpolarizability model. Partly, the same problem is seen here as in paper III, but as it enters only in the many-body effects, it is typically much smaller.

Which of these approaches is most suitable for achieving the goal? Surely this depends on the level of accuracy required and to what computational cost. The first approach will always be the fastest. However, the accuracy may be limited for at least two reasons. First, the coupling between polarization and repulsion will necessarily be rather approximate, and second, most such methods will have to include some type of fitting, which may limit the transferability. The third approach will always be the most accurate one, but to a higher computational cost, allowing much fewer energy evaluations for a given problem. For the second approach to be computationally viable, it is essential to use simpler models for the electric field and the dispersion energy. If this is done, such type of methods could be a suitable compromise between the other two.

What is the accuracy needed? In my view, the best way to address this question is to first use a potential-energy method that is as accurate as possible and then, when agreement with experiment has been confirmed, gradually replace it by computationally cheaper estimates. However, as the result of paper VII shows, this is not a trivial thing to do. Because we are currently limited to methods that treat entropic and solvent effects at a much more approximate level, it is not guaranteed that an improvement is obtained. We are currently investigating the possibility of calculating a rigorous free energy with a standard force field and using perturbations from the standard force field to the quantum-chemical level at the end-points [234], but the preliminary results are disappointing: The geometries obtained with a standard force field are too different from those that would be obtained by a quantum-chemical simulation, thus prohibiting a direct perturbational treatment. However, the use of a better force field for the simulations would possibly solve this problem.

# Bibliography

- [1] J. W. Ponder and D. A. Case, Adv. Protein Chem., 2003, 66, 27.
- [2] P. A. M. Dirac, Proc. R. Soc. A, 1929, 123, 714.
- [3] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry*, Dover Publications, Mineola, 2nd ed., 1989.
- [4] P. W. Atkins and R. S. Friedman, *Molecular quantum mechanics*, Oxford university press, Oxford, 3rd ed., 1997.
- [5] P. Hohenberg and W. Kohn, Phys. Rev. B, 1964, 136, 864.
- [6] W. Kohn and L. J. Sham, Phys. Rev. A, 1965, 140, 1133.
- [7] J. D. Head and M. C. Zerner, Adv. Quantum Chem., 1989, 20, 239.
- [8] G. M. Torrie and J. P. Valleau, J. Chem. Phys., 1977, 23, 187.
- [9] R. H. Swendsen and J. S. Wang, Phys. Rev. Lett., 1986, 57, 2607.
- [10] R. W. Zwanzig, J. Chem. Phys., 1954, 22, 1420.
- [11] A. R. Leach, Molecular Modelling: Principles and Applications, 2nd edition, Pearson Education Limited, Harlow, 2001.
- [12] P. Jurecka, J. Sponer, J. Cerny, and P. Hobza, *Phys. Chem. Chem. Phys.*, 2006, 8, 1985.
- [13] S. Grimme, J. Chem. Phys., 2003, 118, 9095.
- [14] R. A. Bachorz, F. A. Bischoff, S. Höfener, W. Klopper, P. Ottiger, R. Leist, J. A. Frey, and S. Leutwyler, *Phys. Chem. Chem. Phys.*, 2008, **10**, 2758.
- [15] J.-P. Piquemal, A. Marquez, O. Parisel, and C. Giessner-Prettre, J. Comput. Chem., 2005, 26, 1052.
- [16] J. Cerny and P. Hobza, Phys. Chem. Chem. Phys., 2007, 9, 5291.
- [17] C. Diedrich, A. Lüchow, and S. Grimme, J. Chem. Phys., 2005, 123, 184106.
- [18] F. B. van Duijneveldt, J. G. C. M. van Duijneveldt-van de Rijdt, and J. H. van Lenthe, *Chem. Rev.*, 1994, 94, 1873.
- [19] S. F. Boys and F. Bernardi, *Molecular Physics*, 1970, **19**, 553.
- [20] G. Karlström and A. Sadlej, Theor. Chim. Acta, 1982, 61, 1.
