

TECHNISCHE UNIVERSITÄT MÜNCHEN

Fakultät für Chemie
Lehrstuhl für Theoretische Chemie

FIRST-PRINCIPLES MULTI-SCALE
SIMULATIONS OF DYNAMIC CATALYST
SURFACES

CO OXIDATION FROM PALLADIUM SURFACE OXIDE TO
PALLADIUM METAL

MAX J. HOFFMANN

Vollständiger Abdruck der von der Fakultät für Chemie der Technischen Universität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften

genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. Ville Kaila

Prüfer der Dissertation:

1. Univ.-Prof. Dr. Karsten Reuter
2. Univ.-Prof. Dr. Ulrich K. Heiz

Die Dissertation wurde am 26.6.2014 bei der Technischen Universität München eingereicht und durch die Fakultät für Chemie am 16.7.2014 angenommen.

Max J. Hoffmann: *First-principles Multi-scale Simulations of Dynamic Catalyst Surfaces*, CO Oxidation from Palladium surface oxide to Palladium metal, © 2014

Diese Arbeit wurde in der Zeit von Oktober 2010 bis Juni 2014 am Lehrstuhl für Theoretische Chemie der Fakultät für Chemie der Technischen Universität München unter der Aufsicht von Prof. Dr. Karsten Reuter durchgeführt.

ABSTRACT

Heterogeneous catalysts are a keystone technology of the chemical industry. Transition metals represent an important class of materials that display catalytic activity. However, the exact atomistic nature of catalytic activity even for very simple oxidation reactions and elementary catalysts is often still not fully understood. The notion that the catalytically active surface is not a static surface structure but an ever-changing ensemble of different surface terminations that may only emerge under reaction conditions has recently received growing attention. To scrutinize this idea from a theoretical point of view new simulation methods are needed that specifically address complex reaction schemes such as catalytic reactions with ongoing morphological transitions of the catalyst surface.

As a prototypical model system the CO oxidation reaction on Pd(100) is studied. Recent experimental studies which identify either an ultra-thin surface oxide or the metallic fcc(100) termination as the catalytically active state give evidence that the single surface paradigm may not suffice to explain the surface reaction kinetics. The interface between different surface terminations may allow for different elementary processes, and a forming and decomposing surface oxide layer could become the key driver of catalytic activity.

This work employs a first-principles multiscale approach to investigate towards this hypothesis. Density functional theory (DFT) is used to predict the energetics and elementary processes from first-principles. Using the DFT results in combination with methods derived from thermodynamics and statistical mechanics allows to predict the interplay of different surface processes on a meso- and macroscopic time and length scale for a defined temperature and partial pressure conditions.

On this basis a kinetic Monte Carlo (kMC) model for the Pd(100) surface is constructed complementing an existing model for the surface oxide geometry. Comparison of the coverages, predicted by these two models, shows that the stability regimes of these two surfaces overlap significantly around stoichiometric feed conditions. This demonstrates that the temporal development of the surface may indeed be an important factor for the catalytic activity. In a second step an explicit atomistic pathway describing the reduction of the palladium surface oxide exposed to CO is constructed. A straightforward adaption of lattice kMC is devised and implemented into a generic computer package, to describe the complex multi-lattice feature of this reconstruction. A second kMC model is constructed which describes the surface oxide reduction as a result of surface chemical reactions and local reconstructions. In this model the surface oxide reduction kinetics are predicted in very good agreement with experiments unlike in previous single-lattice treatments. The key feature is that reaction steps across the domain boundary between surface oxide and metal surface patches are an essential contributor. Further studies focusing on the surface oxide formation are now needed to fully understand the catalytic activity.

KURZFASSUNG

Heterogene Katalysatoren sind eine Schlüsseltechnologie der chemischen Industrie. Übergangsmetalle sind eine wichtige Materialklasse, die katalytische Aktivität aufweisen. Dennoch ist der genaue Ursprung der katalytischen Aktivität selbst bei sehr einfachen Oxidationsreaktionen und elementaren Katalysatoren oftmals noch nicht vollständig verstanden. Die Vorstellung, dass die katalytisch aktive Oberfläche nicht eine statische Oberflächenstruktur, sondern ein sich ständig veränderndes Ensemble verschiedener Oberflächengeometrien ist, die nur unter den Reaktionsbedingungen entstehen, hat seit Kurzem zunehmend Aufmerksamkeit erhalten. Um diese Idee von einem theoretischen Standpunkt aus zu hinterfragen, sind neue Simulationsmethoden erforderlich, welche komplexe Reaktionsschemata wie Oberflächenreaktionen simultan zu morphologischen Umwandlungen des Katalysators beschreiben können.

Als repräsentatives Beispielsystem wird die CO Oxidationsreaktion auf Pd(100) untersucht. Jüngste experimentelle Studien, die entweder ein ultradünnes Oberflächenoxid oder die metallische fcc(100) Oberfläche als katalytisch aktiven Zustand identifizieren deuten an, dass das Bild einer einzelnen Oberfläche möglicherweise nicht ausreicht, um die Reaktionskinetik zu erklären. Die Phasengrenzfläche zwischen verschiedenen Oberflächendomänen könnte neue Elementarprozesse ermöglichen und eine wachsende und schrumpfende Oberflächenoxidschicht könnte der wichtigste Treiber der katalytischen Aktivität sein.

Die vorliegende Arbeit verfolgt einen ab initio Multiskalenansatz um sich dieser Hypothese zuzuwenden. Dichtefunktionaltheorie (DFT) wird verwendet, um die Energetik und Elementarprozesse vorherzusagen. Mit Hilfe der DFT-Ergebnisse und in Kombination mit Methoden aus der Thermodynamik und der statistischen Mechanik wird das Zusammenspiel verschiedener Oberflächenprozesse und Strukturen auf einer meso- und makroskopischen Zeit- und Längenskala für gegebene Temperatur und Partialdrücke simuliert.

In einem ersten Schritt wurde ein kinetisches Monte Carlo (kMC) Modell für die Pd(100) Oberfläche konstruiert, welches ein bestehendes Modell für das Oberflächenoxid komplementiert. Der Vergleich, der durch diese beiden Modelle vorhergesagten Bedeckungen, zeigt, dass sich die Stabilitätsregime dieser beiden Strukturen deutlich im Bereich stöchiometrischer Bedingungen überlappen. Das bedeutet, dass die zeitliche Entwicklung der Oberflächen ein wichtiger Faktor für die katalytische Aktivität sein kann. In einem zweiten Schritt wird ein atomistischer Reaktionspfad herausgearbeitet, welcher die Reduktion des Palladiumoberflächenoxids durch CO beschreibt. Eine Anpassung der eingittrigen kMC Methode zusammen mit der Implementierung in einem allgemeinen Computerprogramm erlaubt es, diese komplexe mehrgittrige Rekonstruktion effizient zu simulieren. Ein zweites kMC Modell wird konstruiert, das die Oberflächenoxidreduktion als Resultat der fortlaufenden chemischen Reaktionen beschreibt. Die Oberflächenoxid Reduktionskinetik in diesem Modell wird in sehr guter Übereinstimmung mit Experimenten vorhergesagt, im Gegensatz zu vorherigen eingittrigen Beschreibungen. Insbesondere zeigt sich, dass Reaktionsschritte über die Domänengrenze zwischen dem Oberflächenoxid und der Metalloberfläche hinweg in der Tat einen wesentlichen Beitrag dazu liefern. Weitere Studien mit Fokus auf dem Wachstum des Oberflächenoxids werden jetzt benötigt, um die katalytische Aktivität gänzlich zu verstehen.

CONTENTS

Abstract	iii
Kurzfassung	v
1 INTRODUCTION	1
I METHODOLOGICAL BACKGROUND	5
2 BELOW THE ATOMIC SCALE	7
2.1 Electronic Structure Theory	7
2.2 Computing the Electronic Ground-state	12
3 ABOVE THE ELECTRONIC SCALE	17
3.1 Defining the Structural Model	18
3.2 Reactants on Surfaces	19
3.3 Surfaces at Reaction Conditions	21
4 KMOS: A LATTICE KINETIC MONTE CARLO FRAMEWORK	23
4.1 Introduction	23
4.2 Theoretical Background	25
4.2.1 (Surface) Chemistry on a Lattice	25
Site-specific adsorption and lattice mapping	25
Rare-event dynamics and Markovian master equation	26
Locality of elementary reactions	26
Size and structure of the transition matrix	28
4.2.2 Kinetic Monte Carlo	29
4.2.3 Efficient Lattice kMC	31
4.2.4 Sampling of Reaction Rates	33
4.3 The <i>kmos</i> Framework	35
4.3.1 kMC Model Definition	35
4.3.2 Code Generator	39
4.3.3 Kinetic Monte Carlo Solver	43
4.3.4 Simulation Frontend	46
4.4 Performance and Scaling in Practice	46
4.4.1 ZGB Model	46
4.4.2 Literature First-Principles kMC Models	47
4.4.3 Random Models	49
4.5 Summary	51
5 TRANSITION STATE SEARCHES	53
5.1 Theoretical Background	53
5.1.1 (harmonic) Transition State Theory	53
5.1.2 Transition state search	55
5.2 Methods	55
5.2.1 Transition State Finding Methods	55
5.2.1.1 Drag Method	55
5.2.1.2 Climbing Image Nudged Elastic Band	56
5.2.1.3 Freezing string	56

	5.2.1.4	Freezing NEB	57
	5.2.2	Optimization Methods	57
	5.2.3	Benchmark Problem	58
	5.2.4	Implementation/Technical Section	59
5.3		Results	60
	5.3.1	Spring Constant k	60
	5.3.2	Number of Images	60
	5.3.3	Switching from FIRE to LBFGS	63
	5.3.4	Precision of Forces	64
5.4		Conclusions	66
II		CO OXIDATION FROM $\sqrt{5}$ -OXIDE TO Pd(100)	69
6		(BI)STABILITY OF Pd(100) AND $\sqrt{5}$ PHASE	71
	6.1	Introduction	71
	6.2	Methods	72
	6.2.1	1p-kMC simulations	72
	6.2.2	First-principles rate constants	73
	6.2.3	1p-kMC lattice models	75
	6.2.3.1	Pd(100)	75
	6.2.3.2	$\sqrt{5}$ surface oxide	77
	6.2.4	kMC simulation setup	79
	6.3	Results	79
	6.3.1	Kinetic phase diagrams at 600 K	79
	6.3.2	Variation with temperature	81
	6.4	Discussion	83
	6.4.1	Origin and robustness of bistability region	83
	6.4.2	Comparison to experiment	87
	6.5	Conclusions	89
7		REDUCTION OF Pd SURFACE OXIDE BY CO	91
	7.1	Introduction	91
	7.2	Theory	92
	7.2.1	Kinetic Monte Carlo simulations in surface catalysis	92
	7.2.2	Efficient 1p-kMC on one or more lattices	94
	7.2.3	Elementary processes during the reduction of oxidized Pd(100) within unperturbed surfaces	96
	7.2.4	kMC simulation setup	98
	7.3	Results	99
	7.3.1	Single-lattice 1p-kMC simulations of oxide reduction	99
	7.3.2	Atomistic pathway for oxide reduction	101
	7.3.3	Elementary processes at the Pd(100) nucleus under reduction conditions	105
	7.3.4	Multi-lattice 1p-kMC model	110
	7.3.5	Multi-lattice 1p-kMC simulations of oxide reduction	112

7.4	Summary and Conclusions	114
8	SUMMARY & OUTLOOK	117
III	APPENDIX	121
A	CASTEP - A TECHNICAL INTRODUCTION FOR PRACTITIONERS	123
B	CONVERGENCE TESTS	131
C	THE ASE TO CASTEP INTERFACE	137
	PUBLICATIONS	141
	BIBLIOGRAPHY	169
	ACRONYMS	171
	ACKNOWLEDGEMENTS	173

LIST OF FIGURES

Figure 3.1	Regimes and methods of multiscale approach.	17
Figure 3.2	Fcc(100) surface with unit cells.	18
Figure 3.3	Standard surface oxide slab model.	20
Figure 3.4	Fcc(100) high-symmetry sites.	20
Figure 4.1	Schematic lattice kMC reaction network.	25
Figure 4.2	Flowchart of VSSM kMC.	31
Figure 4.3	Flow of information in kmos.	35
Figure 4.4	Lattice representation in kmos.	37
Figure 4.5	Schematic kMC reactions.	38
Figure 4.6	kmos data structures.	44
Figure 4.7	Reference results or ZGB model.	47
Figure 4.8	kmos benchmark for 1p-kMC models.	48
Figure 4.9	kmos benchmark for random models.	49
Figure 5.1	Schematic potential energy surface with TST.	54
Figure 5.2	Transition state searches with different spring constants.	61
Figure 5.3	Transition state searches for different number of images.	62
Figure 5.4	Transition state searches switching the optimizer.	63
Figure 5.5	Transition state searches for different noise.	64
Figure 6.1	Schematic view of Pd(100) and $\sqrt{5}$ -oxide model.	75
Figure 6.2	KMC coverages on Pd(100) and $\sqrt{5}$ -oxide for 600 K.	80
Figure 6.3	Kinetic bistability regime for 400–600 K.	82
Figure 6.4	Bistability in thermodynamic phase diagram.	84
Figure 6.5	Comparison of phase boundaries for different CO lateral interactions.	86
Figure 6.6	Coverages and reaction rates of Pd(100) and $\sqrt{5}$ -oxide model for experimental conditions.	88
Figure 7.1	Schematic view of Pd(100) model.	96
Figure 7.2	Schematic view of $\sqrt{5}$ -oxide model.	97
Figure 7.3	Single-lattice simulations of $\sqrt{5}$ -oxide reduction.	100
Figure 7.4	$\sqrt{5}$ -oxide reduction degrees of freedom.	102
Figure 7.5	Atomistic $\sqrt{5}$ -oxide reduction pathway.	103
Figure 7.6	Adsorption sites near Pd(100)-nucleus.	105
Figure 7.7	Candidate CO oxidation reaction on Pd(100)-nucleus.	109
Figure 7.8	Multi-lattice simulations of $\sqrt{5}$ -oxide reduction.	113
Figure 7.9	Importance of elementary processes towards $\sqrt{5}$ -oxide reduction.	113

Figure 7.10	Multi-lattice simulations of $\sqrt{5}$ -oxide reduction for different domain sizes.	114
Figure B.1	Standard slab geometry.	132
Figure B.2	Converging molecular binding energies w.r.t. cut-off.	132
Figure B.3	Converging bulk energy w.r.t. cut-off.	133
Figure B.4	Converging surface energies w.r.t cut-off and k-point grid.	134
Figure B.5	Converging adsorption energies w.r.t. cut-off.	134
Figure B.6	Converging adsorption energies w.r.t vacuum.	135
Figure B.7	Converging adsorption energy w.r.t slab thickness.	135

LIST OF TABLES

Table 6.1	Activation energies entering single-lattice model.	78
Table 7.1	Binding energies near Pd(100)-nucleus.	106

INTRODUCTION

WHY CATALYSTS ARE IMPORTANT Effective heterogeneous catalysts have an effect on mankind that is difficult to overestimate: the invention of the Haber-Bosch cycle for ammonia synthesis and the three-way catalytic converters in automobiles are just two of the most prominent examples that are enabled by heterogeneous catalysts. They accelerate a chemical reaction by lowering its activation energy.[31, 39] Other industrially important oxidation reactions include the Ostwald process for nitric acid synthesis and the Fischer-Tropsch synthesis to produce e.g. fuels from natural gas. Heterogeneous catalysts are estimated to be responsible for 85% of bulk chemical production.[94] In light of globally diminishing natural resources and in particular fossil fuels, the need for efficient, highly-selective, cheap, and stable catalysts is eminent.[88] In order to meet this need research into catalytic systems and methodologies to develop and predict catalysts is a very active area of research.[1] In the future large scale solutions to problems that mankind is facing such as energy conversion and energy storage will very likely contain a catalyst that is yet to be found as an essential ingredient.

TRADITIONAL CATALYST DEVELOPMENT DOES NOT SCALE Traditionally, the process of discovering and developing catalysts is very slow and requires thousands of trial-and-error experiments, probing many different materials under many different conditions until a reactive, selective, stable, and available catalyst is found. And even if a successful candidate is identified this insight does not necessarily lead to discovery of further catalysts if the success or failure of a catalyst cannot be linked to the underlying chemical reaction network and material properties. This generally slow and cumbersome process has sparked an interest in predicting the functionality of a catalyst from an atomistic understanding[136, 138, 170, 171].

The compelling advantage of such an approach is that nowadays methods are available with which atomistic interactions processes can be predicted without recourse to experimental input and with steadily improving accuracy and efficiency.[138] Such methods predominantly rely on Density Functional Theory to develop an understanding of the functioning of a catalyst from first-principles. If approaches can be established that predict the reactive behavior of a catalyst for any desired reaction and external condition the tantalizing vision of rational design of optimal catalysts from all elements in the periodic table may be realized.

ENTER MODEL CATALYSTS Unfortunately this vision is quite a daunting challenge. In efforts to understand and describe existing catalysts alone it has become clear that the surface geometry is often not known and the full catalyst system may consist of a combination of different structures and different combinations of materials, defects, and various contaminations. Compared with the existing understanding derived from well-defined crystals or individual molecules there is obviously a non-trivial gap in complexity. There are essentially two complementary routes to handle this gap[136]: on the one hand there are high-speed screening approaches[125, 177] that identify atomistic descriptors that are hoped to directly correlate with the catalytic function on the macro-scale. With such a screening approach one can screen the periodic table of elements for optimal materials. On the other hand there is the multi-scale approach which links suitable simulation techniques for each length and time-scale in a bottom-up manner to predict macroscopic behavior of one particular catalyst material or surface. For the latter model catalysts, which include the study of single-crystal surfaces, provide for an essential stepping stone between experiment and theory to gain an integrated understanding. For this very reason the growing field of model catalysis has emerged which focuses on reduced structures (*i.e.* using single-crystals as a starting point) and studies the kinetics of a catalytic reaction.[14, 61]

THE MULTI-SCALE APPROACH Heterogeneous catalysis is inherently a multi-scale problem. From the subatomic scale where the electronic structure determines the rules for making and breaking of chemical bonds, over the interplay between many bond-breaking events, to the continuum scale where individual atoms no longer appear and heat- and mass-transport become the governing paradigm. To combine these different regimes in a consistent and rigorous manner we follow a so-called bottom-up handshake approach[140] where we calculate a multitude of significant features on the smaller scale which then serve as input to the next larger scale and coarse-grained description. By prudent combination of different levels of description up to ten orders of magnitude in time and length can be linked which underlines its uniqueness among theoretical approaches to condensed matter.

THE CURRENT CHALLENGES The study of model catalysts has led to the identification of several gaps or challenges common to theory and experiment between the atomistic surface science kind of picture and the technological application which can be grouped in to the following three areas: the pressure and temperature gap, the multi-scale challenge, and the complexity gap. The temperature and pressure gaps generally describe the fact that the most precise atomistic understanding is deduced via experiments in ultra-high vacuum and up to

room temperature conditions while most catalysts operate at ambient pressures or above and at elevated temperatures. Under those conditions completely different surface structures and compositions often arise and entirely different processes can happen.[195] Furthermore the heat and mass transport might become the limiting factor in the overall reaction.[121] A problem that researchers have only recently started to address by use of suitable continuous phase measurement techniques[19, 200] and coupling of surface models to continuous fluid dynamics calculations[101, 119, 121–123, 127, 173, 192]. The complexity or materials gap is admittedly the fuzziest area of those three as complexity can extend into many directions such as the complexity of the surface structure or the reaction network between intermediate adsorbates. To single out one example: consider the fact that state-of-the-art multi-scale modeling of catalysis is limited to describe the chemical reactions on a fixed single-crystal surface.[154] However, at increasing pressures and temperatures many transition-metal based catalysts are suspected to undergo phase transitions such as the formation of oxides, carbides, or sulfates. In particular the formation of an atomically thin surface oxide has been established as a unique surface phase of relevance in oxidation catalysis.[118]

Taking into account such reconstruction on an atomistic scale and include it in a multi-scale framework was up to now not feasible. The efficient, yet, predictive first-principles lattice kinetic Monte Carlo (kMC) approach[68, 153, 157], which allows for a reliable microkinetic description including concomitant spatial heterogeneities at the surface, did not offer the flexibility to include such reconstructions while simultaneously describing catalytic reactions. Extending the tool-chain to account for lattice reconstructions in an efficient kMC approach is the primary motivation for the present project and a solution is proposed in chapter 7.

PALLADIUM (100) AS A SHOWCASE FOR INCREASING COMPLEXITY Palladium as a catalyst for CO oxidation is more complex than other commonly studied catalysts such as RuO₂ because it is not as straightforward to single out one particular oxidation state that is stable and present under the full range of operating reaction conditions. This is illustrated by the fact that a metallic surface[26, 206], the $(\sqrt{5} \times \sqrt{5})R27^\circ$ surface oxide[10, 63, 71, 188] and even a bulk oxide surface[80] have been suggested as highly reactive surface terminations. The ab-initio thermodynamic phase diagram by Rogal *et al.*[162] established that a $(\sqrt{5} \times \sqrt{5})R27^\circ$ palladium surface oxide (henceforth abbreviated as $\sqrt{5}$ -oxide) is stable for stoichiometric conditions as also confirmed experimentally by van Rijn *et al.*[188]. However both experiment and theory also show that the range of relevant reaction conditions is not too far in phase space from a thermodynamically stable CO covered Pd(100) phase and a significant contribution from

metallic patches within the surface oxide or the interface between metallic surface and surface oxide cannot be ruled out by any of the studies so far.

There has been a strong debate about the most active phase in favor of either the metallic surface on the one hand[53, 54, 56, 126] or a surface oxide covered palladium on the other[70, 71, 74, 189]. In recent years significant progress in experimental techniques allows for a direct observation of the surface state under *in-situ* conditions. Resulting from this Hendriksen *et al.*[73] using *in-situ* Scanning Tunneling Microscope (STM) and Fernandes *et al.*[45] using High Resolution X-ray Photoelectron Spectroscopy (HRXPS) have suggested that surface steps and domain boundaries play a significant role in the conversion between oxidation states by CO. To specifically address the micro-kinetics of such dynamically changing surface terminations in a first-principles multi-scale framework new tools and methodologies are needed.

OUTLINE The remaining chapters are structured as follows: Chapter 4 serves both as an introduction to the kMC method and the efficient kMC approach via the simulation software package *kmos* which was developed and extensively used in the course of this thesis. In Chapter 5 we will study the challenge of the much sought-after transition states and discuss modifications to existing techniques. Chapter 6 uses the single-lattice kMC approach to study the kinetic stability regimes of the Pd(100) and the $(\sqrt{5} \times \sqrt{5})R27^\circ$ surface termination in a CO oxidizing environment. Chapter 7 introduces the multi-lattice kinetic Monte Carlo approach and an atomistic pathway for the transition from the palladium surface oxide to the metallic Pd(100) surface. Finally, Chapter 8 will summarize our findings and discuss the next steps necessary to also describe the reverse transition from the metallic phase to the surface oxide phase.

The main findings of this thesis are that the Pd(100) surface-oxide termination can coexist with a CO covered metallic termination under near-ambient stoichiometric feed conditions. This bistability goes beyond a softened phase boundary at higher temperatures but has a clear overlap that is due to the surface kinetics. The reduction of palladium surface oxide on Pd(100) in a pure CO environment is driven by elementary processes that occur on the surface oxide for temperatures close to room temperature. For higher temperatures it is driven by elementary processes between the metal and the surface oxide. These results suggest that a single surface or even single site picture does not suffice to describe the true nature of heterogeneous catalytic activity even for model catalysts. Further methodological developments and detailed theoretical as well as experimental studies are necessary to form a more complete view of the active surface under reaction conditions.

Part I

METHODOLOGICAL BACKGROUND

2.1 ELECTRONIC STRUCTURE THEORY

In this chapter we will provide an overview over the most important concepts necessary for the first-principles calculations in this treatise. The approach most suitable for the problem at hand are plane wave Density Functional Theory (DFT) calculations. Therefore we will summarize the underlying concepts of how energies and forces for a given arrangement of ionic cores in a periodic geometry are efficiently calculated.

THE SCHRÖDINGER EQUATION The smallest scale we will consider in this multi-scale treatment is the electronic scale. The physical system is described in terms of atomic nuclei and electrons whose stationary states are described by the time-independent non-relativistic Schrödinger equation[175]

$$\hat{H}\Psi = E\Psi \quad (2.1)$$

which allows in principle to predict all ground-state properties of any system under study. In the case of a system consisting of N_N nuclei and N_e electrons in the absence of an external field we can describe with the following Hamiltonian in spatial representation

$$\begin{aligned} \hat{H} &= \hat{T}_N + \hat{T}_e + \hat{V}_{NN} + \hat{V}_{ee} + \hat{V}_{eN} \\ \hat{T}_N &= \sum_{i=1}^{N_N} \frac{-\hbar^2 \nabla_i^2}{2M_i} \\ \hat{T}_e &= \sum_{j=1}^{N_e} \frac{-\hbar^2 \nabla_j^2}{2m_e} \\ \hat{V}_{NN} &= \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^{N_N} \frac{Z_i Z_j}{|\mathbf{R}_i - \mathbf{R}_j|} \\ \hat{V}_{ee} &= \frac{1}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^{N_e} \frac{e^2}{|\mathbf{r}_i - \mathbf{r}_j|} \\ \hat{V}_{eN} &= - \sum_{i=1}^{N_N} \sum_{j=1}^{N_e} \frac{Z_i e}{|\mathbf{r}_j - \mathbf{R}_i|} \end{aligned} \quad (2.2)$$

which accounts for the nuclei kinetic energy \hat{T}_N , the electron kinetic energy \hat{T}_e , the nuclei-nuclei electrostatic interaction \hat{V}_{NN} , the electron-

electron electrostatic interaction \hat{V}_{ee} , and the electron-nuclei electrostatic interaction \hat{V}_{eN} . The other symbols used are the charge of the i th nucleus Z_i , the mass of the i th nucleus, the location of the i th nucleus \mathbf{R}_i , the location of the i th electron \mathbf{r}_i , and the electron mass m_e . We are using SI units throughout this chapter. Though formally exact in practice it cannot be solved except for extremely small systems. The main problem is the fact that the full many-body wave function as described by the Schrödinger equation is defined on a $3(N_e + N_N)$ -dimensional domain. Thus we will recapitulate in the following the historic development of abstractions and approximations to the representation of the Schrödinger equation. As a result one is able to calculate the resulting electronic wave function and resulting properties numerically to a sufficiently high precision. For a more detailed introduction to quantum mechanics the reader is referred to e.g. Sakurai[167].

ADIABATIC APPROXIMATION A first crucial step towards the ability to solve the Schrödinger equation was suggested by Born and Oppenheimer[11]. The key idea is motivated by the fact that electrons are approximately 1800 times lighter than the lightest nucleus and therefore the typical time scale associated with the movement of electrons should correspondingly be much shorter for electrons than for nuclei. Therefore it is assumed that the total wave function can be factorized as

$$\Psi(\mathbf{R}, \mathbf{r}, t) = \sum_i \Phi_i(\mathbf{R}, \mathbf{r}) \eta_i(\mathbf{R}, t) \quad (2.3)$$

into an electronic wave function where the nuclear wave function $\eta_i(\mathbf{R}, t)$ depends on the nuclear coordinates \mathbf{R} and on the time t , whereas the electronic wave function Φ_i only depends on the electronic degrees of freedom and on the nuclear coordinates \mathbf{R} . From this *Ansatz* one can derive an effective Schrödinger equation for the electronic wave function neglecting coupling between the electronic and the nuclear wave function

$$\hat{H}_e \Phi_n(\mathbf{R}, \mathbf{r}) = \epsilon_n \Phi_n(\mathbf{R}, \mathbf{r}) \quad (2.4)$$

where

$$H_e = \hat{T}_e + V_{ee} + V_{eN}. \quad (2.5)$$

The other terms \hat{T}_N and \hat{V}_{NN} have been separated out and the resulting equation now only depends parametrically on those. This electronic Schrödinger equation can now be solved independently for each configuration of nuclei.

CLASSICAL NUCLEI APPROXIMATION Justified by the fact that the spatial distribution of nuclear wave packets at room temperature

is confined to a length that is much smaller than the typical sizes of atoms, the nuclei can be described as classical point charges at the center of their respective wave packet. As a result we have separated from a full quantum mechanical treatment the kinetic energy of the atomic nuclei and the electrostatic interaction between nuclei becomes merely an additive parameter of the system's total energy.

There are two ways to view this result: on the one hand the position of the ionic nuclei are clamped, so that they represent a constant external potential that influences the electronic wave function. On the other hand the motion of the ionic cores can be viewed to evolve on a so-called potential energy surface (PES) manifested by the surrounding electrons which are always in their ground state. The forces acting on the nuclei due to the PES can be calculated as the expectation value of the derivative of the Hamiltonian.

However the electronic wave function still depends on the $3N_e$ spatial electronic coordinates (and possibly spin degrees of freedom) and therefore exponential scaling of the size of the equation quickly renders any attempt to solve it in the general case unfeasible.

HELLMANN-FEYNMAN FORCES As noted independently by Güttinger[64], Hellmann[29], and Feynman[46] a derivative with respect to an external parameter λ on which the wave function does not directly depend can be calculated as the expectation value of the derivative of the Hamiltonian operator

$$\frac{\partial \epsilon_i}{\partial \lambda} = \left\langle \Phi_i \left| \frac{\partial \hat{H}_e}{\partial \lambda} \right| \Phi_i \right\rangle + \left\langle \frac{\partial \Phi_i}{\partial \lambda} \left| \hat{H}_e \right| \Phi_i \right\rangle + \left\langle \Phi_i \left| \hat{H}_e \right| \frac{\partial \Phi_i}{\partial \lambda} \right\rangle \quad (2.6)$$

$$= \left\langle \Phi_i \left| \frac{\partial \hat{H}_e}{\partial \lambda} \right| \Phi_i \right\rangle + \epsilon_n \frac{\partial}{\partial \lambda} \langle \Phi_i | \Phi_i \rangle \quad (2.7)$$

The result is valid for any hermitian operator and can be used to efficiently calculate forces using $\lambda = R_I$

$$\frac{dE_n(\mathbf{R}_I)}{d\mathbf{R}} = \left\langle \Phi_n(\mathbf{R}) \left| \frac{\partial \hat{H}_e}{\partial \mathbf{R}_i} \right| \Phi_n(\mathbf{R}) \right\rangle. \quad (2.8)$$

However in practice when the Hellmann-Feynman theorem is applied to a wave function expanded into a basis set (see below) special corrections need to be made if the basis set is not spatially independent of the nuclear coordinates.[151]

DENSITY FUNCTIONAL THEORY Already one year after the original publication of the Schrödinger equation Thomas[183] and Fermi[44] put forward a quite different approach to describing the electronic structure of condensed matter. In their model the electronic density

$$\rho(\mathbf{r}_1) = N_e \int d\mathbf{r}_2 \dots d\mathbf{r}_{N_e} |\Phi_n(\mathbf{R}, \mathbf{r})|^2 \quad (2.9)$$

is the main protagonist and the total energy is to be calculated as[33]

$$E_{\text{TFD}}[\rho] = C_k \int d\mathbf{r} \rho(\mathbf{r})^{5/3} + \int d\mathbf{r} \rho(\mathbf{r}) v_{\text{ext}}(\mathbf{r}) + \frac{1}{2} \iint d\mathbf{r} d\mathbf{r}' \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \quad (2.10)$$

$$- C_X \int d\mathbf{r} \rho^{4/3}(\mathbf{r}) + E_C[\rho]. \quad (2.11)$$

Here C_k and C_X are suitably chosen constants and E_C is another functional describing the correlation between electrons. This approach has the crucial advantage that instead of dealing with the $3N_e$ -dimensional function $\Phi_i(\mathbf{r})$ one would only deal with a 3-dimensional object $\rho(\mathbf{r}_1)$. Therefore this approach received considerable attention[196]. Though mostly due to the insufficient prediction of the kinetic energy[176] the electronic density picture was never used to describe chemical interactions.

HOHENBERG-KOHN THEOREM A new era of development was ushered in by the seminal paper of Hohenberg and Kohn[87] that proved two things. On the one hand they showed that the ground-state electronic density unequivocally determines the external potential up to an additive constant. In effect it means that the total electronic energy can be calculated in principle using a functional

$$E_V[\rho] = F[\rho] + \int d\mathbf{r} \rho(\mathbf{r})V(\mathbf{r}) \quad (2.12)$$

where

$$F[\rho] = \langle \Phi[\rho] | \hat{T}_e + V_{ee} | \Phi[\rho] \rangle. \quad (2.13)$$

A corollary of this theorem is that the ground-state density unequivocally determines any observable of the quantum mechanical system. This is easy to see since the ground-state density maps unequivocally to an external potential which also determines the full electronic wave function.

On the other hand they showed that the density functional is minimized by the ground state density.

KOHN-SHAM APPROACH Even though impressive for its clarity and generality the Hohenberg-Kohn (HK) theorems give essentially no advice of what the functional mapping from a ground state density to its potential and ultimately to the total energy looks like. The by far most successful approach exploiting the HK-theorems was suggested by Kohn and Sham[103]. They proposed to describe the system in terms of non-interacting single-particle wave functions ϕ_i^{KS} .

If a potential can be found which accounts for all the exchange- and correlation effects of the fully interacting system in terms of a single-particle wave function then the HK theorem holds equally for

the non-interacting system and the density calculated from the single-particle wave functions

$$\rho(\mathbf{r}) = \sum_i |\phi_i^{\text{KS}}(\mathbf{r})|^2 \quad (2.14)$$

and its energy functional will minimize for the same ground state-density as the interacting system. This system is much more accessible because it subdivides the full electronic problem into N_e coupled equations which each only depend on the coordinates of one orbital. This approach has the main advantage that the kinetic energy contributions calculated from such non-interacting orbitals

$$T_S = \frac{-\hbar^2}{2m_e} \langle \phi_i^{\text{KS}} | \nabla^2 | \phi_i^{\text{KS}} \rangle \quad (2.15)$$

matches the exact kinetic energy much better than the earlier density-based approximations and allows to write the full Kohn-Sham energy functional in the form

$$E_{\text{KS}} = T_S[\rho] + \int d\mathbf{r} \rho(\mathbf{r})V(\mathbf{r}) + \frac{1}{2} \iint \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} + E_{\text{xc}}[\rho]. \quad (2.16)$$

By taking the functional derivative of eq. 2.16 with respect to all single-particle wave functions one arrives at a set of eigenvalue equations which are referred to as Kohn-Sham equations

$$\left(-\frac{\hbar^2}{2m_e} \nabla^2 + V_{\text{KS}} \right) |\phi_i^{\text{KS}}(\mathbf{r})\rangle = \epsilon_i |\phi_i^{\text{KS}}(\mathbf{r})\rangle$$

where $V_{\text{KS}} = v_{\text{ext}}(\mathbf{r}) + \int \frac{\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}' + V_{\text{XC}}[\rho]$

(2.17)

has to be satisfied for every single-particle orbital ϕ_i^{KS} which is the main objective of DFT solvers.

EXCHANGE-CORRELATION FUNCTIONALS The Kohn-Sham approach is quite successful at describing the largest energy contributions accurately, *i.e.* the electrostatic electron-electron interaction, the electron kinetic energy, the electrostatic nucleus-electron interaction. However it also left the most complex terms such as the exchange and correlation energy in an extra term, the so called exchange-correlation functional $V_{\text{xc}}[\rho]$. Ever since the arrival of the Kohn-Sham equations have large amounts of time been spent to find efficient and accurate expressions for this V_{xc} . Two approaches have emerged here to arrive at approximations: One is to start with a sufficiently general expansion into different functional forms and fit its coefficients to known results ("result driven") while the other is to use various principles as guidelines that the exact xc-functional has to fulfill to obtain ("principle driven") approximate expressions. Since the inception of the KS-equations various functionals have been proposed. The Jacob's ladder

scheme proposed by Perdew *et al.*[143, 145] has recently become very popular as a tool to group various functionals by complexity.

The lowest rung (local density approximation) only considers the density at each volume element such that the functional can be explicitly written as

$$E_{xc}^{LDA}[\rho] = \int d\mathbf{r} \rho(\mathbf{r}) F_{xc}^{LDA}[\rho(\mathbf{r})]. \quad (2.18)$$

On the second rung (generalized gradient approximation) are all functionals that can be explicitly written in the form

$$E_{xc}^{GGA}[\rho] = \int d\mathbf{r} \rho(\mathbf{r}) F_{xc}^{GGA}[\rho(\mathbf{r}), \nabla\rho(\mathbf{r})]. \quad (2.19)$$

Functionals belonging to the third rung contain second derivatives or the Kohn-Sham orbital kinetic energy density and this rung is called the meta-GGA family. Each functional within this set can be written as

$$E_{xc}^{mGGA}[\rho] = \int d\mathbf{r} \rho(\mathbf{r}) F_{xc}^{mGGA}[\rho(\mathbf{r}), \nabla\rho(\mathbf{r}), \nabla^2\rho(\mathbf{r}), \tau]. \quad (2.20)$$

where

$$\tau = \frac{1}{2} \sum_i^{\text{occ.}} |\nabla\phi_i^{\text{KS}}|^2 \quad (2.21)$$

describes the kinetic energy density of all occupied orbitals. The fourth rung goes a crucial step further by including the exact electron exchange of occupied orbitals. The fifth rung finally uses all available information available within Kohn-Sham theory by also including non-occupied orbitals. It should be noted that by virtue of the Hohenberg-Kohn theorem functionals on the third to fifth rung are still density functionals since the Kohn-Sham orbitals itself are functionals of the density. However with each rung the computational effort to solve the corresponding Kohn-Sham equation also increases significantly.

For the purpose of predicting material properties as well as chemical interactions between atoms already the GGA level in particular as manifested in the form put forward by Perdew, Burke, and Ernzerhof[144] has been quite successful and is the functional used throughout this thesis. The development of new exchange-correlation functionals is ongoing research aiming towards higher accuracy and reliability.

2.2 COMPUTING THE ELECTRONIC GROUND-STATE

SELF-CONSISTENT FIELD PROCEDURE The astute reader will have noticed that eq. 2.17 contains the density which it aims to calculate inside the operator that is acting on the density. Thus this equation

cannot be solved simply by calculating the eigenstates and ϵ -values of the operator. Instead the equation needs to be solved self-consistently. An iterative approach can be used to achieve this. Starting from a suitable initial guess for the density ρ_{in} the diagonalization of the KS equation determines a new density ρ_{out} . This resulting density has to be used as the new initial density $\rho_{\text{in}} := \rho_{\text{out}}$, from which a new KS operator is constructed and the corresponding KS equation has to be solved. ρ_{out} is generally different from ρ_{in} . Thus the cycle has to be repeated until $\rho_{\text{out}} \approx \rho_{\text{in}}$ up to a predetermined tolerance metric. The procedure resulting from this is accordingly coined Self-Consistent Field (SCF) procedure.

INTRODUCTION OF A BASIS So far we were only concerned with solving the Kohn-Sham equations in a formal sense. Little thought has been spent on how to translate it into a form that a computer can be instructed to solve. The first thing to realize is that there needs to be a general and flexible way to represent any wave function in a computer's memory. To this end the most common method today is to first linearly expand the full wave function into a sum of known functions

$$\phi_i = \sum_{\alpha} c_{i\alpha} \varphi_{\alpha} \quad (2.22)$$

where ϕ_i is the wave function one aims to represent, each φ_{α} is a fixed known function and $c_{i\alpha}$ are real coefficients. Expanding the wave function into a basis set effectively turns the differential equation into an algebraic equation for which very efficient numerical methods exist. In principle any complete set of functions is sufficient to reproduce any piecewise continuous function to any desired precision. In practice of course ideally a function system is chosen which is able to reproduce the wave functions that are expected to occur in the solution with great accuracy using as few basis vectors as possible.

BLOCH-THEOREM AND BRILLOUIN ZONE SAMPLING Being interested mostly in periodic systems (surfaces can be considered a truncated crystal lattice) a very practical simplification arises due to an observation by Bloch. This assumption is motivated by the structure of many real-world solid systems, where the arrangement of atoms can be described in terms of 3 lattice vectors k_x, k_y, k_z and a finite arrangement of atoms which is called basis in this context.[2] Felix Bloch showed[8] that the wave function on a periodic symmetry can always be expressed as

$$\psi_{j\mathbf{k}} = \exp(i\mathbf{k} \cdot \mathbf{x}) u_{j\mathbf{k}}(\mathbf{x}). \quad (2.23)$$

That is the wave function factorizes into a global phase factor and a periodic part that does not depend on the unit cell. This reduces the

effort significantly as the periodic function $u_{j\mathbf{k}}(\mathbf{x})$ only depends parametrically on the point in the unit cell which can be solved separately for each \mathbf{k} -point almost independently. Different \mathbf{k} -point calculations only couple via the real space density. Thus in the SCF procedure each calculation for each \mathbf{k} -point can be executed separately and only after each SCF step does the current (non self-consistent) density have to be synchronized. A systematic scheme for generating \mathbf{k} -points in the Brillouin zone has been proposed by Monkhorst and Pack[132]. The question remains how many \mathbf{k} -points have to be considered for a faithful representation of the full 3-dimensional band structure. This cannot be answered in general but has to be tested for every system by systematically increasing the grid-point density until quantities of interest do not change anymore.

PLANE WAVES AS BASIS VECTORS Being primarily concerned with a description of solid-state surfaces we will introduce and justify our choice of plane-waves (PWs) as basis-vectors. Many catalytically relevant materials such as transition metals condense in a crystalline structure that can be described by Bravais lattices[2]. If we therefore assume that this lattice symmetry will occur in most systems under study it is evident to choose the basis set most suited for such a geometry. Plane waves in the form (cf. eq. 2.22)

$$\varphi_{\mathbf{G}}^{(j\mathbf{k})} = \sum_{\mathbf{G}'} \frac{c_{\mathbf{G}'}^{(j\mathbf{k})}}{\sqrt{\Omega}} \exp(i(\mathbf{G} + \mathbf{k}) \cdot \mathbf{r}) \quad (2.24)$$

have the main advantage that if one transforms the Kohn-Sham equation into Fourier space

$$\sum_{\mathbf{G}'} \left[\frac{\hbar^2}{2m_e} |\mathbf{k} + \mathbf{G}|^2 \delta_{\mathbf{G},\mathbf{G}'} + V_{\text{KS}}[\rho](\mathbf{G} - \mathbf{G}') \right] c_{\mathbf{G}'}^{(j\mathbf{k})} = \epsilon^{(j\mathbf{k})} c_{\mathbf{G}}^{(j\mathbf{k})} \quad (2.25)$$

the kinetic energy term becomes diagonal and therefore computationally very efficient to evaluate. The potential term is diagonal in real-space (at least before the introduction of pseudo-potentials) and fast conversion between these two representations is facilitated by the use of fast Fourier transformations[49, 150]. The precision of the plane wave basis can be easily controlled by imposing a cut-off energy E_{cut} :

$$|\mathbf{G}| < \sqrt{\frac{2m_e E_{\text{cut}}}{\hbar^2}} \quad (2.26)$$

has set only a single parameter that determines the size or resolution of it which makes it straightforward to converge in practical calculations.