- [21] F. B. van Duijneveldt in *Molecular Interactions*, ed. S. Scheiner; Wiley, Chichester, 1997; p. 81.
- [22] O. Marchetti and H.-J. Werner, Phys. Chem. Chem. Phys., 2008, 10, 3400.
- [23] A. Halkier, W. Klopper, T. Helgaker, P. Jorgensen, and P. R. Taylor, J. Chem. Phys., 1999, 111, 9157.
- [24] A. Halkier, T. Helgaker, P. Jørgensen, W. Klopper, H. Koch, J. Olsen, and A. K. Wilson, *Chem. Phys. Letters*, 1998, **286**, 243.
- [25] F.-M. Tao and Y.-K. Pan, J. Chem. Phys., 1992, 97, 4989.
- [26] W. Klopper and W. Kutzelnigg, Chem. Phys. Letters, 1987, 134, 17.
- [27] H. Fliegl, W. Klopper, and C. Hättig, J. Chem. Phys., 2005, 122, 084107.
- [28] T. B. Adler, G. Knizia, and H.-J. Werner, J. Chem. Phys., 2007, 127, 221106.
- [29] K. E. Riley and P. Hobza, J. Phys. Chem. A, 2007, 111, 8257.
- [30] A. D. Buckingham, Adv. Chem. Phys., 1967, 12, 107.
- [31] M. Ziólkowski, S. J. Grabowski, and J. Leszczynski, J. Phys. Chem. A, 2006, 110, 6514.
- [32] A. J. Stone, The Theory of Intermolecular Forces, Oxford University Press, Oxford, 1st ed., 1996.
- [33] B. Jeziorski, R. Moszynski, and K. Szalewicz, Chem. Rev., 1994, 94, 1887.
- [34] I. C. Hayes and A. J. Stone, Mol. Phys., 1984, 53, 69.
- [35] K. Szalewicz, K. Patkowski, and B. Jeziorski, Struct. Bond., 2005, 116, 43.
- [36] M. Gutowski, G. Chalasinski, and J. van Duijneveldt-van de Rijdt, Int. J. Quantum Chem., 1984, 26, 971.
- [37] H. L. Williams, E. M. Mas, K. Szalewicz, and B. Jeziorski, J. Chem. Phys., 1995, 103, 7374.

- [38] K. Kitaura and K. Morokuma, Int. J. Quantum Chem., 1976, 10, 325.
- [39] P. S. Bagus, K. Hermann, and C. W. J. Bauschlicher, J. Chem. Phys., 1984, 80, 4378.
- [40] W. J. Stevens and W. H. Fink, Chem. Phys. Letters, 1987, 139, 15.
- [41] W. Chen and M. S. Gordon, J. Phys. Chem., 1996, 100, 14316.
- [42] R. F. Frey and E. R. Davidson, J. Chem. Phys., 1989, 90, 5555.
- [43] W. Cornell, P. Cieplak, C. Bayly, I. Gould, K. M. Merz, D. Ferguson, D. Spellmeyer, T. Fox, J. Caldwell, and P. Kollman, J. Am. Chem. Soc., 1995, 117, 5179.
- [44] S. W. Rick and S. J. Stuart in *Reviews in Computational Chemistry*, ed. K. B. Lipowitz and D. B. Boyd, Vol. 18; Wiley-VCH, New York, 2002; p. 89.
- [45] D. P. Geerke and W. F. van Gunsteren, J. Phys. Chem. B, 2007, 111, 6425.
- [46] N. Gresh, G. A. Cisneros, T. A. Darden, and J.-P. Piquemal, J. Chem. Theory. Comput., 2007, 3, 1960.
- [47] M. S. Gordon, M. A. Freitag, P. Bandyopadhyay, J. H. Jensen, V. Kairys, and W. J. Stevens, J. Phys. Chem. A, 2001, 105, 293.
- [48] A. Wallqvist and G. Karlström, Chem. Scr. A., 1989, 29, 131.
- [49] O. Engkvist, P.-O. Åstrand, and G. Karlström, Chem. Rev., 2000, 100, 4087.