PSEUDOPOTENTIALS The main limitation of plane wave basis sets is that in order to resolve the quickly oscillating features of a typical single-particle orbital near an atomic nucleus very high frequency

components are required. High-frequency components translate into large matrices that have to be diagonalized and Fourier transformed. In order to amend this limitation one essentially divides the electronic states involved into those directly involved in chemical bonding and those that are not: valence electronic states and core electronic states. This approach was originally proposed by Kleinman and Philips[102]. They found that it is advantageous to construct modified valence states ϕ'_v from core states ϕ_c and valence states ϕ_v in the following way

$$|\phi'_v\rangle = |\phi_v\rangle + \sum_c \langle \phi_c | \phi'_v \rangle |\phi_c\rangle. \quad (2.27)$$

If one then defines a new potential $V_{PS} := \sum_c (E_c - E) \langle \phi | \phi'_v \rangle \langle \phi'_v | \phi_c \rangle / |\phi_v\rangle$, the valence states observe a modified Schrödinger equation

$$(\hat{H} + V_{PS})|\phi_v\rangle = E|\phi_v\rangle. \quad (2.28)$$

In this way only those parts of the electronic system are calculated that change as a response to the chemical environment. The electronic core part which usually requires the largest plane-wave basis set to resolve can be calculated once and reused for a number of nuclear configurations.

Being able to calculate total energies and forces is only one part of predicting relevant properties of materials for heterogeneous catalysis. Another important step is to link electronic structure features to geometries and reaction rates observable in experiment under specific conditions. This is the objective of Chapter 3. For more low-level details of the here deployed CASTEP DFT package[28] the reader is referred to Appendix A.

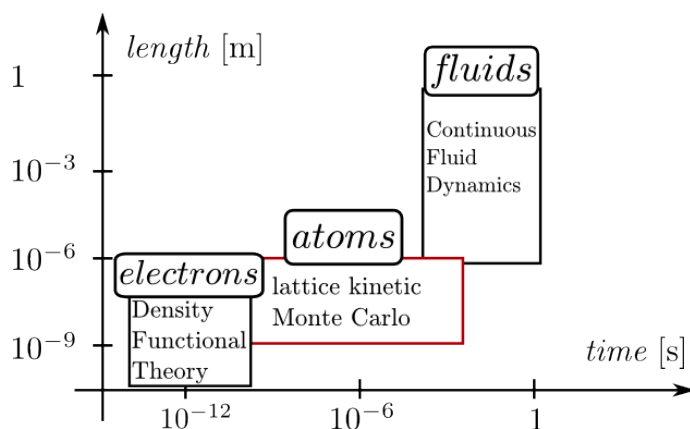


Figure 3.1: Overview over the different regimes in length and time scale that have to be considered to rigorously describe the functionality of heterogeneous catalysts based on the main object of calculation. The scope of this thesis is to make complex morphological transition such as the reduction of a surface oxide accessible to the atomistic description via lattice kinetic Monte Carlo.

Being equipped with methods to calculate total energies and forces of atomic arrangements we can in principle obtain all information to predict the functioning of a catalyst. However, to make timely predictions which can be directly related to experiment under specific conditions additional steps need to be made. Fig. 3.1 shows the distinctly different regimes based on the main physical object of interest and the methods deemed suitable for each. This thesis focuses on the first two stages (electronic and atomistic level) by enhancing the capabilities of the atomistic description.

The quest for predicting catalytic activity on an atomistic level for a given material can be broken down into two questions: What surface termination is present for given conditions? And how can chemical reactions take place on this termination? A very approximate answer to the first question can be obtained within the so-called *ab initio* thermodynamics approach, while the answer to the second question requires full-blown (micro)kinetic modeling. This chapter will introduce the fundamental concepts underlying the thermodynamic predictions while Chapter 4 explains lattice kinetic Monte Carlo as the deployed approach for making kinetic predictions.




Figure 3.2: The fcc(100) surface illustrating the primitive (1×1) unit cell together with a (2×2) , a $(\sqrt{2} \times \sqrt{2})$, and a $(\sqrt{5} \times \sqrt{5})$ unit cell.

3.1 DEFINING THE STRUCTURAL MODEL

First, we will focus on defining a suitable structural model for heterogeneous catalysis. The defining characteristic of heterogeneous catalysis is that the catalytic activity occurs at the interface between two aggregate states such as a solid-liquid interface or a solid-gas interface. For the CO oxidation reaction studied here we focus on solid-gas interfaces. On an atomic scale various features such as terraces, steps, defects, or domain boundaries exist even at a single-crystal model catalyst surface. However, in this thesis we will take a first step and concentrate on the actual terraces as an important structural primitive that constitutes a majority of the overall surface area. These single-crystal surfaces consist largely of low Miller index [131] single crystal surfaces which is the central structural paradigm for this thesis. Fig. 3.2 shows the Pd(100) surface as an example illustrating the so-called primitive (1×1) unit cell next to the (2×2) , the $(\sqrt{2} \times \sqrt{2})$, and the $(\sqrt{5} \times \sqrt{5})$ unit cell. The $(\sqrt{5} \times \sqrt{5})$ unit cell plays an important role in this thesis since it is the smallest unit cell in which the PdO(101) surface oxide can be placed in a commensurate way.

To represent such a surface in a DFT calculation it is crucial to use a geometrical model which is as small as possible to due the high computational cost of DFT, yet, contains the essential features of the physical system. A common approach is to use a supercell geometry which consists of several layers of bulk unit cells divided by an empty space to represent the vacuum above the surface. One example of such a supercell is depicted Fig. 3.3 which is coined a *slab*. We will generally use a Cartesian coordinate system in which the z-axis

is perpendicular to the surface, while the x - and y -axis are parallel to the surface. In the physical system the influence of the reduced symmetry at the surface will attenuate towards the inner parts of the bulk crystal. To describe this the atoms in central layers removed from the modeled surface are constrained to their bulk positions while the layers closer to the surface are allowed to move to minimize the total energy of the system. One either relaxes surface layers on both sides of the surface slab while keeping layer(s) in the center fixed (symmetric) or only relaxes the atoms on one side while keeping atoms on the other side fixed (asymmetric). In this thesis asymmetric slabs are used throughout to minimize the computational cost. In all cases numerical tests have to be performed to ensure that the total number of layers in the slab as well as the number of fixed layers does not introduce deviations to the evaluated quantities compared to the infinite geometry that is larger than a given threshold.

The commonly deployed plane-wave basis set enforces periodic boundary conditions in all 3 direction for the supercell. This is typically desirable for the x - and y -direction since it allows to represent a macroscopic system with only a few atoms. However in the z -direction it may cause artificial effects that do not occur in the system under study. Therefore the gap between slabs (or vacuum thickness) needs to be sufficiently large to prevent interaction between the periodic images of wave functions. Detailed tests for the present system with respect to these model parameters are presented in Appendix B.

3.2 REACTANTS ON SURFACES

The next object of interest towards heterogeneous catalysis is the interaction of reactants with surfaces. On crystalline surfaces molecules typically adsorb centered on high-symmetry sites. This can be easily rationalized by the fact that the potential energy surface at locations away from high-symmetry sites typically features some gradient parallel to the surface. Fig. 3.4 shows the high-symmetry surface sites of the fcc(100) surface. Optimizing the slab with a molecule m at such a high-symmetry site allows to calculate the binding energy

$$E_{m,i}^{\text{bind}} = E_{m@i}^{\text{slab}} - (E^{\text{slab}} + E_m^{\text{gas}}) \quad (3.1)$$

where $E_{m@i}^{\text{slab}}$ is the total energy of the slab with the adsorbed molecule, E^{slab} is the total energy of only the slab, and E_m^{gas} is the energy of the molecule in gas phase. To calculate the gas phase energy of a molecule using periodic plane-wave calculations it has to be placed in an otherwise empty simulation cell that is large enough to prevent interaction between periodic images.

The evaluated binding energy can already give crucial insight if a molecule adsorbs at all on a given surface and what is the relative ordering of the most stable adsorption sites.




Figure 3.3: The slab model used in this thesis. The supercell is indicated by black lines. For illustrative purposes the central unit cell is surrounded by one periodic copy along each crystal axis. Palladium atoms belonging to the metallic bulk lattice are drawn blue, palladium atoms belonging to the surface oxide are drawn green, oxygen atoms are drawn red and carbon atoms are drawn in dark grey.




Figure 3.4: Top view of one (1×1) unit-cell of a fcc(100) surface with the three high-symmetry sites hollow, bridge, and top. These sites are good candidate locations for stable adsorption sites.

The choice of the surface unit cell is also important here. If one constructs a surface slab with a primitive (1×1) surface unit cell and places a reactant on it, this represents the adsorption situation of one adsorbate per surface atom. We will refer to this as one mono-layer coverage (1 ML), the number of adsorbates per surface atoms. If one places one molecule on a (2×2) this corresponds to one adsorbate per four surface atoms or 0.25 ML. A ($\sqrt{2} \times \sqrt{2}$) unit cell can be used to model a 0.5 ML coverage. In general evaluating low coverage binding energies is computationally more expensive than high density coverages due to the increasingly large unit cell size.

Furthermore combinations of different adsorbates can be used to calculate the combined binding energy or lateral interaction between coadsorbates. The magnitude of lateral interaction can be systematically calculated by deploying an expansion of adsorbates grouped in clusters starting from singlets to doublets, triplets and so on.[174, 204] This route will however not be deployed in this thesis since the main focus lays in the influence of a morphological transition of the surface itself.

3.3 SURFACES AT REACTION CONDITIONS

In order to relate and compare the stability of different surfaces under specific temperatures and partial pressures the thermodynamics of the combined bulk, surface, and gas phase system need to be considered. In a purely thermodynamic treatment one first aims to establish the thermodynamically stable surface termination. For a fixed temperature and fixed partial pressures the Gibbs free energy $G(T, \{p_i\})$ is the pertinent thermodynamic potential. Being exposed to both the surface needs to be in equilibrium with the underlying bulk crystal and the gas phase above.

The surface with the lowest Gibbs free surface energy

$$\gamma(T, \{p_i\}) = \frac{1}{A} \left(G(T, p_i) - \sum_i N_i \mu_i(T, p_i) \right) \quad (3.2)$$

is stable[156]. Here G is the Gibbs free energy of the solid including the surface, μ_i the gas phase chemical potential of the i th species, N_i the number of adsorbates of species i , and A the surface area of the surface under study.

The bulk Gibbs free energy can be decomposed further into individual contributions

$$G = E_{\text{total}} + F_{\text{vib}} + F_{\text{config}} + pV. \quad (3.3)$$

Here E_{total} is the internal energy, F_{vib} the vibrational free energy, F_{conf} the configurational free energy, p the total pressure, and V the bulk volume.[161] The total energy contribution is typically the largest contribution in solid state surfaces and differences between total energies

can be calculated quite well using DFT. The gas phase chemical potential can be calculated in a similar fashion from first-principles by calculating the gas phase total energy, translational, rotational, vibrational, and configurational free energy.

By evaluating the Gibbs free surface energy over a range of conditions for a set of candidate surfaces one can readily generate a so-called *ab initio* phase diagram, *i.e.* a diagram depicting the stability range in temperature and partial pressures of various surface phases. We will see an application of this in Chapter 6.

KMOS: A LATTICE KINETIC MONTE CARLO FRAMEWORK

This chapter gives an in-depth introduction to the lattice kinetic Monte Carlo method and focuses specifically on efficiency aspects of lattice kMC implementations. The accompanying software framework `kmos` developed during this thesis is available as open-source[81] and the work of this chapter has been published in *Computer Physics Communications*[85].

4.1 INTRODUCTION

The pressing demands for ever more energy- and resource-efficient processing reinforce the long-standing quest towards a detailed mechanistic understanding of heterogeneous catalysis. At best down to the atomic level, such understanding would pave the way for a rational design of improved catalysts, which ultimately will be tailored to the nanoscale. Quantitative theory increasingly contributes to this quest with refined kinetic models that meanwhile allow to accurately predict the activity of model catalysts of increasing complexity (from single-crystal surfaces up to nanoparticles at planar supports) without any recourse to experimental data. [60, 68, 69, 89, 93, 158, 166, 190] Such models have to span a range of scales in length and time, starting with the making and breaking of the individual chemical bonds at the electronic structure level, over the mesoscopic interplay of the various elementary reactions in the reaction network, to the heat and mass transport at the macroscopic (reactor) scale. [90, 97, 134, 137, 165, 168, 172]

To achieve this, state-of-the-art multi-scale models resort to a hierarchical combination of different methodology. The current framework for the mesoscopic level are microkinetic approaches evaluating a (Markovian) master equation (*vide infra*). [27, 58, 165] Using as input kinetic parameters for all elementary reactions (e.g. provided from first-principles electronic structure theory calculations as described in the preceding chapters), such microkinetic models determine for given operation conditions at the surface (e.g. temperature T and partial pressures $\{p_i\}$ of all reactants i) not only the catalytic activity (typically measured as turn-over frequency, TOF, in units of products per active site and time) but also other important information such as surface composition, the occurrence of individual reaction steps in the network, or in particular the presence of a dominant reaction mechanism as well as rate-determining steps therein.[128, 182] Aver-

aged over a sufficiently large catalyst surface area the TOF output can then for example be used as input for macroscale simulations of heat and mass transport in a given reactor geometry.[101, 119, 121–123, 127, 173, 192]

The traditional and still prevalent microkinetic approach employs a mean-field approximation to solve the master equation, and then only accounts for average surface coverages of the different reaction intermediates at the active surface. In case of heterogeneous arrangement of active sites, strong lateral interactions among the adsorbed species, or diffusion limitations, this approximation is known to break down and lead to qualitatively wrong results [124, 182, 198]. This has contributed to the recent rise of alternative kinetic Monte Carlo (kMC) simulations, which do not need to rely on the mean-field approximation and therefore provide a faithful account of the detailed spatial distributions of species at the catalyst surface, as well as their correlations and fluctuations.[25, 154, 193] In contrast to effective rate equation based models for which a definition of some abstract active site (type) is often sufficient, kMC thus needs as input detailed information about the microscopic arrangement of the true active sites of the crystal surface. In return, it then provides comprehensive information about the (time-resolved) arrangement of chemicals at all these active sites during catalyst operation. Apart from a wealth of mechanistic information, e.g. about correlations in the occupation of neighboring sites at the surface, this allows to obtain proper (not erroneous mean-field) mesoscopic averages of quantities like TOFs that ultimately are required for reactor level modeling.

Due to the inherent methodological simplicity of kMC, seminal works in the surface catalysis context typically relied on specialized code written from scratch.[68, 69, 157, 158, 163, 164] Even though kMC models are used in the field with increasing frequency this practice has largely prevailed and only few general kMC packages have been put forward to date.[34, 57, 106, 179] This stands in stark contrast to the manifold of established and powerful software packages employed in the multi-scale framework for either the underlying electronic structure calculations [197] or the continuum mechanics reactor scale simulations[20]. Noting this as an obstacle to a further, wide-spread use of the kMC approach to surface catalysis has been the motivation for the here presented *kmos* package, which as its core objective aims at a most user-friendly and efficient implementation, execution, and evaluation of increasingly complex lattice kMC models in the surface catalysis context.




Figure 4.1: (Color online) Schematic representation of a typical model catalyst surface, showing the top layer atoms of a metal(100) facet as large spheres. The periodic surface consists of repeating unit cells, each containing one or more active sites (here indicated by circles and squares for hollow and bridge adsorption sites, respectively). Possible elementary reactions in the context of CO oxidation (dissociative adsorption, diffusion, reaction and desorption) are indicated by yellow arrows.

4.2 THEORETICAL BACKGROUND

4.2.1 (Surface) Chemistry on a Lattice

In terms of microkinetic modeling, the atomistic evolution proceeding during surface catalytic reactions is quite representative for a wider class of problems including crystal growth, initial corrosion, diffusion in crystalline (battery) materials or surface self-assembly. These problems feature a range of common characteristics, which motivate a so-called lattice approach to kMC that also underlies the `kmos` package. In the following we use a survey over these characteristics to briefly introduce this lattice approach to kMC and clearly define terminology henceforth employed. For a more detailed account of general kMC methodology we refer to existing reviews.[25, 154, 193] Even though the following introduction is done within the surface catalysis context, the generalization to the other problems mentioned is self-evident.

SITE-SPECIFIC ADSORPTION AND LATTICE MAPPING The first defining characteristic is that the surface adsorption of all reactants and reaction intermediates is assumed to be site-specific, *i.e.* it always occurs in well-defined so-called active *sites* offered by the crystalline surface. Due to the periodicity of the latter this generally leads to a lattice with each lattice point representing one such site.[2] The actual kMC simulations only consider this lattice, which allows to encompass a wide range of system geometries within this framework. Most straightforward are extended low-index single-crystal surfaces,

where the lattice is simply composed of multiple identical surface unit-cells and then continued through the use of periodic boundary conditions. Figure 4.1 illustrates[152] this for a fcc(100) model catalyst surface exhibiting two types of active sites. More complex geometries like entire nanoparticles are accessed through the thoughtful use of non-primitive unit-cells and/or pseudo reaction intermediates (e.g. declaring different active sites, effective species and kinetic parameters for every facet).

At each site an integer *occupation* value represents one of several possible *reaction intermediates* binding to this site (e.g. 1 for adsorbed O, 2 for adsorbed CO in the prominent CO oxidation context), including a reaction intermediate *empty* (e.g. occupation value 0) and also including the possibility that a (larger) reaction intermediate extends over more than one site (in the lattice context simply achieved by additional constraints linking the occupation of neighboring sites). One specific set of occupation values on the entire lattice is called a *configuration* (denoted by small Latin letters u, v, \dots), and a transition from one configuration to another proceeds through the occurrence of an *event* (denoted by small Greek letters α, β, \dots). An event thus changes the occupation of one or more sites.

RARE-EVENT DYNAMICS AND MARKOVIAN MASTER EQUATION

The second defining characteristic is that the time evolution is characterized by a so-called rare-event dynamics[194]. Due to activation barriers well exceeding thermal energies, the reaction intermediates reside most of the time in their adsorption (lattice) sites, and the events in form of the actual *elementary reactions* (adsorption, diffusion, reaction, desorption) happen comparably fast in between. Exploiting this separation of time scales, prevalent microkinetic theory [13, 58] generally assumes that any such event occurs independent of all preceding ones, *i.e.* it applies a *Markov* approximation. The time evolution of the system (in this case the transitions from configuration to configuration through the consecutive occurrence of events) is then described by a Markovian master equation

$$\dot{\rho}_u(t) = \sum_v (w_{uv}\rho_v(t) - w_{vu}\rho_u(t)) \quad , \quad (4.1)$$


where $\rho_u(t)$ is the probability for the system to be in configuration u at time t , and w_{vu} is the transition rate (in units of time^{-1}) at which configuration u changes to configuration v .

LOCALITY OF ELEMENTARY REACTIONS The third defining characteristic of the systems mentioned initially is that changes in configuration due to an event are typically geometrically narrowly confined to as few as ~ 1 -10 sites. Due to this locality it is possible and convenient to uniquely define any elementary reaction a in terms of the local *educt* E_a lattice configuration before and the local *product* P_a

lattice configuration after the event, as well as the concomitant rate constant k_a ,



Obviously, these local lattice configurations have to extend at least over all sites that actually change occupation due to the occurring elementary step. For a simple uni-molecular CO adsorption step the local lattice configuration must e.g. contain the very site involved that changes its occupation from 0 (empty) to 2 (CO). For a dissociative adsorption step of O_2 the minimum local lattice configuration must in turn extend over the two neighboring sites that change their occupation from 0 (empty) to 1 (O), and for more complex reactions involving reaction intermediates covering multiple sites the minimum local lattice configurations span even larger lattice areas. In cases the local lattice configurations may need to include further nearby lattice sites, which do not change their actual occupation from educt to product configuration, but occupation value of which is a necessary information to determine the elementary reaction. This is prominently the case in the presence of lateral interactions. In order to properly capture such interactions the local lattice configuration needs to include all lattice sites within the interaction radius to uniquely define the local adsorbate environment. Imagine for the case of the aforementioned unimolecular CO adsorption that this depends on whether or not a site neighboring the actual adsorption site is also occupied with CO. In this situation the local educt and product lattice configuration need to include the actual adsorption site i (which changes its occupation) and the neighboring site j to uniquely define two distinct elementary reactions:



and



In the presence of periodicity in the employed lattice there can be a large number of events that in fact all represent the same elementary reaction, just occurring at different lattice sites. The definition through the local lattice configurations allows to efficiently achieve this classification by first checking if local educt and product lattice configurations can be transformed into each other through a lateral lattice translation vector. Since an elementary reaction is not affected by any lattice configuration difference outside the local educt and product configuration this grouping correctly includes many events which only differ by the (non-changing) occupations outside these local configurations. To illustrate this consider again CO adsorption on an empty periodic surface featuring one type of active site. Given that

adsorption into any of these sites is equivalent, adsorption events on sites i and j are different events in terms of the overall lattice configuration, yet they would both be grouped to the same elementary reaction by their identical local educt and product lattice configurations. Similarly, adsorption on site i with another adsorbate present on site k is again a different event, but falls still into the same elementary reaction class if there are no lateral interactions between sites i and k , and k is correspondingly outside the local lattice configuration. Notwithstanding, it is important to realize that identical local educt and product lattice configurations are only a necessary, but not a sufficient condition for the same elementary reaction. In the surface catalysis context, this is notably exemplified by Eley-Rideal type reaction events, where an adsorbed reaction intermediate is reacted off in a gas-phase scattering reaction. The local educt and product lattice configurations for such an event are identical to those describing a mere desorption process of the reaction intermediate. Yet, these are two distinct elementary reactions, which in the lattice framework is accounted for through two different rate constants.

SIZE AND STRUCTURE OF THE TRANSITION MATRIX The considerations about locality provide important insight into the structure of the overall transition matrix \mathbf{w} in Eq. (4.1). First, it can be decomposed into a sum of elementary reaction matrices as

$$w_{vu} = \sum_a w_{vu}^a, \quad (4.3)$$

where

$$w_{vu}^a = \begin{cases} k_a & (\mathbf{u} \rightarrow \mathbf{v}) \in a \\ 0 & \text{else} \end{cases} \quad (4.4)$$

and k_a is the reaction rate constant of elementary reaction a . The total number of different matrix elements is thus given by the number of inequivalent elementary reaction steps in the model.

Second, with respect to the structure of \mathbf{w} it is useful to define the set of *available events* $\sigma_{\mathbf{u}}$ for any configuration \mathbf{u} as the set of all events α_{vu} that lead from configuration \mathbf{u} to any other configuration \mathbf{v}

$$\sigma_{\mathbf{u}} = \{\alpha_{vu} | w_{vu} \neq 0\}, \quad (4.5)$$

i.e. $\sigma_{\mathbf{u}}$ is formed by the non-zero elements in the \mathbf{u} th column of \mathbf{w} . The locality of the elementary reactions implies that there are no events that connect largely differing lattice configurations. As such, $\sigma_{\mathbf{u}}$ is much smaller than the total size of \mathbf{w} , *i.e.* the transition matrix is sparse. For the later kMC efficiency discussion we further note that given an event α_{vu} , all events in $\sigma_{\mathbf{v}} \setminus \sigma_{\mathbf{u}}$ are said to be *enabled* by α_{vu} , while all events in $\sigma_{\mathbf{u}} \setminus \sigma_{\mathbf{v}}$ are said to be *disabled* by α_{vu} . As any event is local and thus affects much less sites than the total number of sites

in the lattice, for every event $\alpha_{v,u}$ the number of enabled or disabled events is again much smaller than the number of available events in both u and v , or

$$|\sigma_u \cap \sigma_v| \gg |(\sigma_u \cup \sigma_v) \setminus (\sigma_u \cap \sigma_v)|. \quad (4.6)$$

This difference in size will become the more pronounced the larger the lattice that is to be simulated.

These insights into the structure of the transition matrix are important as the formal simplicity of the master equation (4.1) easily disguises the complexity in solving it in practice in the surface catalytic context. This is due to the sheer size of the space of all possible lattice configurations. To illustrate this, let us assume a simple surface system that exhibits only one type of active site per unit cell and allows for a minimum number of two different reaction intermediates at this site (again in the context of CO oxidation this could be adsorbed O and CO). Together with the possibility of an active site being empty, this yields three possible occupations of every site. In order to properly capture the ensemble characteristics of the system like the average TOF we typically need to explicitly simulate at least a surface area of (10×10) surface unit-cells that is then continued by periodic boundary conditions. The total number of configurations possible on this lattice is $3^{100} \approx 10^{47}$, and the $(3^{100} \times 3^{100})$ transition matrix \mathbf{w} features $(3^{100})^2 \approx 10^{95}$ matrix entries $w_{v,u}$. As discussed this matrix is sparse though, as there are no events that connect largely differing lattice configurations, and σ_u for every configuration u will be much smaller than 3^{100} . Nevertheless, \mathbf{w} still contains a total number of non-zero elements that is too large to be stored directly. On the other hand, due to the translational symmetry of the lattice the total number of inequivalent matrix entries $w_{v,u}$ is typically rather small and determined by the total number of inequivalent elementary reactions α in the reaction network, cf. Eq. (4.4). For a CO oxidation model in the described simple surface system this total number can be as low as seven: dissociative adsorption of O_2 , associative desorption of two adsorbed O, CO adsorption, CO desorption, O diffusion, CO diffusion, and CO+O reaction.

4.2.2 Kinetic Monte Carlo

It is clearly hopeless to explicitly solve such a high-dimensional master equation directly or even just aim to store the entire probability density in a lattice representation. The idea behind kinetic Monte Carlo (kMC) simulations is instead to achieve a numerical solution by generating an ensemble of trajectories of the underlying Markov process, where each trajectory propagates the system correctly from configuration to configuration in the sense that the average over the entire ensemble of trajectories yields the probability densities $\rho_u(t)$

of Eq. (4.1).[12, 25, 47, 59, 154, 193] Analysis of any single (stochastic) kMC trajectory is correspondingly meaningless, unless a stationary system state (in the catalysis context: steady-state operation) allows to replace the ensemble average by a time average over one trajectory. The actual objective for a kMC computer algorithm (and in turn for a software package like `kmoss`) is therefore to generate such kMC trajectories. For this, the kMC code generally only needs to store the (evolving) occupation values on the lattice, as well as the inequivalent rate constants of all elementary reactions. On-the-fly it then generates and focuses on those transition rates $w_{v,u}$ that are actually required to propagate the trajectory. More specifically `kmoss` aims to optimize the execution of previously defined models. While advanced model generation schemes (like self-learning kMC) could be implemented on top of the `kmoss` framework, they are currently out of the scope of general first-principles methodology[154, 165] and not included in this manuscript.

Differences between kMC solvers arise in the way how the event sequence is chosen and the concomitant way how the elapsed system time is determined. For the latter, one generally exploits that waiting times for uncorrelated events are Poisson distributed.[12, 47, 59] This means that given a rate constant k for an event, the probability that n events occur in an interval Δt is

$$p_n(k, \Delta t) = (k\Delta t)^n e^{-k\Delta t} / n! \quad . \quad (4.7)$$

The waiting time between two events is then simply given by the case that no events occur

$$p_0(k, \Delta t) = e^{-k\Delta t} ,$$

for which a suitably distributed random number can be directly computed from a uniformly distributed random number $r \in]0, 1]$ [150] as

$$\Delta t = \frac{-\ln(r)}{k} \quad . \quad (4.8)$$

Lukkien *et al.*[115] proposed a unified scheme consisting of three categories that classify existing kMC solvers: The first reaction method (FRM), the variable step-size method (VSSM), and the random selection method (RSM). We now briefly describe the essential features of each category, primarily to contrast the conceptual differences. The equivalence of all three approaches has also been shown by Lukkien *et al.*[115], such that a preference for one or the other emerges only out of efficiency considerations as discussed in the next section.

FRM At every kMC step the FRM updates the sequence of available events σ_u and their corresponding rate constants k_u . From this it calculates a sequence of time increments $\tau = -\ln(\mathbf{r})/k_u$, where $\mathbf{r} \in$

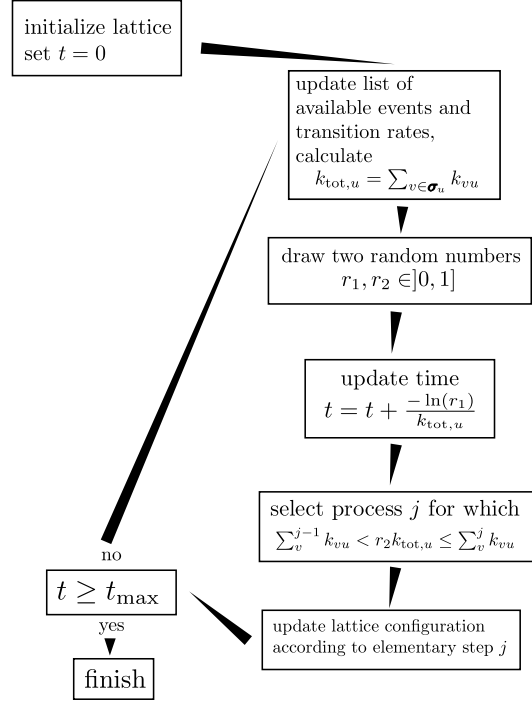


Figure 4.2: The basic steps of a VSSM kinetic Monte Carlo algorithm.

$]0, 1]$ is a sequence of uniformly distributed random numbers. The smallest element of τ is selected, the elapsed time is advanced by the corresponding time increment, and the corresponding event is executed by updating the system configuration accordingly.

vssm At every kMC step the VSSM updates the sequence of available events σ_u and calculates the total rate $k_{\text{tot},u} = \sum_{v \in \sigma_u} k_{vu}$. The time is advanced by $-\ln(r)/k_{\text{tot},u}$, where $r \in]0, 1]$ is a uniformly distributed random number, and one of the available events is selected for execution with a probability weighted by its rate constant. Since VSSM is the algorithm underlying *kmos*, Fig. 4.2 further illustrates these steps in form of a flow chart.

rsm At every kMC step the RSM calculates $K = \sum_a k_a$, where the sum runs over all elementary reactions a . One elementary reaction is chosen with a probability weighted by its rate constant, and one of the N_{sites} sites in the lattice is randomly selected. The time is increased by $-\ln(r)/N_{\text{sites}}K$, where $r \in]0, 1]$ is a uniformly distributed random number. The lattice is updated if the event is available at the selected site.

4.2.3 Efficient Lattice kMC

For small lattice sizes and simple kMC models containing a limited number of elementary reactions efficiency of the kMC code is gen-

erally not an issue. In particular in the context of multi-scale modeling approaches the computational costs of kMC simulations are completely negligible compared to typical costs of first-principles electronic structure theory calculations. With increasing lattice sizes and complexity of the kMC model this situation changes, in particular when considering that likely larger numbers of kMC simulations have to be run to cover different gas-phase operation conditions or when performing sensitivity analyses. In this situation, efficiency of the lattice kMC simulations becomes paramount. For the RSM algorithm a major limitation to efficiency might arise out of a diminishing probability for successful events. In the surface catalysis context this commonly arises in situations of almost fully occupied lattices and the presence of a very fast diffusion process. Such events are then predominantly chosen, but essentially never successful. For the rejection-free FRM and VSSM algorithms the main bottleneck arises instead out of the necessity to update the sequence of available events at every kMC step. A naïve approach that is straightforward to implement (and accordingly chosen in many ‘from scratch’ programs) is to determine σ_u through iteration over all lattice sites. With the number of lattice sites possibly going up to tens of thousands for entire nanoparticle simulations any retraction to such a sequential operation on the full lattice ($O(N_{\text{sites}})$) will then drastically impede the overall performance.

Lukkien *et al.* [115] have systematically analyzed the efficiency of the three kMC approaches and concluded on VSSM as most promising method. In line with this analysis, VSSM has also been chosen as basic algorithm underlying kmos. In contrast to the FRM algorithm VSSM requires only two random numbers per kMC step, cf. Fig. 4.2. As such its main computational burden lies in the repeating update of the set of available events and total reaction rate $k_{\text{tot},u}$. In contrast to the naïve $O(N_{\text{sites}})$ approach, exploitation of the locality of the elementary reactions allows to largely reduce the scaling of both these calculation steps. With respect to the set of available events this is achieved through local update procedures, thereby taking into account eq. (4.6). Rather than building this set anew at every kMC step, these local updates merely determine a new σ_v from the previous set of available events σ_u by strictly removing all disabled events ($\sigma_u \setminus \sigma_v$) and adding all enabled events ($\sigma_v \setminus \sigma_u$), or formally

$$\sigma_v = (\sigma_u \setminus (\sigma_u \setminus \sigma_v)) \cup (\sigma_v \setminus \sigma_u) \quad .$$

From the new set of available events σ_v , its corresponding total rate constant $k_{\text{tot},v} = \sum_{w \in \sigma_v} k_{wv}$ can also be calculated without iterating over the full size of the set (which would generally also scale as $O(N_{\text{sites}})$). For this, not only the contained events directly, but also the number of events $N_{a,v}^{\text{avail}}$ that belong to the same elementary reaction a are stored. This way, if an event α_{wv} belonging to an elementary

reaction a is added to the set of available events, the corresponding counter $N_{a,v}^{\text{avail}}$ is simply increased by 1, whereas if $\alpha_{w,v}$ was removed, the counter is decreased. As a result one can quickly calculate $k_{\text{tot},v}$ as

$$k_{\text{tot},v} = \sum_a k_a N_{a,v}^{\text{avail}} \quad , \quad (4.9)$$

where the sum does not iterate over the elements in σ_v , but only over the much smaller set of elementary reactions ($O(N_{\text{react}})$). Further, if the previous summation is carried out by filling an array of accumulated rates

$$k_{1,v}^{\text{acc}} = 0 \quad (4.10)$$

$$k_{a,v}^{\text{acc}} = k_{a-1,v}^{\text{acc}} + k_a N_{a,v}^{\text{avail}} \quad , \quad (4.11)$$

the next event $\alpha_{w,v}$ can also be selected without retracting to $O(N_{\text{sites}})$ operations by using a random number $r \in]0, 1]$ and selecting the elementary reaction b for which

$$k_{b-1,v}^{\text{acc}} < r k_{\text{tot},v} \leq k_{b,v}^{\text{acc}} \quad , \quad (4.12)$$

through a binary search ($O(\log(N_{\text{react}}))$), and then selecting randomly one of the $N_{b,v}^{\text{avail}}$ available events belonging to elementary reaction b ($O(1)$).

As this analysis shows every required task of a VSSM lattice kMC solver can thus be carried out with a computational effort that is independent of the number of sites in the system.

4.2.4 Sampling of Reaction Rates

A central capability of kMC simulations in the context of heterogeneous catalysis is the calculation of reaction rates. Normalized to active site or surface area, corresponding TOFs yield the occurrence of any elementary reaction per time. If this elementary reaction yields a final product, then its TOF measures the overall catalytic activity with respect to this product. If there are several elementary reactions leading to different products, then the ratios of their TOFs additionally provide the selectivities.

In the context of kMC simulations a straightforward definition of the TOF per active site of any elementary reaction a at any time t is

$$\text{TOF}^a(t) = \frac{\langle N^a(t) \rangle}{N_{\text{sites}}} \quad , \quad (4.13)$$

where $N^a(t)$ is the number of times that reaction a has occurred at time t , and the average $\langle \rangle$ is over a sufficiently large ensemble of kMC trajectories. Realizing that the actual occurrence of an event is

given by the probability for the system to actually be in a configuration where the event is enabled times its rate constant, an equivalent definition is

$$\text{TOF}^{\alpha}(t) = \frac{\sum_{\mathbf{u}} \rho_{\mathbf{u}}(t) \sum_{\mathbf{v}} w_{\mathbf{v}\mathbf{u}}^{\alpha}}{N_{\text{sites}}} , \quad (4.14)$$

where the sums run over all configurations \mathbf{u} and \mathbf{v} , $\rho_{\mathbf{u}}(t)$ is the probability for the system to be in configuration \mathbf{u} at time t , and $w_{\mathbf{v}\mathbf{u}}^{\alpha}$ are as defined in eq. (4.4) the transition rates of all events $\alpha_{\mathbf{v}\mathbf{u}}$ that correspond to the elementary reaction α .

In catalytic applications the primary focus is typically on steady-state operation. Even if time-dependent operation conditions and consequently time-dependent TOFs are of interest, the changes generally occur over at least mesoscopic times, and can therefore be captured through appropriate binning in constant-condition time windows. In a corresponding stationary situation the ensemble averages in eqs. (4.13) and (4.14) can be replaced by time averages. For the first definition this leads numerically to

$$\text{TOF}^{\alpha} \approx \frac{\int_0^{t_{\text{final}}} N^{\alpha}(t) dt}{N_{\text{sites}} t_{\text{final}}} = \frac{N_{t_{\text{final}}}^{\alpha}}{N_{\text{sites}}} , \quad (4.15)$$

where $N_{t_{\text{final}}}^{\alpha}$ is the total number of times elementary reaction α took place in a time span t_{final} . A corresponding straightforward counting to determine TOFs is what is primarily implemented in simple 'from scratch' kMC codes. This approach becomes highly inefficient though, if small TOFs are to be measured. Due to the irregular and rare occurrence of the corresponding elementary reaction long time spans need to be simulated to sufficiently converge the TOF. In this situation it is advantageous to resort to the second TOF definition in eq. (4.14),

$$\begin{aligned} \text{TOF}^{\alpha} &= \frac{\sum_{\mathbf{u}} \bar{\rho}_{\mathbf{u}} \sum_{\mathbf{v}} w_{\mathbf{v}\mathbf{u}}^{\alpha}}{N_{\text{sites}}} \\ &\approx \frac{\sum_{i=1}^{N_{\text{final}}} \sum_{\mathbf{v}} w_{\mathbf{v}\mathbf{u}_i}^{\alpha} \Delta t_i}{N_{\text{sites}} t_{\text{final}}} \\ &= \frac{\sum_{i=1}^{N_{\text{final}}} k_{\alpha} N_{\mathbf{a},\mathbf{u}_i}^{\text{avail}} \Delta t_i}{N_{\text{sites}} t_{\text{final}}} , \end{aligned}$$

where N_{final} are the number of kMC steps in the time span t_{final} , \mathbf{u}_i is the configuration occupied at the beginning of kMC step i and Δt_i is its duration, that is the time until the simulation jumps out of \mathbf{u}_i . The second equality demonstrates the efficiency of this approach, which adds to the convergence of the TOF with every kMC step even if k_{α} is very small, and which furthermore comes at negligible overhead

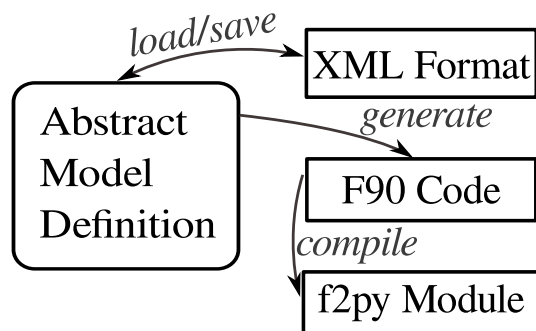


Figure 4.3: Scheme for the flow of information in the `kmos` framework. An abstract kMC model is defined by using the `kmos` API. Thus, either the Python code using the `kmos` API itself can serve as definitive specification of the kMC model, or the model can be stored in an XML scheme for archiving and exchange. From the abstract model definition `kmos` generates tailored Fortran90 source code that performs the actual kMC simulations. This code can be compiled and exposed to Python again using `f2py`.

as N_{a,u_i}^{avail} is calculated at every kMC step i anyway. For the determination of low TOFs this approach can therefore significantly reduce the time required for a converged sampling and is correspondingly implemented in `kmos` by default.

4.3 THE kmos FRAMEWORK

The essential idea of the `kmos` approach to kMC modeling is to use a code generator to produce highly efficient code from an abstract definition of a kMC model. As further detailed below `kmos` thus avoids a static and in full generality cumbersome hard-coding of the complex conditional dependencies between arbitrary events. Instead it custom tailors the code on the basis of a defined model, which in particular allows for most efficient local updates of enabled and disabled events. The general flow of information in the `kmos` framework is illustrated in Fig. 4.3. The following three subsections consecutively describe the three main parts apparent from this scheme: The specification of the kMC model, the code generation from the model, and how the generated code implements the VSSM kMC algorithm.

4.3.1 kMC Model Definition

Since this part of the `kmos` framework is Python based, this subsection will borrow a subset of object-oriented terminology to describe its structure: Essentially, a kMC model is a hierarchy of objects with attributes.

The information necessary to define a kMC model generally falls into two related, but distinct categories. On the one hand, there is the

information required for the actual kMC simulations. This is information on the sites and lattice structure, on the reaction intermediates (code-internally called species), on general parameters like temperature or partial pressures that can be used to internally compute the rate constants, as well as all possible elementary reactions. On the other hand, there is additional information required for the analysis, in particular for an atomistic visualization of the generated kMC trajectories. This is prominently any explicit geometric information (size and shape of unit-cell, Cartesian coordinates of sites within the unit-cell, representation of substrate and reaction intermediates). In summary, this leads to the following schematic structure of the model definition:

- model
 - lattice (geometry): unit cell, [sites]
 - sites: name, position
 - reaction intermediates (species): name, representation
 - parameters: name, value
 - elementary reactions: name, [conditions], [actions], rate constant

Within this basic skeleton the user has to define the model specific parts. For this, one could envision some configuration file-like format. However, in particular with respect to the sequence of elementary reactions it turns out that one would have to type many very similar statements. For instance, if the same elementary reaction can be executed in several different directions due to the symmetry of the lattice. `kmoss` therefore offers an application programming interface (API) that allows to create each object in the model by one constructor call (in terms of object-oriented programming). This has the benefit of offering a fairly straightforward syntax, while at the same time allowing for all the flexibility and expressiveness of a high-level programming language and its control constructs such as for-loops and if-statements.

LATTICE DEFINITION The system is represented as a finite lattice with periodic boundary conditions. At present `kmoss` only supports one global Bravais lattice; a limitation that we intend to overcome in future work. Multiple active sites within the unit cell are accounted for through a basis. Each site is defined through a unique name. Internally every lattice point can thus be represented with a four-tuple $n.(x, y, z)$, where x, y, z are the integer coordinates of the unit cell, and n goes over the different active sites within the unit cell as illustrated in Fig. 4.4. Naturally this scheme can describe one-, two-, or three-dimensional lattices by setting 2, 1, or 0 entries to zero respectively. By default `kmoss` enforces periodic boundary conditions by

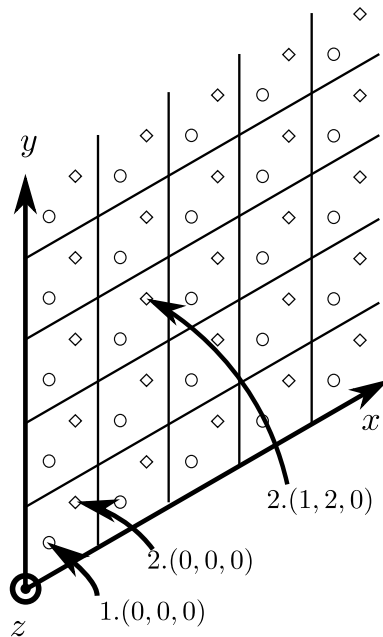


Figure 4.4: Illustration of the lattice representation using the four-tuple $n \cdot (x, y, z)$, where x, y, z are the integer coordinates of the unit cell, and n goes over the different active sites within the unit cell.

internally expanding the lattice by one unit cell along each lattice axis. When interested in modeling a finite lattice, this feature can be blocked by defining an inactive dummy reaction intermediate and initializing the edges of the simulated geometry with it. For visualization the shape and size of the unit cell can be specified, as well as the fractional Cartesian coordinates of all active sites within the unit cell.

REACTION INTERMEDIATE DEFINITION Reaction intermediates are specified through a unique name, and internally get assigned an integer value. A species empty needs to be explicitly defined. Site blocking, e.g. in multidentate adsorption or to mimic infinitely repulsive lateral interactions, can be achieved through the definition of dummy species (say `A_blocked` as additional dummy for a reaction intermediate `A` covering multiple sites). In the specification of the elementary reaction, the blocked sites are then occupied with the dummy, which thus prevents them from being empty for other elementary reactions. Special boundary conditions such as a source or a drain that continuously inserts or removes surface intermediates at the edges of the lattice can similarly be modeled by using such special intermediates and corresponding elementary reactions. For the purpose of visualization it is possible to enter a string in the atoms-object-constructor form as understood by the Atomic Simulation Environment (ASE).[3]




Figure 4.5: Graphical representations illustrating the defining characteristics of elementary reactions: Dissociative O_2 adsorption (top panel), CO diffusion (middle panel) and CO oxidation (bottom panel).

ELEMENTARY REACTION DEFINITION The elementary reactions are defined in terms of the occupations in the local educt and local product lattice configuration, as well as the corresponding rate constant. Some sample definitions are depicted in Fig. 4.5. For elementary reactions involving more than one site, other involved sites are specified by relative vectors in the four-tuple representation. If the central site used to define the elementary reaction is e.g. a bridge site and another bridge site in the unit cell in the positive direction of the first lattice vector is involved, then this additional site is referred to as `bridge.(1,0,0)`. `kmos` makes no effort to recognize symmetries in the lattice (e.g. to deduce that a diffusion from `bridge.(1,0,0)` implies the possible equivalent diffusion to `bridge.(-1,0,0)`). However, the `kmos` API allows to select pairs of sites based on type and geometrical distance making the inclusion of such equivalences straightforward. Since each element (site and occupation) of the local educt lattice configuration acts as a requirement that the elementary step can be executed, it is coined `Condition`. The definition of an elementary reaction can in principle contain an arbitrary number of such `Conditions`, though there are limits on the number that the compiler can process as discussed in the next section. Nevertheless, this allows to describe fairly complex elementary reactions involving lateral interactions, concerted processes, bystander adsorbates, multidentate adsorption, and even some reconstruction of the underlying lattice structure. The provided primitives (`Conditions` and rate-constant expressions) allow in principle for any sophisticated level of included lateral interaction in the sense of a surface cluster expansion [133, 203]. Yet, the analytical dependence of the rate constant on these lateral interactions has to be explicitly provided by the user. Each element (site and occupation) of the local product lattice configuration describes a

change induced by the elementary reaction and is thus coined `Action`. Internally, `Condition` and `Action` are identical data types, but for sake of clarity different class names are used.

RATE CONSTANT EXPRESSIONS AND PARAMETERS `kmos` accepts hard-coded values for the rate constants of the individual elementary reactions. However, in the context of surface catalysis the rate constants are often calculated using expressions such as

$$k = \frac{k_B T}{h} \exp(-\beta \Delta G) \quad ,$$

for activated surface processes or

$$k = \frac{p_i A}{\sqrt{2\pi m_i k_B T}}$$

for adsorption processes of ideal-gas particles (see e.g. ref. [153], which also includes the definition of the various parameters appearing in these expressions). Since it is convenient to quickly iterate external parameters (like temperature T and partial pressures p_i) or directly change activation barriers ΔG for example in a sensitivity analysis study, `kmos` also allows to directly enter such mathematical expressions for the rate constants as strings, which are later evaluated at runtime for the parameters currently present. Since these evaluations are quite expensive, they are only updated if any of the parameters change though.

4.3.2 Code Generator

As discussed in section 4.2.3 the main efficiency driver of a VSSM-based kMC code is the local update procedure, with its concomitant determination of disabled and enabled events. This local update procedure is also the only heavily model-dependent part of any kMC program, whereas as detailed in the next section all other parts of the actual kMC algorithm can be written in a generic way. Complicating matters, practical kMC work typically involves frequent changes of the kMC model (refinement through addition of new elementary reaction processes, consideration of further sites and reaction intermediates etc.). These changes require modifications of the code in typically as many locations as there are elementary reactions, since each new reaction might be affected by all existing elementary reactions while it can also affect every existing elementary reaction (*vide infra*). Doing these modifications by hand (as in the early 'from scratch' codes) is therefore not only highly cumbersome, but also extremely prone to human error. The modifications concern furthermore precisely that part of the code that determines the overall efficiency and should therefore not be implemented in an unoptimized way. `kmos`' answer

to this situation is to fully automatize this aspect of the work by outsourcing it to a secondary code generator program. On the basis of a defined kMC model this code generator writes the required update procedure in a compilable program language (Fortran90), which consecutively can be included in the remaining kMC program. This way, one gets the best of two worlds: A flexible high-level interface for defining kMC models and at the same time an optimized low-level implementation of the model.

The key to design an efficient code generator is to carefully reflect every logical dependence in the model to infer as many decisions as possible during the generation step and minimize the number of computational steps at runtime. This is quite different from traditional programming approaches as reflected in the following explanation, since the two levels of algorithmic description are inevitably entangled. As stated above every elementary reaction is defined in terms of Conditions and Actions, and in turn each of these is defined by a relative site coordinate and a concomitant species occupation. All Conditions need to be satisfied for an event representing the elementary reaction to become enabled and only one of the Conditions needs to be dissatisfied for an event to be disabled.

In order to implement the required updates of the set of available events σ_v after an event α_{vu} has been selected, some events have to be added and some have to be removed. Removing events is conceptually and computationally simpler than adding, since removing does not require any inspection of the actual lattice configuration or evaluation if *all* Conditions are satisfied. Instead it can be based on evaluating if any Action of α_{vu} dissatisfies a Condition of an available event in σ_u . One therefore iterates over the Actions (that is configuration changes) due to α_{vu} . For each Action i , which is defined by a *species* and a *site*, one iterates over the sequence of elementary reactions. For each elementary reaction b the first check is if b contains at least one Condition j on the same *site* as in Action i . If this is the case, then such a reaction b could potentially have been affected by the occurrence of event α_{vu} . We correspondingly then also check if the *species* of Action i *does not* match with the *species* of Condition j . If this is also the case then an event β_{wv} corresponding to elementary reaction b at the lattice site where α_{vu} has occurred, has become disabled. Thus compilable code is generated which removes β_{wv} from the available events σ_v , if it was enabled in σ_u .

After the lattice configuration itself is updated (by generating corresponding compilable code to change the occupation entries) the newly enabled events can be added. Again one iterates over all Actions of α_{vu} . For each Action i again defined by a *species* and *site* one iterates over all elementary reactions. For each elementary reaction b , the first check is again if it contains at least one Condition j on the same *site* as Action i . If in addition the *species* of Action i *does* match

with the species of Condition j , the corresponding event β_{wv} of elementary reaction b might have been enabled by the occurrence of event α_{vu} . Other than in the disabling procedure we now have to iterate over all other Conditions of β_{wv} though, which in fact involves inspection of the occupation of all sites contained in the local educt lattice configuration of β_{wv} . Only if *all* Conditions are satisfied, β_{wv} has indeed become enabled through the occurrence of event α_{vu} . Thus compilable code is generated which iterates over all Conditions of event β_{wv} at the corresponding location in the lattice and checks whether in fact the *species* of all Conditions of b match with the *species* present at the relative site. If all Conditions are satisfied, β_{wv} is added to the available events σ_v .

In pseudo-code the combined algorithm developed above can be concisely written as follows. Code executing before compile-time (code generation) is set in roman type, while code executed at runtime is set in monospaced type (*vide infra*). Variable names are set in *italics*. The **for** statement borrows on the Python style syntax (**for** i **in** $x \leftrightarrow$ *block*), which instructs to execute *block* on every element of x and the element will be named i inside *block*.

```

# Update available events for
# elementary reaction a

#Disable events
for  $i = (\textit{species}, \textit{site})$  in actions of  $a$ 
  for  $b$  in elementary reactions involving site
    for  $j$  in conditions of  $b$ 
      if  $i$  contradicts  $j$ 
        disable  $\beta_{wv}$  if enabled

Update lattice configuration

#Enable events
for  $i = (\textit{species}, \textit{site})$  in actions of  $a$ 
  for  $b$  in elementary reactions involving site
    for  $j$  in conditions of  $b$ 
      if  $i$  fulfills  $j$ 
        if all conditions of  $\beta_{wv}$  are met
          enable  $\beta_{wv}$ 

```

A crucial feature of this general update algorithm is hereby that even though it requires four nested loops, the outcome of the Conditions checked in the outermost loops is uniquely determined by the given kMC model. Rather than evaluating these conditions during the actual runtime of the kMC simulation over and over again, the kmos code generator evaluates them beforehand and builds the outcome

directly into the generated code. The parts that need to be executed at runtime consist therefore at most of two nested loops and are by construction optimized for the defined kMC model. In terms of the generated code, the most tricky part is hereby the implementation of the check, if *all* Conditions of a possibly enabled event are satisfied. Checking such interdependencies between different Conditions typically involves many memory reads over all sites of the local educt lattice configuration and thus affects the performance. So, the question arises whether there is an optimal way how to arrange the corresponding queries. The corresponding problem of constructing an optimal binary decision tree has generally been shown to be NP-complete [92]. In the kMC context the corresponding intricacy is given in loose terms by the fact that frequent local lattice configuration motifs, which would be the basis for an optimal construction algorithm, are unknown beforehand and precisely the outcome of a kMC simulation. Accordingly, only two heuristic approaches are presently implemented in `kmos` and can be selected from the command line in the code generation step (`kmos export -b<code-generator>`).

In the first approach the generated code is arranged in such a way that the average number of memory accesses are likely to be minimal for the case that every outcome is equally probable. To this end, during the code generation phase, all required read accesses are collected and sorted by decreasing frequency. All possible outcomes are grouped by the result of the most common read access and accordingly written into different conditional branches. Within each branch this process is recursively repeated. This approach shows exceptional performance at runtime for kMC models without lateral interactions and few species. Though, for more complex models involving more than three different species or considerably far-ranging lateral interactions, it often produces an exceedingly large code tree (on the order of 100MB) and accordingly long compilation time (on the order of hours).

The second approach correspondingly aims at a moderate size of the generated code and for this assumes that the primary source for the existence of multiple Conditions is the existence of lateral interactions extending over many lattice sites. All elementary reactions are then automatically grouped into sets of identical Actions. That is each group contains elementary reactions that are identical in the sites and species that are changed in the execution, and only differ by their Conditions that are not changed by the elementary reaction. The rationale behind this is that models involving lateral interactions contain only a few of these sets. Within each set the present lateral interactions can be determined by as few read accesses as there are lateral interactions, since one specific lateral interaction educt excludes all others within the set. The code resulting from this approach proves

to be much shorter even for models involving as much as five species and up to 40 Conditions (on the order of few MB), and the compilation time stays typically on the order of minutes.

4.3.3 Kinetic Monte Carlo Solver

The generated code is combined with other generic parts to form a VSSM lattice kMC solver that follows the general flow chart shown in Fig. 4.2. To realize the efficiency considerations summarized in Section 4.2.3 this solver operates on a well designed data structure. The base of this data structure is a bijective mapping from the four-tuple $n \cdot (x, y, z)$ lattice representation to a one-dimensional representation, which simply enumerates all lattice points. This mapping can be cached in 1D and 4D arrays which makes it very efficient. Any of the frequently executed core parts are then performed on the 1D representation, and only if explicit inspection of the lattice configuration is required is the trivial inverse mapping applied. As shown below the largest arrays then have a size $(N_{\text{react}} \times N_{\text{sites}})$, where N_{react} is the total number of elementary reactions and N_{sites} is the total number of sites. Even for very large lattices such arrays do not represent any notable memory requirements. kmos correspondingly uses fixed array sizes and avoids dynamic data types which would require continuous memory allocation and deallocation. On this data structure the fundamental data operations to (a) determine the next event and (b) add and delete events to and from the set of available events can be executed independent of the lattice size.

DATA STRUCTURES The deployed kMC solver operates on these 6 arrays, cf. Fig. 4.6:

- The array $\mathbf{L} = \mathbf{L}(N_{\text{sites}})$ stores the current configuration of the system, *i.e.* the integer value of L_x represents the occupation at the x th site.
- The array $\mathbf{k} = \mathbf{k}(N_{\text{react}})$ stores the rate constants for all elementary processes.
- The array $\mathbf{N}^{\text{avail}} = \mathbf{N}^{\text{avail}}(N_{\text{react}})$ stores the number of available sites for all elementary reactions, cf. Eq. (4.9).
- The array $\mathbf{k}^{\text{acc}} = \mathbf{k}^{\text{acc}}(N_{\text{react}})$ stores the accumulated rate constants, cf. Eq. (4.11).
- The array $\mathbf{A} = \mathbf{A}(N_{\text{react}}, N_{\text{sites}})$ stores the available events. Each row of \mathbf{A} represents one elementary reaction and is filled from the left, *i.e.* an element $A_{ai} = x > 0$ tells that site x is currently available for elementary reaction a .




Figure 4.6: Main data structures and event selection process of the kmos VSSM-kMC solver. The array of rate constants \mathbf{k} is usually unchanged during a kMC simulation. $\mathbf{N}^{\text{avail}}$ reflects the number of sites available for each elementary reaction. Using one random number r_2 an elementary reaction a is selected using a binary search on the accumulated rate constants \mathbf{k}^{acc} . For this elementary reaction a one of the available events is selected from \mathbf{A} using the product of a random number r_3 and the number of available sites N_a^{avail} as an index. Note that for each row a in \mathbf{A} the first N_a^{avail} elements are non-zero after which all entries are zero. The array \mathbf{I} stores the position under which the available events are stored in \mathbf{A} so that all necessary updates can be executed on \mathbf{A} without traversing it.

- The array $\mathbf{I} = \mathbf{I}(N_{\text{react}}, N_{\text{sites}})$ allows to retrieve the available events in array \mathbf{A} . For this, if site x is currently available for elementary reaction a and the corresponding event is stored in element A_{ai} , then $I_{ax} = i$. If site x is currently not available, then $I_{ax} = 0$.

DETERMINATION OF THE NEXT EVENT In every kMC step the solver determines the next event and therewith the concomitant elementary reaction and site as illustrated in Fig. 4.6. First, the array of accumulated rate constants \mathbf{k}^{acc} is updated according to the current set of available events. This includes the calculation of the total rate constant as last element $k^{\text{acc}}(N_{\text{react}})$. The elapsed time is updated as $-\ln(r_1)/k_{\text{tot}}$, where $r_1 \in]0, 1]$. Using another uniform random number $r_2 \in]0, 1]$ and a binary search [30] an elementary reaction a is determined for which $k_a^{\text{acc}} < k_{\text{tot}}r_2 \leq k_{a+1}^{\text{acc}}$ by performing a binary search on \mathbf{k}^{acc} . Using a third uniformly distributed random number $r_3 \in]0, 1]$ the concomitant site for the selected event is determined from array \mathbf{A} as the value of element A_{ai} , where $i = \lfloor r_3 N_a^{\text{avail}} \rfloor$.

UPDATE OF THE SET OF AVAILABLE EVENTS After having selected the event, that is elementary reaction and site, the code-generated part takes over to call the required additions and deletions to the set of available events, as well as the update of the lattice configuration. The prior two operations are straightforward but critical primitives of

the generated local update code. In terms of the relevant data structures enabling site x for elementary reaction a consists of the following steps:

1. Increase number of available events:

$$N_a^{\text{avail}} := N_a^{\text{avail}} + 1$$
2. Store site x : $A_{aN_a^{\text{avail}}} := x$
3. Assign address for site x : $I_{ax} := N_a^{\text{avail}}$

Similarly, the deletion of a disabled elementary reaction a at site x proceeds as:

1. Overwrite site x with last site enabled for a : $A_{aI_{ax}} := A_{aN_a^{\text{avail}}}$
2. Empty last site $A_{aN_a^{\text{avail}}} := 0$
3. Reassign address of moved site:

$$I_{aA_{aI_{ax}}} := I_{ax}$$
4. Empty address of deleted site: $I_{ax} := 0$
5. Decrease available events: $N_a^{\text{avail}} := N_a^{\text{avail}} - 1$

As one can see an *enabling* or *disabling* operation requires three or five memory transactions, respectively, and thus does not depend on the total system size or complexity. Only the search time for the next elementary reaction grows logarithmically with the number of elementary reactions N_{react} due to the binary search involved. However, this is not expected to become a bottleneck as this number is generally much smaller than the total number of events that have to be enabled or disabled.

This concludes all required algorithmic work in one kMC step and the next step can follow.

RANDOM NUMBERS As indicated above the kMC solver requires three uniformly distributed random numbers per kMC step. `kmos` relies on the pseudo random number generator (PRNG) provided by the Fortran compiler. Sometimes kMC practitioners are concerned whether such a source of randomness introduces non-physical bias to the generated kMC trajectory. We have not observed any such bias in a kMC simulation so far. In case doubt arises this can be easily tested by changing the PRNG seed conveniently in the configuration file of the compiled kMC model. Furthermore, the currently specified PRNG periods of the most commonly used compilers typically exceed the maximum number of kMC steps during one simulation by several orders of magnitude.

OVERALL CODE LAYOUT Having specified the kMC solver independently of any lattice geometry or specifics of elementary processes means one can reuse this part of the algorithm for virtually any lattice kMC model. This is also reflected in the structure of the overall Fortran code: It is subdivided into the modules `base`, `lattice`, and `proclist`, of which `base` contains the model-independent parts of the VSSM loop that has been described in this Subsection. The module `lattice` replicates the base API in terms of lattice coordinates and implements corresponding information about the model (number of lattice dimensions, numbers of sites per unit cell, and names of sites) for visualization, as well as the central VSSM loop. The third module `proclist` is the one produced by the code-generator and implements how the set of available events is updated after an event has been selected in a kMC step, cf. Section 4.3.2.

4.3.4 *Simulation Frontend*

The complete kMC model is stored in an XML text file by using the `elementtree` XML library. This also allows for easy archiving and exchange of models. A basic graphical user interface (GUI) is provided to visually inspect all aspects of the model definition including the elementary reactions. The generated Fortran code is compiled and exposed as a Python module with the `f2py` [148] interface generator. `kmoss` offers a concise API which allows to control all runtime aspects of a compiled model including setup and evaluation, as a script or interactively using IPython [146] and `numpy` [139], as well as a GUI which visualizes the model geometry using ASE [3] and coverages and turnover frequencies using `matplotlib` [91], while allowing to visually change parameters during the simulation.

4.4 PERFORMANCE AND SCALING IN PRACTICE

4.4.1 *ZGB Model*

We demonstrate the performance and scaling behavior of `kmoss` using a range of kMC models, and start with the seminal model by Ziff, Gulari, and Barshad (ZGB) [208], that has evolved into an influential reference for the development of stochastic approaches to surface catalytic processes. The original ZGB model generically considers CO oxidation at a simple cubic lattice, featuring one active site and only three elementary reactions: irreversible unimolecular adsorption of CO with rate constant y_{CO} , irreversible dissociative adsorption of O_2 at two neighboring sites with rate constant $1 - y_{\text{CO}}$, and instantaneous CO oxidation reaction of directly neighboring adsorbed CO and O. The only free parameter of the model is thus y_{CO} , which is varied in the range $[0,1]$ a.u. In the context of numerical kMC simulations we




Figure 4.7: Coverage dependence and catalytic activity of the ZGB model as implemented using *kmos*. In the idealized ZGB model catalytic activity is defined as the number of CO_2 molecules produced per reactant impingement.[208]

realize this model by approximating the instantaneous CO oxidation reaction with an exceeding rate constant of 10^{15} a.u., and adding unimolecular CO and associative oxygen desorption reactions with negligible rate constants of 10^{-13} a.u. to mimic the irreversible adsorption. Especially the latter is necessary to prevent the system from getting trapped in completely oxygen or CO poisoned configurations, but we validated that neither the obtained results nor runtime performance depends on the particular choice of the finite rate constants chosen for these processes. Figure 4.7 shows the resulting lattice occupations and CO_2 TOF in the relevant range of y_{CO} , perfectly reproducing the two critical y_{CO} values of $y_1 = 0.389$ and $y_2 = 0.527$ that delimit the O and CO coexistence at the surface and the concomitant catalytic activity.[208] The simulations were performed on a lattice containing (200×200) sites, and for this benchmark system *kmos* executed 2.15 million kMC steps per second on a 3.4 GHz Intel Core i7 processor with 16GB RAM. Given that the simulated elapsed time per kMC step varies with every configuration, the corresponding CPU time per million kMC steps (0.47 sec) is the only transferable benchmark property across implementations and somewhat even across models of similar complexity (*vide infra*).

4.4.2 Literature First-Principles kMC Models

As representative examples for modern first-principles based kMC models we consider the CO oxidation model at $\text{RuO}_2(110)$ as put forward by Reuter and Scheffler [157, 158] and the CO oxidation model at a thin $\text{PdO}(101)$ film on top of $\text{Pd}(100)$ as put forward by Rogal, Reuter and Scheffler [163, 164]. The prior model does not include lateral interactions, while the latter model does include pairwise nearest-neighbor lateral interactions at an otherwise comparable number of inequivalent elementary reactions. The comparison of the two models




Figure 4.8: Single-core CPU times required to execute one million kMC steps for the CO oxidation at RuO₂(110) model (solid line) and the CO oxidation at PdO(101) film model (dashed line) as a function of the simulated lattice size (in numbers of unit cells). For both models the kmos performance is essentially independent of the lattice size in the size range relevant for catalytic applications. The simulation time is instead primarily determined by the model complexity. The benchmarks were carried out on a 3.4 GHz Intel Core i7 processor with 16GB RAM.

therefore provides first insight into the performance dependence of kmos on the number of Conditions. Specifically, the CO oxidation at RuO₂(110) model includes two different active sites per surface unit cell, and a total of 26 inequivalent elementary processes (unimolecular CO adsorption and desorption, dissociative adsorption and associative desorption of O₂, CO and O diffusion, as well as CO oxidation and CO₂ decomposition).[157, 158] The PdO(101) model includes the same types of elementary reactions and also two different active sites per unit cell. In addition, it accounts for nearest-neighbor lateral interactions that modify the rate constants of all diffusion, desorption and reaction steps.

Figure 4.8 shows the CPU time required to execute 1 million kMC steps for both models, again calculated on the 3.4 GHz Intel Core i7 with 16GB RAM benchmark system. Summarized is the scaling up to a maximum system size comprising 10⁵ lattice sites, which is already much larger than the 10²–10³ lattice sites on which these models were reliably evaluated in the original publications. In both cases the runtime is practically independent of the lattice size, confirming the scaling considerations made in Section 4.2.3. The moderate increase is presumably due to a less efficient utilization of the processor cache. Memory limitations eventually also determine the maximum system sizes that kmos can currently handle (outside the size range shown). The runtime is instead critically determined by the system complexity, and in particular by the number of Conditions implied by the model. Even though the RuO₂(110) model contains a larger number

of elementary reactions than the ZGB model, the CPU time per million kMC steps is thus almost the same (0.5 sec). In contrast, the pairwise lateral interactions in the PdO(101) model and the concomitant number of Conditions increase this CPU time by a factor of ~ 25 .

4.4.3 Random Models




Figure 4.9: Single-core CPU times required to execute one million kMC steps for various random models on the standard 3.4 GHz Intel Core i7 benchmark processor with 16GB RAM. Each panel shows the dependence along one model parameter (N_{sites} , N_{species} , $N_{\text{Condition}}$ and N_{react}). Continuous lines connect simulation results obtained for random models in which all other parameters are identical, *i.e.* in the N_{react} panel each line represents the runtime dependence on N_{react} for a constant set of N_{sites} , N_{species} and $N_{\text{Condition}}$.

To further investigate the performance dependence on the model complexity we finally consider random models with varying number of active sites per unit cell, number of possible reaction intermediates (species), number of Conditions per elementary reaction, and number of elementary reactions using the moderate-code-size generator. That is, first N_{sites} sites are initialized (`site1`, `site2`, ...). Next,

N_{species} are initialized (`species1`, `species2`, ..). Using these ingredients we construct N_{react} times a pair of elementary reactions: A forward reaction which consists of $N_{\text{Condition}}$ Conditions with the default species empty on $N_{\text{Condition}}$ s random sites within a finite cut-off radius and corresponding Actions on these same sites with random species. The corresponding backward reaction uses the Actions of the forward reaction as Conditions and uses empty on these same sites as Conditions. By creating all elementary reactions in such pairs we automatically prevent dead-lock configurations in which no events are available. All elementary reactions have the same constant rate constant, and in all cases, the simulated lattice size was (20×20) unit cells, as the preceding sections have shown that the performance scaling with model complexity is independent of the system size.

Using each combination of $N_{\text{sites}} \in [1, 5, 10]$, $N_{\text{species}} \in [2, 5, 10]$, $N_{\text{Condition}} \in [1, 2, \dots, 10]$, and $N_{\text{react}} \in [1, 5, 10]$, we evaluate the single CPU time to execute 1 million kMC steps as in the preceding subsections. Figure 4.9 compiles the obtained results, *i.e.* the dependence on each model dimension. To further analyze the obtained dependencies the obtained runtimes are fitted to

$$t \propto (N_{\text{sites}})^a \times (N_{\text{species}})^b \times (N_{\text{react}})^c \times (N_{\text{Condition}})^d \quad ,$$

yielding $a \approx -0.99$, $b \approx -0.07$, $c \approx 1.24$, and $d \approx 2.00$. This shows empirically that the runtime depends approximately quadratically on the number of Conditions per elementary step. Furthermore it demonstrates that the runtime is basically independent of N_{species} and slightly above linear with N_{react} , confirming the observations made above with the first-principles kMC models. Last, it reveals the seemingly paradoxical result that the runtime decreases with N_{sites} . This can be rationalized by the fact that for a fixed number of elementary reactions N_{react} the probability that different events enable or disable each other shrink with increasing content in the unit-cell. This leads on average to fewer add or delete operations to the set of available events and concomitantly to decreasing runtimes. We stress though that this dependence is of little relevance for physically motivated kMC models, since there the number of elementary reactions N_{react} is expected to grow at least linearly with the number of different active sites N_{sites} . In practice, kMC models will also exhibit a different number of Conditions for each elementary reaction. As such, the benchmark results obtained for the random models should not be taken too literally. Nevertheless, they should convey a useful rough orientation for the to-be-expected runtimes of real kMC models featuring corresponding numbers of sites, species, and elementary reactions, as well as average number of Conditions per reaction.

Most centrally, the results obtained with the random models underscore that the number of Conditions is the most critical property in terms of runtime. This is not critical for model complexities currently

addressed, in particular in the context of first-principles kMC simulations of surface catalysis. Notwithstanding, if eventually more than 5–6 reaction intermediates over multiple active sites and with extensive lateral interactions need to be handled, this will change – and the current `moderate_code_size` code generating algorithm might also reach the capabilities of current compilers. Long-term systematic improvements of `kmos` and its efficiency are therefore best spent on this aspect and in particular the binary decision tree to group the queries checking on the interdependencies between different `Conditions`.

4.5 SUMMARY

We have presented the open source [50] package `kmos`, which offers a versatile software framework for efficient lattice kMC simulations, in particular in surface catalysis. `kmos` can handle site-specific reaction networks of arbitrary complexity in one- to three-dimensional lattice systems, involving multiple active sites in periodic or aperiodic arrangements, as well as site-resolved pairwise and higher-order lateral interactions. For the kMC model definition `kmos` offers an extended application programming interface. On the basis of this model definition, a code generator creates an optimized low-level implementation of the main efficiency driver of a VSSM-based kMC code, the local update procedure that determines the disabled and enabled events after the execution of each kMC step. Together with a well designed data structure, this leads to an efficient kMC solver the runtime performance of which is essentially independent of the lattice size. Instead, the runtime sensitively depends on the model complexity and there in particular on the number of `Conditions` implied by the elementary reactions. For the complexity of reaction networks currently perceivable in the surface catalytic context this is not critical. Should higher efficiency eventually be required, improvements to this end and the code generation algorithm either through improved binary decision trees or parallelization strategies could become of interest.

Next to the efficiency, `kmos` other core objective is a most user-friendly implementation, execution, and evaluation of lattice kMC simulations. For this the API allows to control all runtime aspects interactively, through scripts or via a basic graphical user interface. Enhancing the reproducibility and reusability of the kMC models through a standard file format, `kmos` is thus hoped to contribute to a further, wide-spread use of the kMC approach by an extending user community.

TRANSITION STATE SEARCHES

In the previous chapter we have learned about kMC as an efficient method to solve very large Master equations. An essential input that constitutes these Master equations are the (real) matrix elements of the transition matrix w_{uv} (cf. eq.4.1). What we have not covered so far is how to obtain those matrix elements from the physical system under study. The common approach here is presently (harmonic) transition state theory (hTST)[42, 67, 105]. All input required for hTST such as energies and vibrational frequencies can be obtained directly from the potential energy surface of the atomistic model using density functional theory as introduced in Chapter 2.

To apply transition state theory the most decisive parameter is the energy of the so-called transition state (vide infra). Though calculating these transition states for surface systems comes with many practical issues that make transition states expensive to find. In fact experience shows that converging to accurate transition states is computationally so expensive that it can be considered currently the biggest (or thinnest) bottle-neck in pursuing first-principles multi-scale simulations of heterogeneous catalysis.

Therefore this chapter serves both as methodological introduction to the transition state finding methods used throughout this thesis as well as an assessment of current challenges and possible improvements.

5.1 THEORETICAL BACKGROUND

5.1.1 (*harmonic*) Transition State Theory

To calculate the rate constant k_a of an elementary process a (which then constitutes the matrix elements w_{uv} cf. eq.4.1, 4.3, 4.4) from features of the potential energy a set of concepts and approximations named transition state theory is deployed. Transition state theory aims to answer the following question: Given a system is in a local minimum called initial state, how long does it take on average to leave this local minimum towards another local minimum called the final state? A prototypical potential energy surface for this problem is depicted in Fig. 5.1.

Between the initial state and the final state one can define a dividing (hyper-)surface consisting of points r_c that all reaction paths connecting initial and final state have to cross. The dividing surface is constructed in such a way that if the system is on either side of it, it




Figure 5.1: Schematic depiction of a prototypical potential energy surface studied by transition state theory. Let the system be in a local minimum with energy E_{initial} . It then has to overcome a so-called transition state to reach the final state with energy E_{final} . The transition state is defined as the energetically highest point along the transition path connecting initial and final state with the lowest maximum energy. This point along the so-called minimum energy path is by construction a saddle point of the potential energy surface $V(\mathbf{r})$.

would relax back into either the initial state or final state. An expression for the frequency of crossing can be calculated from statistical mechanics using the flux-over-population method[43, 67]

$$k^{\text{TST}} = \frac{\langle \delta(\mathbf{r}_c) \Theta(\dot{\mathbf{r}}) \rangle}{\langle \Theta(\mathbf{r}_c - \mathbf{r}) \rangle} \quad (5.1)$$

which describes the ratio of states crossing the dividing surface away from the initial state to the number of states in the initial basin. Implicitly this assumes that any system that has crossed the dividing surface will not recross the dividing surface before thermalizing in the final state basin. One can evaluate this expression by Taylor expanding the potential energy surface up to quadratic order and after performing some arithmetic one arrives at

$$k^{\text{hTST}} = \frac{\prod_{i=1}^{3N} \omega_i^{\text{initial}}}{\prod_{i=1}^{3N-1} \omega_i^{\text{TST}}} \exp(-\beta(E_{\text{TST}} - E_{\text{initial}})). \quad (5.2)$$

This is the central expression of classical harmonic transition state theory. The $\omega_i^{\text{initial}}$ and ω_i^{final} are the vibrational modes of the system at the initial state and the transition state, respectively, and $\beta = 1/k_B T$ is the inverse temperature. As one can see the Arrhenius like dependence on the energy difference between the transition state and the initial state has a large impact on the result of this expression. Determining the transition state and its energy is therefore quintessential for predicting rate constants. This will be the objective of the remainder of this chapter.

5.1.2 *Transition state search*

At present the Climbing Image Nudged Elastic Band (CI-NEB) method is the predominant approach for transition state searches when both initial state and final state are known. The original Nudged Elastic Band (NEB) method[95] has undergone several evolutionary steps to either cure common failures resulting in sub-optimal convergence[76] or to enable it to find the exact transition state[77].

Despite very widespread use the method contains several key technical parameters, namely the spring constant, the number of images, and the precision of the forces, which have to be chosen carefully in every application. Common practice here is to iterate these parameters on a trial-and-error basis until the result is 'good enough' and thereafter use these parameters for production. Guidance for choosing the parameters arises therefore from experiences in particular cases and a common body of knowledge is only at the beginning of being established.

Current suggestions for improving the performance of NEB include 'tricks-of-the-trade' like the switch between optimizers depending on the degree of convergence or to use a so called two-step procedure[135] which involves switching between alternative methods such as NEB and the dimer method[75].

Here we contribute to these efforts by systematically gathering statistics over a previously established set of reactions to analyze the detailed influence of the technical settings. An emphasis is placed on the influence of finite noise of calculated forces that is inherent to all electronic structure calculations involving an SCF procedure.

Last, we introduce a modified optimization protocol for the CI-NEB algorithm as well as modified force-based convergence criterion which shows small improvements in performance across a wide range of technical settings and could generally speed up expensive transition state searches on the order of 10 %.

5.2 METHODS

5.2.1 *Transition State Finding Methods*

5.2.1.1 *Drag Method*

A very intuitive method for determining a transition state which is referred to under many different names in the literature is the drag method. The idea is to start from an initial geometry and a fixed collective coordinate. The system is then propagated by a small increment along the collective coordinate and the new geometry is later optimized along all other coordinates. During the optimization the collective coordinate is held fixed. Since we start in a minimum of the potential energy surface the total energy of the optimized geometry is

expected to increase between propagation steps until a transition state or saddle point is passed. Therefore the propagation and relaxation cycle is repeated until the resulting energies decrease again. Once such a point has been found the location of the transition state can be refined by repeatedly moving the collective coordinate between the energetically highest optimized points until a given threshold has been identified.[78]

5.2.1.2 *Climbing Image Nudged Elastic Band*

Climbing image NEB (CI-NEB) as introduced by Henkelman *et al.*[77] optimizes an array of atomic configurations approximating the minimum energy path (MEP) between initial and final state.[95] In order to keep these so-called images roughly equidistant along the MEP neighboring images are artificially coupled by a harmonic potential. The forces of the combined system (true atomic forces + spring forces) are projected in two ways. All forces due to springs orthogonal to the NEB string are projected out to minimize corner cutting.[76] All forces due to atomic interaction parallel to the NEB string are projected out to minimize down-sliding of images towards initial and final state. The tangent of the NEB string at each image is estimated by the line segment connecting the image in question and the neighboring with the higher total energy and for the image with the largest total energy the average of the line segments to the two neighboring images is used.[76] The atoms are then moved to minimize these projected forces using a geometry optimization algorithm until a preset threshold value for all projected forces is reached.

Finally in order to converge exactly to the saddle point connecting initial and final state the atomic forces along the NEB string are inverted for the image with the highest energy (the climbing image) such that it converges to the maximum energy configuration along the Minimum Energy Path (MEP).[77]

5.2.1.3 *Freezing string*

The Freezing String (FS) method was introduced more recently by Behn *et al.*[6] as another transition state finding method. FS starts like NEB from an initial and a final configuration differing by a line segment τ . The line segment has a length $\tau = |\tau|$. The transition state is approximated by a series of iteration steps.

Each iteration step consists of an interpolation and an optimization part. For the interpolation a new pair of images is generated by moving the existing configurations closer to one another along the line segment τ by a distance $d = \tau/N$. Here N is a parameter controlling the number of iteration steps and therefore the computational cost and accuracy. For the optimization both images are optimized with all force components parallel to the line segment τ projected out. That

is the images are allowed to relax in all directions perpendicular to the line segment τ . Then τ is updated by calculating the difference between the current configuration of the current pair of nodes.

This iterative procedure is repeated until the distance between the two innermost nodes is less than the interpolation step length d . From there more refined searches can be started either using the FS method with an even smaller step length or for example with the dimer method[75].

5.2.1.4 *Freezing NEB*

Our modification of NEB, coined freezing Nudged Elastic Band (fNEB), starts by linearly interpolating an initial and final state using N images just like NEB. Next though it only optimizes the two outermost images like FS (one step in from the endpoint) using the same forces as in a standard NEB, while keeping all other images in the original position.¹ After one pair of images is optimized the remaining $N - 2$ images are linearly interpolated between the optimized images and the algorithm continues with the next inner neighbors. This procedure is repeated until the center image is reached. Throughout N may be even or odd. The climbing image projection is applied whenever the image to be optimized has a total energy which is higher than the total energy of both neighboring images. Instead of completely converging each image to the preset threshold, it has been found that better total performance is achieved if each image is only optimized until either the threshold or a maximum number of optimization steps (say 20) is reached. Therefore once the center image is reached the 'outside-in' procedure is repeated until all images are below the convergence threshold.

We note that a different "family" of transition state finding methods less often used for surface chemical reactions[21] are the string methods[35] that optimize a chain of states between initial and final, but instead of using spring forces the images are directly redistributed according to some metric. Here the growing string method[147] has emerged as a version that does not start with one complete initial guess but alternately places and optimizes images starting from the outer ends. In this sense fNEB can be seen as the growing string equivalent within the NEB methods.

5.2.2 *Optimization Methods*

The NEB procedure (and fNEB) can in principle be combined with any local optimization algorithm. Among the most common algorithms in surface chemistry are Steepest-descent, Quick-min, Fast In-

¹ One evaluation of the next inner neighbor is necessary to apply the energy guided tangent estimate as in NEB. This evaluation is not lost in the computational sense though as it serves as the first force evaluation in the next iteration step

ertial Relaxation Engine (FIRE)[7], Conjugate Gradient[79], or Limited Memory Broyden-Fletcher-Goldfarb-Shanno (LBFGS)[18, 48] without line search LBFGS(Hess). Their performance in the NEB context has already been benchmarked by Sheppard *et al.*[178], using the same reference system employed below. FIRE and LBFGS(Hess) have been found to be the two fastest overall algorithms, which is why we concentrate our assessment on these two approaches. A short introduction of the Broyden-Fletcher-Goldfarb-Shanno (BFGS) approach is given in Appendix A. Specifically, it has been found that LBFGS(Hess) performs best if it acts on the $[3 \times N \times P]$ dimensional space, where N is the number of images and P the number of atoms per image, rather than acting on N separate $[3 \times P]$ dimensional spaces of the individual images. The optimization method has therefore been coined *global* LBFGS and we consistently use LBFGS in this sense.

5.2.3 Benchmark Problem

The benchmark problem chosen here are 13 low-lying transitions all starting from a close-packed seven atom island at the (111) surface of an fcc crystal. This problem has been devised by Henkelman *et al.*[78] before as a benchmark system, and investigated further by Sheppard *et al.*[178] to compare the performance of different optimization methods. In this chapter we use the identical initial and final state geometries as they are publicly available². As in the two previous studies[78, 178] the (111) surface is modeled by a $[7 \times 8 \times 6]$ slab, *i.e.* 6 layers containing (7×8) atoms in the surface unit-cell each with the three lowest layers fixed and using periodic boundary conditions in the horizontal directions. The atomic interactions of all atoms are described by using a pairwise Morse potential

$$V(r) = D_e [e^{-2\alpha(r-r_0)} - 2e^{-\alpha(r-r_0)}]$$

where $D_e = 0.7102$ eV, $\alpha = 1.6047 \text{ \AA}^{-1}$, and $r_0 = 2.8970 \text{ \AA}$ (supposedly describing the interaction between Pt atoms). The potential is truncated and set to zero at a distance of $r = 9.5 \text{ \AA}$.

Even though the transferability of the insights and optimized settings obtained for this benchmark system is of course unclear, we emphasize that the set of 13 transition states comprises fairly different kinds of rearrangements: some move the entire island in one concerted movement, some move different atoms in opposite directions, and some even move only one or two atoms from the remaining island atoms. We therefore expect at least some transferability to other systems, geometries, and potentials.

All transition state finding calculations converged to identical energy barriers for each transition except for a few cases in the *number of images* benchmark where NEB and fNEB strings using only 2 or

² <http://theory.cm.utexas.edu/henkelman/research/saddle/benchmark.php>

3 images did not converge within on the order of 10^5 steps as will be further described below. All transition states were confirmed by vibrational analysis of the seven atoms of the island yielding exactly one imaginary frequency.

5.2.4 Implementation/Technical Section

In this chapter we use NEB as implemented in the Atomic Simulation Environment (ASE)[3]. The fNEB implementation is built on top of it. We use the implementation of the Morse potential via the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS)[149] interface of ASE with the potential parameters stated above. It has been verified that the energy barriers obtained in this way are identical to use those found by Henkelman *et al.*[78]. For all CI-NEB calculations the climbing image projection was enabled from the first step of the convergence to simplify the protocol. In a scenario with more complicated reaction paths we note that it is common practice to only enable the climbing image protocol after a certain force threshold has been satisfied. The number of images N stated for NEB and fNEB calculations does not include the initial and final state, but only refers to movable images.

Throughout this study the default parameters for the geometry optimizers as given in the ASE framework have been used, which means for LBFGS an initial guess of the inverse Hessian of $0.1 \text{ \AA}^2/\text{eV}$ and maximum step size of 0.04 \AA and for FIRE $N_{\text{min}} = 5$, $f_{\text{inc}} = 1.1$, $f_{\text{dec}} = 0.5$, $\alpha_{\text{start}} = 0.1$, and $f_{\alpha} = 0.99$, which are identical to those settings used by Bitzek *et al.* when introducing the FIRE method.[7]

In every geometry optimization a convergence threshold for which the calculation is considered converged needs to be fixed. The threshold needs to be small enough to make the remaining uncertainty small compared to the uncertainty in the potential, yet not too small to defer convergence due to finite numerical reproducibility of SCF calculations. A common choice for transition state geometries of chemical reactions at surfaces is typically between $0.01\text{-}0.05 \text{ eV/\AA}$. Throughout this chapter a convergence threshold of 0.01 eV/\AA maximum atomic force for each image is used.