- [50] G. A. Kaminski, H. A. Stern, B. J. Berne, and R. A. Friesner, J. Phys. Chem. A, 2004, 108, 621–627.
- [51] A. G. Donchev, V. D. Ozrin, M. V. Subbotin, O. V. Tarasov, and V. I. Tarasov, Proc. Natl. Acad.. Sci., 2005, 102, 7829.
- [52] P. Ren and J. W. Ponder, J. Comput. Chem., 2002, 23, 1497.
- [53] P. Cieplak, J. Caldwell, and P. Kollman, J. Comput. Chem., 2001, 22, 1048.
- [54] S. Patel and C. L. Brooks III, J. Comput. Chem., 2004, 25, 1.
- [55] A. Warshel, M. Kato, and A. V. Pisliakov, J. Chem. Theory. Comput., 2007, 3, 2034.
- [56] T. A. Halgren and W. Damm, Curr. Opin. Struct. Biol., 2001, 11, 236.

- [57] C. J. F. Böttcher, O. C. van Belle, P. Bordewijk, and A. Rip, *Theory of Electric Polarization*, vol. 1, Elsevier Scientific Publishing Company, Amsterdam, 2nd ed., 1973.
- [58] E. Sigfridsson and U. Ryde, J. Comput. Chem., 1998, 19, 377.
- [59] C. Chipot, J. G. Angyán, G. G. Ferenczy, and H. A. Scheraga, J. Phys. Chem., 1993, 97, 6628.
- [60] C. I. Bayly, P. Cieplak, W. D. Cornell, and P. A. Kollman, J. Phys. Chem., 1993, 97, 10269.
- [61] R. C. A., J. W. Essex, and W. G. Richards, J. Am. Chem. Soc., 1992, 114, 9075.
- [62] P. Söderhjelm and U. Ryde, J. Comput. Chem., 2008, p. DOI: 10.1002/jcc.21097.
- [63] D. E. Williams, J. Comput. Chem., 1988, 9, 745.
- [64] R. S. Mulliken, J. Chem. Phys., 1955, 23, 1833.
- [65] P.-O. Löwdin, J. Chem. Phys., 1950, 18, 365.
- [66] A. J. Stone and M. Alderton, Mol. Phys., 1985, 56, 1047.
- [67] G. Karlström in Proceeding of fifth seminar on Computational Methods in Quantum Chemistry, ed. P. T. van Duijnen and W. C. Nieuwpoort; Laboratory of Chemical Physics, University of Groningen, Groningen, The Netherlands, 1981; p. 353.
- [68] F. Vigné-Maeder and P. Claverie, J. Chem. Phys., 1988, 88, 4934.
- [69] A. Reed, L. Curtiss, and F. Weinhold, Chem. Rev., 1988, 88, 899.
- [70] R. F. W. Bader, Atoms in Molecules: A quantum theory, Oxford Press, Oxford, 1990.
- [71] R. F. W. Bader, Chem. Rev., 1991, 91, 893.
- [72] P. L. A. Popelier, Atoms in Molecules: An introduction, Pearson, London, 2000.
- [73] J. Pilmé and J.-P. Piquemal, J. Comput. Chem., 2008, 29, 1440.
- [74] A. J. Stone and M. Alderton, J. Chem. Theory. Comput., 2005, 1, 1128.
- [75] J. Cioslowski, J. Am. Chem. Soc., 1989, 111, 8333.
- [76] L. Gagliardi, R. Lindh, and G. Karlström, J. Chem. Phys., 2004, 121, 4494.

- [77] M. Rafat and P. L. A. Popelier, J. Comput. Chem., 2007, 28, 832.
- [78] G. Náray-Szabó and G. G. Ferenczy, Chem. Rev., 1995, 95, 829.
- [79] G. G. Hall and C. M. Smith, Int. J. Quantum Chem., 1984, 25, 881.
- [80] R. J. Wheatley and J. B. O. Mitchell, J. Comput. Chem., 1994, 15, 1187.
- [81] K. Strasburger, Computers Chem., 1998, 22, 7.
- [82] G. A. Cisneros, J.-P. Piquemal, and T. A. Darden, J. Chem. Phys., 2005, 123, 044109.
- [83] J.-P. Piquemal, G. A. Cisneros, P. Reinhardt, N. Gresh, and T. A. Darden, J. Chem. Phys., 2006, 124, 104101.
- [84] G. A. Cisneros, J.-P. Piquemal, and T. A. Darden, J. Chem. Phys., 2006, 125, 184101.
- [85] G. A. Cisneros, D. Elking, J.-P. Piquemal, and T. A. Darden, J. Phys. Chem. A, 2007, 111, 12049.
- [86] B. I. Dunlap, J. W. D. Connolly, and J. R. Sabin, J. Chem. Phys., 1979, 71, 4993.
- [87] A. Gavezzotti, J. Phys. Chem. B, 2003, 107, 2344.
- [88] Y. G. Ma and P. Politzer, J. Chem. Phys., 2004, 120, 8955.
- [89] A. Volkov, T. S. Koritsanszky, and P. Coppens, *Chem. Phys. Letters*, 2004, 391, 170.
- [90] A. Volkov and P. Coppens, J. Comput. Chem., 2004, 25, 921.
- [91] V. Kairys and J. H. Jensen, Chem. Phys. Letters, 1999, 315, 140.
- [92] M. A. Freitag, M. S. Gordon, J. H. Jensen, and W. J. Stevens, J. Chem. Phys., 2000, 112, 7300.
- [93] J.-P. Piquemal, N. Gresh, and C. Giessner-Prettre, J. Phys. Chem. A, 2003, 107, 10353.
- [94] J. Applequist, J. R. Carl, and K.-K. Fung, J. Am. Chem. Soc., 1972, 94, 2952.
- [95] B. T. Thole, Chem. Phys., 1981, 59, 341.
- [96] M. Masia, M. Probst, and R. Rey, J. Chem. Phys., 2005, 123, 164505.
- [97] A. Warshel and M. Levitt, J. Mol. Biol., 1976, 103, 227.

- [98] S. Nakagawa, P. Mark, and H. Ågren, J. Chem. Theory. Comput., 2007, 3, 1947.
- [99] T. Straatsma and J. McCammon, Molecular Simulation, 1990, 5, 181.
- [100] G. A. Kaminski, R. A. Friesner, and R. Zhou, J. Comput. Chem., 2003, 24, 267.
- [101] O. Engkvist, P.-O. Åstrand, and G. Karlström, J. Phys. Chem., 1996, 100, 6950.
- [102] A. Holt and G. Karlström, J. Comput. Chem., 2008, 29, 1084.
- [103] A. Holt and G. Karlström, J. Comput. Chem., 2008, 29, 1905.
- [104] T. D. Rasmussen, P. Ren, J. W. Ponder, and F. Jensen, Int. J. Quantum Chem., 2007, 107, 1390.
- [105] A. J. Stone, Mol. Phys., 1985, 56, 1065.
- [106] C. R. Le Sueur and A. J. Stone, Mol. Phys., 1994, 83, 293.
- [107] D. R. Garmer and W. J. Stevens, J. Phys. Chem., 1989, 93, 8263.
- [108] A. J. Misquitta and A. J. Stone, J. Chem. Phys., 2006, 124, 024111.
- [109] A. Holt and G. Karlström, Chem. Phys. Letters, 2007, 436, 297.
- [110] N. Celebi, J. G. Ángyán, F. Dehez, C. Millot, and C. Chipot, J. Chem. Phys., 2000, 112, 2709.
- [111] S. Nakagawa and N. Kosugi, Chem. Phys. Letters, 1993, 210, 180.
- [112] G. J. Williams and A. J. Stone, J. Chem. Phys., 2003, 119, 4620.
- [113] C. R. Le Sueur and A. J. Stone, Mol. Phys., 1993, 78, 1267.
- [114] J. G. Angyán, G. Jansen, M. Loos, C. Hättig, and B. A. Heß, Chem. Phys. Letters, 1994, 219, 267.