One additional aspect considered in this chapter is the influence of a random perturbation of the exact potential forces due to not strictly converged forces e.g. in ab-initio calculations. The SCF procedure, being an iterative method, improves approximations to the exact ground state solution of the, e.g. Kohn-Sham equations, in every iteration step. Since the computational effort per step is costly though one usually performs the SCF procedure only until a certain preset threshold of change between successive evaluations has been reached. Therefore the resulting forces contain a bound but random error with respect to the exact result of the potential. Here this ef-

fect is simulated by modifying the Cartesian forces of the analytical potential such that

$$\mathbf{F}_{\text{noisy}} = \mathbf{F}_{\text{exact}} + \nu \mathbf{R},$$

where \mathbf{R} is a matrix with identical shape as \mathbf{F} and each entry is a random number in the half-open interval $[-0.5, 0.5)$ and ν is a control parameter to set the amplitude of the random noise. In other words a random number from a fixed interval is added on every component of the atomic forces in every evaluation.

5.3 RESULTS

5.3.1 Spring Constant k

In NEB and fNEB calculations, the distance between neighboring images is controlled by applying a fictitious spring force between corresponding atoms. Let $\hat{\mathbf{t}}$ be the tangent between two neighboring image configurations. The NEB force $\mathbf{F}_i = -k(|\mathbf{R}_{i+1} - \mathbf{R}_i| - |\mathbf{R}_i - \mathbf{R}_{i-1}|)\hat{\mathbf{t}}$ then contains a spring constant k that has to be chosen empirically. In principle, this spring constant may be varied for different images, though for simplicity we study the force evaluations required as a function of *one global* k only.

Fig. 5.2(a) shows the number of force evaluations averaged over the 13 transitions as a function of the spring constant k using 8 images and no numerical noise. In all cases the force evaluation only depends weakly on k for small k until it reaches about $0.05 \text{ eV}/\text{\AA}^2$, which is when stiffer spring constants start to deteriorate fast convergence. Comparing between the different optimizers LBFGS outperforms FIRE, as has been concluded before by Sheppard *et al.*[178]. For each optimizer the fNEB method outperforms the NEB method by a few steps at least for lower k . In light of this data we choose a spring constant of $0.02 \text{ eV}/\text{\AA}^{-2}$ for the remainder of this chapter, except where explicitly stated otherwise.

5.3.2 Number of Images

Fig. 5.3(a) shows the number of force evaluations per image as a function of the number of images in the string. All methods display unstable behavior for less than 4 images. Some transitions did not even converge after 10^5 iteration steps for both NEB and fNEB, which clearly highlights the risk of using too few images.

Between 5 and 10 images all methods except fNEB/FIRE show more or less similarly stable convergence behavior. Interestingly, too many images take longer even when normalized to the number of images. It is thus worth pointing out that adding more images than

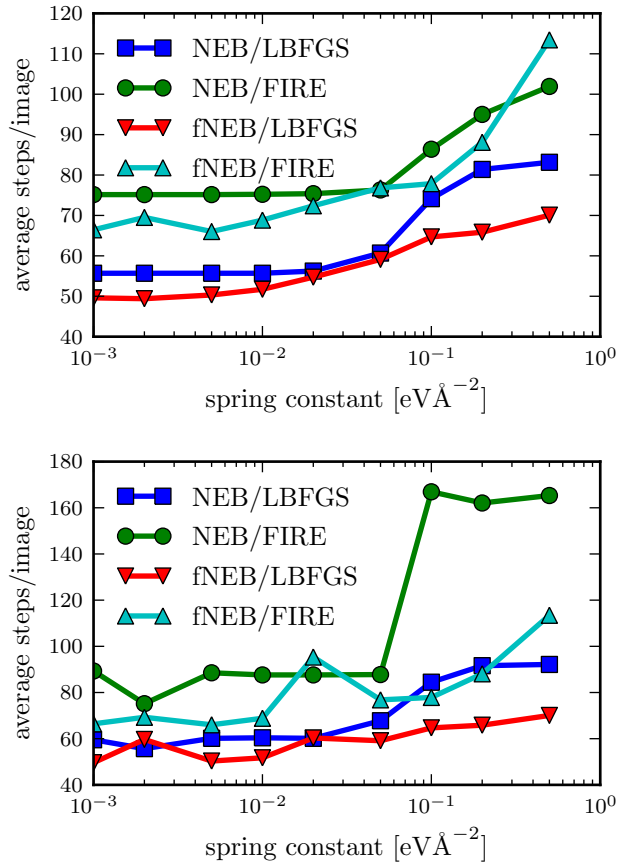


Figure 5.2: The average number of force evaluations per image plotted as a function of the spring constant. The convergence threshold is $0.01 \text{ eV}/\text{\AA}$, the NEB string contains 8 images not counting the initial and final image. No noise is used in the upper figure and a random noise of $0.005 \text{ eV}\text{\AA}^{-2}$ is added in the lower figure. The spike in the fNEB/FIRE curve is due to transition 7. The fNEB converges into slightly different trajectories with yet the same transition state causing slow convergence.

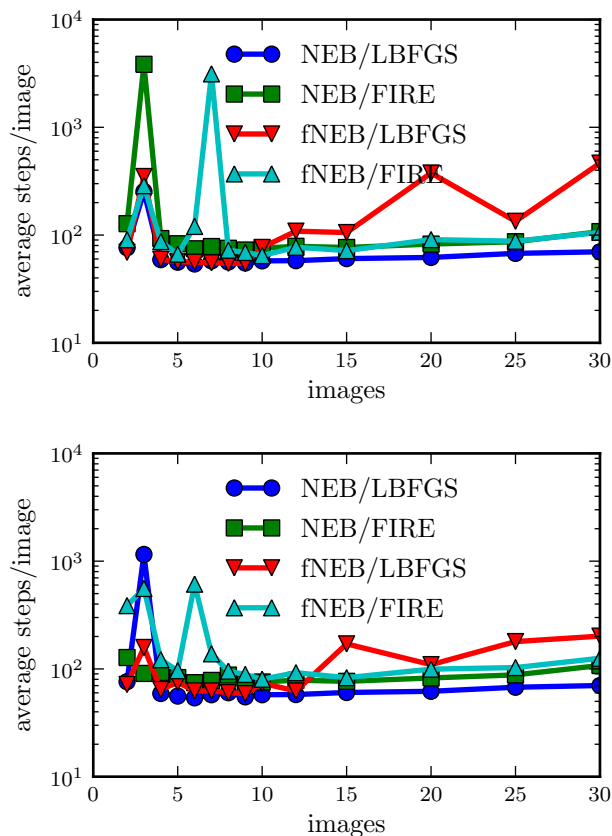


Figure 5.3: Average number of force evaluations as a function of number of images of the NEB string using $k = 0.02 \text{ eV \AA}^{-2}$. No noise is used in the upper figure and a random noise of $0.005 \text{ eV \AA}^{-2}$ in the lower figure. The spike for the fNEB/FIRE combination for 7 images and noiseless forces and 6 images for noisy forces is caused by transition #7. The minimum energy path of this transition involves movement of nearby atoms into opposite directions. Also not all optimizations for 2, 3, and 25 images where converged successfully.

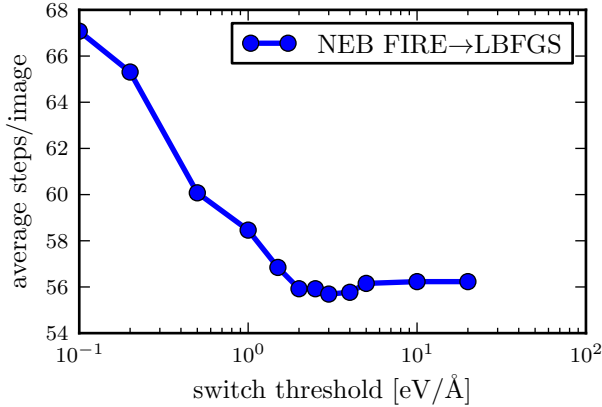


Figure 5.4: The average number of force evaluations per image plotted as a function of the force threshold to switch from the LBFGS to the FIRE optimizer. The final convergence threshold is $0.01 \text{ eV}/\text{\AA}$ and the NEB string contains 8 images not counting the initial and final image. The spring constant is set to $k = 0.02 \text{ eV } \text{\AA}^{-2}$. No noise is added to the Morse potential forces.

required for stable convergence neither improves the overall computational efficiency nor the numerical precision of the resulting transition state.

Critical about this graph is the spike using the fNEB/FIRE combination for 6 and 7 images, respectively. In both cases this slow convergence is caused by transition #7 (cf. Henkelman *et al.*[78]). The MEP of this transition prominently features movement of neighboring atoms in the island into approximately opposite directions. This jumpy behavior shows how crucial it is to choose the correct number of images but that taking more images does not necessarily improve convergence.

5.3.3 Switching from FIRE to LBFGS

The LBFGS optimizer tends to display slower convergence when the target function is far from a local minimum. This can be readily explained through the fact that LBFGS builds upon the assumption that the PES is harmonic along all coordinates. However, if the initial guess is poor this is not necessarily true and therefore one common approach is to use the FIRE optimizer until a certain force threshold has been met and apply the LBFGS optimizer to this partially optimized structure until the final threshold is met. In order to analyze the possible gain of this technique as well as the optimal magnitude of forces to do the switch, we have optimized the same 13 transition states using the NEB method and using the force threshold when to switch from FIRE to LBFGS as a parameter.

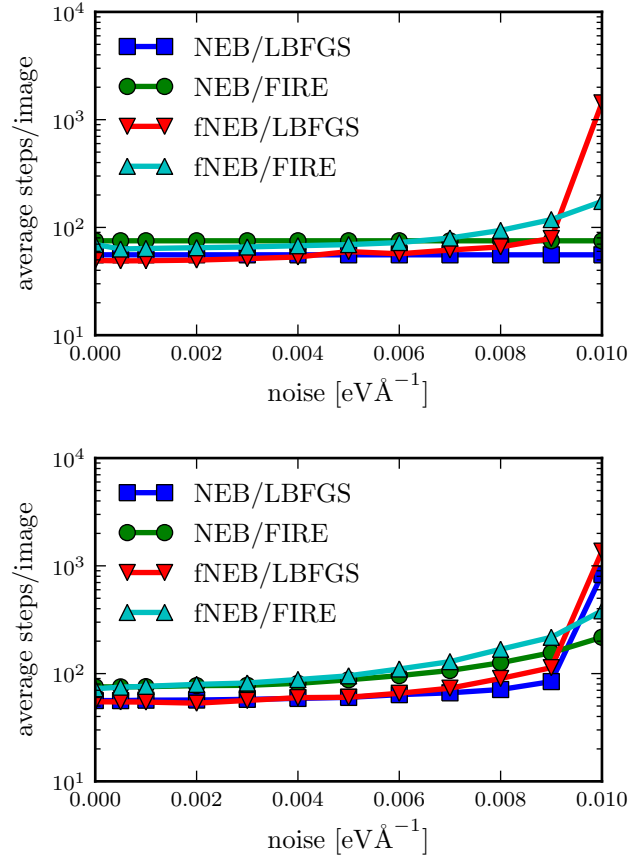


Figure 5.5: Average number of force evaluations as a function of the maximum magnitude of the random noise added onto the evaluated forces. 8 images are inserted between the initial and the final states. The spring constant is $k = 0.002 \text{ eV } \text{Å}^{-2}$ in the upper figure and $k = 0.02 \text{ eV } \text{Å}^{-2}$ in the lower figure.

Fig. 5.4 shows the resulting average number of optimization steps per image as function of the switch threshold for settings that have been identified as optimal in preceding sections III A-C. The curve shows that switching from FIRE to LBFGS below $1 \text{ eV}/\text{Å}$ annihilates any gain and the number of force evaluations quickly approaches the computational effort of applying the FIRE optimizer only. Above $5 \text{ eV}/\text{Å}$ as threshold the curve also levels out since most transitions never display forces above that value even for the linear guess used as starting point here.

5.3.4 Precision of Forces

Due to the SCF procedure used in most ab-initio potentials the exact solution to the eigenvalue problem is only approximated and the remaining uncertainty decreases monotonically with each SCF step. Thus even though in principle the numerical error of forces can be

reduced down to machine precision, the computational effort grows correspondingly. Since the only information extracted from a NEB calculation are the final energies of the converged NEB string and not the energies or forces of intermediate NEB steps, increasing the precision of those forces is only justified as long as it reduces the total time required to find the final NEB configuration. Fig. 5.5(a) shows the number of force evaluations as a function of the random noise added on top of the exact potential forces. For all combinations of methods and optimizers the number of steps increases monotonically but fairly moderately with the magnitude of random noise as long as the noise is small compared to the convergence threshold (here $0.01 \text{ eV}/\text{\AA}$). This changes abruptly when the maximum noise amplitude reaches half of the convergence threshold, which for the fNEB/LBFGS combination means about one order of magnitude more steps to reach convergence. The fNEB/FIRE combination increases roughly by a factor 3. Instead, the combinations NEB/FIRE and NEB/LBFGS are very stable and are basically not affected by this level of noise. Fig. 5.5(b) shows the equivalent data but using a spring constant of $k = 0.02 \text{ eV}/\text{\AA}$. Under these settings all transition state searches are considerably more sensitive to noise though. Once the noise reaches half of the convergence criterion the optimizer becomes more distinguishing than the search method. In all cases the LBFGS based optimizations slow down by about one order of magnitude over the noise level range tested, while the FIRE based optimizations increase only by a factor 3–5. Within each group of optimizers the NEB method copes better with noise than the fNEB method.

In light of the robustness of the NEB method with regards to a finite noise the question whether this could be exploited to save compute time was raised.³ Current *ab initio* codes use a combination of fixed thresholds for change in total energy, eigenvalues, and maximum change in forces between subsequent SCF steps to decide whether the electronic structure is converged. In particular for plane-wave calculations the cost of calculating Hellmann-Feynman forces is almost negligible. Therefore one could relax the total energy criterion and only apply a force based criterion. Furthermore the force criterion could be adapted to the level of NEB convergence. The idea is that during initial NEB steps when the magnitudes of forces are quite large a fixed but random error should have less influence on the next NEB step than if the overall forces are small. We therefore propose a modified convergence criterion: the maximum force of the current NEB iteration should be below a fixed fraction of the maximum force in the previous NEB iteration. In the first NEB iteration no previous forces exists. In order to bootstrap the protocol we set the previous force magnitude initially to $1 \text{ eV}/\text{\AA}$. Experience shows that this is a quite conservative estimate of the initial gradient.

³ Karsten Reuter (*private communication*)

Given the overall cost of first-principles NEB calculations this idea was only tested on a small test problem of oxygen diffusion on Pd(100) between nearest neighbor hollow sites. A NEB calculation was performed using 5 images and a spring constant of $k = 0.1 \text{ eV/\AA}$. A reference calculation was performed with a total energy threshold on each DFT calculation of $1 \times 10^{-6} \text{ eV}$. All calculations combined (including the vibrational analysis) required 1043 SCF steps. The alternative method using the relative force-based criterion with convergence threshold of 3% of the maximum force in the previous steps required (including vibrational analysis) 846 SCF steps or a relatively difference of more than 18%.

5.4 CONCLUSIONS

Thorough comparison of fNEB and NEB using FIRE and LBFGS as an optimizer on a set of 13 benchmark transitions illustrates how critical it is to choose the right technical settings, with performance variations over the employed benchmark problem of the order of 1–2 orders of magnitude. In particular, the optimal number of images can be difficult to predict correctly. Using more than 10 images only increases the computational effort proportional to the number of images, while using less than 4 images dramatically increases the risk of not converging calculations.

The spring constant k has little effect on the computational time below a certain threshold. It is yet to be established how this threshold relates exactly to features of the potential energy surface. A very small spring constant causes slower convergence only in combination with random noise.

A random perturbation on the forces increases the steps required for convergence roughly linearly with the amplitude of the perturbation, when the latter is between approximately 5% and 50% of the convergence threshold. Beyond that the convergence time increases dramatically faster. Smaller perturbations have no measurable effect on the convergence. This observation can be used to further fine-tune the relaxation criteria for the ab initio calculations. In the case of plane-wave calculations and its computationally cheap force evaluations it could be more efficient to define convergence based on forces relative to the maximum force in the previous NEB iteration.

Another direction for improvement (but considerably more complex) could be the observation that in NEB currently intermediate iteration steps serve no other purpose than to produce the next optimization step even though intermediate steps bear the same numerical accuracy and energy as the final converged step. If one could come up with a way to somehow recycle the previous optimization steps that are located increasingly close to the MEP, e.g. by automatically building up a force field, fewer steps might be needed.

The presented fNEB method is a straightforward combination of the NEB using a key advantage of the FS method. For the test set it reduces the number of required force evaluations on the order 10 % if used under optimal settings. Though in particular the benchmark using a different number of images proves that it is no 'silver bullet' solution and some transition paths are particularly prone to steer fNEB into wrong directions. An even larger test set of well documented benchmark reactions is urgently needed to fully evaluate and establish optimal transition state finding methods.

Part II

CO OXIDATION FROM $\sqrt{5}$ -OXIDE TO PD(100)

(BI)STABILITY OF Pd(100) AND $\sqrt{5}$ PHASE

This chapter predicts the stability ranges of both the $\sqrt{5}$ surface oxide termination and the metallic Pd(100) termination by evaluating two separate lattice kinetic Monte Carlo models and finds that the two stable regimes overlap significantly in $(T, \{p_i\})$ phase space. The main work of this chapter has been published *Topics in Catalysis* [83]. The herein established Pd(100) 1p-kMC model has additionally been used to interpret experimental data in a publication in *Physical Review Letters*[10].

6.1 INTRODUCTION

Common use of Pd for catalytic exhaust gas purification [52] has motivated frequent studies on Pd single-crystal model catalysts to elucidate the underlying molecular-level mechanisms. Notwithstanding, despite significant efforts the surface structure and composition at *near-ambient* conditions remains unclear. This is largely due to inherent difficulties in achieving atomic-scale information/resolution in this technologically relevant regime of near-ambient pressures and temperatures of 300–800 K. Generally, the outcome of such experiments seems strongly dependent on the exact preparation conditions and experimental setup. Specifically, there exists a longstanding and controversial debate on whether reactivity is due to the pristine metal surface[56], a formed surface oxide film[189], or even a thicker bulk-like oxide overlayer.[80]

Under ultra-high vacuum (UHV) conditions, different oxidation stages of Pd(100) [206] have been thoroughly characterized both experimentally and theoretically. By using low-energy electron diffraction (LEED) Chang and Thiel first identified five distinct ordered structures before the onset of bulk oxide formation: a $p(2 \times 2)$ oxygen adlayer, a $c(2 \times 2)$ adlayer, a (5×5) , and a $(\sqrt{5} \times \sqrt{5})R27^\circ$ (for brevity henceforth $\sqrt{5}$) reconstruction.[23] Todorova *et al.*[184] established that the latter $\sqrt{5}$ structure corresponds to a single-layer of PdO(101) on top of Pd(100).

At low pressures several experiments indicate that formation of this surface oxide is accompanied with a low activity of the catalyst.[51, 100] This is primarily attributed to the low CO binding energy at the surface oxide [206], which Gao *et al.* estimated to be around 0.5–0.6 eV.[55] At elevated pressures the situation is less clear. Lundgren *et al.*[116] compared structural information from *in situ* surface x-ray diffraction measurements in a pure oxygen environment to a

systematic *ab initio* surface phase diagram [163, 164], and concluded that the formation of bulk oxide on Pd(100) is kinetically severely hindered even at temperatures around 675 K. This suggests that at the (T, p) -conditions relevant for near-ambient CO oxidation the most likely surface terminations are either pristine Pd(100) with some coverage of O adsorbates, or a monolayer-thin film of $\sqrt{5}$. On top of this, reactor scanning tunneling microscopy (STM) measurements by Hendriksen *et al.* in fact suggested bistability in this near-ambient regime. [70, 73, 74] In the understanding of the work by Lundgren *et al.* this would thus translate into a range of (T, p) conditions where both O@Pd(100) and the $\sqrt{5}$ surface oxide are metastable.

Here, we investigate this hypothesis with first-principles kinetic Monte Carlo (1p-kMC) simulations that focus on either of the two surface states, *i.e.* either O@Pd(100) or the $\sqrt{5}$ surface oxide, thereby extending previous 1p-kMC work focused exclusively on the $\sqrt{5}$ phase [162, 163]. We indeed find a range of (T, p) -conditions where both models appear metastable on time scales up to seconds. Moreover, in this bistability regime, both models yield roughly similar turnover frequencies (TOFs) for near-ambient pressures, while consistent with experiment the Pd(100) model is the much more reactive one in the UHV regime. This puts experiments [56, 189] into perspective that claim one surface state to be *the* active one, irrespective of the specific gas-phase conditions.

6.2 METHODS

6.2.1 1p-kMC simulations

Our central goal is to describe the stability ranges of either pristine Pd(100) or the $\sqrt{5}$ surface oxide by analyzing the steady-state coverages of CO and O on them as a function of external feed conditions $(T, p_{\text{CO}}, p_{\text{O}_2})$. For this, two factors indicate that standard microkinetic modeling on the level of mean-field rate equations will not be sufficient: On pristine Pd(100) lateral interactions between adsorbed O and CO are known to be rather strong [5, 111, 112, 204], while on the $\sqrt{5}$ surface oxide it is the essentially one-dimensional trench structure (*vide infra*) that will lead to adlayer inhomogeneities. We thus opt for 1p-kMC simulations as presently only technique that provides the desired microkinetic information while fully accounting for the correlations, fluctuations and explicit spatial distributions of the chemicals at the catalyst surface. [154, 165, 182]

In 1p-kMC the time evolution of the system is coarse-grained to the discrete rare-event dynamics. Relying on a Markov approximation rejection-free 1p-kMC algorithms thus generate configuration-to-configuration trajectories that in their average yield the probability density function $P_i(t)$ to find the system at time t in configuration i

representing the corresponding potential energy surface (PES) basin i . The propagation of this probability density is governed by the Markovian master equation,

$$\frac{dP_i(t)}{dt} = - \sum_{j \neq i} w_{ij} P_i(t) + \sum_{j \neq i} w_{ji} P_j(t) \quad , \quad (6.1)$$

where the sums run over all system configurations j and w_{ij} is the rate constant to go from configuration i to j (cf. eq. 4.4). The central first-principles ingredients required for the 1p-kMC simulation are thus the individual rate constants of the considered elementary processes (adsorption, desorption, diffusion, reaction). To keep this input tractable, *i.e.* arrive at a finite number of inequivalent processes and corresponding rate constants, 1p-kMC simulations are commonly performed on lattice models. In the following sections we will first provide the working equations to determine the first-principles rate constants k_{ij} and then detail the specific lattice models employed in the present work for CO oxidation at pristine Pd(100) and at the $\sqrt{5}$ surface oxide.

For given gas-phase conditions (specifying the adsorption rate constants), the output of 1p-kMC simulations are then the detailed surface composition and occurrence of each individual elementary process at any time. Since the latter comprises the surface reaction events, this also gives the catalytic activity in form of products per surface area and time (*i.e.* TOFs), either time-resolved, e.g. during induction, or time-averaged during steady-state operation.

6.2.2 First-principles rate constants

The first-principles rate constants in this work are evaluated following the approach put forward by Reuter and Scheffler [157]. In brief, this approach relies on (harmonic) transition state theory (TST) for bound to bound processes like diffusion, and kinetic gas theory together with detailed balance to calculate adsorption and desorption rate constants. As the approach and its derivation have been detailed before, we here restrict ourselves to the presentation of the working equations for self-containment.

The adsorption rate constant for species i is given by the rate with which these particles impinge on the unit-cell surface area A_{uc} and the local sticking coefficient $\tilde{S}_{st,i}(T)$, which gives the fraction of the impinging particles that actually stick to a given free site st at temperature T

$$k_{st,i}^{ad}(T, p_i) = \tilde{S}_{st,i}(T) \frac{p_i A_{uc}}{\sqrt{2\pi m_i k_B T}} \quad . \quad (6.2)$$

Here, k_B is the Boltzmann constant, p_i the partial pressure of species i , and m_i the particle mass. In unactivated adsorption events the lo-

cal sticking coefficient merely accounts for the number of inequivalent sites in the surface unit-cell (*vide infra*). For activated adsorption events or Eley-Rideal (ER) type CO oxidation events, it is additionally governed by the adsorption resp. reaction barrier $\Delta E_{i,st,j}$,

$$\tilde{S}_{st,i}(T) = \left(\frac{A_{st,i}}{A_{uc}} \right) \exp \left(-\frac{\Delta E_{i,st,j}}{k_B T} \right) . \quad (6.3)$$

where $A_{st,i}$ is a geometrical factor reflecting the relative share with which molecules impinge on the different inequivalent surface sites st .

Desorption of a particle adsorbed on a surface site st is modeled as the time reversed process of adsorption, and its rate constant thus has to fulfill detailed balance or microscopic reversibility.

$$\begin{aligned} \frac{k_{st,i}^{ad}(T, p_i)}{k_{st,i}^{des}(T)} &= \exp \left(\frac{\Delta G_{st,i}(T, p_i)}{k_B T} \right) \\ &\approx \exp \left(\frac{\mu_{gas,i}(T, p_i) - E_{st,i}^{bind}}{k_B T} \right) \end{aligned} \quad (6.4)$$

where $\Delta G_{st,i}(T, p_i)$ is the difference in Gibbs free energy between the particle adsorbed at the surface state and in the gas phase. This is approximated by the difference between the gas-phase chemical potential $\mu_{gas,i}(T, p_i)$ [155] and the binding energy of the particle in the adsorbed state $E_{st,i}^{bind}$. Following the procedure detailed in Ref. [155] we interpolate tabulated values [24] to determine the gas-phase chemical potentials at any gas-phase condition.

The diffusion rate constant of an adsorbate from one surface site st to another site st' is approximated as

$$k_{st,st'i}^{diff}(T) \approx \left(\frac{k_B T}{h} \right) \exp \left(-\frac{\Delta E_{st,st'i}^{diff}}{k_B T} \right) , \quad (6.5)$$

where $\Delta E_{st,st'i}^{diff}$ is the diffusion barrier. Langmuir-Hinshelwood (LH) type CO oxidation reactions are described with an equivalent expression containing the reaction barrier. They (as well as the ER reactions) are generally treated as associative desorption events though, *i.e.* the formed CO_2 immediately desorbs from the surface and thereby creates two (one) vacant surface site(s) in case of LH (ER).

Within this approach the required first-principles input to determine the rate constants boils down to the binding energies of the species at the surface sites (for the desorption rate constant), as well as to their diffusion and reaction barriers. For the latter barriers the transition states were approximately identified through scans along suitable reaction coordinates. For diffusion barriers the specific reaction coordinate employed was the lateral coordinate along a straight line connecting the known initial and final state. For CO oxidation reactions the distance between the C atom of the CO molecule and




Figure 6.1: (Color online) Top view of Pd(100) and the $\sqrt{5}$ surface oxide, illustrating the employed lattice models. Hollow sites are depicted by yellow circles and bridge sites by green rectangles (see text).

the O adsorbate was employed. We carefully checked for hysteresis effects and estimate the uncertainty in the determined barriers to be of the order of ± 0.1 eV. The effect of this energetic uncertainty on the presented 1p-kMC results will be critically discussed below.

All required total energies were obtained from density-functional theory (DFT) with the generalized gradient approximation functional by Perdew, Becke and Ernzerhof (PBE) [144] to treat electronic exchange and correlation. Using the plane-wave code CASTEP [28] with standard library ultrasoft pseudopotentials systematic convergence tests showed that the quantities of interest (binding energies, diffusion and reaction barriers) are converged to within 30 meV at the employed energy cut-off of 400 eV and k-point density of 0.4 \AA^{-1} . The surfaces were modeled in periodic supercell geometries, containing vacuum separations of more than 10 \AA . For the $\sqrt{5}$ surface oxide the calculations were done within a (1×1) surface unit-cell of surface oxide on top of four layers of Pd(100). For pristine Pd(100) we employed a (2×2) surface unit-cell and also four metal layer slabs. In both cases, all geometries were fully relaxed to residual forces below 50 meV/\AA , while keeping the bottom two slab layers fixed to their bulk positions.

6.2.3 1p-kMC lattice models

6.2.3.1 Pd(100)

In setting up the lattice model for Pd(100) we exploit the limited coverage range, for which the 1p-kMC simulations need to provide a faithful representation. Detailed experimental work [23] indicates the formation of surface oxides at around a critical O coverage of 0.5

monolayer (ML, defined as the ratio of O to top layer Pd atoms). As the purpose of the model is to derive the gas-phase conditions under which *no* surface oxide is formed, the developed model only aims to faithfully reproduce this coverage range $\theta < 0.5$ ML and the elementary processes occurring in it. Similarly, CO coverages of more than 0.5 ML have only been characterized for strongly overstoichiometric CO pressures and low temperatures[5, 180, 186, 187]. As this regime is outside the catalytic context, only CO coverages below 0.5 ML are of interest here, too.

In this dilute coverage range O atoms bind preferably to the four-fold hollow sites at Pd(100) [159, 204], while CO binds to twofold bridge sites [15, 164]. In consequence we set up the lattice model shown in Fig. 6.1, in which O can exclusively occupy hollow sites and CO bridge sites. Lateral interactions between oxygen adatoms have been systematically calculated by Zhang, Blum and Reuter and were found to be strongly repulsive at nearest and moderately repulsive at next-nearest hollow-hollow distance [204]. For the targeted dilute coverage regime, this is taken into account in the 1p-kMC model by blocking adsorption or diffusion events into configurations that would result in O atoms in nearest-neighbor positions. For dissociative adsorption this implies a required motif of eight empty sites (two next-nearest neighbor sites for the actual adsorption and their immediately adjacent six nearest-neighbor sites) known as the 8-site rule[16, 22, 110, 112] that is required for an adsorption event to take place. Similarly repulsive interactions between adsorbed CO molecules are indicated by the experimentally characterized superstructures [5, 180, 186, 187] and DFT calculations [163]. Accordingly and accounting for the shorter bridge-bridge distances, we also block CO adsorption and diffusion events into configurations resulting in CO molecules up to next nearest-neighbor bridge-bridge distance. In this respect, this default model, henceforth denoted as Pd(100)-2NN, contains less geometric detail than the one motivated recently by Liu and Evans [110]. However, by varying the extent of lateral interaction in our model below we show that these details are not central to the bistability question addressed in this work. Specifically, this will be done with a model representing only shorter ranged repulsive interactions, in which CO-CO coadsorption is blocked exclusively at nearest bridge-bridge distance (Pd(100)-1NN).

Finally, the absence of mixed O-CO adsorbate structures at low temperatures [180] suggests O-CO lateral interactions to also be repulsive. Unfortunately, no systematic first-principles studies characterizing these interactions have been performed to date. In this situation we note that the nearest-neighbor hollow-bridge distance at Pd(100) is smaller than the O-CO distance in the CO oxidation transition states described below. O-CO lateral interactions are thus crudely

modeled by blocking adsorption and diffusion events into configurations involving O and CO at such close distance.

Within this lattice model and site-blocking rules we consider all non-concerted O and CO adsorption, desorption, as well as nearest-neighbor site diffusion and reaction events. Oxygen adsorption is modeled as dissociative process into two neighboring vacant hollow sites. CO adsorption occurs unimolecularly into one free bridge site. Both adsorption processes are treated as non-activated, *i.e.* using a local sticking coefficient of 1 for O₂ and 0.5 for CO, the latter accounting for the two bridge sites per surface unit-cell. For oxygen adsorption in the low-coverage range this is supported by explicit DFT sticking coefficient calculations [129], while for CO we simply verified that lifting the CO molecule vertically up from the bridge adsorption site yields a path without adsorption barrier. This together with the observation that a CO lowered above an adsorbed oxygen atom is efficiently steered into a neighboring vacant bridge site, supports the assumption that all CO molecules impinging in the vicinity of a vacant bridge site will stick.

Neglecting any lateral interactions beyond the short-range blocking rules, the required O and CO binding energies, diffusion and reaction barriers are independent of the local adsorbate environment. The corresponding values computed within our DFT setup are summarized in Table I and are generally in very good agreement with previous computations [163, 184, 204]. The obtained LH reaction barrier of 0.9 eV in particular is in good agreement with previous DFT calculations at varying coverages (0.76-1.05 eV) [36, 65, 201]. Attempts to calculate an ER type reaction path over a (2 × 2) oxygen adlayer at Pd(100) showed that CO and O coadsorption is energetically more favorable and thus an ER reaction is excluded from the Pd(100) model. As apparent from Table I the O and CO diffusion barriers are very low. This high mobility severely limits the numerical efficiency of the 1p-kMC simulations, which are completely dominated by frequent executions of diffusion events at minute time increments. In order to speed up the 1p-kMC simulations we thus artificially raised the diffusion barriers by as much as 0.5 eV without observing any effect on the computed coverages or TOFs.

6.2.3.2 $\sqrt{5}$ surface oxide

For the $\sqrt{5}$ surface oxide we employ the 1p-kMC model established and detailed by Rogal, Reuter and Scheffler [163, 164]. In brief, this model considers two non-equivalent sites named bridge and hollow as depicted on the right of Fig. 6.1. Since there is no process involving sites from adjacent bridge-hollow trenches, the lattice may be viewed as quasi one-dimensional. In analogy to the just described Pd(100) model, the elementary process list consists of all non-concerted adsorption, desorption, diffusion and LH reaction processes involving

Table 6.1: Summary of DFT binding energies, diffusion and reaction barriers (LH and ER) used in the Pd(100) and $\sqrt{5}$ 1p-kMC models. All values are in eV.

PdO $\sqrt{5}$ desorption barriers:

E^0 on-site energy, V nearest-neighbor lateral interaction¹

$E_{O,bridge}^0$	$E_{O,hollow}^0$	$E_{CO,bridge}^0$	$E_{CO,hollow}^0$
-0.51^1	-1.95^1	-1.40^1	-1.92^1
$V_{O-O,br-br}$	$V_{O-O,hol-hol}$	$V_{O-O,br-hol}$	
0.08^1	0.07^1	0.08^1	
$V_{CO-CO,br-br}$	$V_{CO-CO,hol-hol}$	$V_{CO-CO,br-hol}$	
0.08^1	0.13^1	0.14^1	
$V_{O-CO,br-br}$	$V_{O-CO,hol-hol}$	$V_{O-CO,br-hol}$	$V_{O-CO,hol-br}$
0.06^1	0.11^1	0.13^1	0.12^1

PdO $\sqrt{5}$ diffusion barriers

$CO,br \rightarrow br$	$CO,hol \rightarrow hol$	$CO,br \rightarrow hol$
0.4^1	0.6^1	0.3^1
$O,br \rightarrow br$	$O,hol \rightarrow hol$	$O,br \rightarrow hol$
1.2^1	1.4^1	0.1^1

PdO $\sqrt{5}$ reaction barriers (LH)

$E_{Obr,CObr}$	$E_{Ohol,COhol}$	$E_{Obr,COhol}$	$E_{Ohol,CObr}$
1.0^1	1.6^1	0.5^1	0.9^1

PdO $\sqrt{5}$ reaction barriers (ER)

E_{Ohol1}	E_{Ohol2}	E_{Obr1}	E_{Obr2}
0.8	0.5	0.0	0.0

Pd(100) desorption barriers

$E_{O,hollow}^0$	$E_{CO,bridge}^0$
-1.25	-1.93

Pd(100) diffusion barriers

$CO,br \rightarrow br$	$O,hol \rightarrow hol$
0.14	0.28

Pd(100) reaction barriers

$E_{CObr,Ohol}$
0.9

these sites. Dissociative O₂ adsorption requires two neighboring sites, and is only hindered by a sizable adsorption barrier of 1.9 eV in the case of adsorption into two bridge sites. CO adsorption is unimolecular and not hindered by adsorption barriers.

As only modification of the Rogal model we additionally consider an ER reaction mechanism, as recently suggested by Hirvi *et al.* for bulk PdO(101) [80]. Reaction path calculations vertically impinging a CO molecule over O atoms adsorbed in the different surface sites indeed also yield rather low reaction barriers over the $\sqrt{5}$ surface oxide, namely about 0.7 eV over hollow and essentially zero over bridge. The latter rather astonishing result has very little consequences for the surface oxide stability and the catalytic activity in near stoichiometric feeds though, as under corresponding gas-phase conditions the O coverage of bridge sites is negligible.

6.2.4 kMC simulation setup

Both 1p-kMC models were implemented using the kmos framework[81]. All simulations were performed in simulation cells containing (20 × 20) unit cells and using periodic boundary conditions. Systematic checks showed that the quantities of interest here, *i.e.* the average steady-state coverages and TOFs, are perfectly converged at these cell sizes. For defined gas-phase feed conditions (T, p_{CO}, p_{O₂}) the 1p-kMC simulations eventually reach a steady state, with constant TOF and average surface coverages $\bar{\Theta}_{i,st}$ of species i on site st,

$$\bar{\Theta}_{i,st} = \frac{\sum_n \Theta_{i,st,n} \Delta t_n}{\sum_n \Delta t_n} , \quad (6.6)$$

where n denotes the value at the nth kMC step after steady state has been reached. Starting from different initial adsorbate distributions and using different random number seeds we validated that this steady state is well defined, *i.e.* we never observed multiple steady-states in this system. We also carefully varied the initial conditions to ensure that the unique steady state has been reached in all simulations. In the kinetic phase diagrams shown in Fig. 6.2 below the obtained results are summarized in form of contour plots. All contour plots were interpolated using radial basis functions in order to reduce numerical noise.

6.3 RESULTS

6.3.1 Kinetic phase diagrams at 600 K

Predicted oxygen and CO coverages on the $\sqrt{5}$, as well as on the Pd(100) surface are shown in Fig. 6.2 for a temperature of 600 K. For the Pd(100) surface the coverages refer to CO in bridge sites and O




Figure 6.2: (Color online) 1p-kMC computed O and CO coverage as a function of partial pressures for $T = 600$ K at Pd(100) (upper panels) and at the $\sqrt{5}$ surface oxide (lower panels). At Pd(100) the coverage refers to O in hollow sites and CO in bridge sites, at the $\sqrt{5}$ the coverage refers to O and CO in upper hollow sites (see text). The thick black line in the left panels indicates the boundary where each phase is expected to be stable: For Pd(100) this corresponds to the $p(2 \times 2)$ oxygen adlayer, for PdO $\sqrt{5}$ this is a coverage > 0.9 ML.

in hollow sites, as these are the only available sites for these species in the employed model, respectively. For the $\sqrt{5}$ surface only the O and CO coverage in the upper hollow sites is shown, as this is the expected critical descriptor for the stability of the surface oxide layer [163]. For Pd(100) a coverage of 1 monolayer (ML) corresponds to one adsorbate (O or CO) per surface Pd atom. For the $\sqrt{5}$ surface a coverage of 1 ML corresponds to a complete occupation of the upper hollow rows by the respective species.

Both models yield the intuitively expected coverage variations with partial pressure. Starting with pure O gas conditions (going horizontally along the bottom of the panels in Fig. 6.2), both models yield the correct surface coverages in the respective stable regimes by construction: For Pd(100) increasing O content in the gas phase leads to a gradual coverage increase starting from the clean surface. For the $\sqrt{5}$ the full ML O coverage in the bottom right part of the panel in turn reflects a fully intact layer of the surface oxide in which all upper hollow sites are fully covered by oxygen. Increasing CO pressure also influences the surface coverages as intuitively expected: An increasing CO population at the surface defers the stabilization of surface O at Pd(100) to higher O_2 pressures, while it leads to a quicker depletion of upper hollow O atoms at the $\sqrt{5}$ at decreasing O_2 pressures. For highly overstoichiometric CO pressures (upper horizontal line

in the panels) the Pd(100) model displays CO coverages of 0.5 ML. This is slightly less than the highest CO concentration characterized experimentally, namely 0.75 ML [186]. This deviation results from the employed simplified lateral interaction model, which likely features too repulsive CO-CO interactions. We will scrutinize this with the less repulsive Pd-1NN model below, but also note that these CO-rich gas-phase conditions (and dense CO adlayers) are not the focus of our present interest. The same holds for the incorrect limit of a purely CO-covered $\sqrt{5}$ surface in the upper left part of the respective panel (where the surface oxide would in reality be reduced away), as well as for the $c(2 \times 2)$ 0.5 ML O coverage in the lower right part of the Pd(100) panel (where instead a surface or bulk oxide would be formed).

The very transition between CO- and O-covered regimes requires closer inspection. In the case of the Pd(100) surface a CO-poisoned surface would be catalytically inactive, but certainly stable. This surface is expected to be stable at least until the oxygen concentration does not rise above ~ 0.25 ML as this corresponds to a coverage regime representative for the experimentally characterized $p(2 \times 2)$ overlayer [116, 207]. In Fig. 6.2 we therefore denote the stability boundary of the Pd(100) surface at 0.25 ML coverage, *i.e.* we would expect a stable Pd(100) surface for any gas-phase conditions to the upper left of this line. Note that due to the steep coverage rise in the transition region, this stability boundary would be barely affected on the scale of Fig. 6.2 if we had e.g. chosen 0.4 ML coverage as the stability criterion. In the same spirit Rogal *et al.* have used a coverage exceeding 0.9 ML as stability criterion for the $\sqrt{5}$ surface oxide before [164]. The corresponding stability boundary for the surface oxide is also drawn in Fig. 6.2, and we would expect the surface oxide to be stable anywhere to the bottom right of this line. Again, on the scale of Fig. 6.2 it would make little difference, if a stability criterion of e.g. 0.95 ML or 0.99 ML coverage had been used. With the two stability regions thus quite narrowly defined, the central and intriguing feature of Fig. 6.2 is that there is a finite range of partial pressures where both models are predicted to be stable. This is highlighted again in Fig. 6.3, where the resulting bistability region is marked in green.

6.3.2 Variation with temperature

The results presented in the preceding section were obtained for $T = 600$ K, a temperature that falls in the middle of the temperature range from ~ 300 – 800 K that is generally most relevant for CO oxidation catalysis. Not least to make contact with the dedicated reactor STM experiments performed by Hendriksen *et al.* at 408 – 443 K [74], an important next step is to assess the variation of our findings, in particular the existence of a bistability region, with temperature. Correspond-




Figure 6.3: (Color online) Deduced region of bistability, *i.e.* gas-phase conditions where the 1p-kMC models would predict the simultaneous stability of pristine Pd(100) and the $\sqrt{5}$ surface oxide (see text). All diagram show the same range of gas phase chemical potentials.

ingly and following the same protocol as before, the two lower panels of Fig. 6.3 summarize our findings for $T = 500$ K and $T = 400$ K. Both panels show the same range of gas-phase chemical potentials as the upper panel for 600 K, allowing to ascribe all differences between the three panels directly to kinetic effects, that is deviations from thermodynamic scaling. Intriguingly, these deviations are rather small, *i.e.* the differences for the three temperatures amount primarily to the expected scaling of the pressure axes [157]. This leaves the central result in form of a bistability over a finite range of gas-phase pressures essentially untouched within the temperature range 400–600 K. Particularly intriguing is that certainly for $T = 600$ K the catalytically most relevant near-ambient conditions at stoichiometric feed fall right within the region of bistability. For the lower temperatures the bistability region shifts increasingly towards O-rich conditions, such that e.g. the conditions of the Hendriksen experiments fall just at the border of this region.

6.4 DISCUSSION

6.4.1 Origin and robustness of bistability region

An immediate concern with the observed bistability region is that it is rather narrow in $(p_{\text{O}_2}, p_{\text{CO}})$ -space. Considering the range of uncertainties underlying the 1p-kMC models, prominently the semi-local DFT energetics and lateral interaction model, this raises doubts as to the robustness of this finding. Fortunately, in the present case an analysis of the atomic-scale reason behind the bistability shows that its actual existence emerges rather independently from these uncertainties, which is why we lead these two discussions jointly in this section.

A first important step towards an understanding of the atomic-scale origin of the bistability comes from the observed almost perfect thermodynamic scaling of the location and extent of the bistability region in $(p_{\text{O}_2}, p_{\text{CO}})$ -space in the temperature range 400–600 K, cf. Fig. 6.3. This suggests that the actual reaction kinetics, and the uncertainties in the concomitant reaction barriers, is not central to its existence. This view is confirmed by the fact that an equivalent bistability region is already obtained within a "constrained" *ab initio* thermodynamics approach [156, 156], *i.e.* an approach that neglects the reaction kinetics completely. Figure 6.4 shows the corresponding phase diagram for the same range of chemical potentials also underlying the panels in Fig. 6.3. For a comparison to the 1p-kMC results, exactly the same DFT energetics and only ordered structures consistent with the 1p-kMC lateral interaction models are used. This explains small differences with respect to the corresponding phase diagram published before by Rogal, Reuter and Scheffler [163], which e.g. con-

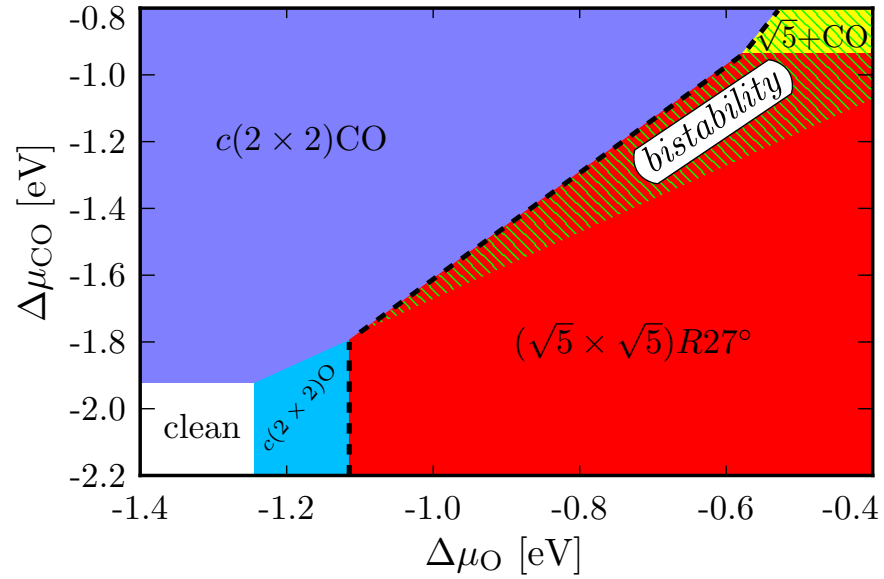


Figure 6.4: (Color online) "Constrained" *ab initio* thermodynamic phase diagram of the Pd(100)/ $\sqrt{5}$ system, considering only ordered structures consistent with the 1p-kMC lateral interaction models (see text). The thick dashed line denotes the stability boundary between adsorption phases on Pd(100) (upper left part) and adsorption phases on the $\sqrt{5}$ surface oxide (bottom right part). Additionally shown as hatched area is the stability region of Pd(100) phases with O coverage below 0.5 ML. This is the thermodynamic equivalent to the bistability region identified in the 1p-kMC simulations, if the low O coverage defers formation of the (thermodynamically preferred) surface oxide.

sidered additional experimentally characterized CO ordered phases (not of relevance for the present discussion). Denoted by the thick dashed line is the stability boundary between adsorption phases on the pristine Pd(100) surface (upper left part) and adsorption phases on the $\sqrt{5}$ surface oxide (lower right part). For all catalytically relevant gas-phase conditions, this boundary runs in fact between CO-covered Pd(100) and different adsorption phases on the $\sqrt{5}$, *i.e.* with increasing O-content in the gas phase constrained thermodynamics predicts an abrupt phase transition from a CO-poisoned surface directly to the surface oxide. Knowing that a critical local coverage of about ~ 0.5 ML O is necessary to induce the formation of the surface oxide, it is therefore interesting to also include in the phase diagram the stability region of O-containing Pd(100) phases with less than 0.5 ML O coverage, *i.e.* phases where (coming from the pristine metal side) oxide formation would not yet start. Intriguingly, the resulting region shown in Fig. 6.4 extends over similar gas-phase conditions as the bistability region deduced from the 1p-kMC simulations.

As such we ascribe the bistability region to gas-phase conditions, where the $\sqrt{5}$ surface oxide is thermodynamically more stable, but where a too low O coverage at Pd(100) would not readily induce the formation of the oxide. When crossing the bistability region from O-rich to CO-rich gas-phase conditions the system will thus prevail in the (thermodynamically preferred) oxidized state, while in the opposite direction it will prevail in the metal state as kinetic limitations to stabilize enough oxygen at the surface prevent the formation of the surface oxide. With this understanding of the atomic-scale origin, it is primarily the adsorbate binding energies and the lateral interaction model that are crucial to the robustness of the bistability region – as they govern the stabilization of oxygen at the surface. More specifically, it is the relative binding energy differences at the metal and the $\sqrt{5}$ that will primarily affect the extension of the bistability region, while the absolute binding energetics will rather shift its location in $(p_{\text{O}_2}, p_{\text{CO}})$ -space. We would expect uncertainties in the employed DFT exchange-correlation functional to rather affect the absolute binding energetics (*i.e.* systematic over- or underbinding) and only to a lesser extent the binding energy differences of O and CO at the two surfaces. This view is confirmed by systematic tests, in which we increased or decreased all binding energies by 0.2 eV and in both cases found only small differences in the extension of the bistability region at more pronounced changes in its position in $(p_{\text{O}_2}, p_{\text{CO}})$ -space. Correspondingly, we do not expect that the predicted gas-phase conditions for the bistability are accurate to better than some orders of magnitude in pressure. Its actual existence, however, should be very robust with respect to the uncertainties of present-day DFT functionals.

This leaves as final important aspect the employed lateral interaction model on the Pd(100) surface. Already the comparison to the




Figure 6.5: (Color online) Comparison of the bistability regions arising from 1p-kMC simulations with different lateral interaction models. Pd(100)-2NN is the default model employed so far, where configurations are excluded that would result in CO adsorbed in nearest-neighbor and next-nearest neighbor positions (green solid area). Pd(100)-1NN mimics less repulsive interactions and only excludes configurations resulting in CO adsorbed in nearest-neighbor positions (hatched area).

experimentally determined maximum CO coverage (0.75 ML [186] versus 0.5 ML in Fig. 6.2) indicated the site-blocking rules in the default Pd(100)-2NN model to be too repulsive. We assess the consequences for the bistability region by determining this region with the less repulsive Pd(100)-1NN model, where now only configurations that would result in CO adsorbed in nearest-neighbor positions are excluded. The resulting 1p-kMC bistability region at 600 K is shown in Fig. 6.5 and compared to the corresponding bistability region deduced before within the default Pd(100)-2NN model. Obviously the eased stabilization of CO at the Pd(100) surface (maximum coverage now 1 ML) blocks the O adsorption even more effectively and consequently leads to a much increased range of gas-phase conditions where both surface states are predicted to be stable. If we take the comparison to the maximum CO coverage from experiment as measure, we would expect the true CO-CO repulsion to be somewhere between the one represented within the Pd(100)-1NN and Pd(100)-2NN site blocking models. Correspondingly, we would estimate the true extension of the bistability region to be somewhere between those predicted within the two models. While thus not fully quantified in its extension, the actual existence of a bistability region *per se* is robust against the uncertainty in the employed lateral interaction model.

6.4.2 Comparison to experiment

The finding of a finite bistability region in our calculations agrees nicely with the bistability and oscillations reported experimentally by Hendriksen and coworkers [70, 74]. Unfortunately, we can not directly compare the temperatures and pressures where this bistability is found. On the theoretical side this is due to the aforesaid uncertainty with respect to the location of this bistability region in $(p_{\text{O}_2}, p_{\text{CO}})$ -space caused by the approximate first-principles parameters entering the kinetic models. On the experimental side this is due to the suspected non-negligible mass-transfer effects in the employed reactor STM setup [121–123, 188]. Nevertheless, the deduced pressure ranges are intriguingly close, as is the extension of the bistability region. At the experimental $T \sim 400$ K and $p_{\text{O}_2} = 1$ atm, the measured width of the bistability region in CO pressure is 0.02 atm [74], while it is ≈ 0.5 atm in the simulations. Also, the finding that at constant p_{O_2} , the bistability region shifts with increasing temperature to increasing p_{CO} agrees with the experimental observations [70]. As such our interpretation is that the here obtained bistability between metallic Pd(100) and the $\sqrt{5}$ surface oxide provides an atomic-scale model for the oscillations of the Hendriksen experiments. This is close to the interpretation arrived at in the experimental study, with the slight modification that they suspected the metallic phase to correspond to predominantly O-covered Pd(100), whereas our simulations show that this is predominantly CO-covered Pd(100).

The atomic-scale insight into the surface composition provided by our 1p-kMC simulations also allows to further qualify the reaction mechanism. An intriguing observation made by Hendriksen and coworkers was that in contrast to the metal surface, the reaction rate on the oxide did not follow traditional LH kinetics. They proposed this as a signature of low oxide stability and a concomitant Mars-van-Krevelen type mechanism, in which the oxide is continually consumed and reformed. Figure 6.6 shows that this can instead be rationalized within a much simpler picture that does not involve the oxide stability. Shown are the surface coverages on the Pd(100) metal and $\sqrt{5}$ surface oxide over the p_{CO} pressure range, where we obtain bistability at fixed $T = 400$ K and $p_{\text{O}_2} = 1$ atm, *i.e.* for gas-phase conditions comparable to the Hendriksen experiments. At the Pd(100) surface, O and CO adsorb in inequivalent sites, hollow and bridge, respectively. Due to the strongly repulsive interactions, mimicked in our simulations by mutual site blocking, this nevertheless leads effectively to a competition for adsorption sites between the two adsorbates and in consequence to a LH-type kinetics. In contrast, on the $\sqrt{5}$ surface oxide, O and CO also adsorb in inequivalent sites at these gas-phase conditions, namely O almost exclusively on the hollow sites and CO almost exclusively on the bridge sites, cf. Fig. 6.6. At the surface oxide




Figure 6.6: (Color online) 1p-kMC surface coverages and turn-over-frequencies (TOFs) at gas-phase conditions comparable to the reactor STM experiments by Hendriksen *et al.* [74]. Shown is the bistability region in p_{CO} at fixed $T = 400 \text{ K}$ and $p_{\text{O}_2} = 1 \text{ atm}$. Upper panel: O and CO coverage on the Pd(100) surface; middle panel: O and CO coverage on the $\sqrt{5}$ surface oxide; lower panel: comparison of the intrinsic TOFs of the two system states.

these sites are located further away from each other and adsorbates even in nearest-neighbor sites experience only very small lateral interactions, cf. Table I. Consequently, the occupation of these two site types occurs almost independently of each other at varying gas-phase conditions and gives rise to a TOF proportional in CO pressure, cf. Fig. 6.6, exactly as observed in the experiment. Similarly, we find the TOF to vary only little when changing the oxygen partial pressure away from the conditions shown in Fig. 6.6 – again fully consistent with the experimental findings.

The understanding that the actual reaction mechanism over the $\sqrt{5}$ proceeds by adsorption on the otherwise intact surface oxide suggests that the continued roughening of the surface observed in the experiments is a by-product and not essential to the catalytic activity. With respect to the total activity Fig. 6.6 shows that over the bistability region both surface states, Pd(100) and $\sqrt{5}$, exhibit rather similar intrinsic TOFs to within a factor of three. While the mass-transfer limitations present in the reactor STM measurements prevent a direct comparison, this finding alone sets the controversial discussion concerning *the* active state of the surface into perspective. At low temperatures and UHV conditions, our 1p-kMC models indeed yield a significantly higher activity of the Pd(100) surface. This arises predominantly from the comparatively weak CO binding at the $\sqrt{5}$ surface oxide and the concomitant limitations in stabilizing it at the surface. However, at technological gas-phase conditions, which in our simulations do fall within the bistability region, the TOF differences between the two models are not large enough and furthermore change sensitively with the detailed gas-phase conditions, cf. Fig. 6.6, to support a detailed discussion as to which state is the more active one. In this respect, the Pd(100) surface differs qualitatively from the equally prominently discussed Ru/RuO₂ system [142, 154]. On ruthenium the catalytic activity can be clearly attributed to the oxidized surface, while in the present system two competing surface states can equally contribute to the catalytic activity, which thus also explains why a clear signature of palladium (surface) oxide formation may in some circumstances be more difficult to find.

6.5 CONCLUSIONS

Detailed 1p-kMC simulations on either the Pd(100) surface or the $\sqrt{5}$ surface oxide yield a finite range of CO oxidation gas-phase conditions, where both surface states appear stable. This finding of a bistability region is robust against variations in the detailed criteria used to assess the stability of either state, as well as against the uncertainties arising from the approximate first-principles energetics and lateral interaction models. Notwithstanding the latter uncertainties do affect the position where in (T, p)-space the bistability region is

found and we expect this to translate in uncertainties up to the order of 100 K and some orders of magnitude in pressure.

The bistability region arises from limitations in stabilizing oxygen at the Pd(100) surface, which extends the stability region of the metallic surface beyond that predicted by thermodynamics. In some respect this is therefore reminiscent of the generic Langmuir-Hinshelwood kinetics that lead to bistability between a CO-poisoned state that effectively blocks O₂ adsorption and an O-rich state, *i.e.* the bistability of the reaction rate originates essentially from the inequivalence of the adsorption of carbon monoxide and oxygen.[4, 17, 37, 169, 205]. The difference is that in these classical models the two stable solutions were either supposed to be different adsorbate phases on the same substrate or did correspond to an active metal and an inactive oxide state. Instead, in the present system we have a CO-rich Pd(100) state that blocks oxygen adsorption and the surface oxide as O-rich state. Furthermore, the bistability does comprise technological conditions and conditions comparable to the dedicated reactor STM experiments performed by Hendriksen *et al.*[73, 74]. Under these conditions, both surface states do show similar intrinsic activity, which sets preceding discussions with respect to *the* active state into perspective.

Within the uncertainties of the exact location of the bistability region in (T, p)-space we suggest the here obtained bistability between metallic Pd(100) and the $\sqrt{5}$ surface oxide as an atomic-scale model for the oscillations of the Hendriksen experiments. Likewise, our data lends further support to the view that oxide formation plays an important role in understanding the catalytic activity of Pd catalysts. If it is even the oscillations themselves, catalytic activity would again be rationalized as a kinetic phase transition phenomenon, *i.e.* the catalyst surface being close to an instability[40, 156, 208]. At present our simulations performed separately on the two intact surface states, corresponding to ideal terraces of Pd(100) and the $\sqrt{5}$ surface oxide, can not address this notion directly. This holds equally for the experimental interpretation that the continued roughening of the oxidic surface during reaction, as well as the formation of steps are crucial ingredients to the oscillatory behavior. The following chapter extends the 1p-kMC capabilities towards the actual oxide decomposition, which is a necessary first step to start scrutinizing these points with predictive-quality theory.

REDUCTION OF PD SURFACE OXIDE BY CO

After exploring the stability boundaries of the Pd(100) and the $\sqrt{5}$ -oxide using static lattice kinetic Monte Carlo this chapter will address the explicit reduction of the $\sqrt{5}$ -oxide towards the Pd(100) surface, *i.e.* the explicit morphological transition from one lattice to the other. By doing so we will first construct an explicit atomistic reconstruction pathway which is then implemented using a suitable extension of the lattice kMC method. The work of this chapter is currently finalized for publication in *The Journal of Chemical Physics*[84].

7.1 INTRODUCTION

In recent years first-principles (1p) kinetic Monte Carlo (kMC) simulations have established a new standard for microkinetic modeling in heterogeneous catalysis.[165] On the one hand, this approach offers material-specific predictive quality by relying on first-principles rate constants for the involved elementary processes, typically computed through density-functional theory (DFT) and transition-state theory (TST). On the other hand and in comparison to still prevalent mean-field rate equation based microkinetic models, it explicitly resolves the involved active sites, and correspondingly accounts correctly for possible spatial heterogeneities, correlations and fluctuations in the adsorbate layer. The downside of the approach is the still significant computational cost connected with the DFT-TST calculation of the rate constants, which generally involves for each required rate constant a transition state search for the activation energy, and – depending on the level of approximation employed – further calculations to determine the pre-factor. In practice, 1p-kMC simulations in heterogeneous catalysis are therefore predominantly performed within the lattice kMC approach, where the translational symmetry of the crystalline lattice leads to a dramatic reduction of the number of inequivalent elementary processes and correspondingly required number of different 1p rate constants [86, 154]. For simple reaction networks like for the multiply studied CO oxidation reaction, in the absence of lateral interactions, and for a small number of different active site types as e.g. offered by a low-index single-crystal model catalyst surface, this number of 1p rate constants can be as low as 10–20, comprising adsorption, desorption, diffusion and reaction steps at and between the different active sites.[158, 160, 163]

A major limitation of the lattice 1p-kMC approach is that it cannot directly be applied to (surface) morphological rearrangements, *i.e.*

any type of phase transition that includes changes in the considered lattice arrangement of the active sites. In the catalytic context such rearrangements may prominently occur in response to the reactive environment. For the (late) transition metal catalysts typically employed in oxidation catalysis, a possible formation of (surface) oxides has e.g. been controversially discussed, with as complicated scenarios perceivable as the continuous formation and reduction of such oxidic films being a major driver for the observed catalytic activity.[9, 72–74, 116, 118, 130, 141, 188]

As a first step to make such systems and phenomena accessible to scrutiny by 1p-kMC we here introduce a multi-lattice 1p-kMC approach, which is applicable to morphological transitions between commensurate crystalline lattices. Such multi-lattice kMC simulations simultaneously consider all involved lattices and their respective lattice sites. In any spatial region describing one lattice the sites of all other lattice types are deactivated, while suitable elementary processes describe local transitions between the lattices. We illustrate this multi-lattice 1p-kMC idea by modeling recent surface oxide reduction experiments performed on Pd(100) [45]. Exposing an initially prepared ($\sqrt{5} \times \sqrt{5}$)R27° PdO(101) surface oxide film (henceforth coined $\sqrt{5}$ -oxide for brevity)[104, 184] to a 5×10^{-11} bar CO atmosphere, significant changes in the rate of oxide reduction were observed in the temperature range from 300–400 K. Analysis of the data using fitted mean-field kinetic models suggested that the reduction is controlled by boundary processes between surface oxide and metal domains [45]. This renders this system an ideal test case for a quantitative description using multi-lattice 1p-kMC, which then has to include a microscopic pathway for the actual oxide reduction mechanism.

The remainder of this paper is structured as follows: After a brief recap of the lattice 1p-kMC approach and a summary of the previously developed single-lattice 1p-kMC models for pristine Pd(100) and the $\sqrt{5}$ -oxide, we first demonstrate that single-lattice 1p-kMC cannot account for the experimentally observed temperature dependence of the reduction time. Introducing a 1p-guided atomistic pathway for the oxide reduction process we then show that corresponding multi-lattice 1p-kMC simulations perfectly reproduce the observed trends, and that indeed the critical steps in the reduction of the surface oxide correspond to cross-reactions between the metal domains and the surface oxide.

7.2 THEORY

7.2.1 Kinetic Monte Carlo simulations in surface catalysis

Thermally driven surface chemical reactions involving site-specific binding of the reaction intermediates generally follow a so-called *rare-*

event time evolution. That is, on microscopic time scales (that would e.g. be reachable by typical molecular dynamics simulations) the system's state can for most of the time be described as reaction intermediates vibrating around the active sites to which they are bound. Only rarely, *i.e.* each time after long time intervals, an elementary process changes this population at the active sites, where the actual duration of this process itself is quite short. A diffusion process brings a reaction intermediate to another adjacent active site, adsorption and desorption fill and empty active sites at the surface, surface reactions convert reaction intermediates into another etc. During the long time intervals between such individual elementary processes the reaction intermediates equilibrate fully and thereby are assumed to forget through which elementary process they actually reached their current site or state.

Viewing the entire population at all active sites as one possible configuration, individual events in form of elementary processes thus advance the system's state from one configuration to another. In each configuration reaction intermediates have forgotten about their past, *i.e.* in this Markov approximation events occur uncorrelated from previous ones. The time evolution of the system can then be described through a Markovian master equation

$$\dot{\rho}_u(t) = \sum_v w_{uv} \rho_v(t) - w_{vu} \rho_u(t) \quad , \quad (7.1)$$

where $\rho_u(t)$ is the probability for the system to be in configuration u at time t , and w_{vu} is the transition rate (in units of time^{-1}) at which configuration u changes to configuration v .

For surface catalytic applications, this equation is in practice unsolvable by any direct, even numerical approach. The total number of matrix elements in eq. (7.1) unequivocally outgrows storage capabilities of available hardware already for rather small models. This is readily illustrated considering a simple catalyst surface model that contains 100 active sites and considers O and CO as possible reaction intermediates. Including the possibility of empty sites, the corresponding number of configurations (possible different populations of the active sites) is 3^{100} and the transition matrix describing the connection of every in principle possible configuration with every other one accordingly has $(3^{100})^2 \approx 2.65 \times 10^{95}$ matrix elements.

The kMC approach circumvents this problem by generating configuration-to-configuration trajectories in such a way that an average over an ensemble of such trajectories yields the correct probability density $\rho_u(t)$ of eq. (7.1). [86, 154] For steady-state applications, this ensemble average over many kMC trajectories may alternatively be replaced by a time average over one sufficiently long trajectory. Instead of evaluating and storing the entire transition matrix, a kMC code thus only focuses on-the-fly on those transition matrix elements w_{vu} that are actually required to propagate the trajectory. Existing

kMC algorithms differ in the way they propagate this trajectory by randomly choosing events out of all events possible for a given configuration and in the way they determine the elapsed system time.[115] In this work we specifically employ the variable step size method as implemented in the kMC package `kmos` [86].

7.2.2 *Efficient 1p-kMC on one or more lattices*

Even if a kMC code only focuses on those w_{vu} that are actually required to propagate the trajectory, this can still be an intractable task, if the rate constants determining them are to be provided through 1p-calculations. To appropriately choose one possible event to propagate the kMC trajectory, a kMC algorithm needs to know all possible (or for approximate procedures at least all energetically low-lying) escape paths out of the current system configuration and their corresponding rate constants. In the prevalent DFT-TST approach to obtain these rate constants in 1p-kMC this equates to determining all low-lying transition states on the potential energy surface (PES) that lead out of the current PES minimum (plus ensuing calculations for the kinetic prefactors). While various approaches to this problem have been put forward and feature under names like self-learning or adaptive kMC [38, 96, 185, 199], a tremendously more efficient route is possible, if the active sites involved in the surface catalysis form a static lattice.

In corresponding lattice kMC the entirety of all possible escape paths breaks down into groups of individual elementary processes that can possibly take place at each given active site [86, 154]. Due to the translational symmetry of the lattice, each equivalent active site type furthermore features the same group of in principle possible elementary processes, and these elementary processes typically depend only on the occupation of nearby sites. This locality then boils down the number of actually required inequivalent 1p rate constants to a manageable number, at least for not too complex reaction networks. In lattice 1p-kMC these 1p rate constants are in practice precomputed and stored, to be looked up on demand in the evolving kMC simulation. For a given system configuration, expressed as the population of each active site in the lattice (empty or occupied with a certain reaction intermediate), a naive lattice kMC code correspondingly loops over all active sites in the lattice. Depending on the occupation of the site and possibly the occupation of nearby sites, it determines the currently available events and looks up their corresponding rate constants. Out of all of these available events at all sites, it executes one as determined by the kMC algorithm, updates the lattice configuration and elapsed time, and proceeds from thereon. Modern and efficient lattice kMC codes like `kmos` [86] operate on the same con-

ceptual recipe, but e.g. resort to local updates of the list of available events instead of full lattice loops.

Assuming a static arrangement of active sites, lattice 1p-kMC can generally not be applied to systems undergoing a morphological transition. In the present context, this would be the reduction of the initially oxidized surface to a metallic state, where the oxide overlayer and metal substrate exhibit different crystal lattices. With the here suggested generalization to multi-lattice 1p-kMC this situation is changed in as much, as transitions between commensurable lattices become accessible. The essence of the approach is hereby to choose as basic unit cell of the lattice employed in the kMC simulations a suitable, large enough commensurate cell that embraces all primitive unit cells of the different lattices to be treated. This commensurate cell (and the lattice built up from it) then contains as active sites simultaneously all active sites of all lattices. The fact that the lattice involved need to fit into a common commensurate unit cell may appear as a disadvantage of the multi-lattice approach at presented here. On the other hand if there is no such commensurability this also means that the number of possible interface combinations and therefore the number of required inequivalent 1p-calculations grow excessively which means that any lattice kMC approach loses its main advantage.

The trick to efficiently distinguish during a simulation in which lattice state a given local region is currently in, is to introduce an additional artificial species (canonically named `null`) to represent inactive sites. If say, a certain region is currently in one lattice structure, then `null` species are placed on all sites corresponding to the other lattices in this region. A site "occupied" with a `null` species is hereby to be distinguished from an empty site. No elementary processes are available at sites occupied by a `null` species, *i.e.* elementary processes that require an empty active site explicitly look for the sites to be empty and not "occupied" by `null`. In a corresponding region where all sites belonging to the other lattices are deactivated through `null` species, multi-lattice 1p-kMC thus reduces to a regular lattice 1p-kMC simulation and only the elementary processes corresponding to the 1p-kMC model of this lattice can possibly occur. In this setup, local transitions between the different lattices are finally enabled by introducing new elementary processes which involve the removal or placement of `null` species at sites belonging to the lattices involved. Trivially, for the oxide reduction example this could be as simple as an elementary process placing `null` species on one or more active sites of the oxide lattice and removing `null` species from nearby active sites of the metal lattice. In reality, local lattice transitions are unlikely that simple, but follow complex, multi-stepped reaction paths. Nevertheless, as demonstrated before [86], the computational cost of additional sites and species is very low in the `kmos` code. The true challenge of setting up and performing a multi-lattice 1p-kMC simulation is

therefore to devise atomistic pathways for the lattice transitions and implement them in form of corresponding elementary processes.

7.2.3 Elementary processes during the reduction of oxidized Pd(100) within unperturbed surfaces

The multi-lattice 1p-kMC simulations describing the reduction of the Pd(100) surface initially covered with a $\sqrt{5}$ -surface oxide film build on existing 1p-kMC models for CO oxidation at the $\sqrt{5}$ -oxide and for CO oxidation at Pd(100). These single-lattice models and the elementary processes considered in them have been detailed before [82, 164]. Here, we only briefly recapitulate their essentials for completeness and as a means to introduce the two lattice types and their active sites involved.




Figure 7.1: Illustration of the lattice employed in the 1p-kMC model of CO oxidation at Pd(100). Shown is a top view of the Pd(100) surface with a primitive unit cell containing one hollow site and two bridge sites. CO molecules can adsorb at bridge sites and oxygen molecules adsorb (dissociatively) at hollow sites (Pd atoms are depicted as large green spheres, O atoms as small red spheres, and C as small grey spheres).

The CO oxidation model on Pd(100)[82] contains three sites per unit cell: one hollow site and two bridge sites, cf. Fig. 7.1. Oxygen molecules can adsorb dissociatively on two empty next-nearest neighbor hollow sites to comply with the so-called 8-site rule.[16, 22, 110, 112] CO can adsorb on an empty bridge site. Desorption processes are time-reversals of these non-activated adsorption processes, with rate constants fulfilling detailed balance. Oxygen diffusion occurs between nearest-neighbor hollow sites, CO diffusion between nearest-neighbor bridge sites. The repulsive lateral interactions between the adsorbed reaction intermediates [111, 204] are modeled through site-blocking rules that prevent processes leading to O-O pairs at nearest-neighbor hollow-hollow distances, to CO-CO pairs closer or at next-nearest-neighbor bridge-bridge distance, and O-CO pairs at nearest-

neighbor hollow-bridge distance. CO oxidation proceeds via a Langmuir-Hinshelwood mechanism that involves adjacent adsorbed O and CO in next-nearest-neighbor hollow-bridge positions and leads to an instantaneously desorbing CO₂ molecule. The rate constants for all these elementary processes have been computed with DFT-TST, specifically employing the PBE exchange-correlation (xc) functional [144]. As in preceding work [82], the very low O and CO diffusion barriers have been artificially raised by 0.5 eV to increase the numerical efficiency of the 1p-kMC simulations. We validated that this has a negligible effect on the 1p-kMC simulation results reported below, as diffusion is then still fast enough to achieve the equilibration of the ad-layer without dominating the simulations in form of frequent executions of diffusion events at minute time increments.




Figure 7.2: Illustration of the lattice employed in the 1p-kMC model of CO oxidation at the $\sqrt{5}$ -oxide. Shown is a top view of the $(\sqrt{5} \times \sqrt{5})R27^\circ$ PdO(101) surface oxide reconstruction with a primitive unit cell containing one hollow and one bridge site. O and CO can adsorb at both site types (Pd atoms are depicted as large green spheres, O atoms as small red spheres, and C as small grey spheres; Pd atoms of the underlying Pd(100) substrate are darkened).

The CO oxidation model on the $\sqrt{5}$ -oxide [164] considers two non-equivalent sites per unit cell, bridge and hollow, as illustrated in Fig. 7.2. O and CO can adsorb at both sites, otherwise the model includes an equivalent list of non-concerted adsorption, desorption, diffusion and Langmuir-Hinshelwood reaction processes as just detailed for the Pd(100) 1p-kMC model. An additional Eley-Rideal reaction mechanism is possible for adsorbed O atoms in both sites, *i.e.* they can be picked up by an impinging CO molecule to form CO₂. Short-ranged lateral interactions are accounted for through an *ab initio* parametrized lattice gas Hamiltonian approach. Adsorption is non-activated, apart from dissociative adsorption of O₂ into neighboring bridge sites, which is hindered by a sizable adsorption barrier of 1.9 eV. As for the Pd(100) 1p-kMC model, all 1p rate constants of

this $\sqrt{5}$ -oxide model have been computed with DFT-TST, employing the PBE xc functional [144]. Due to the geometric structure of the surface oxide, the elementary processes only connect one-dimensional rows of nearest-neighbor bridge and hollow sites. A lattice 1p-kMC simulation of the intact $\sqrt{5}$ -oxide therefore corresponds to a parallel simulation of many separate one-dimensional sub-systems.

The multi-lattice 1p-kMC simulations of the oxide reduction process consider additional elementary processes to model the atomistic pathway from oxide to metal lattice. This pathway and the involved elementary processes are detailed in Sections 7.3.2-7.3.4 below. Their 1p rate constants have been calculated with DFT employing the same TST-based expressions [157] and the same PBE xc functional [144] as used in the Pd(100) and $\sqrt{5}$ -oxide 1p-kMC models. Employing the same computational setup as detailed in ref. [82], the DFT calculations have been performed with the CASTEP package [28] via the Atomic Simulation Environment[3]. Geometry optimizations and transition-state searches were done with supercell geometries containing asymmetric slabs containing one (frozen) layer of Pd(100) and one layer of PdO(101), separated by at least 10 Å vacuum. Geometry optimizations were performed using the BFGS method with a force threshold of 0.05 eV/Å. Transition state searches have been performed with the climbing-image nudged elastic band method [77] or simple drag procedures, with the nature of the transition state confirmed through frequency calculations for those cases where an accurate barrier height was required. The energetics of the thus obtained intermediate geometries were refined in single-point calculations with further Pd(100) layers added to the slab. The gas-phase calculations for O₂ and CO were performed in a cubic supercell of 15 Å box length and Γ -point sampling. At a plane-wave cutoff of 400 eV and a k-point density of 0.4 Å⁻¹ along periodic directions in the slab calculations convergence tests demonstrate a numerical convergence of the obtained DFT barriers and binding energies to within 0.1 eV.

7.2.4 kMC simulation setup

All 1p-kMC simulations have been performed using the kmos framework [86]. The periodic boundary simulation cell contains (40 × 20) $\sqrt{5}$ -unit cells in case of the single-lattice 1p-kMC simulations on the surface oxide, and (20 × 20) commensurate $\sqrt{5}$ -unit cells each containing a total of 19 active sites (from both the metal and the oxide lattice) in case of the multi-lattice 1p-kMC simulations of the oxide reduction process. For the latter, one row of $\sqrt{5}$ -surface oxide was removed from the otherwise fully intact surface oxide film to act as a seed for the reduction process as further detailed in Section 7.3.4. Averages over 10 kMC trajectories starting with different random number seeds were taken to obtain the correct transient quantities.

7.3 RESULTS

7.3.1 Single-lattice 1p-kMC simulations of oxide reduction

The reduction of the $\sqrt{5}$ -oxide by CO was recently experimentally studied in detail using high-resolution X-ray photoelectron spectroscopy (HRXPS) by Fernandes *et al.* [45]. An initially prepared $\sqrt{5}$ -oxide was exposed to a pure CO atmosphere of 5×10^{-11} bar and at three different constant temperatures of 303 K, 343 K, and 393 K. The presence of the surface oxide and adsorbed CO was thereby monitored as a function of time through the corresponding peaks in the XPS spectra. At each temperature studied, the surface oxide coverage θ_{ox} was found to decrease monotonically after a short induction period of 1-3 min, while simultaneously the CO coverage increases monotonically essentially as $\theta_{\text{CO}} = 1 - \theta_{\text{ox}}$. This suggests that CO almost immediately covers any surface area that is not covered by surface oxide. Rather large variations were observed for the total time required to fully reduce the surface oxide at the three temperatures: At 303 K the oxide is reduced after approximately 90 minutes, while the same process takes only approximately 12 min at 343 K and less than 4 min at 393 K.

Circumventing the need to explicitly account for the oxide deconstruction, the coverage of O in the upper hollow sites of the $\sqrt{5}$ -oxide was previously proposed as an indicator for the oxide stability in single-lattice 1p-kMC simulations [164]. Using this criterion 1p-kMC simulations based on the $\sqrt{5}$ -oxide model detailed in subsection 7.2.3 fail qualitatively to reproduce the experimental findings. Even at the highest temperature of 393 K no reduction of the oxide is observed at all after 120 mins, *i.e.* the oxygen coverage remains essentially at 100% for the entire time. The latter can hereby readily be explained by the lack of CO on the surface, *i.e.* at an $\approx 0\%$ CO coverage in the simulations efficient removal of oxygen via CO oxidation is simply not possible.

An obvious reason for this discrepancy could be inaccuracies in the 1p rate constants employed in the 1p-kMC simulations. An underestimation of the CO adsorption energy by the approximate PBE xc functional could e.g. be one rationalization for the low CO coverage in the simulations. Alternatively, even at the low CO coverage a smaller CO oxidation barrier as predicted by PBE could enhance the oxide reduction rate. We explore this possibility by rerunning the 1p-kMC simulations using accordingly modified rate constants. Specifically, we systematically vary the CO oxidation barrier in 0.05 eV steps and the CO adsorption energy on the bridge and hollow sites in 0.1 eV steps. Within any reasonable bounds for these values, no thus generated parameter set leads to kMC simulation results that would reconcile the discrepancy with the experimental data: Either the time scale for

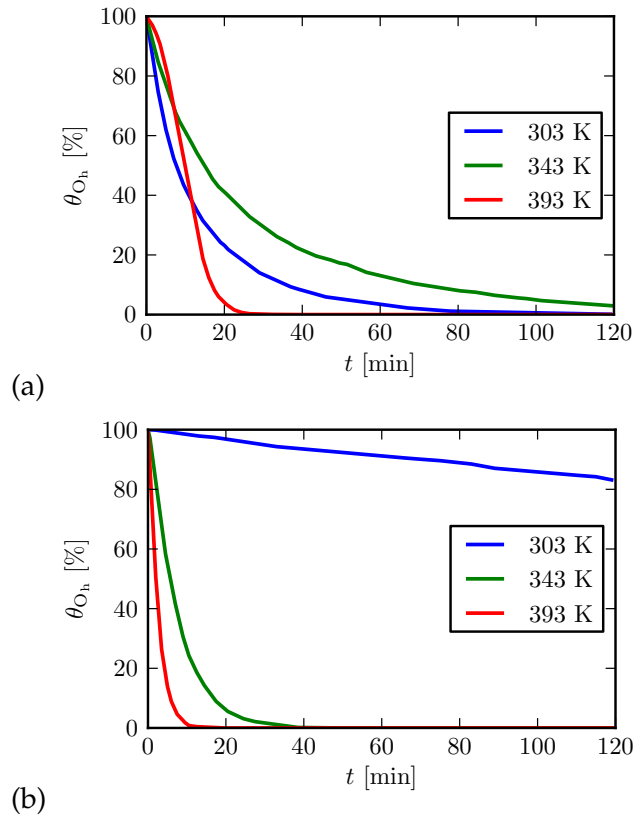


Figure 7.3: Single-lattice kMC simulations for $\sqrt{5}$ -oxide reduction at the three experimentally employed temperatures, using the O coverage at the hollow sites (θ_{O_h}) as indicator for the progressing reduction of the surface oxide (see text). The 1p rate constants are modified to achieve at least some agreement with the experimental findings: (a) CO adsorption energy lowered by 0.2 eV. The reduction times for all three temperatures are in the right order of magnitude, but the ordering with temperature is wrong. (b) CO adsorption energy lowered by 0.5 eV; CO oxidation barrier increased by 0.15 eV. The oxide reduction time correctly decreases with temperature, but the reduction time at 303 K disagrees qualitatively with experiment.

reduction is too long or too short for all three temperatures, or if the simulations exhibit a correct time scale for at least one of the temperatures, then the trend with temperature is wrong, *i.e.* the reduction time does not decrease for each higher temperature as found in the measurements. We illustrate this in Fig. 7.3 for two cases: Once the CO adsorption energy is lowered by 0.2 eV, and once the CO adsorption energy is lowered by 0.5 eV, while the CO oxidation barrier is simultaneously increased by 0.15 eV. In the first case, reasonable reduction time scales are obtained for all three temperatures, but the ordering is wrong: At the intermediate temperature of 343 K reduction is slower than at the lower temperature of 303 K. In the second case, this ordering is correct, but the time scale for oxide reduction at 303 K is much too long.

It is difficult to generally rule out that also varying the rate constants of all other elementary processes would not coincidentally yield a parameter set that would reproduce the experimental reduction times. Nevertheless, the failure to obtain such a parameter set when modifying the two processes that most obviously relate to the reduction time, let us believe that a shortcoming of the employed 1p rate constants is unlikely the reason for the qualitative discrepancy to the measurements. Looking for another reason we recall the original data analysis performed by Fernandes *et al.* [45], who concluded on the crucial role of elementary processes taking place at the boundary between oxide and metal domains. Scrutinizing this possibility requires on the one hand to perform multi-lattice 1p-kMC simulations allowing to toggle between oxide and metal lattice. More importantly though, it also requires an atomistic pathway for the actual oxide reduction process, which shall be established in the next subsection.

7.3.2 Atomistic pathway for oxide reduction

As a natural starting point to establish an atomistic pathway from the intact $\sqrt{5}$ -oxide to reduced Pd(100) we investigate the structural evolution of the surface oxide film when progressively removing O atoms from the oxide lattice. These studies are consistently performed with large supercell models spanning four $\sqrt{5}$ -oxide surface unit cells in a (2×2) arrangement as shown in Fig. 7.4. This model is large enough to study the effects of O removal on one $\sqrt{5}$ cell that is still surrounded by intact $\sqrt{5}$ -oxide units in every direction. For these studies, an important structural aspect of the commensurate $\sqrt{5}$ -oxide/Pd(100) interface is that in every $\sqrt{5}$ -oxide unit cell two of the four Pd atoms in the oxide layer are located in fourfold hollow sites of the underlying Pd(100) substrate, with the other two being rather close to bridge sites, cf. Fig. 7.4. After complete reduction, the Pd atoms of the surface oxide will form domains of a new Pd(100) layer,[117, 184] in which they are also located in the fourfold hollow sites. This means




Figure 7.4: Top view of the supercell model employed in the oxide reduction DFT calculations, spanning four $\sqrt{5}$ -oxide unit cells in a (2×2) arrangement. Each intact $\sqrt{5}$ -oxide unit cell contains four O atoms, labeled from A to D, and four Pd atoms, two of which (marked through hatching) are located in hollow sites of the underlying Pd(100) substrate. The arrows in the upper right three unit-cells indicate the directions in which the two Pd atoms can be shifted to generate a minimal Pd(100) nucleus: This defines the three cases I (upper left unit-cell), II (upper right unit-cell) and III (lower right unit-cell) referred to in the text. Pd atoms are depicted as large green spheres, O atoms as small red spheres; Pd atoms of the underlying Pd(100) substrate are darkened.

that during the local oxide decomposition two out of four Pd atoms in the oxide layer do not necessarily have to significantly change their lateral position.

In the stoichiometric $\sqrt{5}$ -oxide termination expected to initially be present in the reduction experiments there are four O atoms per $\sqrt{5}$ unit cell, two above the Pd atoms of the oxide layer and two below, as shown in Fig. 7.4. In agreement to previous DFT studies performed in smaller supercells [163], we find no significant structural changes upon removal of any single of these four O atoms in one $\sqrt{5}$ unit cell of our model. In particular, none of the oxide Pd atoms change their lateral registry. The same is observed, when removing pairs of O atoms in all possible combinations. Apparently neither single O vacancies, nor O divacancies lead to an unactivated destabilization of the $\sqrt{5}$ -oxide structure.




Figure 7.5: Top view of intermediate structures along the suggested oxide reduction pathway. (a) Divacancy created by removal of two neighboring upper O atoms in one $\sqrt{5}$ -oxide unit-cell (labeled A and B in Fig. 7.4). (b) Intermediate structure formed from (a) through an essentially unactivated shift of the two bridge-site Pd atoms as in case I of Fig. 7.4. This structure is by 0.4 eV more stable than structure (a). (c) Pd(100) nucleus formed from (b) after the activated diffusion of the sub-surface O labeled D in Fig. 7.4 to the fourfold on-surface site in the center of the thus formed Pd(100) nucleus (see text).