- [115] F. Dehez, C. Chipot, C. Millot, and J. G. Ángyán, *Chem. Phys. Letters*, 2001, **338**, 180.
- [116] M. Ferraro, M. Caputo, and P. Lazzeretti, J. Chem. Phys., 1988, 109, 2987.
- [117] G. Karlström, Theor. Chim. Acta, 1982, 60, 535.
- [118] A. Serr and R. R. Netz, Int. J. Quantum Chem., 2006, 106, 2960.
- [119] A. Holt and G. Karlström, J. Comput. Chem., 2008, 29, 2033.

- [120] J.-P. Piquemal, R. Chelli, P. Procacci, and N. Gresh, J. Phys. Chem. A, 2007, 111, 8170.
- [121] K. T. Tang and Toennies, J. Chem. Phys., 1984, 80, 3726.
- [122] M. P. Hodges and A. J. Stone, Mol. Phys., 2000, 98, 275.
- [123] I. Adamovic and M. S. Gordon, Mol. Phys., 2005, 103, 379.
- [124] G. J. Williams and A. J. Stone, J. Chem. Phys., 2003, 119, 4620.
- [125] A. J. Misquitta, B. Jeziorski, and K. Szalewicz, Phys. Rev. Lett., 2003, 91, 033201.
- [126] A. D. Becke and E. R. Johnson, J. Chem. Phys., 2005, 122, 154104.
- [127] E. R. Johnson and A. D. Becke, J. Chem. Phys., 2005, 123, 024101.
- [128] A. Öhrn and G. Karlström, J. Phys. Chem. A, 2007, 111, 10468.
- [129] S. Brdarski and G. Karlström, J. Phys. Chem. A, 1998, 102, 8182–8192.
- [130] E. Fraschini and A. J. Stone, J. Comput. Chem., 1998, 19, 847.
- [131] E. J. Baerends in *Cluster Models for Surface and Bulk Phenomena*, ed. G. Pacchione, P. S. Bagus, and F. Parmigiani; Plenum, New York, 1992.
- [132] A. J. Stone and S. L. Price, J. Phys. Chem., 1988, 92, 3325.
- [133] H. Margenau and N. R. Kestner, Theory of Intermolecular Forces, Pergamon, Oxford, 1969.
- [134] B. Jeziorski, M. Bulski, and L. Piela, Int. J. Quantum Chem., 1976, 10, 281.
- [135] J. H. Jensen, J. Chem. Phys., 1996, 104, 7795.
- [136] J. Jensen and M. S. Gordon, Mol. Phys., 1996, 89, 1313.
- [137] A. Wallqvist and G. Karlström, Chem. Scripta, 1989, 29A, 131.
- [138] N. Gresh, P. Claverie, and A. Pullman, Int. J. Quantum Chem., 1986, 29, 101.
- [139] J. N. Murrell and J. J. N. Teixeiras-Dias, Mol. Phys., 1970, 19, 521.
- [140] N. Gresh, J.-P. Piquemal, and M. Krauss, J. Comput. Chem., 2005, 26, 1113.
- [141] Y. S. Kim, S. K. Kim, and W. D. Lee, Chem. Phys. Letters, 1981, 80, 574.

- [142] I. Nobeli, S. L. Price, and R. J. Wheatley, Mol. Phys., 1998, 95, 525.
- [143] A. J. Stone and C.-S. Tong, J. Comput. Chem., 1994, 15, 1377.
- [144] R. Moszynski, P. E. S. Wormer, B. Jeziorski, and A. van der Avoird, J. Chem. Phys., 1995, 103, 8058.
- [145] E. M. Mas, R. Bukowski, and K. Szalewicz, J. Chem. Phys., 2003, 118, 4386.
- [146] J. H. Jensen, J. Chem. Phys., 2001, 114, 8775.
- [147] M. Gutowski, M. Kakol, and L. Piela, Int. J. Quantum Chem., 1983, 23, 1843–1853.
- [148] Z. Barandiarán and L. Seijo, J. Chem. Phys., 1988, 89, 5739–5746.