In this situation, we focus on a straightforward way to generate a minimal Pd(100) nucleus within the surface oxide layer and explore

if this (together with an increasing number of O vacancies) does not constitute a pathway with at least a very low activation barrier. This pathway centers on the above mentioned realization that two of the four Pd atoms in the $\sqrt{5}$ -oxide unit cell are already located in fourfold hollow sites with respect to the underlying Pd(100) substrate. A direct way to convert one entire $\sqrt{5}$ cell into a Pd(100)-type nucleus would therefore be to simply shift the other two Pd atoms from their close to bridge positions to adjacent hollow sites. As illustrated in Fig. 7.4 there are three ways to realize this: Either both Pd atoms move to hollow sites in the same direction (cases I and II), or the two Pd atoms move to hollow sites in opposite directions (case III). If we enforce such lateral shifts on the Pd atoms of one unit cell for the intact surface oxide, they relax back to their original positions upon geometry optimization. The same holds, if this is done with one O vacancy in the $\sqrt{5}$ -oxide unit cell, or when moving only one of the two Pd atoms. In contrast, metastable structures result when the same movements are applied in the presence of an O divacancy, specifically when the two upper O atoms, labeled A and B in Fig. 7.4, that are directly coordinated to one of the moving Pd atoms are removed. For the three cases defined in Fig. 7.4 the relative energies of the resulting metastable structures compared to the original divacancy structure with the Pd atoms in their regular positions are -0.40 eV (I), -0.25 eV (II) and $+0.17$ eV (III). Cases I and II lead therefore to more stable intermediate structures, and in particular for the most stable case I we furthermore find only a negligibly small activation barrier below 0.1 eV.

This suggests the following initial steps in the oxide reduction pathway: Upon local formation of an O divacancy as shown in Fig. 7.5a, the two bridge-site Pd atoms of the corresponding $\sqrt{5}$ -oxide cell shift laterally and essentially unactivated to yield the intermediate structure shown in Fig. 7.5b. This metastable structure, more stable than the divacancy by 0.4 eV, already features a minimal Pd(100) nucleus with four Pd atoms in hollow sites, but still contains the two O atoms of the $\sqrt{5}$ -oxide cell that were originally below the Pd layer. One of these two, labeled C in Fig. 7.4, automatically moves to a position above the Pd atoms of the Pd(100) nucleus along with the lateral shift of the bridge-site Pd atoms. The other O atom D remains in a sub-surface position as depicted in Fig. 7.5c. We calculate a barrier of 0.66 eV for the diffusion process that lets this O atom pop up to the on-surface site E in the center of the Pd(100) nucleus, cf. Fig. 7.5c, where it is by -0.11 eV more stable than in its original sub-surface site. From there and as further detailed below, the on-surface O atom can be reacted off through Langmuir-Hinshelwood-type CO oxidation or diffuse away, leaving a bare minimal Pd(100) nucleus as the end product of the here suggested oxide reduction pathway. Continued reduction of the oxide layer will then lead to an increasing num-

ber of such Pd(100) nuclei. As the Pd density in the intact surface oxide is 20% lower than in a Pd(100) layer, Pd vacancies will emerge along with the formation of these nuclei. The diffusion of these Pd vacancies will be very fast, with a calculated activation barrier of only ~ 0.2 eV, similar to the diffusion of Pd on Pd(100).[41, 98, 99, 109, 113] Quickly, this will lead to the coalescence of the individual Pd(100) nuclei and with that to the formation of extended Pd(100) islands as seen in experiment.[41, 72–74].

7.3.3 Elementary processes at the Pd(100) nucleus under reduction conditions




Figure 7.6: Top view of the Pd(100) nucleus as formed during the oxide reduction process. Shown are all high-symmetry binding sites explored to assess the binding properties of the nucleus itself (sites A-E) and at its boundary (sites F-K). Pd atoms are depicted as large green spheres, O atoms as small red spheres; Pd atoms of the underlying Pd(100) substrate and of the formed nucleus are darkened.

The Pd(100) nucleus is the central new feature within the obtained oxide reduction pathway. During the oxide reduction process it may emerge upon sufficient local oxygen depletion of the $\sqrt{5}$ -oxide, and may then enable new elementary processes specific to either the nucleus itself or the boundary between nucleus and surrounding oxide. As the next step we therefore proceed to a systematic identification and description of these new processes enabled in the vicinity of a nucleus. As starting point for this we first consider the binding at the different high-symmetry sites illustrated in Fig. 7.6, which represent sites on top of the nucleus and at its boundary. Specifically, sites labeled A-D and site E correspond to Pd(100)-type bridge and

Table 7.1: Calculated DFT binding energies for CO and O at high-symmetry sites on top of the Pd(100) nucleus and at its boundary. Shown are results for the situation with the sub-surface O atom still beneath the nucleus (labeled "w/ sub-surf O") and without the sub-surface O atom (labeled "w/o sub-surf O"). Sites A-K are defined in Fig. 7.6. Fields marked with a letter indicate that the adsorbate moved to the respective site during geometry optimization. Fields marked "-" indicate sites for which no binding was obtained. Binding in sites marked with an asterisk ended up in a threefold coordination after geometry optimization. All energies are in eV.

	w/ sub-surf O		w/o sub-surf O	
	CO	O	CO	O
A	-1.72	E	-1.85	E
B	—	—	-1.73	I
C	-1.40	D	-1.62	-0.98*
D	-2.06	-0.96	-2.06	-1.00
E	D	-1.22	-1.90*	-1.25
F	-1.49	-0.96	-1.72	-1.25
G	-1.27	-0.40	-1.31	-0.47
H	-0.83	J	-1.56	J
I	—	—	B	-0.96
J	H	-0.39	H	-0.75
K	-1.03	-0.30	-1.10	-0.37

hollow sites on the nucleus, respectively. Sites F-J are sites directly at the nucleus/oxide boundary, whereas site K corresponds to a $\sqrt{5}$ -oxide bridge site in the immediate vicinity of the nucleus. Table 7.1 summarizes the computed binding energies at these sites for the two "oxidation" stages of the nucleus along the oxide reduction pathway: With the sub-surface oxygen atom still beneath it, cf. Fig. 7.5b, and without it. The latter situation results as a consequence of the activated diffusion of the O atom to the on-surface hollow site of the nucleus and subsequent reaction or diffusion, *vide infra*, and is the situation depicted in Fig. 7.6.

With the exception of site B directly adjacent to the sub-surface O atom, the effect of the latter atom on the binding properties of the sites on top of the nucleus is rather small. Quite similar binding energies are obtained for these sites for the two cases. *Cum grano salis* these binding energies on the nucleus sites are also quite similar to those computed for the extended Pd(100) surface. Roughly we obtain a preferred O binding at the hollow site E of about -1.25 eV, to be compared to -1.25 eV at the corresponding site on Pd(100), and a preferred CO binding at the bridge sites A-D of about ~ -1.4 - 2.1 eV,

to be compared to -1.93 eV at bridge sites of Pd(100). Without the presence of the sub-surface O atom this also extends to the bridge sites F and H at the interface between the nucleus and the $\sqrt{5}$ -oxide, which again exhibit CO binding properties very similar to regular bridge sites at extended Pd(100). Furthermore considering the close geometric and energetic vicinity of these different sites we therefore also expect similarly low diffusion barriers as on Pd(100) and correspondingly a predominant presence of the adsorbates in the most favorable sites, *i.e.* CO in the bridge site D and O in the hollow site E at the temperatures employed in the oxide reduction experiments.

Multiple adsorption on one nucleus in form of CO adsorption at several bridge sites, or O and CO coadsorption on hollow and bridge site, respectively, is not possible. Also in this respect, adsorption at the Pd(100) nucleus seems very similar to adsorption at extended Pd(100). Likewise the disturbance of the $\sqrt{5}$ -oxide created by the Pd(100) nucleus seems quite short-ranged. Already at the oxide bridge site K directly adjacent to the nucleus, O and CO binding is essentially identical to the binding at the bridge site on the stoichiometric $\sqrt{5}$ termination, *i.e.* the CO binding is calculated as -1.03 eV and -1.10 eV at site K (with and without sub-surface O underneath the adjacent Pd(100) nucleus) and as -1.08 eV on the intact $\sqrt{5}$ -oxide, whereas the O binding energy is 0.30 eV, -0.37 eV and -0.19 eV for the three cases, respectively. Finally, the binding at the interfacial sites F, G, H and I in the presence of the sub-surface oxygen atom and at the interfacial site H without the sub-surface O atom is intermediate to the one obtained on the nucleus and on the intact $\sqrt{5}$ -oxide.

In light of these results, we can now systematically analyze the new elementary processes that may arise due to the nucleus under the reducing conditions of the experiment and assess their qualitative relevance for the oxide reduction process. Regardless of whether the sub-surface O atom is still present underneath the nucleus or not, CO adsorption and desorption processes can take place at the nucleus bridge sites, with a maximum coverage of one CO molecule on a nucleus. CO diffusion on the nucleus itself will be fast, and will lead to a preferential population of bridge site D at the temperatures employed in the reduction experiments. After coalescence of nuclei to Pd(100) domains, equivalent processes (CO adsorption, desorption, diffusion) will occur, with properties as on extended Pd(100). In principle, CO diffusion could also occur from nucleus (or Pd(100) domain) to oxide or vice versa. An important point to realize here, however, is that CO binding at the nucleus (and at Pd(100)) is about 1 eV more stable than at the hollow or bridge sites of the $\sqrt{5}$ -oxide. This already imposes a quite large mere thermodynamic barrier for CO diffusion onto the oxide. Simultaneously CO diffusion processes from the oxide to the nucleus will be rare, owing to the near zero CO coverage on

the surface oxide under the conditions of the reduction experiments, cf. Section 7.3.1.

If all sites A-E of the nucleus are unoccupied, the sub-surface O atom can pop up into the on-surface fourfold hollow site E of the nucleus, with a calculated barrier of 0.66 eV. The reverse process is possible with a corresponding barrier of 0.77 eV. From the nucleus site E the oxygen atom can diffuse to an adjacent vacant hollow site of the surrounding oxide, or vice versa. Since O binding is about 0.4 eV more stable at the oxide hollow sites, this diffusion will preferentially take oxygen away from the nucleus (or coalesced larger Pd(100) domains). For O diffusion between the oxide hollow sites directly adjacent to the nucleus (those above site K in Fig. 7.6) we find a slightly lowered diffusion barrier of 1.1 eV as compared to the 1.4 eV for the diffusion over the intact surface oxide. This lowering can be rationalized with the higher geometric flexibility of the underlying Pd atoms in the partially reduced surface oxide environment. Oxygen diffusion from site E to adjacent bridge sites of the $\sqrt{5}$ -oxide face a similarly large thermodynamic barrier as the CO diffusion processes due to the much weaker O binding at these oxide sites. As these bridge sites are furthermore essentially all unoccupied under the conditions of the reduction experiments, the reverse diffusion processes bringing O atoms from adjacent oxide bridge sites to the nucleus will also occur at a negligibly low rate. Associative oxygen desorption requires two nearby O atoms and thus can only take place either after coalescence of nuclei (then from next-nearest neighbor Pd(100)-type hollow sites) or in case of an individual nucleus by involving one O atom from the nucleus site E and one O atom from nearest neighbor O hollow sites of the oxide. Due to the similar binding energies of O on both metal and oxide, such processes are as rare as they were in the single-lattice 1p-kMC simulations focusing only on the surface oxide. Correspondingly, such processes are unlikely to have a relevant effect on the reduction kinetics.

Finally, new CO oxidation reactions can be enabled at the nucleus or its boundary. One possibility is a reaction of a CO molecule adsorbed on the nucleus with the sub-surface O atom. For a CO molecule occupying the preferred bridge site D we calculate a barrier of 1.27 eV for this process, which is rather high for this process to play a significant role at the temperatures employed in the experiments. Alternatively, the CO molecule could react with O atoms of the surrounding $\sqrt{5}$ -oxide. Focusing on the preferred binding site D Fig. 7.7 illustrates the four corresponding reactions to the closest such O atoms and compiles the calculated barriers for these processes. Only one of these processes exhibits a barrier below 1 eV and should therewith exhibit a similar reactivity as the reaction processes occurring on either Pd(100) or on the $\sqrt{5}$ -oxide. As last possibility, an O atom occupying the on-surface site E of the nucleus could react with a nearby ad-




Figure 7.7: Top view of the Pd(100) nucleus illustrating possible CO oxidation reactions across the nucleus/oxide boundary. The CO molecule is adsorbed at the most stable bridge site D of the Pd(100)-nucleus. The calculated barriers for the four reaction processes are: NW 1.63 eV, SW 1.46 eV, SE 1.54 eV, and E 0.95 eV. Pd atoms are depicted as large green spheres, O atoms as small red spheres; Pd atoms of the underlying Pd(100) substrate and of the formed nucleus are darkened.

sorbed CO molecule. Due to the near zero CO coverage on the oxide during the reduction experiments, such a process is unlikely to occur with a CO molecule adsorbed on the $\sqrt{5}$ though. At the same time, the strong repulsive interactions prevent CO adsorption at the bridge sites A-D of the nucleus, if an O atom occupies the hollow site E. This leaves as only possibility for such a cross-reaction, the reaction with a CO molecule adsorbed at the bridge site F, cf. Fig. 7.6. Both sites, E and F, feature binding properties highly comparable to those of corresponding sites at extended Pd(100). We therefore also expect such a reaction process to exhibit a barrier similar to the corresponding process at Pd(100). In general, such a reaction process can nevertheless not have a strong influence on the oxide reduction process. As O diffusion from the $\sqrt{5}$ -oxide onto the nucleus is rare, this process will predominantly occur after the sub-surface O atom has popped up to the surface. The process then takes place once per formed nucleus and is thereafter unlikely to reoccur.

Summarizing these insights, the qualitatively new feature emerging from the presence of a formed Pd(100) nucleus during the oxide reduction process is the significantly enhanced stabilization of CO at the nucleus as compared to the $\sqrt{5}$ -oxide. Whereas the CO coverage on the latter is close to zero under the reducing conditions of the experiment, a much higher CO coverage can thus be expected at formed nuclei. An oxidation reaction of such CO molecules across the nucle-

us/oxide interface with O atoms of the surface oxide then provides a much more efficient means to deplete the surface oxygen compared to the two mechanisms possible at the $\sqrt{5}$ -oxide alone. There the associative O desorption is a rare process due to the strong O binding energy on the oxide, and the CO oxidation is a rare process due to the low CO coverage at the surface. In contrast, the new reduction process enabled at the oxide/nucleus interface only depends on the CO adsorption, which is comparatively frequent in the reducing environment of the experiment, and on the replenishment of the created O vacancies at the oxide sites at the interface. The latter occurs via O diffusion on the surface oxide, which is not a very fast process, but still faster than the oxide-only reaction mechanism, considering in particular that the number of formed Pd(100) nuclei scales with the increasing number of concomitantly created O vacancies in the oxide layer.

7.3.4 Multi-lattice 1p-kMC model

We now proceed by casting the conceptual picture of the reduction process via the formation of Pd(100) nuclei into an explicit multi-lattice 1p-kMC model. A 1p-kMC model hereby means a list of elementary processes, where each elementary process is defined by a set of conditions that have to be met so that the process can be executed, a set of actions that occur due to the elementary process, and a corresponding 1p rate constant.[86] For regions of Pd(100) and $\sqrt{5}$ -oxide the multi-lattice 1p-kMC model will employ exactly the same 1p-kMC models as used for the single-lattice simulations above. The rate constants for all new elementary processes arising in the context of the Pd(100) nuclei are derived from identical TST expressions as before, *i.e.* following the recipe detailed in Ref. 157.

Once an oxygen divacancy involving two neighboring upper hollow sites on the $\sqrt{5}$ -oxide as illustrated in Fig. 7.5 is formed during the 1p-kMC simulation, the nucleus formation is modeled as a non-activated process, *i.e.* it automatically follows the formation of the divacancy. In practice we realized this in the kmos implementation by assigning this process a very high arbitrary rate constant. In terms of the multi-lattice representation the nucleus formation corresponds to deactivating the two bridge and two upper hollow sites associated with the corresponding $\sqrt{5}$ -oxide unit cell by changing their occupation status to null. In their place, Pd(100) lattice sites are activated. Specifically, this concerns Pd(100) bridge sites at sites labeled A, B, C, D, F, H in Fig. 7.6, the occupation of which is changed from null to empty. Furthermore, a Pd(100) hollow site is created empty at the site labeled E in Fig. 7.6. The quenched-in sub-surface oxygen atom is modeled to remain at its lower hollow $\sqrt{5}$ -oxide site. However, now that the Pd(100) hollow site labeled E in Fig. 7.6 is created, this sub-surface O

atom can hop there with a barrier of $E_a = 0.66$ eV, if the Pd(100) hollow site and its four nearest-neighbor Pd(100) bridge sites are empty in accordance with the site-blocking rules used in the Pd(100) model [82]. From there, the oxygen atom can reverse the diffusion with a barrier $E_a = 0.77$ eV, or, if there are already other immediately adjacent Pd(100) nuclei, it can also diffuse to other nearest-neighbor Pd(100) hollow sites. O diffusion from the nucleus to surrounding $\sqrt{5}$ -oxide (and the reverse back-diffusion process) seems not to play a role for the reduction conditions addressed here. We tentatively included such processes in the 1p-kMC model with a range of barriers fulfilling the thermodynamic reversibility constraint of a 0.4 eV stronger binding at the $\sqrt{5}$ -oxide. The effects on the simulated reduction times reported below were insignificant, which is why we did not explicitly compute the DFT barriers for this pair of processes. The accelerated hollow-hollow diffusion of oxygen on the $\sqrt{5}$ -oxide next to the Pd(100)-nucleus with a barrier of +1.1 eV is accounted for by making the presence of one of the Pd atoms in a Pd(100) lattice position an explicit condition for the elementary process.

The second oxygen atom that is originally in a lower $\sqrt{5}$ -hollow site and pops up to the surface during the formation of the Pd(100) nucleus, cf. Fig. 7.5b, has a binding energy of -1.42 eV, irrespective of whether the other quenched-in sub-surface O atom is still present or not. Even though thus exhibiting a somewhat stronger binding by 0.17 eV, we still approximate this as a regular O atom in a Pd(100) hollow site and correspondingly activate another Pd(100) hollow lattice site at the location that is obtained by going from site E one lattice constant to the left in Fig. 7.6 and directly occupy it with the popped-up O atom. Simultaneously, the lower hollow- $\sqrt{5}$ site in which the oxygen atom resided before the nucleus formation is deactivated by setting its occupation to null.

CO diffusion from possibly existing neighboring Pd(100) sites onto bridge sites of the new Pd(100) nucleus is also allowed, if the nearest-neighbor hollow as well as all up to next-nearest-neighbor bridge sites are empty – again in accordance with the site-blocking rules for CO diffusion in the Pd(100) model.[82] If all Pd(100) nucleus bridge and hollow sites are empty, non-activated CO adsorption is possible onto a bridge site. CO desorption is modeled with a desorption barrier of 1.93 eV and obeys detailed balance with the adsorption process as in the Pd(100) model. As additional reaction processes we only consider the most likely process labeled E in Fig. 7.7 with a barrier of 0.95 eV, involving a CO molecule adsorbed at site D of the nucleus and an O atom adsorbed at a directly adjacent upper hollow $\sqrt{5}$ -oxide site.

With these prescriptions and new processes the multi-lattice kMC simulation will allow for the generation of Pd(100) nuclei, which if generated next to each other will increasingly lead to larger Pd(100)-type patches in the course of a reduction condition simulation. As

a Pd(100) unit-cell of the commensurate $\sqrt{5}$ -oxide/Pd(100) interface exhibits only 4/5 of the Pd atom density, these patches will not correspond to dense Pd(100) terraces though, but contain a regular array of Pd vacancies. In reality, Pd self-diffusion will lead to the coalescence into closely packed Pd(100) islands, and connected with it to an entire class of additional elementary processes. However, in the here proposed atomistic pathway, cf. Fig. 7.5, the Pd(100)-nucleus formation step shifts the two mobile Pd atoms in such a way that the emerging Pd(100) vacancy is located furthest from the relevant $\sqrt{5}$ -oxide/Pd(100) domain boundary, where the cross-reaction E of Fig. 7.7 can take place. As such we anticipate that the coalescence of Pd vacancies through Pd self-diffusion occurs predominantly displaced from the reduction front and then has no effect on the actual reduction kinetics. Within this view, we neglect Pd self-diffusion in the multi-lattice kMC model, such that the end product of a complete reduction of the $\sqrt{5}$ -oxide film corresponds to a Pd(100) surface with a regular array of Pd vacancies.

7.3.5 Multi-lattice 1p-kMC simulations of oxide reduction

With the multi-lattice 1p-kMC model set up, we return to the oxide reduction experiments of Fernandes *et al.* described in Section 7.3.1. The qualitatively new feature introduced by the explicit account of reduced oxide patches in the multi-lattice simulations is the possibility of reduction through cross-reactions at $\sqrt{5}$ -oxide/Pd(100) domain boundaries. In order to enable an efficient realization of such processes we employ kMC simulation cells, in which initially one row of $\sqrt{5}$ -oxide is already removed. For the coordinate system used in this model the reduced row follows along the [010]-direction as illustrated in Fig. 7.8. In the lattice coordinate system of bulk PdO the surface oxide is a stretched layer of Pd(011) (which is equivalent to (101)) and the reduced row follows the [0 $\bar{1}$ 1] direction. Any chosen simulation cell size then mimics initially present domain sizes in the experiments due to surface imperfections like steps, dislocations etc. To assess the sensitivity on this point, we conduct the multi-lattice 1p-kMC simulations using (10 × 20), (20 × 20) and (40 × 20) cells. As demonstrated in Fig. 7.9 we obtain only small variations that do not affect at all the conclusions put forward below. The further analysis of the obtained data is therefore only presented for the (20 × 20) simulation cell.

Figure 7.9 shows the surface oxide coverage as a function of time using the experimental conditions of 5×10^{-11} bar CO pressure, and temperatures of 303 K, 343 K, and 393 K. In contrast to the single-lattice 1p-kMC simulations discussed in Section 7.3.1, very good agreement with the experimentally observed times and trends is obtained – without the need to modify any of the first-principles rate constants. In order to better understand how the new cross-reaction feature




Figure 7.8: Illustration of the initial configuration used in the multi-lattice 1p-kMC simulations. For clarity, only a small (4×8) commensurate $\sqrt{5}$ -oxide cell area of the larger simulation cell is shown. One row of $\sqrt{5}$ -oxide cells has been reduced to Pd(100) nuclei along the (100) direction of the multi-lattice kMC model. Shown in parenthesis are the corresponding crystallographic axes of the PdO crystal.




Figure 7.9: Multi-lattice 1p-kMC simulations of the $\sqrt{5}$ -oxide reduction at the three experimentally employed temperatures. Shown is the fraction of the surface covered by the surface oxide as a function of time using a (20×20) (solid, bold lines), a (10×20) (dashed, thin lines), and a (40×20) (dotted, thin lines) simulation cell. The slight difference in the initial $\sqrt{5}$ -oxide fraction for the different cell sizes results from the fact that in every cell initially one row of oxide is removed.




Figure 7.10: Relative contribution of elementary processes leading to the creation of divacancies and therewith to the ongoing oxide reduction through the formation of Pd(100) nuclei at the surface. "Oxide reaction" denotes divacancy creation through a CO oxidation reaction within the $\sqrt{5}$ -oxide domain, "cross-reaction" denotes divacancy creation through a cross-reaction over the $\sqrt{5}$ -oxide/Pd(100) domain boundary, and "O diffusion" denotes divacancy creation through diffusion of O atoms on the $\sqrt{5}$ -oxide.

leads to this dramatic improvement, we analyze which elementary process immediately precedes the formation of an appropriate divacancy and therewith triggers the creation of a Pd(100) nucleus in the simulations. Figure 7.10 shows these relative contributions and reveals that at lower temperatures reduction is predominately caused by CO oxidation processes within the surface oxide domain, while at the higher temperatures the ongoing reduction is primarily driven by cross-reactions over $\sqrt{5}$ -oxide/Pd(100) domain boundaries. The latter is a natural consequence of the enhanced CO binding at the Pd(100) nuclei as compared to the $\sqrt{5}$ -oxide. At the higher temperatures CO is then predominantly stabilized on the Pd(100) patches at the surface, concomitantly enhancing the cross-reaction compared to the regular CO oxidation within the surface oxide domain. This change in the mechanism behind the oxide reduction leads to a temperature dependence of the reduction time that is impossible to grasp within a single-lattice model, regardless of how one tries to tweak the underlying rate constants.

7.4 SUMMARY AND CONCLUSIONS

We presented a general multi-lattice first-principles kinetic Monte Carlo approach that allows to describe surface morphological transitions within lattice kMC, as long as these transitions occur between structures exhibiting commensurate lattices. This approach allows to evaluate long-time kinetic behavior subject to corresponding transi-

tions and therewith allows to scrutinize their influence. The necessary requirement for the approach is a detailed atomistic pathway for the transition, which then needs to be mapped into the 1p-kMC model in form of stable intermediates at the involved lattice sites and new elementary processes involving these intermediate structures.

As a showcase we illustrated the approach by addressing recent oxide reduction experiments at a Pd(100) model catalyst by Fernandes *et al.* [45]. We developed the required atomistic pathway for the concomitant transition from the $\sqrt{5}$ -surface oxide film to reduced Pd(100) by inspecting surface geometries upon O vacancy creation, systematically identifying stable intermediates by means of geometry optimizations and transition state calculations to connect intermediate geometries. An important rational in this development was the realization that several oxide metal atoms do not necessarily have to change their geometric position during the transition from oxide to pristine metal. The choice to only consider pathways where corresponding position changes are not taking place then leads to a drastic reduction in the space of possibly stable intermediates and could also be a useful strategy to conceive complex reconstruction mechanisms in other systems.

Single-lattice 1p-kMC simulations could not account for the temperature dependence of the reduction time observed in the experiments, not even when abandoning the first-principles character of the simulations and optimizing the rate constants of relevant elementary processes to fit the experimental data. In contrast, very good agreement with the measurements is obtained with the multi-lattice 1p-kMC simulations on the basis of the proposed oxide reduction pathway. The crucial new feature leading to this improvement is the possibility for CO oxidation reactions across $\sqrt{5}$ -oxide/Pd(100) domain boundaries introduced by the explicit account of already reduced Pd(100) nuclei in the simulations. These cross-reactions constitute a largely varying contribution to the overall oxide reduction within the temperature range covered by the experiments, and therewith generate a temperature-dependence of the reduction rate that cannot be effectively described by a single-lattice model.

In the catalysis context, a multi-lattice 1p-kMC model explicitly accounting for a domain boundary between two phases offers in general advanced possibilities to scrutinize the (sometimes emphasized) role of the latter for the catalytic function. Such an interface controlled reactivity has for instance been put forward by Li *et al.* [108], which can now be revisited and explored in detail and with a predictive character on the basis not only of trends on specific reaction barriers, but in the framework of a rigorous multiscale methodology. The present work already contributes to this by highlighting that even though reaction barriers may be quite similar on the metal, the surface oxide and across the metal/oxide domain boundary, the latter

could still be the dominant reaction mechanism due to the steady-state coverages around the metal/oxide interface. For the present model to directly address steady-state CO oxidation activity primarily additional elementary processes are required that describe surface oxidation, in particular dissociative oxygen adsorption at highly O-precovered surfaces e.g. featuring Pd(100) nuclei surrounded by $\sqrt{5}$ -oxide.

Summary

The discovery of materials and materials functionality will in the future without a doubt be accelerated by powerful computational tools that allow to predict the properties and functionalities from an atomistic level. Heterogeneous catalysis is one application area where multiscale computer simulations have already been particularly fruitful to gain a new level of understanding since here the making and breaking of bonds which is governed by quantum mechanics inevitably affects the performance at a macroscopic scale.

Almost 90 years after Taylor introduced the concept of the catalytic surface[181] the exact nature of the catalytic activity even for very simple reactions on elementary surfaces is today still a highly active area of research. The focus of this thesis was the CO oxidation on Pd(100). The stability boundary between the metallic Pd(100) surface termination and a surface covered by an atomistically thin surface oxide (characterized as a $(\sqrt{5} \times \sqrt{5})R27^\circ$ overstructure and henceforth abbreviated as $\sqrt{5}$ -oxide) falls into the range of catalytically relevant regions around stoichiometric feed conditions. Such a transition regime makes the relation between structure and activity particularly intricate. One conjectured scenario is the continuous formation and reduction of such oxidic films as a major driver for the observed catalytic activity. A necessary requirement for making this hypothesis verifiable is a methodology that can account for such complex atomistic detail while simultaneously being able to predict microkinetic observables such as surface coverages and reaction rates.

In this thesis a first-principles multiscale approach has been deployed to address this issue. Density functional theory (DFT) at the generalized gradient approximation level is used to describe slab geometries that model different intermediate stages from the $\sqrt{5}$ -oxide to the Pd(100) surface. DFT is then combined with lattice kinetic Monte Carlo (kMC) to connect information from the electronic structure with phenomena at the mesoscale and beyond. An extension to the lattice kMC method is developed which allows to describe a catalytic surface with ongoing morphological transformations. Emerging schemes which couple the outcome of such atomistic simulations in the form of a boundary condition to continuum scale fluid simulations are hereby an important factor to specifically pursue the computationally highly efficient lattice kMC approach.

In a first step a kMC model for CO oxidation on Pd(100) was constructed to complement an existing model for the $\sqrt{5}$ -oxide termination. The two kMC models for Pd(100) and $\sqrt{5}$ -oxide were then used to separately predict the stability boundaries of the respective termination based on oxygen coverage criteria. Contrary to expectations it was found that the two boundaries do not coincide in phase space. Instead the stability regimes overlap over a wide range of feed conditions between 400-600 K and around stoichiometric reactant ratios at near-ambient pressures. It was found that the tendency of adsorbed CO molecules to hinder oxygen adsorption on Pd(100) extends the stability range of the metallic termination beyond its thermodynamically stable regime. This bistability therefore only exists in the presence of both CO and oxygen and matches quite well with the bistability regime found experimentally by Hendriksen *et al.*[74] on the same surface which is accompanied by oscillatory catalytic activity. The two models also reproduce the experimentally observed kinetics and the relative activity of the two surfaces correctly. The hindered oxygen adsorption model offers hereby an alternative explanation for the origin of the bistability which does not require the complementary roughening of the surface required in the originally proposed rationalization of the oscillatory behavior.

In the next step an atomistic description of the reduction of the $\sqrt{5}$ -oxide by CO was pursued. The essential insight towards a reconstruction mechanism was that half of the palladium atoms in the $\sqrt{5}$ -oxide do not necessarily have to move significantly during the reconstruction. An atomistic pathway was obtained via systematic probing of intermediate geometries using DFT calculations. The resulting pathway is initiated by oxygen divacancies and proceeds via the formation of Pd(100)-nuclei consisting of four palladium atoms. The Pd(100)-nucleus geometry allows for an efficient Langmuir-Hinshelwood CO oxidation process that crosses the Pd(100) to $\sqrt{5}$ -oxide domain boundary.

To account for this reconstruction which obviously deviates from a single crystal structure an extension to the lattice kMC approach had to be devised. One necessary step was to choose a unit cell that is large enough to represent all active sites of both surface terminations simultaneously. Second, in the language of kMC an artificial species was used to deactivate the sites belonging to the currently non-present lattice. In this way lattice reconstruction mechanisms of almost arbitrary sophistication can be described with the advantageous efficiency of lattice kMC.

The developed atomistic pathway and concomitant elementary processes were implemented using this multi-lattice kMC approach. The resulting oxide decay function under a pure CO atmosphere of 5×10^{-11} bar and 303-393 K was found to be in excellent agreement with those found in XPS experiments by Fernandes *et al.*[45]. The

same temperature dependency could not be explained using only the $\sqrt{5}$ -oxide model. Analysis of the resulting kMC trajectories showed that towards the upper end of simulated temperatures the $\sqrt{5}$ -oxide is in fact driven by the cross-boundary reaction as originally conjectured from the experiments.

Outlook

The atomistic $\sqrt{5}$ -oxide reduction pathway was constructed with a focus on CO oxidation. However, the resulting mechanism is independent of the reducing agent itself as long as it is able to efficiently remove oxygen atoms from the surface oxide. Thus, many of the intermediate steps will be useful when constructing similar models with other reducing agents such as H_2 or CH_4 . The guiding protocol for constructing the surface oxide reduction pathway can be abstractly formulated in the following way: Identify which atoms will have to move during the reduction to arrive at stable sites of the underlying metal substrate. Remove locally an increasing number of oxygen atoms and test by use of geometry optimizations for all reasonable combinations of final sites whether the reconstructed geometry is energetically more favorable than the oxidic geometry. Finally calculate activation energies for transition paths connecting the initial configuration with the most stable reconstruction. This protocol could be applied similarly to other elementary transition metal surface oxides. Surface oxides have been found to exist on Pd(111), Rh(111), Rh(100), Rh(110), Pt(110), Ag(111), and Ag(100).[118]

Having established a first model for the reduction of Palladium $\sqrt{5}$ -oxide the complementary next step is to construct a model for the formation of $\sqrt{5}$ -oxide. Compared to the modeling of the reduction the modeling of surface oxide formation is complicated by two factors. First, the $\sqrt{5}$ -oxide surface contains 20% less palladium atoms per unit area compared to the Pd(100) surface. Thus during the formation of $\sqrt{5}$ -oxide inevitably every fifth palladium atom would have to be expelled from the closed packed Pd(100) surface. The additional diffusion of these excess atoms either on top of the metallic surface or even on top of the surface oxide and subsequent accumulation at existing terraces or islands leads to a further complication of the already non-trivial model. Second is the fact that the highest thermodynamically stable oxygen concentration on Pd(100) of 0.25 ML corresponds to the $p(2 \times 2)$ -O adlayer. In comparison the $\sqrt{5}$ -oxide corresponds to an oxygen concentration of 0.8 ML. So far in the Pd(100) model oxygen is not allowed to adsorb above 0.5 ML coverage which is motivated using simple geometrical rules. However, $\sqrt{5}$ -oxide formation will most likely require oxygen to adsorb either on top of a kinetically stabilized $c(2 \times 2)$ -O or near an existing Pd(100) to $\sqrt{5}$ -oxide phase boundary to create the critical oxygen density required for $\sqrt{5}$ -oxide formation. To

obtain predictive rate constants for these processes one or more oxygen sticking coefficients for these oxygen covered geometries will be needed. Since the first-principles calculation of sticking coefficients for pre-covered surfaces is a very active area of research[62, 114], calculation of oxygen adsorption at higher coverages will most likely require further methodological developments.

Detailed experimental studies can be very instrumental for validating first-principles guided multi-scale models as has been shown for the case of $\sqrt{5}$ -oxide reduction in this thesis. Similar studies on the temperature dependent formation of $\sqrt{5}$ -oxide could prove equally useful to scrutinize a corresponding multi-lattice kMC model. Finally, after both $\sqrt{5}$ -oxide reduction and formation is accounted for one can make quantitative predictions about the dynamic surface composition of Pd(100) for CO oxidizing conditions and its corresponding catalytic activity.

Part III

APPENDIX

CASTEP - A TECHNICAL INTRODUCTION FOR PRACTITIONERS

Plane-wave DFT Calculations

The core functionality of any Density Functional Theory (DFT) package is to solve the Kohn-Sham equation for a given system

$$\hat{\mathbf{H}}_{\text{KS}}(\mathbf{r}) \phi_i(\mathbf{r}) = \varepsilon_i \phi_i(\mathbf{r})$$

which in physical terms is a single-particle equation, where each $\phi_i(\mathbf{r})$ denotes the i th single particle state. As indicated already we will work in spatial representation and drop explicit representation of spin for sake of simplicity. The idea behind DFT is that $\hat{\mathbf{H}}_{\text{KS}}$ is constructed ideally in such a way that the solution ϕ_i yields a density

$$\rho(\mathbf{r}) := \sum_i |\phi_i(\mathbf{r})|^2$$

that is identical to the true ground-state density of the many-particle electronic wavefunction of our system. The Kohn-Sham operator can be written as

$$\begin{aligned} \hat{\mathbf{H}}_{\text{KS}}(\mathbf{r}) &\hat{=} -\frac{\hbar^2 \nabla^2}{2m_e} + V_{\text{eff}}(\mathbf{r}) \\ &= -\frac{\hbar^2 \nabla^2}{2m_e} + \underbrace{\sum_j \int d\mathbf{r}' \frac{\langle \phi_j(\mathbf{r}') | \phi_j(\mathbf{r}') \rangle}{|\mathbf{r} - \mathbf{r}'|}}_{V_{\text{Hartree}}[\rho](\mathbf{r})} + V_{\text{ext}}(\mathbf{r}) + V_{\text{xc}}[\rho](\mathbf{r}) \end{aligned}$$

where V_{ext} subsumes the potential acting on the electronic density due to the ionic cores or an external field, and the static potential energy between ions, V_{Hartree} , represents the electrostatic potential interacting between the electrons. Finally, the so called exchange-correlation functional V_{xc} constructed such that it corrects all differences between the single-particle wavefunctions and the many-body wavefunction. The construction of such a functional is a topic for a group of experts and as a practitioner I will not be concerned with it here.¹ Instead I will view this as given and highlight practical issues involved in solving the Kohn-Sham equations on a computer using a plane-wave basis set with the CASTEP package[28]. It is crucial to note that V_{xc} is a functional of the density and therefore depends on

¹ Throughout this summary I will use the GGA-PBE[144] as exchange correlation functional

the single-particle wave function. This means that for each wavefunction found from solving the Kohn-Sham equation we get a new Kohn-Sham operator. Thus one keeps solving and constructing Kohn-Sham operators until the differences between successive wavefunctions are below a certain threshold (which will be explained below) and one speaks of a self-consistent solution to the Kohn-Sham equations.

Computationally it is efficient to represent the single-particle wavefunction by introducing a not necessarily orthogonal basis²

$$\phi_i(\mathbf{r}) = \sum_{\alpha} c_{\alpha}^i \varphi_{\alpha}(\mathbf{r})$$

which lets us write the Kohn-Sham equation as a matrix equation that corresponds to what is known in the Hartree-Fock context as Roothaan-Hall³ equation⁴

$$\mathbf{H}_{\alpha\beta}^i \mathbf{c}_{\beta}^i = \varepsilon^i \mathbf{S}_{\alpha\beta}^i \mathbf{c}_{\beta}^i \quad (\text{A.1})$$

for which efficient numeric methods exist. The matrix elements are defined as

$$\mathbf{H}_{\alpha\beta} = \int d^3r \varphi_{\alpha}^*(\mathbf{r}) \left[-\frac{\hbar^2 \nabla^2}{2m_e} + V_{\text{eff}}(\mathbf{r}) \right] \varphi_{\beta}(\mathbf{r}) \quad (\text{A.2})$$

and

$$\mathbf{S}_{\alpha\beta} = \int d^3r \varphi_{\alpha}^*(\mathbf{r}) \varphi_{\beta}(\mathbf{r}). \quad (\text{A.3})$$

The wavefunction in a crystal lattice needs to obey Bloch's theorem (switching to a collective index $i \rightarrow i, \mathbf{k}$)

$$\phi_{i\mathbf{k}}(\mathbf{r} + \mathbf{R}) = e^{i\mathbf{k} \cdot \mathbf{R}} \phi_{i\mathbf{k}}(\mathbf{r})$$

which through the wave vector quantum number \mathbf{k} alone represents an infinite number of possible states.

For solid-state and various related systems plane-waves are a natural choice for several reasons. If we cast them in the form

$$\phi_{i\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{G}} c_{i\mathbf{k},\mathbf{G}} \varphi_{i\mathbf{k},\mathbf{G}} = \sum_{\mathbf{G}} c_{i\mathbf{k},\mathbf{G}} \frac{e^{i\mathbf{r} \cdot (\mathbf{k} + \mathbf{G})}}{\sqrt{\Omega}} \quad \Omega : \text{volume of the simulated cell}$$

² in this section a superscript is merely a convenient place for indices

³ For the closed shell case to be specific. But we have to solve the same type of equation for the restricted open shell case.

⁴ using the sum convention

the Bloch theorem is automatically satisfied if \mathbf{G} is a lattice vector of the reciprocal lattice defined by the simulated cell. Then the matrix elements in (A.1) become particularly simple, namely:

$$\begin{aligned} \langle \phi_{\mathbf{k},\mathbf{G}} | T | \phi_{\mathbf{k},\mathbf{G}'} \rangle &= \frac{1}{\Omega} \int d^3\mathbf{r} e^{-i(\mathbf{k}+\mathbf{G})\cdot\mathbf{r}} \frac{-\hbar^2 \nabla^2}{2m_e} e^{i(\mathbf{k}+\mathbf{G}')\cdot\mathbf{r}} \\ &= \frac{1}{\Omega} \int d^3\mathbf{r} e^{-i(\mathbf{k}+\mathbf{G})\cdot\mathbf{r}} \frac{\hbar^2 (\mathbf{k} + \mathbf{G})^2}{2m_e} e^{i(\mathbf{k}+\mathbf{G}')\cdot\mathbf{r}} \\ &= \frac{1}{\Omega} \frac{\hbar^2 (\mathbf{k} + \mathbf{G})^2}{2m_e} \int d^3\mathbf{r} e^{i\mathbf{r}\cdot(\mathbf{G}-\mathbf{G}')} = \frac{\hbar^2 (\mathbf{k} + \mathbf{G})^2}{2m_e} \delta_{\mathbf{G},\mathbf{G}'} \end{aligned}$$

$$\begin{aligned} \langle \phi_{\mathbf{k},\mathbf{G}} | V_{\text{eff}}(\mathbf{r}) | \phi_{\mathbf{k},\mathbf{G}'} \rangle &= \frac{1}{\Omega} \int d^3\mathbf{r} e^{-i(\mathbf{k}+\mathbf{G})\cdot\mathbf{r}} V_{\text{eff}}(\mathbf{r}) e^{i(\mathbf{k}+\mathbf{G}')\cdot\mathbf{r}} \\ &= \frac{1}{\Omega} \int d^3\mathbf{r} e^{-i(\mathbf{G}-\mathbf{G}')\cdot\mathbf{r}} V_{\text{eff}}(\mathbf{r}) \\ &= \tilde{V}_{\text{eff}}(\mathbf{G} - \mathbf{G}') \end{aligned}$$

In the language of linear algebra this means: in the plane-wave basis the full Hamiltonian for one SCF step has block diagonal form and each block belongs to exactly one \mathbf{k} -point.

So the Kohn-Sham equation that needs to be solved for each point in the first Brillouin zone simply becomes

$$\sum_{\mathbf{G}'} \tilde{V}_{\text{eff}}(\mathbf{G} - \mathbf{G}') c_{\mathbf{k},\mathbf{G}'} = \left[\varepsilon_{\mathbf{k}} - \frac{\hbar^2 (\mathbf{k} + \mathbf{G})^2}{2m_e} \right] c_{\mathbf{k},\mathbf{G}}. \quad (\text{A.4})$$

After each iteration step the density and the electronic energy can be calculated according to

$$\begin{aligned} E_{\text{tot}} &= \frac{1}{V_{\text{BZ}}} \sum_i^{\text{occ.}} \int_{\text{BZ}} d^3\mathbf{k} \varepsilon_{\mathbf{k}} \\ \rho(\mathbf{r}) &= \frac{1}{V_{\text{BZ}}} \sum_i^{\text{occ.}} \int_{\text{BZ}} d^3\mathbf{k} |\phi_{\mathbf{k}}(\mathbf{r})|^2. \end{aligned}$$

While theoretically speaking this is still impossible to do in finite time on any computer, since there is still a summation of infinitely many \mathbf{k} -points and the wavefunction of each band at each \mathbf{k} -point consists of infinitely many basis functions, experience shows that a finite number of each is sufficient to compute well-converged quantities.