- [149] S. Huzinaga, J. Mol. Struc. (Theochem), 1991, 234, 51–73.
- [150] I. Panas, Chem. Phys. Letters, 1993, 201, 255–260.
- [151] A. Öhrn and G. Karlström, Mol. Phys., 2006, 104, 3087–3099.
- [152] H. M. Senn and W. Thiel, Curr. Opin. Chem. Biol., 2007, 11, 182.
- [153] E. Valderrama and R. J. Wheatley, J. Comput. Chem., 2003, 24, 2075.
- [154] G. A. Cisneros, J.-P. Piquemal, and T. A. Darden, J. Chem. Phys., 2006, 125, 184101.
- [155] J. Neugebauer, M. J. Louwerse, E. J. Baerends, and T. A. Wesolowski, J. Chem. Phys., 2005, **122**, 094115.
- [156] H. Li, M. S. Gordon, and J. H. Jensen, J. Chem. Phys., 2006, 124, 214108.
- [157] A. J. Stone, Chem. Phys. Letters, 1993, 211, 101.
- [158] D. Hagberg, G. Karlström, B. O. Roos, and L. Gagliardi, J. Am. Chem. Soc., 2005, 127, 14250.
- [159] H. Gohlke and G. Klebe, Angew. Chem. Int. Ed., 2002, 41, 2644.
- [160] P. Ferrara, H. Gohlke, D. J. Price, G. Klebe, and C. L. Brooks, III, J. Med. Chem., 2004, 47, 3032.
- [161] M. K. Gilson, J. A. Given, B. L. Bush, and J. A. McCammon, *Biophys. Jour.*, 1997, 72, 1047.
- [162] C.-E. Chang and M. K. Gilson, J. Am. Chem. Soc., 2004, 126, 13156.
- [163] M. K. Gilson and H.-X. Zhou, Annu. Rev. Biophys. Biomol. Struct., 2007, 36, 21.

- [164] B. L. Tembe and J. A. McCammon, Comput. Chem., 1984, 8, 281.
- [165] D. A. Pearlman and P. S. Charifson, J. Med. Chem., 2001, 44, 3417.
- [166] C. Oostenbrink and W. F. van Gunsteren, Proteins: Struct. Funct. Bioinformatics, 2004, 54, 237.
- [167] J. W. Essex, D. L. Severance, J. Tirado-Rives, and W. L. Jorgensen, J. Phys. Chem. B, 1997, 101, 9663.
- [168] W. L. Jorgensen, J. K. Buckner, S. Boudon, and J. Tirado-Rives, J. Chem. Phys., 1988, 89, 3742.
- [169] H. Fujitani, Y. Tanida, M. Ito, G. Jayachandran, C. D. Snow, M. R. Shirts, E. J. Sorin, and V. S. Pande, J. Chem. Phys., 2005, 123, 084108.
- [170] S. Boresch, F. Tettinger, M. Leitgeb, and M. Karplus, J. Phys. Chem. B, 2003, 107, 9535.
- [171] J. Wang, Y. Deng, and B. Roux, Biophys. Jour., 2006, 91, 2798.
- [172] V. Helms and R. C. Wade, Proteins: Struct. Funct. Genet., 1998, **32**, 381.
- [173] W. Hyang-June and B. Roux, Proc. Natl. Acad. Sci., 2005, 102, 6825.
- [174] M. S. Lee and M. A. Olson, *Biophys. Jour.*, 2006, 90, 864.
- [175] T. Rodinger and R. Pomès, Curr. Opin. Struct. Biol., 2005, 15, 164.
- [176] C. Oostenbrink and W. F. van Gunsteren, Proteins: Struct. Funct. Bioinformatics, 2004, 54, 237.
- [177] C. Oostenbrink and W. F. van Gunsteren, Proc. Natl. Acad. Sci., 2005, 102, 6750.
- [178] R. J. Radmer and P. A. Kollman, J. Comput.-Aid. Mol. Design, 1998, 12, 215.