Well-converged is a tricky attribute here, because the meaning very much depends on the context. I think one should aim to reach convergence that is at least within the accuracy of corresponding experimental measurement. However if the calculated property is used to calculate other quantities it may be quite hard to extrapolate how a small error in these technical settings affects the final quantity.

Fortunately it is usually sufficient to sample the Brillouin zone at a very finite number of \mathbf{k} -points and each calculation only couples loosely in computational terms through the density in the Hamiltonian so it can be parallelized over different processors efficiently. The standard method to use as few points to sample the Brillouin zone as possible is the one by Monkhorst and Pack[132] which produces a uniform grid in \mathbf{k} -space.

FEATURES OF PLANE-WAVE BASIS SETS A good general discussion of plane-wave basis sets can be found in the book by Martin[120]. Here I will merely look at some practical issues. A great advantage of plane-waves is the fact that convergence of the basis set is particularly straightforward. One includes in the plane-wave expansion all \mathbf{G} vectors with

$$|\mathbf{k} + \mathbf{G}| < \sqrt{\frac{2m_e E_{\text{cut}}}{\hbar^2}}$$

and increases the cut-off energy E_{cut} until the quantity of interested does not change beyond a desired tolerance window.

To be more specific I estimate the number of plane-waves used in my rectangular simulation cell for the Γ -point ($\mathbf{k} = 0$). So the question is: how many $\mathbf{n} \in \mathbb{N}^3$ exist such that

$$|\mathbf{n} \cdot \mathbf{A}^{-1}| < \sqrt{\frac{2m_e E_{\text{cut}}}{\hbar^2}}$$

where \mathbf{A}^{-1} is 2π times the inverse of the unit cell. The volume of an ellipsoid is $V = \frac{4}{3}\pi abc$, where a, b, c are the radii along three different axes. Casting this in the form of the commonly used units we get

$$\begin{aligned} N_{\text{pw}} &= \frac{4}{3}\pi \sqrt{\frac{2m_e eV^3}{\hbar^2}} \left(\frac{\text{\AA}^3}{2\pi}\right)^3 a b c (E_{\text{cut}})^{3/2} \\ &\approx 2.3 \cdot 10^{-3} a b c E_{\text{cut}}^{3/2} \end{aligned}$$

Using a simulation cell of $6.24\text{\AA} \times 6.24\text{\AA} \times 30.19\text{\AA}$ and a cut-off energy of 450 eV we find roughly 25,500 plane waves.

Using the plane-wave basis requires discrete Fourier transformations (DFT⁵) between real space and Fourier space (often called \mathbf{G} space) on at several stages in the calculation. This has been made economical by using a numerical technique known as fast Fourier transform (FFT). The key idea behind FFT is to reduce the number of required mathematical operations by decomposing the Fourier transform

$$X[k] = \sum_{j=1}^n X[j] \omega_n^{jk}, \quad \omega_n = \exp(2\pi i/n)$$

⁵ no to be confused with Density Functional Theory

into blocks of smaller discrete Fourier transformation. The results from each block can be reused in the overall Fourier transformation without the need to recalculate them. FFTW₃ is considered the fastest implementation of this idea, which works best if n can be decomposed into many small integer factors (*i.e.* some power of 2 is ideal) and worst if n is prime. FFTW₃ is designed to use an optimal decomposition for each case of n and each computer architecture.[49]

A clear disadvantage of a plane wave basis set is that they are extremely inefficient at representing a quickly oscillating wavefunction such as typically found close to the atomic nucleus. Therefore the use of pseudo-potentials is intimately connected to a plane wave basis set. A pseudo-potential essentially divides the problem of solving the full Kohn-Sham equation into states participating in chemical bonding (valence electrons) and those not participating in chemical bonding (non-valence or core electrons). Then one first solves the Kohn-Sham equation once per atomic species and extracts from the solution an effective potential as felt by the valence electrons due to the ionic core surrounded by non-valence electrons. The construction of pseudo-potentials is out the scope of this summary. In general they should be constructed such that the pseudo-potential is a *soft* as possible thus allowing for a smaller basis set for the remaining electrons but at the same time *transferable*, which means it works equally well for different environments. One should realize that this idea becomes more difficult to apply to transition- or late-transition metals since here the distinction between valence and non-valence electrons is less clear.

In the present case of Palladium there are therefore two candidate pseudo-potentials to choose from: one treating 18 electrons explicitly which we will call *on-the-fly* (OTF) pseudo-potentials and one treating 10 electrons explicitly which we will call *Materials Studio* (MS) pseudo-potentials for the lack of a better name. From the computational standpoint the latter is very attractive since the total number of bands enters the runtime at least linearly but only the former has proven reliable enough to calculate CO binding energies on Pd(100) in good agreement with reference calculations.

GRIDS IN PLANE-WAVE CALCULATIONS Even though a plane-wave code solves the Kohn-Sham equations in reciprocal space, a real space density has to be calculated to apply the Kohn-Sham Hamiltonian and usually to update the exchange-correlation potential which is a functional of the density. To Fourier transform the density a \mathbf{k} -space grid that has twice the size of the corresponding basis set cut-off is fully sufficient since the product of two wavefunctions has a \mathbf{k} vector with a modulus bound by twice the maximum \mathbf{k}

$$\exp(ix\mathbf{k}_1) \exp(ix\mathbf{k}_2) = \exp(ix(\mathbf{k}_1 + \mathbf{k}_2))$$

On the other hand one generally needs a finer grid, when Fourier transforming the charge due to the pseudo-potential. For better efficiency CASTEP performs two FFTs one two different grids. The grid size is set to $1.75|\mathbf{k}_{\text{cut}}|$ instead of $2|\mathbf{k}_{\text{cut}}|$ since it is generally considered to be “equivalent” but this empirical setting needs to be checked to each system. The grid for transforming the so-called augmented charges is usually chosen to be slightly larger. This technique of *double-gridding* is described by Laasonen *et al.*[107].

Geometry Optimization

As a last topic one I want to touch on the BFGS algorithm used by CASTEP to perform geometry optimizations. BFGS is an iterative procedure that can be motivated from a quadratic expansion to the potential energy surface $f(\mathbf{x})$

$$f(\mathbf{x}) \approx f(\mathbf{x}_i) + (\mathbf{x} - \mathbf{x}_i) \nabla f(\mathbf{x}_i) + \frac{1}{2} (\mathbf{x} - \mathbf{x}_i) \mathbf{A} (\mathbf{x} - \mathbf{x}_i)$$

where \mathbf{A} is called the Hessian of f and \mathbf{x}_i represents the configuration of the system at hand at the i th step. The goal is to find a configuration \mathbf{x} such that $\nabla f(\mathbf{x}) = 0$. Thus

$$0 = \nabla f(\mathbf{x}_i) + \mathbf{x} \mathbf{A} \rightarrow \mathbf{x}_{i+1} = -\mathbf{A}^{-1} \nabla f(\mathbf{x}_i)$$

which simply describes the Newton algorithm. Since in a DFT calculation ∇f is simple to calculate due to the Hellmann-Feynman theorem but \mathbf{A} is prohibitively expensive many methods have been developed to circumvent this of which BFGS is widely accepted. The idea is to find increasingly better approximations \mathbf{H}_i to \mathbf{A}^{-1} using the following formula to update \mathbf{H}

$$\mathbf{H}_{i+1} = \mathbf{H}_i + \left(1 + \frac{\gamma \mathbf{H}_i \gamma}{\delta \gamma} \right) \frac{\delta \delta}{\delta \gamma} - \left(\frac{\delta \gamma \mathbf{H}_i + \mathbf{H}_i \gamma \delta}{\delta \gamma} \right)$$

where $\delta = \mathbf{x}_{i+1} - \mathbf{x}_i$ and $\gamma = \nabla f(\mathbf{x}_{i+1}) - \nabla f(\mathbf{x}_i)$ [48].

CASTEP FLOWCHART To summarize where different parameters enter the calculation and how some of them affect the runtime, I will give a short flowchart. The runtime information has been obtained partially from the CASTEP workshop slides as well as studying the CASTEP code itself.

assumptions: we are using density mixing without spin-polarization and check for force, energy, and eigenvalue convergence.

symbols: N_b : number of bands $\sim \frac{1}{2}$ number of valence electrons, N_{pw} : number of planewaves, E_{cut} : cut-off energy, s_g : grid scale, s_{fg} : fine grid scale

1. initialize wavefunction using random numbers.

$$N_{\text{pw}} \propto V_{\text{cell}} (E_{\text{cut}})^{3/2} N_b$$

2. **start loop:** for each **k**-point construct Hamiltonian $\mathbf{H}_{\mathbf{k}S}^{(k)}$ which includes
 - a) calculate real space density, $O(N_b s_g N_{pw} \ln(s_g N_{pw}))$
 - b) augment density with pseudo charges, $O(N_b s_{fg} N_{pw} \ln(s_{fg} N_{pw}))$
3. iteratively diagonalize $\mathbf{H}_{\mathbf{k}S}^{(k)}$ until eigenvalue tolerance is met, $O(N_{pw})$, orthogonalize band $O(N_b^2 N_{pw})$
4. find occupancy of bands using *smearing width*
5. calculate forces using Hellmann-Feynman (if checked), $O(N_b N_{atoms}^2)$ for non-local potential
6. if force and total energy threshold is met exit, otherwise jump to 2.

CONVERGENCE TESTS

This part of the appendix serves to scrutinize the accuracy of the technical settings chosen when using CASTEP to calculate the input required for constructing the kMC models. The energy differences of partially reconstructed surface geometries and relative binding energies of O and CO on defect geometries are used as the primary property for guiding the construction of the atomistic pathway of the $\sqrt{5}$ -oxide. Therefore the numerical convergence of these quantities and their contributions is analyzed in detail in the following.

The absolute accuracy of these calculated properties is determined by the numerical settings of the DFT software package besides the exchange-correlation functional. The PBE functional is used throughout this thesis. The optimal settings in the DFT package are a trade-off between accuracy and computational speed. Experience shows that errors due to the GGA xc functional can be estimated on the order of 0.1–0.2 eV. This serves as a benchmark for all other settings. The goal is therefore to not worsen the accuracy due to technical settings beyond the error of the xc-functional with the smallest possible computational effort. We therefore set the target for the numerical error to be 0.03 eV or less due to any single numerical parameter.

The CASTEP package is a plane-wave pseudo potential code and therefore uses the cut-off energy as the single parameter to control the size and accuracy of the basis set. Plane wave basis sets are usually used in combination with pseudo potentials to reduce the required number of plane wave primitives as explained before. Throughout this thesis we rely on the ultra-soft pseudo-potentials[191] using the default on-the-fly (OTF) settings provided by CASTEP version 5.5.

Furthermore the density of the deployed k-point grid determines the quality of the Brillouin zone sampling in periodic calculations. Therefore the cut-off energy and the k-point grid are the two central technical settings which need to be chosen for reliable accuracy for the system under study. These settings are then used consistently throughout this thesis. To improve the convergence of the metallic system a smearing scheme is required to produce a slightly smoothed transition from occupied to unoccupied states. The Gaussian smearing scheme with a conservative width of 0.05 eV was used throughout this thesis.

Asymmetric slab geometries are used as the physical model to calculate binding energies and reaction barriers. To accurately decouple slabs in the direction perpendicular to the surface a sufficiently large vacuum has to be placed between slabs to avoid spurious non-




Figure B.1: The prototypical surface slab model consisting of a layer $\sqrt{5}$ -oxide and one layer of Pd(100) metal atoms fixed at the bulk positions. The solid black lines indicate the unit cell placing at least 10 \AA of vacuum between two surfaces.

physical interactions across periodic boundary conditions. The necessary minimal vacuum thickness will be determined, too. Lastly the slab geometries required to represent the $\sqrt{5}$ -oxide to Pd(100) reconstructions need to be fairly large containing up to 20 Pd atoms per fcc(100) layer. Therefore the slab geometry will be chosen quite thin (the $\sqrt{5}$ -oxide layer on top of one frozen Pd(100) layer). We will furthermore estimate the magnitude of error that is introduced due to the finite slab thickness.

Molecular Binding Energies




Figure B.2: Molecular binding energies for the PBE xc functional as a function of the cut-off energy. The basis set is more accurate the larger the basis set

The binding energy of oxygen was calculated in a rectangular cell of size $14\text{\AA} \times 15\text{\AA} \times 16\text{\AA}$ to avoid interaction between periodic images. The binding energy for oxygen was calculated using the formula

$$E_{\text{O}_2}^{\text{bind}} = E_{\text{O}_2} - 2E_{\text{O}} \quad (\text{B.1})$$

and for CO

$$E_{\text{CO}_2}^{\text{bind}} = E_{\text{CO}} - (E_{\text{O}} + E_{\text{C}}). \quad (\text{B.2})$$

The binding energy for oxygen (cf. Fig. B.2) can reasonably reproduce reference pseudo-potential calculations such as -5.81 eV[66]. The value for an energy cut-off at 400 eV of -5.86 eV deviates only slightly from the most accurate value at 500 eV of -5.87 eV. The binding energy for CO using the same setting is calculated as -11.57 eV at 400 eV compared to -11.58 eV at 500 eV cut-off energy.

Bulk energies




Figure B.3: Bulk cohesive energy of fcc Pd with respect to cut-off and k-point grid using an MP k-point grid of $[14 \times 14 \times 14]$ and energy cut-off of 400 eV respectively.

The bulk fcc structure of Pd was optimized using the tensor based optimization included in CASTEP. The reference calculation for Pd was again performed in a periodic box of size $14\text{\AA} \times 15\text{\AA} \times 16\text{\AA}$ and the cohesive energy calculated using the formula

$$E_{\text{coh}} = E_{\text{Pd,bulk}} - E_{\text{Pd,Atom}}. \quad (\text{B.3})$$

For a periodic structure such as bulk palladium Brillouin zone sampling becomes necessary to capture the band structure of the material. To find reasonably converged settings for each the cut-off energy was first tested using a moderate Monkhorst-Pack k-point grid[132] of $[5 \times 5 \times 5]$ using the conventional unit cell. The resulting cohesive energy can be seen in Fig. B.3. From this it was decided that a energy cut-off of 400 eV would suffice here as well. Using the cut-off energy of 400 eV the setting for the optimal k-point grid is tested next. The resulting graph shows that a k-point grid of $[10 \times 10 \times 10]$ suffices to obtain the desired accuracy. Using these settings the lattice constant was calculated as 3.938 eV which is used to construct the geometries throughout this thesis. [32, 202]

Adsorption Energies




Figure B.4: Surface formation energy with respect to cut-off for fcc Pd using a MP k-point grid of $[10 \times 10 \times 1]$.

Absolute surface formation energies cannot be calculated with the asymmetric slabs used in this thesis. However relative total energies are plotted with respect to the most precise value in Fig. B.4. Again they confirm that an energy cut-off of 400 eV and a k-point grid of $[10 \times 10 \times 1]$ is sufficient for this geometry. Only one k-point is used in the z direction perpendicular to the slab surface as no periodic band structure is expected in this geometry. The adsorption energy is calculated using the formula

$$E_{\text{mol}}^{\text{ads}} = E_{\text{mol@slab}} - (E_{\text{mol}} + E_{\text{slab}}). \quad (\text{B.4})$$

All geometry optimization were performed using the BFGS optimizer with a force threshold of $0.05 \text{ eV}/\text{\AA}$.




Figure B.5: CO and O adsorption energy for fcc Pd and PdO using a plane-wave energy cut-off of 400 eV. CO is adsorbed the bridge site of a (1×1) slab and O is adsorbed to the hollow site of a (2×2) slab.

The binding energies for O and CO as a function of cut-off energy are plotted in Fig. B.5. Again the cut-off energy proves sufficient for the required accuracy.

As a final test the effect of the slab thickness is tested. Due to the fairly large surface model needed for describing the surface oxide reconstruction a very thin slab model of consisting of the $\sqrt{5}$ -oxide layer and one layer of Pd(100) atoms fixed at its bulk positions is used throughout this thesis as depicted in Fig. B.1. This physical model must have an effect on the absolute binding energies on top of this




Figure B.6: Surface energy for fcc Pd(100) and CO adsorption energy as a function of the vacuum separation between slab using a MP k-point grid of $[4 \times 4 \times 4]$ for the commensurate $\sqrt{5}$ -oxide surface unit cell and a plane-wave cut-off energy of 400 eV.




Figure B.7: Binding energy of a CO molecule in the bridge position of the $\sqrt{5}$ -oxide as a function of the number of Pd(100) layers below. The lowest layer of Pd(100) atoms are frozen at their bulk positions.

surface. A typical deviation can be estimated by inspecting Fig. B.3. We see that this may introduce error in binding energies on the order of 0.05-0.1 eV. Therefore special care needs to be taken to interpret differences in binding energies if they are on this order of magnitude. However in particular for relative changes between similar geometries and reaction barriers the error is expected to be smaller since the deviation due to the thin slab is contained in both cases (error cancellation).

THE ASE TO CASTEP INTERFACE

A calculator module or interface between the Atomic Simulation Environment (ASE) and CASTEP has been developed and used throughout this thesis. This Appendix gives an overview over its capabilities and its design guidelines.

The community of researchers employing density functional theory software is segmented across different implementations such as Wien2k, VASP, CASTEP, FHI-Aims, GPAW, QuantumEspresso, CP2K, Gaussian, and many more. The difference between these software packages are usually given by the chosen basis-function and (pseudo-)potential approach and a different feature sets on top of the core functionality and besides non-functional characteristics such as license and computer language. Each code typically has some unique feature or functionality but all in all none of them is suitable for all intends and purposes of first-principles calculations. The fact that each package typically comes with its unique input file format and command set creates a serious barrier to the flow of information, verification, and cross-use and therefore synergies between codes. To this end the Atomic Simulation Environment (ASE) pioneered by Bahn and Jacobsen [3] pragmatically mitigates the problem by developing an application programming interface (API) with which in principle all those different packages can be used via one unified interface and for one calculation job several different codes can be used easily. In the context of ASE each first-principles code is accessed via a so called Calculator module which translates the unified description of the calculation job into input files and command line calls suitable for the corresponding package. This has the nice benefit that higher-level functionality such as geometry optimization, molecular dynamics, vibrational analysis, QM/MM embedding, geometry generation, transition state searches, phonon calculations and many more have to be implemented only once and are readily available for all connected packages and thereby extend the feature set of all the supported electronic structure packages simultaneously.

WHAT MAKES THE CASTEP ASE CALCULATOR UNIQUE? A problem that all ASE calculators face is how to stay up to date with changes in the underlying quantum chemical software. Fortunately new features are continuously added to electronic structure code packages which may reflect in new keywords in input file. In order to allow users to access to the new features the interfaces typically need to be updated as well to know the correct syntax or at least loca-

tion of the new keywords. Many electronic structure codes use more than one input file. Therefore even if the order in the input file is irrelevant it needs to be written to the correct input file. The CASTEP code is also continuously developed with new version releases approximately once per year. The ASE-CASTEP calculator is unique in this sense as it generates all known keywords from the compiled binary installed on the user's system. The cost of maintaining this interface between regular version updates is therefore marginal. This auto maintenance feature is nicely facilitated by the online `-help` function of CASTEP. For instance issuing the command `castep -help task` on the command line returns concise information on the task keyword as reproduced below:

```
Help information on PARAMETERS keywords:
      TASK      *! type of calculation!*
Type: String                                     Level: Basic

Defines the type of calculation to be performed. Allowed values:
SINGLEPOINT, BANDSTRUCTURE, GEOMETRYOPTIMIZATION, MOLECULARDYNAMICS, OPTICS,
TRANSITIONSTATESEARCH, PHONON, EFIELD, PHONON+EFIELD, THERMODYNAMICS, WANNIER,
MAGRES, ELNES, GENETICALGOR Default value : SINGLEPOINT
```

From this string alone we can readily extract a brief description, the format of the expected arguments (String, Float, Integer, Vector, ...), the location of the arguments (`.cell` or `.param` file and a brief description. Upon first invocation a keyword file is generated on the user's computer containing exactly the available keywords for the installed version. When the CASTEP calculator is invoked for the first time in a fresh ASE installation, the interface will try to automatically fetch the available keywords in the installed binary by inspecting the string returned by the command `castep -help all`. The interface will then parse the result of the online documentation to determine its format, the file location, and the online description. It will then write the collected information to a file named `castep_keywords.py` in the current directory. The gathered information is stored in the form of valid python code so that it can be easily accessed at runtime. This `castep_keywords.py` has to be copied by the user into the calculators directory of the ASE installation. Afterwards and subsequent times the CASTEP calculator is invoked all available keyword are directly obtained from the keywords file. The available options can then be used in a straightforward and pythonic way, e.g. by setting

```
calc.param.task = 'GeometryOptimization'
calc.param.cut_off_energy = 400
calc.cell.kpoints_mp_grid = '4 4 1'
...
```

in an ASE deploying python script.

This way the ASE-CASTEP calculator readily obtains a full support of all CASTEP keywords and minimizes the amount of required maintenance work as the CASTEP code itself evolves. The ASE-CASTEP

has been submitted to ASE code base and already found a sizeable and growing user base among CASTEP users.

LIST OF PUBLICATIONS

- [10] S. Blomberg, M. J. Hoffmann, J. Gustafson, N. M. Martin, V. R. Fernandes, A. Borg, Z. Liu, R. Chang, S. Matera, K. Reuter, and E. Lundgren. In situ x-ray photoelectron spectroscopy of model catalysts: At the edge of the gap. *Phys. Rev. Lett.*, 110:117601, Mar 2013. doi: 10.1103/PhysRevLett.110.117601. URL <http://link.aps.org/doi/10.1103/PhysRevLett.110.117601>. (pages 3 and 71.)
- [83] Max J. Hoffmann and Karsten Reuter. Co oxidation on $\text{pd}(100)$ versus $\text{pd}(101)$ - $(\sqrt{5} \times \sqrt{5})r27^\circ$: First-principles kinetic phase diagrams and bistability conditions. *Topics in Catalysis*, 57(1-4):159–170, 2014. ISSN 1022-5528. doi: 10.1007/s11244-013-0172-5. URL <http://dx.doi.org/10.1007/s11244-013-0172-5>. (page 71.)
- [84] Max J Hoffmann and Karsten Reuter. Detailed kinetics from an atomistic picture: Multi-lattice first-principles kinetic monte carlo simulations of $\text{pd}(100)$ surface oxide reduction, 2014. (in preparation). (page 91.)
- [85] Max J Hoffmann, S Matera, and K Reuter. kmos: A lattice kinetic monte carlo framework. *Computational Physics Communications*, 185(7):2138–2150, 2014. doi: 10.1016/j.cpc.2014.04.003. URL <http://dx.doi.org/10.1016/j.cpc.2014.04.003>. (page 23.)

BIBLIOGRAPHY

- [1] Hironori Arakawa, Michele Aresta, John N. Armor, Mark A. Barteau, Eric J. Beckman, Alexis T. Bell, John E. Bercaw, Carol Creutz, Eckhard Dinjus, David A. Dixon, Kazunari Domen, Daniel L. DuBois, Juergen Eckert, Etsuko Fujita, Dorothy H. Gibson, William A. Goddard, D. Wayne Goodman, Jay Keller, Gregory J. Kubas, Harold H. Kung, James E. Lyons, Leo E. Manzer, Tobin J. Marks, Keiji Morokuma, Kenneth M. Nicholas, Roy Periana, Lawrence Que, Jens Rostrup-Nielson, Wolfgang M. H. Sachtler, Lanny D. Schmidt, Ayusman Sen, Gabor A. Somorjai, Peter C. Stair, B. Ray Stults, and William Tumas. Catalysis research of relevance to carbon management: Progress, challenges, and opportunities. *Chemical Reviews*, 101(4):953–996, April 2001. ISSN 0009-2665. doi: 10.1021/cro00018s. URL <http://dx.doi.org/10.1021/cr000018s>. (page 1.)
- [2] Neil W. Ashcroft and N. David Mermin. *Solid State Physics*. Brooks Cole, 1 edition, January 1976. ISBN 0030839939. (pages 13, 14, and 25.)
- [3] Sune R. Bahn and Karsten W. Jacobsen. An object-oriented scripting interface to a legacy electronic structure code. *Computing in Science and Engineering*, 4(3):56–66, 2002. ISSN 1521-9615. (pages 37, 46, 59, 98, and 137.)
- [4] M. Bär, Ch Zülicke, M. Eiswirth, and G. Ertl. Theoretical modeling of spatiotemporal self-organization in a surface catalyzed reaction exhibiting bistable kinetics. *The Journal of Chemical Physics*, 96(11):8595–8604, June 1992. ISSN 00219606. doi: doi:10.1063/1.462312. URL http://jcp.aip.org/resource/1/jcpsa6/v96/i11/p8595_s1. (page 90.)
- [5] R. J. Behm, K. Christmann, G. Ertl, and M. A. Van Hove. Adsorption of CO on pd(100). *The Journal of Chemical Physics*, 73(6):2984, 1980. ISSN 00219606. doi: 10.1063/1.440430. URL <http://link.aip.org/link/JCPSA6/v73/i6/p2984/s1&Agg=doi>. (pages 72 and 76.)
- [6] Andrew Behn, Paul M. Zimmerman, Alexis T. Bell, and Martin Head-Gordon. Efficient exploration of reaction paths via a freezing string method. *The Journal of Chemical Physics*, 135(22):224108–224108–9, December 2011. ISSN 00219606. doi: doi:10.1063/1.3664901. URL http://jcp.aip.org/resource/1/jcpsa6/v135/i22/p224108_s1. (page 56.)

- [7] Erik Bitzek, Pekka Koskinen, Franz Gähler, Michael Moseler, and Peter Gumbsch. Structural relaxation made simple. *Physical Review Letters*, 97(17):170201, October 2006. doi: 10.1103/PhysRevLett.97.170201. URL <http://link.aps.org/doi/10.1103/PhysRevLett.97.170201>. (pages 58 and 59.)
- [8] Felix Bloch. Über die quantenmechanik der elektronen in kristallgittern. *Zeitschrift für Physik*, 52(7-8):555–600, July 1929. ISSN 0044-3328. doi: 10.1007/BF01339455. URL <http://link.springer.com/article/10.1007/BF01339455>. (page 13.)
- [9] S. Blomberg, M. J. Hoffmann, J. Gustafson, N. M. Martin, V. R. Fernandes, A. Borg, Z. Liu, R. Chang, S. Matera, K. Reuter, and E. Lundgren. In situ x-ray photoelectron spectroscopy of model catalysts: At the edge of the gap. *Physical Review Letters*, 110(11):117601, March 2013. doi: 10.1103/PhysRevLett.110.117601. URL <http://link.aps.org/doi/10.1103/PhysRevLett.110.117601>. (page 92.)
- [10] S. Blomberg, M. J. Hoffmann, J. Gustafson, N. M. Martin, V. R. Fernandes, A. Borg, Z. Liu, R. Chang, S. Matera, K. Reuter, and E. Lundgren. In situ x-ray photoelectron spectroscopy of model catalysts: At the edge of the gap. *Phys. Rev. Lett.*, 110:117601, Mar 2013. doi: 10.1103/PhysRevLett.110.117601. URL <http://link.aps.org/doi/10.1103/PhysRevLett.110.117601>. (pages 3 and 71.)
- [11] M. Born and R. Oppenheimer. Zur quantentheorie der molekeln. *Annalen der Physik*, 389(20):457–484, January 1927. ISSN 1521-3889. doi: 10.1002/andp.19273892002. URL <http://onlinelibrary.wiley.com/doi/10.1002/andp.19273892002/abstract>. (page 8.)
- [12] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz. A new algorithm for monte carlo simulation of ising spin systems. *Journal of Computational Physics*, 17(1):10–18, January 1975. ISSN 0021-9991. doi: 10.1016/0021-9991(75)90060-1. URL <http://www.sciencedirect.com/science/article/B6WHY-4DDR37H-9W/2/c5e217e509b6a99a10c20c683754ce68>. (page 30.)
- [13] M. Boudart. From the century of the rate equation to the century of the rate constants: a revolution in catalytic kinetics and assisted catalyst design. *Catalysis Letters*, 65(1-3):1–3, March 2000. ISSN 1011-372X, 1572-879X. doi: 10.1023/A:1019057002970. URL <http://link.springer.com/article/10.1023/A%3A1019057002970>. (page 26.)
- [14] M. Boudart. Model catalysts: reductionism for understanding. *Topics in Catalysis*, 13(1-2):147–149, July 2000. ISSN 1022-5528,

- 1572-9028. doi: 10.1023/A:1009080821550. URL <http://link.springer.com/article/10.1023/A%3A1009080821550>. (page 2.)
- [15] A. M. Bradshaw and F. M. Hoffmann. The chemisorption of carbon monoxide on palladium single crystal surfaces: IR spectroscopic evidence for localised site adsorption. *Surface Science*, 72(3):513–535, April 1978. ISSN 0039-6028. doi: 10.1016/0039-6028(78)90367-9. URL <http://www.sciencedirect.com/science/article/B6TVX-46SXTGJ-G8/2/723598f47e6f16ac6a500cf5ae0b12d5>. (page 76.)
- [16] C. R. Brundle, J. Behm, and J. A. Barker. Summary abstract: How many metal atoms are involved in dissociative chemisorption? *Journal of Vacuum Science Technology*, 2:1038–1039, April 1984. URL <http://adsabs.harvard.edu/abs/1984JVST...2.1038B>. (pages 76 and 96.)
- [17] V.I. Bykov, G.S. Yablonskii, and V.I. Elokhin. Steady state multiplicity of the kinetic model of CO oxidation reaction. *Surface Science Letters*, 107(1):L334–L338, May 1981. ISSN 0167-2584. doi: 10.1016/0167-2584(81)90536-3. URL <http://www.sciencedirect.com/science/article/pii/0167258481905363>. (page 90.)
- [18] R. Byrd, P. Lu, J. Nocedal, and C. Zhu. A limited memory algorithm for bound constrained optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, September 1995. ISSN 1064-8275. doi: 10.1137/0916069. URL <http://epubs.siam.org/doi/abs/10.1137/0916069>. (page 58.)
- [19] Paul T. Callaghan. Rheo-NMR and velocity imaging. *Current Opinion in Colloid & Interface Science*, 11(1):13–18, April 2006. ISSN 1359-0294. doi: 10.1016/j.cocis.2005.10.003. URL <http://www.sciencedirect.com/science/article/pii/S135902940500097X>. (page 3.)
- [20] cfd online. Codes – CFD-Wiki, the free CFD reference, 2013. URL <http://www.cfd-online.com/Wiki/Codes>. (page 24.)
- [21] Hugh Chaffey-Millar, Astrid Nikodem, Alexei V. Matveev, Sven Krüger, and Notker Rösch. Improving upon string methods for transition state discovery. *Journal of Chemical Theory and Computation*, 8(2):777–786, February 2012. ISSN 1549-9618. doi: 10.1021/ct200639w. URL <http://dx.doi.org/10.1021/ct200639w>. (page 57.)
- [22] S. L. Chang and P. A. Thiel. Formation of a metastable ordered surface phase due to competitive diffusion and adsorption kinetics: Oxygen on pd(100). *Physical Review Letters*, 59(3):296, July

1987. doi: 10.1103/PhysRevLett.59.296. URL <http://link.aps.org/doi/10.1103/PhysRevLett.59.296>. (pages 76 and 96.)
- [23] S.-L. Chang and P. A. Thiel. Oxygen on pd(100): Order, reconstruction, and desorption. *The Journal of Chemical Physics*, 88(3):2071, 1988. ISSN 00219606. doi: 10.1063/1.454084. URL <http://link.aip.org/link/JCPSA6/v88/i3/p2071/s1&Agg=doi>. (pages 71 and 75.)
- [24] M. W. Chase Jr. NIST-JANAF thermochemical tables. *J. Phys. Chem. Ref. Data*, 9:1, 1998. (page 74.)
- [25] Abhijit Chatterjee and Dionisios G. Vlachos. An overview of spatial microscopic and accelerated kinetic monte carlo methods. *Journal of Computer-Aided Materials Design*, 14(2):253–308, February 2007. ISSN 0928-1045. doi: 10.1007/s10820-006-9042-9. URL http://apps.isiknowledge.com/full_record.do?product=UA&search_mode=Refine&qid=8&SID=V2Dj8EMnbkBelc6NjG3&page=1&doc=10&colname=WOS. (pages 24, 25, and 30.)
- [26] M Chen, Y Cai, Z Yan, K Gath, S Axnanda, and D Goodman. Highly active surfaces for CO oxidation on rh, pd, and pt. *Surface Science*, 601(23):5326–5331, December 2007. ISSN 00396028. doi: 10.1016/j.susc.2007.08.019. URL http://apps.isiknowledge.com/full_record.do?product=UA&search_mode=Refine&qid=2&SID=Z29c52PGFe0J0ppI074&page=1&doc=8&colname=WOS. (page 3.)
- [27] I. Chorkendorff and J. W. Niemantsverdriet. *Concepts of Modern Catalysis and Kinetics*. John Wiley & Sons, March 2006. ISBN 9783527605644. (page 23.)
- [28] Stewart J. Clark, Matthew D. Segall, Chris J. Pickard, Phil J. Hasnip, Matt I. J. Probert, Keith Refson, and Mike C. Payne. First principles methods using CASTEP. *Zeitschrift für Kristallographie*, 220(5-6-2005):567–570, May 2005. ISSN 0044-2968. doi: 10.1524/zkri.220.5.567.65075. URL <http://www.oldenbourg-link.com/doi/abs/10.1524/zkri.220.5.567.65075>. (pages 15, 75, 98, and 123.)
- [29] Clusius. Einführung in die quantenchemie. von h. hellmann. 350 s., 43 abb., 35 tab. franz deuticke, leipzig u. wien 1937. pr. geh. RM. 20,-. geb. RM. 22,-. *Angewandte Chemie*, 54(11-12):156–156, March 1941. ISSN 1521-3757. doi: 10.1002/ange.19410541109. URL <http://onlinelibrary.wiley.com/doi/10.1002/ange.19410541109/abstract>. (page 9.)
- [30] Thomas H Cormen. *Introduction to algorithms*. The MIT Press, Cambridge, Massachusetts; London, 2009. ISBN 9780262033848 0262033844 9780262533058 0262533057. (page 44.)

- [31] Boy Cornils, Wolfgang A. Herrmann, Robert Schlögl, and Chi-Huey Wong. *Catalysis from A to Z: A Concise Encyclopedia*. Wiley-VCH, Weinheim, auflage: 2 edition, 2003. ISBN 9783527303731. (page 1.)
- [32] Juarez L. F. Da Silva, Catherine Stampfl, and Matthias Scheffler. Converged properties of clean metal surfaces by all-electron first-principles calculations. *Surface Science*, 600(3):703–715, February 2006. ISSN 0039-6028. doi: 10.1016/j.susc.2005.12.008. URL <http://www.sciencedirect.com/science/article/pii/S0039602805013154>. (page 133.)
- [33] P. a. M. Dirac. Note on exchange phenomena in the thomas atom. *Mathematical Proceedings of the Cambridge Philosophical Society*, 26(03):376–385, 1930. doi: 10.1017/S0305004100016108. (page 10.)
- [34] David J. Dooling and Linda J. Broadbelt. Generic monte carlo tool for kinetic modeling. *Industrial & Engineering Chemistry Research*, 40(2):522–529, January 2001. doi: 10.1021/ie000310q. URL <http://dx.doi.org/10.1021/ie000310q>. (page 24.)
- [35] Weinan E and Eric Vanden-Eijnden. Transition-path theory and path-finding algorithms for the study of rare events. *Annual Review of Physical Chemistry*, 61(1):391–420, 2010. doi: 10.1146/annurev.physchem.040808.090412. URL <http://www.annualreviews.org/doi/abs/10.1146/annurev.physchem.040808.090412>. (page 57.)
- [36] Andreas Eichler. CO oxidation on transition metal surfaces: reaction rates from first principles. *Surface Science*, 498(3):314–320, February 2002. ISSN 0039-6028. doi: 10.1016/S0039-6028(01)01805-2. URL <http://www.sciencedirect.com/science/article/pii/S0039602801018052>. (page 77.)
- [37] M. Eiswirth, J. Bürger, P. Strasser, and G. Ertl. Oscillating Langmuir-Hinshelwood mechanisms. *The Journal of Physical Chemistry*, 100(49):19118–19123, January 1996. ISSN 0022-3654. doi: 10.1021/jp961688y. URL <http://dx.doi.org/10.1021/jp961688y>. (page 90.)
- [38] Fedwa El-Mellouhi, Normand Mousseau, and Laurent J. Lewis. Kinetic activation-relaxation technique: An off-lattice self-learning kinetic monte carlo algorithm. *Physical Review B*, 78(15):153202, October 2008. doi: 10.1103/PhysRevB.78.153202. URL <http://link.aps.org/doi/10.1103/PhysRevB.78.153202>. (page 94.)
- [39] Gerhard Ertl, Helmut Knözinger, Ferdi Schüth, and Jens Weitkamp. *Handbook of Heterogeneous Catalysis: 8 Volumes*. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, auflage: 2.

- vollst. überarb. u. erw. auflage edition, February 2008. ISBN 9783527312412. (page 1.)
- [40] J. W. Evans. Kinetic phase transitions in catalytic reaction models. *Langmuir*, 7(11):2514–2519, November 1991. ISSN 0743-7463. doi: 10.1021/la00059a020. URL <http://dx.doi.org/10.1021/la00059a020>. (page 90.)
- [41] J. W. Evans, D. K. Flynn-Sanders, and P. A. Thiel. Surface self-diffusion barrier of pd(100) from low-energy electron diffraction. *Surface Science*, 298(2–3):378–383, December 1993. ISSN 0039-6028. doi: 10.1016/0039-6028(93)90051-K. URL <http://www.sciencedirect.com/science/article/pii/003960289390051K>. (page 105.)
- [42] Henry Eyring. The activated complex in chemical reactions. *The Journal of Chemical Physics*, 3(2):107–115, February 1935. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.1749604. URL <http://scitation.aip.org/content/aip/journal/jcp/3/2/10.1063/1.1749604>. (page 53.)
- [43] L. Farkas. Keimbildungsgeschwindigkeit in übersättigten dämpfen. *Z. phys. Chem.*, 125:236–242, 1927. (page 54.)
- [44] Enrico Fermi. A statistical method for the determination of some atomic properties and the application of this method to the theory of the periodic system of elements. *Z. Phys*, 48:73–79, 1928. (page 9.)
- [45] V. R. Fernandes, J. Gustafson, I. H. Svenum, M. H. Farstad, L. E. Walle, S. Blomberg, E. Lundgren, and A. Borg. Reduction behavior of oxidized pd(100) and Pd₇₅Ag₂₅(100) surfaces using CO. *Surface Science*, 621:31–39, March 2014. ISSN 0039-6028. doi: 10.1016/j.susc.2013.10.018. URL <http://www.sciencedirect.com/science/article/pii/S0039602813003154>. (pages 4, 92, 99, 101, 115, and 118.)
- [46] R. P. Feynman. Forces in molecules. *Physical Review*, 56(4):340–343, August 1939. doi: 10.1103/PhysRev.56.340. URL <http://link.aps.org/doi/10.1103/PhysRev.56.340>. (page 9.)
- [47] Kristen A. Fichthorn and W. H. Weinberg. Theoretical foundations of dynamical monte carlo simulations. *The Journal of Chemical Physics*, 95(2):1090–1096, July 1991. doi: 10.1063/1.461138. URL <http://link.aip.org/link/?JCP/95/1090/1>. (page 30.)
- [48] R. Fletcher. *Practical Methods of Optimization*. Wiley, 2 edition, May 2000. ISBN 0471494631. (pages 58 and 128.)

- [49] M. Frigo and S.G. Johnson. The design and implementation of FFTW3. *Proceedings of the IEEE*, 93(2):216–231, 2005. ISSN 0018-9219. doi: 10.1109/JPROC.2004.840301. (pages 14 and 127.)
- [50] FSF. gnu.org, June 2007. URL <http://www.gnu.org/licenses/gpl.html>. (page 51.)
- [51] Harald Gabasch, Axel Knop-Gericke, Robert Schlögl, Marta Borasio, Christian Weilach, Günther Rupprechter, Simon Penner, Bernd Jenewein, Konrad Hayek, and Bernhard Klötzer. Comparison of the reactivity of different PdO species in CO oxidation. *Physical Chemistry Chemical Physics*, 9(4):533, 2007. ISSN 1463-9076, 1463-9084. doi: 10.1039/b610719b. URL <http://pubs.rsc.org/en/Content/ArticleLanding/2007/CP/b610719b>. (page 71.)
- [52] H. S. Gandhi, G. W. Graham, and R. W. McCabe. Automotive exhaust catalysis. *Journal of Catalysis*, 216(1-2):433–442, May 2003. ISSN 0021-9517. doi: 10.1016/S0021-9517(02)00067-2. URL <http://www.sciencedirect.com/science/article/B6WHJ-484SH74-4/2/7ef50950d709764519ead5c9d66aeb5d>. (page 71.)
- [53] F. Gao, S.M. McClure, Y. Cai, K.K. Gath, Y. Wang, M.S. Chen, Q.L. Guo, and D.W. Goodman. CO oxidation trends on pt-group metals from ultrahigh vacuum to near atmospheric pressures: A combined in situ PM-IRAS and reaction kinetics study. *Surface Science*, 603(1):65–70, January 2009. ISSN 0039-6028. doi: 10.1016/j.susc.2008.10.031. URL <http://www.sciencedirect.com/science/article/B6TVX-4TVHSGT-3/2/6d2f019b76644eaad1f1c2fca51be49a>. (page 4.)
- [54] F. Gao, Y. Wang, Y. Cai, and D. W. Goodman. CO oxidation on pt-group metals from ultrahigh vacuum to near atmospheric pressures. 2. palladium and platinum. *The Journal of Physical Chemistry C*, 113(1):174–181, January 2009. doi: 10.1021/jp8077985. URL <http://dx.doi.org/10.1021/jp8077985>. (page 4.)
- [55] Feng Gao, Matthew Lundwall, and D. Wayne Goodman. Infrared reflection absorption spectroscopy study of CO adsorption and reaction on oxidized pd(100). *J. Phys. Chem. C*, 112(15):6057–6064, 2008. ISSN 1932-7447. doi: 10.1021/jp710713r. URL <http://dx.doi.org/10.1021/jp710713r>. (page 71.)
- [56] Feng Gao, Yilin Wang, and D. Wayne Goodman. Reply to “Comment on ‘CO oxidation on pt-group metals from ultrahigh vacuum to near atmospheric pressures. 2. palladium and platinum’”. *The Journal of Physical Chemistry C*, 114(14):6874, April 2010. doi: 10.1021/jp100134e. URL <http://dx.doi.org/10.1021/jp100134e>. (pages 4, 71, and 72.)