- [179] F. S. Lee, Z.-T. Chu, M. B. Bolger, and A. Warshel, Protein Engineering, 1992, 5, 215.
- [180] Y. Y. Sham, Z. T. Chu, H. Tao, and A. Warshel, *Proteins: Struct. Funct. Genet.*, 2000, **39**, 393.
- [181] J. Åqvist, C. Medina, and J.-E. Samuelsson, Protein Engineering, 1994, 7, 385.
- [182] J. Åqvist and T. Hansson, J. Phys. Chem., 1996, 100, 9512.
- [183] J. Aqvist, V. B. Luzhkov, and B. O. Brandsdal, Acc. Chem. Res., 2002, 35, 358.

- [184] M. Almlöf, B. O. Brandsdal, and J. Åqvist, J. Comput. Chem., 2004, 25, 1242.
- [185] R. Zhou, R. A. Friesner, A. Ghosh, R. C. Rizzo, W. L. Jorgensen, and R. M. Levy, J. Phys. Chem. B, 2001, 105, 10388.
- [186] M. N. J. Carlsson, M. Andér and J. Åqvist, J. Phys. Chem. B, 2006, 110, 12034.
- [187] P. A. Kollman, I. Massova, C. Reyes, B. Kuhn, S. Huo, L. Chong, M. Lee, T. Lee, Y. Duan, W. Wang, O. Donini, P. Cieplak, J. Srinivasan, D. A. Case, and T. E. Cheatham, Acc. Chem. Res., 2000, 33, 889.
- [188] J. Wang, P. Morin, W. Wang, and P. A. Kollman, J. Am. Chem. Soc., 2001, 123, 5221.
- [189] A. Weis, K. Katebzadeh, P. Söderhjelm, I. Nilsson, and U. Ryde, J. Med. Chem., 2006, 49, 6596.
- [190] F. Gräter, S. M. Schwarzl, A. Dejaegere, S. Fischer, and J. C. Smith, J. Phys. Chem. B, 2005, 109, 10474.
- [191] M. Wang and C. F. Wong, J. Chem. Phys., 2007, 126, 026101.
- [192] M. Kaukonen, P. Söderhjelm, J. Heimdal, and U. Ryde, J. Phys. Chem. B, 2008, 112, 12537.
- [193] J. Kongsted and U. Ryde, J. Comput.-Aid. Mol. Design, 2008, DOI: 10.1007/s10822-008-9238-z.
- [194] C. A. Chang, W. Chen, and M. K. Gilson, Proc. Natl. Acad. Sci., 2007, 104, 1534.
- [195] J. M. J. Swanson, R. H. Henchman, and J. A. McCammon, *Biophys. Jour.*, 2004, 86, 67.
- [196] D. A. Pearlman, J. Med. Chem., 2005, 48, 7796.
- [197] B. Kuhn, P. Gerber, T. Schulz-Gasch, and M. Stahl, J. Med. Chem., 2005, 48, 4040.
- [198] J. M. Hermida-Ramón, S. Brdarski, G. Karlström, and U. Berg, J. Comput. Chem., 2002, 24, 161.
- [199] M. S. Lee and M. A. Olson, J. Phys. Chem. B, 2005, 109, 5223.
- [200] M. R. Shirts, J. W. Pitera, W. C. Swope, and V. S. Pande, J. Chem. Phys., 2003, 119, 5740.
- [201] L. Y. Zhang, E. Gallicchio, and R. A. Friesner, J. Comput. Chem., 2001, 22, 591.

- [202] D. Beglov and B. Roux, J. Phys. Chem. B, 1997, 101, 7821.
- [203] B. Roux and T. Simonson, J. Chem. Phys., 2006, 124, 084905.
- [204] D. Bashford and D. A. Case, Annu. Rev. Phys. Chem., 2000, 51, 129.
- [205] S. Miertus, E. Scrocco, and J. Tomasi, Chem. Phys., 1981, 55, 117.
- [206] H. Li, C. S. Pomelli, and J. H. Jensen, Theor. Chem. Acc., 2003, 109, 71.
- [207] S. T. Russell and A. Warshel, J. Mol. Biol., 1985, 185, 389.
- [208] A. Papazyan and A. Warshel, J. Phys. Chem. B, 1997, 101, 11254.
- [209] L. Sandberg, R. Casemyr, and O. Edholm, J. Phys. Chem. B, 2002, 106, 7889.
- [210] F. J. Luque, C. Curutchet, J. Munoz-Muriedas, A. Bidon-Chanal, I. Soteras, A. Morreale, J. L. Gelpi, and M. Orozco, *Phys. Chem. Chem. Phys.*, 2003, 5, 3827.
- [211] N. Huang, C. Kalyanaraman, K. Bernacki, and M. P. Jacobson, *Phys. Chem. Chem. Phys.*, 2006, 8, 5166.
- [212] J. Antony, J.-P. Piquemal, and N. Gresh, J. Comput. Chem., 2005, 26, 1131.
- [213] C. Roux, N. Gresh, L. E. Perera, J.-P. Piquemal, and L. J. Salmon, J. Comput. Chem., 2007, 28, 938.
- [214] D. Jiao, P. A. Golubkov, T. A. Darden, and P. Ren, Proc. Natl. Acad.. Sci., 2008, 105, 6290.
- [215] J. R. Maple, Y. Cao, W. Damm, T. A. Halgren, G. A. Kaminski, L. Y. Zhang, and R. A. Friesner, J. Chem. Theory. Comput., 2005, 1, 694.
- [216] N. Gresh, A. Pullman, and P. Claverie, *Theor. Chim. Acta*, 1985, 67, 11.
- [217] N. Gresh, P. Claverie, and A. Pullman, Theor. Chim. Acta, 1984, 66, 1.
- [218] P. A. Molina, H. Li, and J. H. Jensen, J. Comput. Chem., 2003, 24, 1971.
- [219] K. Raha, M. B. Peters, B. Wang, N. Yu, A. M. Wollacott, L. M. Westerhoff, and K. M. Merz, *Drug Discovery Today*, 2007, **12**, 731.
- [220] A. E. Cho, V. Guallar, B. J. Berne, and R. Friesner, J. Comput. Chem., 2005, 26, 915.
- [221] A. Khandelwal, V. Lukacova, D. Comez, D. M. Kroll, S. Raha, and S. Balaz, J. Med. Chem., 2005, 48, 5437.

- [222] W. Yang, Phys. Rev. Lett., 1991, 66, 1438.
- [223] K. Raha and K. M. Merz, J. Am. Chem. Soc., 2004, 126, 1020.
- [224] K. Kitaura, E. Ikeo, T. Asada, T. Nakano, and M. Uebayasi, Chem. Phys. Letters, 1999, 313, 701.
- [225] K. Fukuzawa, Y. Mochizuki, S. Tanaka, K. Kitaura, and T. Nakano, J. Phys. Chem. B, 2006, 110, 16102.
- [226] I. Nakanishi, D. G. Fedorov, and K. Kitaura, Proteins: Struct. Funct. Bioinformatics, 2007, 68, 145.
- [227] D. G. Fedorov and K. Kitaura, J. Phys. Chem. A, 2007, 111, 6904.
- [228] D. W. Zhang and J. Z. H. Zhang, J. Chem. Phys., 2003, 119, 3599.
- [229] D. W. Zhang, Y. Xiang, M. Gao, and J. Z. H. Zhang, J. Chem. Phys., 2004, 120, 1145.
- [230] D. W. Zhang, Y. Xiang, and J. Z. H. Zhang, J. Phys. Chem. B, 2003, 107, 12039.
- [231] R. P. A. Bettens and A. M. Lee, J. Phys. Chem. A, 2006, 110, 8777.
- [232] R. P. A. Bettens and A. M. Lee, Chem. Phys. Letters, 2007, 449, 341.
- [233] E. E. Dahlke and D. G. Truhlar, J. Chem. Theory. Comput., 2007, 3, 46.
- [234] T. H. Rod and U. Ryde, Phys. Rev. Lett., 2005, 94, 138302.