- [57] Cristina Garcia Cardona, Veena Tikare, and Steven J. Plimpton. Parallel simulation of 3D sintering. *International Journal of Computational Materials Science and Surface Engineering*, 4(1): 37–54, January 2011. doi: 10.1504/IJCMSSE.2011.037351. URL <http://dx.doi.org/10.1504/IJCMSSE.2011.037351>. (page 24.)
- [58] Crispin Gardiner. *Handbook of Stochastic Methods: for Physics, Chemistry and the Natural Sciences*. Springer, 3rd edition, April 2004. ISBN 3540208828. (pages 23 and 26.)
- [59] Daniel T Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22(4):403–434, December 1976. ISSN 0021-9991. doi: 10.1016/0021-9991(76)90041-3. URL <http://www.sciencedirect.com/science/article/pii/0021999176900413>. (page 30.)
- [60] Amit A. Gokhale, James A. Dumesic, and Manos Mavrikakis. On the mechanism of low-temperature water gas shift reaction on copper. *Journal of the American Chemical Society*, 130(4):1402–1414, January 2008. ISSN 0002-7863. doi: 10.1021/ja0768237. URL <http://dx.doi.org/10.1021/ja0768237>. (page 23.)
- [61] D. W. Goodman. Model catalysts: from imagining to imaging a working surface. *Journal of Catalysis*, 216(1–2): 213–222, May 2003. ISSN 0021-9517. doi: 10.1016/S0021-9517(02)00112-4. URL <http://www.sciencedirect.com/science/article/pii/S0021951702001124>. (page 2.)
- [62] Axel Groß. Coverage effects in the adsorption of h₂ on pd(100) studied by ab initio molecular dynamics simulations. *The Journal of Chemical Physics*, 135(17):174707, November 2011. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.3656765. URL <http://scitation.aip.org/content/aip/journal/jcp/135/17/10.1063/1.3656765>. (page 120.)
- [63] J. Gustafson, M. Shipilin, C. Zhang, A. Stierle, U. Hejral, U. Ruett, O. Gutowski, P.-A. Carlsson, M. Skoglundh, and E. Lundgren. High-energy surface x-ray diffraction for fast surface structure determination. *Science*, 343(6172):758–761, February 2014. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1246834. URL <http://www.sciencemag.org/content/343/6172/758>. (page 3.)
- [64] P. Güttinger. Das verhalten von atomen im magnetischen drehfeld. *Zeitschrift für Physik*, 73(3-4):169–184, March 1932. ISSN 1434-6001, 1434-601X. doi: 10.1007/BF01351211. URL <http://link.springer.com/article/10.1007/BF01351211>. (page 9.)
- [65] B. Hammer. The NO+CO reaction catalyzed by flat, stepped, and edged pd surfaces. *Journal of Catalysis*, 199

- (2):171–176, April 2001. ISSN 0021-9517. doi: 06/jcat.2000.3147. URL <http://www.sciencedirect.com/science/article/pii/S0021951700931476>. (page 77.)
- [66] B. Hammer, L. B. Hansen, and J. K. Nørskov. Improved adsorption energetics within density-functional theory using revised perdedw-burke-ernzerhof functionals. *Physical Review B*, 59(11):7413, March 1999. doi: 10.1103/PhysRevB.59.7413. URL <http://link.aps.org/doi/10.1103/PhysRevB.59.7413>. (page 133.)
- [67] Peter Hänggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after kramers. *Reviews of Modern Physics*, 62(2):251, April 1990. doi: 10.1103/RevModPhys.62.251. URL <http://link.aps.org/doi/10.1103/RevModPhys.62.251>. (pages 53 and 54.)
- [68] Eric W Hansen and Matthew Neurock. First-principles-based monte carlo methodology applied to O/Rh(100). *Surface Science*, 464(2–3):91–107, October 2000. ISSN 0039-6028. doi: 10.1016/S0039-6028(00)00598-7. URL <http://www.sciencedirect.com/science/article/pii/S0039602800005987>. (pages 3, 23, and 24.)
- [69] Eric W. Hansen and Matthew Neurock. First-principles-based monte carlo simulation of ethylene hydrogenation kinetics on pd. *Journal of Catalysis*, 196(2):241–252, December 2000. ISSN 0021-9517. doi: 10.1006/jcat.2000.3018. URL <http://www.sciencedirect.com/science/article/pii/S0021951700930185>. (pages 23 and 24.)
- [70] B Hendriksen, S Bobaru, and J Frenken. Oscillatory CO oxidation on pd(100) studied with in situ scanning tunneling microscopy. *Surface Science*, 552(1-3):229–242, March 2004. ISSN 00396028. doi: 10.1016/j.susc.2004.01.025. URL http://apps.isiknowledge.com/full_record.do?product=UA&search_mode=GeneralSearch&qid=4&SID=N28eH9fMMkLm8P4LNJo&page=1&doc=3&colname=WOS. (pages 4, 72, and 87.)
- [71] B. L. M. Hendriksen and J. W. M. Frenken. CO oxidation on pt(110): Scanning tunneling microscopy inside a high-pressure flow reactor. *Physical Review Letters*, 89(4):046101, July 2002. doi: 10.1103/PhysRevLett.89.046101. URL <http://link.aps.org/doi/10.1103/PhysRevLett.89.046101>. (pages 3 and 4.)
- [72] Bas Hendriksen, Stefania Bobaru, and Joost Frenken. Looking at heterogeneous catalysis at atmospheric pressure using tunnel vision. *Topics in Catalysis*, 36(1):43–54, 2005. doi: 10.1007/s11244-005-7861-7. URL <http://dx.doi.org/10.1007/s11244-005-7861-7>. (pages 92 and 105.)

- [73] Bas L. M. Hendriksen, Marcelo D. Ackermann, Richard van Rijn, Dunja Stoltz, Ioana Popa, Olivier Balmes, Andrea Resta, Didier Wermeille, Roberto Felici, Salvador Ferrer, and FrenkenJoost W. M. The role of steps in surface catalysis and reaction oscillations. *Nat Chem*, 2(9):730–734, 2010. ISSN 1755-4330. doi: 10.1038/nchem.728. URL <http://dx.doi.org/10.1038/nchem.728>. (pages 4, 72, and 90.)
- [74] B.L.M. Hendriksen, S.C. Bobaru, and J.W.M. Frenken. Bistability and oscillations in CO oxidation studied with scanning tunnelling microscopy inside a reactor. *Catalysis Today*, 105(2): 234–243, July 2005. ISSN 0920-5861. doi: 10.1016/j.cattod.2005.02.041. URL <http://www.sciencedirect.com/science/article/pii/S0920586105002427>. (pages 4, 72, 81, 87, 88, 90, 92, 105, and 118.)
- [75] Graeme Henkelman and Hannes Jónsson. A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. *The Journal of Chemical Physics*, 111(15):7010, 1999. ISSN 00219606. doi: 10.1063/1.480097. URL <http://link.aip.org/link/JCPSA6/v111/i15/p7010/s1&Agg=doi>. (pages 55 and 57.)
- [76] Graeme Henkelman and Hannes Jónsson. Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. *The Journal of Chemical Physics*, 113(22):9978, 2000. ISSN 00219606. doi: 10.1063/1.1323224. URL <http://adsabs.harvard.edu/abs/2000JChPh.113.9978H>. (pages 55 and 56.)
- [77] Graeme Henkelman, Blas P. Uberuaga, and Hannes Jónsson. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *The Journal of Chemical Physics*, 113(22):9901, 2000. ISSN 00219606. doi: 10.1063/1.1329672. URL <http://link.aip.org/link/JCPSA6/v113/i22/p9901/s1&Agg=doi>. (pages 55, 56, and 98.)
- [78] Graeme Henkelman, Gísli Jóhannesson, and Hannes Jónsson. Methods for finding saddle points and minimum energy paths. In Steven D. Schwartz, editor, *Theoretical Methods in Condensed Phase Chemistry*, number 5 in Progress in Theoretical Chemistry and Physics, pages 269–302. Springer Netherlands, January 2002. ISBN 978-0-7923-6687-4, 978-0-306-46949-7. URL http://link.springer.com/chapter/10.1007/0-306-46949-9_10. (pages 56, 58, 59, and 63.)
- [79] M. R. Hestenes and E. Stiefel. *Methods of conjugate gradients for solving linear systems*. National Bureau of Standards, 1952. URL <http://archive.org/details/jresv49n6p409>. (page 58.)

- [80] Janne T. Hirvi, Toni-Jani J. Kinnunen, Mika Suvanto, Tapani A. Pakkanen, and Jens K. Nørskov. CO oxidation on PdO surfaces. *The Journal of Chemical Physics*, 133(8):084704, 2010. ISSN 00219606. doi: 10.1063/1.3464481. URL <http://link.aip.org/link/JCPSA6/v133/i8/p084704/s1&Agg=doi>. (pages 3, 71, and 79.)
- [81] Max J Hoffmann. kmos, 2012. URL <http://mhoffman.github.com/kmos/>. (pages 23 and 79.)
- [82] Max J. Hoffmann and Karsten Reuter. CO oxidation on pd(100) versus PdO(101)-($\sqrt{5}$) \times ($\sqrt{5}$)R27 $^\circ$: first-principles kinetic phase diagrams and bistability conditions. *Topics in Catalysis*, 57(1-4):159–170, February 2014. ISSN 1022-5528, 1572-9028. doi: 10.1007/s11244-013-0172-5. URL <http://link.springer.com/article/10.1007/s11244-013-0172-5>. (pages 96, 97, 98, and 111.)
- [83] Max J. Hoffmann and Karsten Reuter. Co oxidation on pd(100) versus pdo(101)-($\sqrt{5} \times \sqrt{5}$)r27 $^\circ$: First-principles kinetic phase diagrams and bistability conditions. *Topics in Catalysis*, 57(1-4):159–170, 2014. ISSN 1022-5528. doi: 10.1007/s11244-013-0172-5. URL <http://dx.doi.org/10.1007/s11244-013-0172-5>. (page 71.)
- [84] Max J Hoffmann and Karsten Reuter. Detailed kinetics from an atomistic picture: Multi-lattice first-principles kinetic monte carlo simulations of pd(100) surface oxide reduction, 2014. (in preparation). (page 91.)
- [85] Max J Hoffmann, S Matera, and K Reuter. kmos: A lattice kinetic monte carlo framework. *Computational Physics Communications*, 185(7):2138–2150, 2014. doi: 10.1016/j.cpc.2014.04.003. URL <http://dx.doi.org/10.1016/j.cpc.2014.04.003>. (page 23.)
- [86] Max J. Hoffmann, Sebastian Matera, and Karsten Reuter. kmos: A lattice kinetic monte carlo framework. *Computer Physics Communications*, 185(7):2138–2150, April 2014. ISSN 0010-4655. doi: 10.1016/j.cpc.2014.04.003. URL <http://www.sciencedirect.com/science/article/pii/S001046551400126X>. (pages 91, 93, 94, 95, 98, and 110.)
- [87] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Physical Review*, 136(3B):B864–B871, November 1964. doi: 10.1103/PhysRev.136.B864. URL <http://link.aps.org/doi/10.1103/PhysRev.136.B864>. (page 10.)
- [88] John Holdren. Federal energy research and development for the challenges of the 21st century. *Lewis M. Branscomb and James H. Keller*, 1999. URL http://belfercenter.hks.harvard.edu/publication/2055/federal_energy_research_and_

- development_for_the_challenges_of_the_21st_century.html.
(page 1.)
- [89] K. Honkala, A. Hellman, I. N. Remediakis, A. Logadottir, A. Carlsson, S. Dahl, C. H. Christensen, and J. K. Nørskov. Ammonia synthesis from first-principles calculations. *Science*, 307(5709):555–558, January 2005. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1106435. URL <http://www.sciencemag.org/content/307/5709/555>. (page 23.)
- [90] K. Honkala, Z. Łodziana, I. N. Remediakis, and N. Lopez. Expanding and reducing complexity in materials science models with relevance in catalysis and energy. *Topics in Catalysis*, pages 1–11, February 2014. ISSN 1022-5528, 1572-9028. doi: 10.1007/s11244-013-0158-3. URL <http://link.springer.com/article/10.1007/s11244-013-0158-3>. (page 23.)
- [91] J.D. Hunter. Matplotlib: A 2D graphics environment. *Computing in Science Engineering*, 9(3):90–95, 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.55. (page 46.)
- [92] Laurent Hyafil and RL Rivest. Constructing optimal binary decision trees is NP-complete. *Information Processing Letters*, 5(1): 15–17, 1976. (page 42.)
- [93] Oliver R. Inderwildi, Stephen J. Jenkins, and David A. King. Fischer-Tropsch mechanism revisited: Alternative pathways for the production of higher hydrocarbons from synthesis gas. *The Journal of Physical Chemistry C*, 112(5):1305–1307, February 2008. ISSN 1932-7447. doi: 10.1021/jp710674q. URL <http://dx.doi.org/10.1021/jp710674q>. (page 23.)
- [94] Andreas Jess and Peter Wasserscheid. *Chemical Technology: An Integral Textbook*. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, auflage: 1. auflage edition, February 2013. ISBN 9783527304462. (page 1.)
- [95] Hannes Jónsson, Greg Mills, and Karsten W. Jacobsen. Nudged elastic band method for finding minimum energy paths of transitions. In *Classical and Quantum Dynamics in Condensed Phase Simulations*. World Scientific, 385. (pages 55 and 56.)
- [96] Abdelkader Kara, Oleg Trushin, Handan Yildirim, and Talat S. Rahman. Off-lattice self-learning kinetic monte carlo: application to 2D cluster diffusion on the fcc(111) surface. *Journal of Physics: Condensed Matter*, 21(8):084213, February 2009. ISSN 0953-8984. doi: 10.1088/0953-8984/21/8/084213. URL <http://iopscience.iop.org/0953-8984/21/8/084213>. (page 94.)

- [97] F. J. Keil. Multiscale modelling in computational heterogeneous catalysis. In Barbara Kirchner and Jadran Vrabec, editors, *Multiscale Molecular Methods in Applied Chemistry*, number 307 in Topics in Current Chemistry, pages 69–107. Springer Berlin Heidelberg, January 2012. ISBN 978-3-642-24967-9, 978-3-642-24968-6. URL http://link.springer.com/chapter/10.1007/128_2011_128. (page 23.)
- [98] Sung Youb Kim, In-Ho Lee, and Sukky Jun. Transition-pathway models of atomic diffusion on fcc metal surfaces. II. stepped surfaces. *Physical Review B (Condensed Matter and Materials Physics)*, 76(24):245408–20, December 2007. doi: 10.1103/PhysRevB.76.245408. URL <http://link.aps.org/abstract/PRB/v76/e245408>. (page 105.)
- [99] Sung Youb Kim, In-Ho Lee, and Sukky Jun. Transition-pathway models of atomic diffusion on fcc metal surfaces. i. flat surfaces. *Physical Review B (Condensed Matter and Materials Physics)*, 76(24):245407–15, December 2007. doi: 10.1103/PhysRevB.76.245407. URL <http://link.aps.org/abstract/PRB/v76/e245407>. (page 105.)
- [100] Janos T. Kiss and Richard D. Gonzalez. Deactivation of selected noble metal catalysts during carbon monoxide oxidation: an in situ IR and kinetic study. *Industrial & Engineering Chemistry Product Research and Development*, 24(2):216–219, 1985. ISSN 0196-4321. doi: 10.1021/i300018a008. URL <http://dx.doi.org/10.1021/i300018a008>. (page 71.)
- [101] R. Kissel-Osterrieder, F. Behrendt, and J. Warnatz. Dynamic monte carlo simulations of catalytic surface reactions. *Proceedings of the Combustion Institute*, 28(1):1323–1330, 2000. ISSN 1540-7489. doi: 10.1016/S0082-0784(00)80346-X. URL <http://www.sciencedirect.com/science/article/pii/S008207840080346X>. (pages 3 and 24.)
- [102] Leonard Kleinman and James C. Phillips. Crystal potential and energy bands of semiconductors. III. self-consistent calculations for silicon. *Physical Review*, 118(5):1153–1167, June 1960. doi: 10.1103/PhysRev.118.1153. URL <http://link.aps.org/doi/10.1103/PhysRev.118.1153>. (page 15.)
- [103] W. Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):A1133, November 1965. doi: 10.1103/PhysRev.140.A1133. URL <http://link.aps.org/doi/10.1103/PhysRev.140.A1133>. (page 10.)
- [104] Petr Kostelník, Nicola Seriani, Georg Kresse, Anders Mikkelsen, Edvin Lundgren, Volker Blum, Tomáš Šikola, Peter Varga, and

- Michael Schmid. The $\text{pd}(100)\text{-(5}\times\text{5)}\text{R}27^\circ\text{-O}$ surface oxide: A LEED, DFT and STM study. *Surface Science*, 601(6):1574–1581, March 2007. ISSN 00396028. doi: 10.1016/j.susc.2007.01.026. URL http://apps.isiknowledge.com/full_record.do?product=UA&colname=WOS&search_mode=CitingArticles&qid=24&SID=N11ILgijLnFDGk18Mim&page=2&doc=20. (page 92.)
- [105] H. A. Kramers. Brownian motion in a field of force and the diffusion model of chemical reactions. *Physica*, 7(4):284–304, April 1940. ISSN 0031-8914. doi: 10.1016/S0031-8914(40)90098-2. URL <http://www.sciencedirect.com/science/article/pii/S0031891440900982>. (page 53.)
- [106] Lothar Kunz. Entwicklung eines computerprogramms zur kinetischen monte carlo simulation von oberflächenreaktionen auf nanopartikeln. Diploma thesis, Universität Karlsruhe, Karlsruhe, 2007. URL http://www.itcp.kit.edu/deutschmann/img/content/06_LKunz_DA_UniKA.pdf. (page 24.)
- [107] Kari Laasonen, Alfredo Pasquarello, Roberto Car, Changyol Lee, and David Vanderbilt. Car-parrinello molecular dynamics with vanderbilt ultrasoft pseudopotentials. *Physical Review B*, 47(16):10142, April 1993. doi: 10.1103/PhysRevB.47.10142. URL <http://link.aps.org/doi/10.1103/PhysRevB.47.10142>. (page 128.)
- [108] W Li and B Hammer. Reactivity of a gas/metal/metal-oxide three-phase boundary: CO oxidation at the $\text{pt}(111)\text{-c}(4\times 2)\text{-}2\text{CO}/\alpha\text{-PtO}_2$ phase boundary. *Chemical Physics Letters*, 409(1-3):1–7, June 2005. ISSN 00092614. doi: 10.1016/j.cplett.2005.04.064. URL http://apps.isiknowledge.com/full_record.do?product=UA&search_mode=Refine&qid=4&SID=S2mj7Lmemg55h0dNFn0&page=1&doc=2&colname=WOS. (page 115.)
- [109] C. L. Liu, J. M. Cohen, J. B. Adams, and A. F. Voter. EAM study of surface self-diffusion of single adatoms of fcc metals ni, cu, al, ag, au, pd, and pt. *Surface Science*, 253(1–3):334–344, August 1991. ISSN 0039-6028. doi: 10.1016/0039-6028(91)90604-Q. URL <http://www.sciencedirect.com/science/article/pii/003960289190604Q>. (page 105.)
- [110] Da-Jiang Liu and J. W. Evans. Atomistic lattice-gas modeling of CO oxidation on $\text{pd}(100)$: Temperature-programed spectroscopy and steady-state behavior. *The Journal of Chemical Physics*, 124(15):154705, 2006. ISSN 00219606. doi: 10.1063/1.2186314. URL <http://link.aip.org/link/JCPSA6/v124/i15/p154705/s1&Agg=doi>. (pages 76 and 96.)

- [111] Da-Jiang Liu and J. W. Evans. Chemical diffusion of CO in mixed CO+O adlayers and reaction-front propagation in CO oxidation on pd(100). *The Journal of Chemical Physics*, 125(5):054709, 2006. ISSN 00219606. doi: 10.1063/1.2221690. URL <http://link.aip.org/link/JCPSA6/v125/i5/p054709/s1&Agg=doi>. (pages 72 and 96.)
- [112] Da-Jiang Liu and J.W. Evans. Atomistic and multiscale modeling of CO-oxidation on pd(1 0 0) and rh(1 0 0): From nanoscale fluctuations to mesoscale reaction fronts. *Surface Science*, 603(10-12):1706–1716, June 2009. ISSN 0039-6028. doi: 10.1016/j.susc.2008.10.058. URL <http://www.sciencedirect.com/science/article/pii/S0039602809000673>. (pages 72, 76, and 96.)
- [113] Fusheng Liu, Wangyu Hu, Yifeng Chen, Huiqiu Deng, Han Chen, Xiyuan Yang, and Wenhua Luo. Atomic self-diffusion behaviors relevant to 2D homoepitaxy growth on stepped pd(001) surface. *Surface Science*, 624:89–94, June 2014. ISSN 0039-6028. doi: 10.1016/j.susc.2014.01.011. URL <http://www.sciencedirect.com/science/article/pii/S0039602814000144>. (page 105.)
- [114] A. Lozano, A. Groß, and H. F. Busnengo. Molecular dynamics study of h₂ dissociation on h-covered pd(100). *Physical Review B*, 81(12):121402, March 2010. doi: 10.1103/PhysRevB.81.121402. URL <http://link.aps.org/doi/10.1103/PhysRevB.81.121402>. (page 120.)
- [115] J. J. Lukkien, J. P. L. Segers, P. A. J. Hilbers, R. J. Gelten, and A. P. J. Jansen. Efficient monte carlo methods for the simulation of catalytic surface reactions. *Physical Review E*, 58(2):2598–2610, August 1998. doi: 10.1103/PhysRevE.58.2598. URL <http://link.aps.org/doi/10.1103/PhysRevE.58.2598>. (pages 30, 32, and 94.)
- [116] E. Lundgren, J. Gustafson, A. Mikkelsen, J. Andersen, A. Stierle, H. Dosch, M. Todorova, J. Rogal, K. Reuter, and M. Scheffler. Kinetic hindrance during the initial oxidation of pd(100) at ambient pressures. *Physical Review Letters*, 92(4):046101, January 2004. ISSN 0031-9007. doi: 10.1103/PhysRevLett.92.046101. URL http://apps.isiknowledge.com/full_record.do?product=UA&search_mode=Refine&qid=3&SID=W111374Cm9FDcC113Ag&page=1&doc=5&colname=WOS. (pages 71, 81, and 92.)
- [117] Edvin Lundgren. private communication, May 2014. (page 101.)
- [118] Edvin Lundgren, Anders Mikkelsen, Jesper N. Andersen, Georg Kresse, Michael Schmid, and Peter Varga. Surface oxides on close-packed surfaces of late transition metals. *Journal of Physics: Condensed Matter*, 18(30):R481–R499, 2006. ISSN 0953-8984. URL <http://www.iop.org/EJ/abstract/-search=65751872.1/0953-8984/18/30/R01>. (pages 3, 92, and 119.)

- [119] Debarshi Majumder and Linda J. Broadbelt. A multiscale scheme for modeling catalytic flow reactors. *AIChE Journal*, 52(12):4214–4228, 2006. ISSN 1547-5905. doi: 10.1002/aic.11030. URL <http://onlinelibrary.wiley.com/doi/10.1002/aic.11030/abstract>. (pages 3 and 24.)
- [120] Richard M. Martin. *Electronic Structure: Basic Theory and Practical Methods*. Cambridge University Press, April 2004. ISBN 0521782856. (page 126.)
- [121] Sebastian Matera and Karsten Reuter. First-principles approach to heat and mass transfer effects in model catalyst studies. *Catalysis Letters*, 133(1-2):156–159, September 2009. ISSN 1011-372X, 1572-879X. doi: 10.1007/s10562-009-0168-8. URL <http://www.springerlink.com/content/2527g41g2357742j/>. (pages 3, 24, and 87.)
- [122] Sebastian Matera and Karsten Reuter. Transport limitations and bistability for in situ CO oxidation at RuO₂(110) : First-principles based multiscale modeling. *Physical Review B*, 82(8):085446, 2010. doi: 10.1103/PhysRevB.82.085446. URL <http://link.aps.org/doi/10.1103/PhysRevB.82.085446>.
- [123] Sebastian Matera and Karsten Reuter. When atomic-scale resolution is not enough: Spatial effects on in situ model catalyst studies. *Journal of Catalysis*, 295:261–268, November 2012. ISSN 0021-9517. doi: 10.1016/j.jcat.2012.08.020. URL <http://www.sciencedirect.com/science/article/pii/S0021951712002771>. (pages 3, 24, and 87.)
- [124] Sebastian Matera, Hakim Meskine, and Karsten Reuter. Ad-layer inhomogeneity without lateral interactions: Rationalizing correlation effects in CO oxidation at RuO₂(110) with first-principles kinetic monte carlo. *The Journal of Chemical Physics*, 134(6):064713, 2011. ISSN 00219606. doi: 10.1063/1.3553258. URL <http://link.aip.org/link/JCPSA6/v134/i6/p064713/s1&Agg=doi>. (page 24.)
- [125] Ian E. Maxwell. Combinatorial chemistry: Connecting with catalysis. *Nature*, 394(6691):325–326, July 1998. ISSN 0028-0836. doi: 10.1038/28510. URL <http://www.nature.com/nature/journal/v394/n6691/full/394325a0.html>. (page 2.)
- [126] Sean M. McClure and D. Wayne Goodman. New insights into catalytic CO oxidation on pt-group metals at elevated pressures. *Chemical Physics Letters*, 469(1–3):1–13, February 2009. ISSN 0009-2614. doi: 10.1016/j.cplett.2008.12.066. URL <http://www.sciencedirect.com/science/article/pii/S000926140801720X>. (page 4.)

- [127] Donghai Mei and Guang Lin. Effects of heat and mass transfer on the kinetics of CO oxidation over RuO₂(1 1 0) catalyst. *Catalysis Today*, 165(1):56–63, May 2011. ISSN 0920-5861. doi: 10.1016/j.cattod.2010.11.041. URL <http://www.sciencedirect.com/science/article/pii/S0920586110007820>. (pages 3 and 24.)
- [128] Hakim Meskine, Sebastian Matera, Matthias Scheffler, Karsten Reuter, and Horia Metiu. Examination of the concept of degree of rate control by first-principles kinetic monte carlo simulations. *Surface Science*, 603(10-12):1724–1730, June 2009. ISSN 0039-6028. doi: 10.1016/j.susc.2008.08.036. URL <http://www.sciencedirect.com/science/article/B6TVX-4VCH6T7-M/2/13208a32b93e687c21ce6aa40dac7ac7>. (page 23.)
- [129] Jörg Meyer and Karsten Reuter. First-principles initial sticking of O₂ at Pd(100). to be published, 2011. (page 77.)
- [130] Angelos Michaelides, Karsten Reuter, and Matthias Scheffler. When seeing is not believing: Oxygen on Ag(111), a simple adsorption system? *Journal of Vacuum Science & Technology A*, 23(6):1487–1497, October 2005. ISSN 0734-2101, 1520-8559. doi: 10.1116/1.2049302. URL <http://scitation.aip.org/content/avs/journal/jvsta/23/6/10.1116/1.2049302>. (page 92.)
- [131] William Hallows Miller. *A Treatise on Crystallography*. Deighton, 1839. (page 18.)
- [132] Hendrik J. Monkhorst and James D. Pack. Special points for Brillouin-zone integrations. *Physical Review B*, 13(12):5188, June 1976. doi: 10.1103/PhysRevB.13.5188. URL <http://link.aps.org/doi/10.1103/PhysRevB.13.5188>. (pages 14, 126, and 133.)
- [133] S. Müller. Bulk and surface ordering phenomena in binary metal alloys. *Journal of Physics: Condensed Matter*, 15(34):R1429–R1500, 2003. (page 38.)
- [134] Matthew Neurock. Engineering molecular transformations for sustainable energy conversion. *Industrial & Engineering Chemistry Research*, 49(21):10183–10199, November 2010. ISSN 0888-5885. doi: 10.1021/ie101300c. URL <http://dx.doi.org/10.1021/ie101300c>. (page 23.)
- [135] Astrid Nikodem, Alexei V. Matveev, Bo-Xiao Zheng, and Notker Rösch. Efficient two-step procedures for locating transition states of surface reactions. *Journal of Chemical Theory and Computation*, 9(1):588–599, January 2013. ISSN 1549-9618. doi: 10.1021/ct300728a. URL <http://dx.doi.org/10.1021/ct300728a>. (page 55.)

- [136] J. K. Norskov, T. Bligaard, J. Rossmeisl, and C. H. Christensen. Towards the computational design of solid catalysts. *Nat Chem*, 1(1):37–46, April 2009. ISSN 1755-4330. doi: 10.1038/nchem.121. URL <http://dx.doi.org/10.1038/nchem.121>. (pages 1 and 2.)
- [137] Jens K. Nørskov, Frank Abild-Pedersen, Felix Studt, and Thomas Bligaard. Density functional theory in surface chemistry and catalysis. *Proceedings of the National Academy of Sciences*, 108(3):937–943, January 2011. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1006652108. URL <http://www.pnas.org/content/108/3/937>. (page 23.)
- [138] J.k. Nørskov, M. Scheffler, and H. Toulhoat. Density functional theory in surface science and heterogeneous catalysis. *MRS Bulletin*, 31(09):669–674, 2006. doi: 10.1557/mrs2006.175. (page 1.)
- [139] Travis E. Oliphant. Python for scientific computing. *Computing in Science & Engineering*, 9(3):10–20, 2007. URL <http://link.aip.org/link/?CSX/9/10/1>. (page 46.)
- [140] M. Ortiz, A.m. Cuitiño, J. Knap, and M. Koslowski. Mixed Atomistic–Continuum models of material behavior: The art of transcending atomistics and informing continua. *MRS Bulletin*, 26(03):216–221, 2001. doi: 10.1557/mrs2001.45. (page 2.)
- [141] Herbert Over. Surface chemistry of ruthenium dioxide in heterogeneous catalysis and electrocatalysis: From fundamental to applied research. *Chemical Reviews*, 112(6):3356–3426, June 2012. ISSN 0009-2665. doi: 10.1021/cr200247n. URL <http://dx.doi.org/10.1021/cr200247n>. (page 92.)
- [142] Herbert Over and Martin Muhler. Catalytic CO oxidation over ruthenium—bridging the pressure gap. *Progress in Surface Science*, 72(1–4):3–17, June 2003. ISSN 0079-6816. doi: 10.1016/S0079-6816(03)00011-X. URL <http://www.sciencedirect.com/science/article/pii/S007968160300011X>. (page 89.)
- [143] John P. Perdew and Karla Schmidt. Jacob’s ladder of density functional approximations for the exchange-correlation energy. In *AIP Conference Proceedings*, volume 577, pages 1–20, Antwerpen (Belgium), July 2001. ISBN 0094-243X. doi: 10.1063/1.1390175. URL <http://adsabs.harvard.edu/abs/2001AIPC..577...1P>. (page 12.)
- [144] John P. Perdew, Kieron Burke, and Matthias Ernzerhof. Generalized gradient approximation made simple. *Physical Review Letters*, 77(18):3865, October 1996. doi: 10.1103/PhysRevLett.77.3865. URL <http://link.aps.org/doi/10.1103/PhysRevLett.77.3865>. (pages 12, 75, 97, 98, and 123.)

- [145] John P. Perdew, Adrienn Ruzsinszky, Jianmin Tao, Viktor N. Staroverov, Gustavo E. Scuseria, and Gábor I. Csonka. Prescription for the design and selection of density functional approximations: More constraint satisfaction with fewer fits. *The Journal of Chemical Physics*, 123(6):062201, August 2005. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.1904565. URL <http://scitation.aip.org/content/aip/journal/jcp/123/6/10.1063/1.1904565>. (page 12.)
- [146] Fernando Perez and Brian E. Granger. IPython: a system for interactive scientific computing. *Computing in Science & Engineering*, 9(3):21–29, 2007. ISSN 1521-9615. doi: 10.1109/MCSE.2007.53. URL <http://ipython.org/citing.html>. (page 46.)
- [147] Baron Peters, Andreas Heyden, Alexis T. Bell, and Arup Chakraborty. A growing string method for determining transition states: Comparison to the nudged elastic band and string methods. *The Journal of Chemical Physics*, 120(17):7877–7886, May 2004. ISSN 00219606. doi: doi:10.1063/1.1691018. URL http://jcp.aip.org/resource/1/jcpsa6/v120/i17/p7877_s1. (page 57.)
- [148] Pearu Peterson. F2PY: a tool for connecting fortran and python programs. *International Journal of Computational Science and Engineering*, 4(4):296, 2009. ISSN 1742-7185, 1742-7193. doi: 10.1504/IJCSE.2009.029165. URL <http://www.inderscience.com/info/inarticle.php?artid=29165>. (page 46.)
- [149] Steve Plimpton. Fast parallel algorithms for short-range molecular dynamics. *Journal of Computational Physics*, 117(1): 1–19, March 1995. ISSN 0021-9991. doi: 10.1006/jcph.1995.1039. URL <http://www.sciencedirect.com/science/article/pii/S002199918571039X>. (page 59.)
- [150] William Press, Saul Teukolsky, William Vetterling, and Brian Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, 3 edition, September 2007. ISBN 0521880688. (pages 14 and 30.)
- [151] P. Pulay. Ab initio calculation of force constants and equilibrium geometries in polyatomic molecules. *Molecular Physics*, 17(2): 197–204, 1969. ISSN 0026-8976. doi: 10.1080/00268976900100941. URL <http://dx.doi.org/10.1080/00268976900100941>. (page 9.)
- [152] P. Ramachandran and G. Varoquaux. Mayavi: 3D visualization of scientific data. *Computing in Science Engineering*, 13(2):40–51, 2011. ISSN 1521-9615. doi: 10.1109/MCSE.2011.35. (page 26.)
- [153] Karsten Reuter. First-principles kinetic monte carlo simulations for heterogeneous catalysis: Concepts, status and frontiers, 2009. URL <http://edoc.mpg.de/397332>. (pages 3 and 39.)

- [154] Karsten Reuter. First-principles kinetic monte carlo simulations for heterogeneous catalysis: Concepts, status, and frontiers. In *Modeling and Simulation of Heterogeneous Catalytic Reactions: From the Molecular Process to the Technical System*, pages 71–112. Wiley-VCH, Weinheim, 2012. editor: Deutschmann, Olaf. (pages 3, 24, 25, 30, 72, 89, 91, 93, and 94.)
- [155] Karsten Reuter and Matthias Scheffler. Composition, structure, and stability of RuO₂(110) as a function of oxygen pressure. *Physical Review B*, 65(3):035406, December 2001. doi: 10.1103/PhysRevB.65.035406. URL <http://link.aps.org/doi/10.1103/PhysRevB.65.035406>. (page 74.)
- [156] Karsten Reuter and Matthias Scheffler. First-principles atomistic thermodynamics for oxidation catalysis: Surface phase diagrams and catalytically interesting regions. *Physical Review Letters*, 90(4):046103, January 2003. doi: 10.1103/PhysRevLett.90.046103. URL <http://link.aps.org/doi/10.1103/PhysRevLett.90.046103>. (pages 21, 83, and 90.)
- [157] Karsten Reuter and Matthias Scheffler. First-principles kinetic monte carlo simulations for heterogeneous catalysis: Application to the CO oxidation at RuO₂(110). *Physical Review B*, 73(4):045433, January 2006. doi: 10.1103/PhysRevB.73.045433. URL <http://link.aps.org/doi/10.1103/PhysRevB.73.045433>. (pages 3, 24, 47, 48, 73, 83, 98, and 110.)
- [158] Karsten Reuter, Daan Frenkel, and Matthias Scheffler. The steady state of heterogeneous catalysis, studied by first-principles statistical mechanics. *Physical Review Letters*, 93(11):116105, 2004. doi: 10.1103/PhysRevLett.93.116105. URL <http://link.aps.org/doi/10.1103/PhysRevLett.93.116105>. (pages 23, 24, 47, 48, and 91.)
- [159] K.H. Rieder and W. Stocker. A helium diffraction study of the p(2X2) phase of oxygen on pd(100). *Surface Science Letters*, 150(1):L66–L70, February 1985. ISSN 0167-2584. doi: 10.1016/0167-2584(85)91213-7. URL <http://www.sciencedirect.com/science/article/pii/0167258485912137>. (page 76.)
- [160] Michael Rieger, Jutta Rogal, and Karsten Reuter. Effect of surface nanostructure on temperature programmed reaction spectroscopy: First-principles kinetic monte carlo simulations of CO oxidation at RuO₂(110). *Physical Review Letters*, 100(1):016105, January 2008. doi: 10.1103/PhysRevLett.100.016105. URL <http://link.aps.org/doi/10.1103/PhysRevLett.100.016105>. (page 91.)
- [161] Jutta Rogal and Karsten Reuter. Ab initio atomistic thermodynamics for surfaces: A primer. Technical report, February 2006.

- URL <http://stinet.dtic.mil/oai/oai?&verb=getRecord&metadataPrefix=html&identifier=ADA476575>. (page 21.)
- [162] Jutta Rogal, Karsten Reuter, and Matthias Scheffler. CO oxidation at pd(100): A first-principles constrained thermodynamics study. *Physical Review B*, 75(20), May 2007. ISSN 1098-0121. doi: 10.1103/PhysRevB.75.205433. URL <http://prb.aps.org/abstract/PRB/v75/i20/e205433>. (pages 3 and 72.)
- [163] Jutta Rogal, Karsten Reuter, and Matthias Scheffler. First-principles statistical mechanics study of the stability of a subnanometer thin surface oxide in reactive environments: CO oxidation at pd(100). *Physical Review Letters*, 98(4):046101, January 2007. doi: 10.1103/PhysRevLett.98.046101. URL <http://link.aps.org/doi/10.1103/PhysRevLett.98.046101>. (pages 24, 47, 72, 76, 77, 80, 83, 91, and 103.)
- [164] Jutta Rogal, Karsten Reuter, and Matthias Scheffler. CO oxidation on pd(100) at technologically relevant pressure conditions: First-principles kinetic monte carlo study. *Physical Review B (Condensed Matter and Materials Physics)*, 77(15):155410–12, April 2008. doi: 10.1103/PhysRevB.77.155410. URL <http://link.aps.org/abstract/PRB/v77/e155410>. (pages 24, 47, 72, 76, 77, 81, 96, 97, and 99.)
- [165] Maarten K. Sabbe, Marie-Françoise Reyniers, and Karsten Reuter. First-principles kinetic modeling in heterogeneous catalysis: an industrial perspective on best-practice, gaps and needs. *Catalysis Science & Technology*, 2(10):2010, 2012. ISSN 2044-4753, 2044-4761. doi: 10.1039/c2cy20261a. URL <http://pubs.rsc.org/en/content/articlehtml/2012/cy/c2cy20261a>. (pages 23, 30, 72, and 91.)
- [166] Mark Saeys, M.-F. Reyniers, M. Neurock, and G. B. Marin. Ab initio reaction path analysis of benzene hydrogenation to cyclohexane on pt(111)†. *The Journal of Physical Chemistry B*, 109(6):2064–2073, February 2005. ISSN 1520-6106. doi: 10.1021/jp049421j. URL <http://dx.doi.org/10.1021/jp049421j>. (page 23.)
- [167] J. J. Sakurai. *Modern Quantum Mechanics*. Addison Wesley, Reading, Mass, 1 edition edition, September 1993. ISBN 9780201539295. (page 8.)
- [168] M. Saliccioli, M. Stamatakis, S. Caratzoulas, and D.G. Vlachos. A review of multiscale modeling of metal-catalyzed reactions: Mechanism development for complexity and emergent behavior. *Chemical Engineering Science*, 66(19):4319–4355, October 2011. ISSN 0009-2509. doi: 10.1016/j.ces.2011.05.050. URL <http://www.sciencedirect.com/science/article/pii/S000925091100368X>. (page 23.)

- [169] B.C. Sales, J.E. Turner, and M.B. Maple. Oscillatory oxidation of CO over pt, pd and ir catalysts: Theory. *Surface Science*, 114(2-3):381-394, February 1982. ISSN 0039-6028. doi: 10.1016/0039-6028(82)90692-6. URL <http://www.sciencedirect.com/science/article/pii/0039602882906926>. (page 90.)
- [170] Rutger A. van Santen and Matthew Neurock. *Molecular Heterogeneous Catalysis: A Conceptual and Computational Approach*. John Wiley & Sons, June 2009. ISBN 9783527608348. (page 1.)
- [171] Rutger A. van Santen and Philippe Sautet. *Computational Methods in Catalysis and Materials Science*. Wiley-VCH, 1 edition edition, March 2009. ISBN 9783527320325. (page 1.)
- [172] Philippe Sautet and Françoise Delbecq. Catalysis and surface organometallic chemistry: a view from theory and simulations. *Chemical reviews*, 110(3):1788-1806, March 2010. ISSN 1520-6890. doi: 10.1021/cr900295b. (page 23.)
- [173] C. Schaefer and A. P. J. Jansen. Coupling of kinetic monte carlo simulations of surface reactions to transport in a fluid for heterogeneous catalytic reactor modeling. *The Journal of Chemical Physics*, 138(5):054102-054102-9, February 2013. ISSN 00219606. doi: doi:10.1063/1.4789419. URL http://jcp.aip.org/resource/1/jcpsa6/v138/i5/p054102_s1. (pages 3 and 24.)
- [174] David J. Schmidt, Wei Chen, C. Wolverton, and William F. Schneider. Performance of cluster expansions of coverage-dependent adsorption of atomic oxygen on pt(111). *Journal of Chemical Theory and Computation*, 8(1):264-273, January 2012. ISSN 1549-9618. doi: 10.1021/ct200659c. URL <http://dx.doi.org/10.1021/ct200659c>. (page 21.)
- [175] E. Schrödinger. An undulatory theory of the mechanics of atoms and molecules. *Physical Review*, 28(6):1049-1070, December 1926. doi: 10.1103/PhysRev.28.1049. URL <http://link.aps.org/doi/10.1103/PhysRev.28.1049>. (page 7.)
- [176] J.M.C. Scott. LXXXII. the binding energy of the thomas-fermi atom. *Philosophical Magazine Series 7*, 43(343):859-867, 1952. ISSN 1941-5982. doi: 10.1080/14786440808520234. URL <http://www.tandfonline.com/doi/abs/10.1080/14786440808520234>. (page 10.)
- [177] Selim M. Senkan. High-throughput screening of solid-state catalyst libraries. *Nature*, 394(6691):350-353, July 1998. ISSN 0028-0836. doi: 10.1038/28575. URL <http://www.nature.com/nature/journal/v394/n6691/abs/394350a0.html>. (page 2.)

- [178] Daniel Sheppard, Rye Terrell, and Graeme Henkelman. Optimization methods for finding minimum energy paths. *The Journal of Chemical Physics*, 128(13):134106–134106–10, April 2008. ISSN 00219606. doi: doi:10.1063/1.2841941. URL http://jcp.aip.org/resource/1/jcpsa6/v128/i13/p134106_s1. (pages 58 and 60.)
- [179] Michail Stamatakis and Dionisios G. Vlachos. A graph-theoretical kinetic monte carlo framework for on-lattice chemical kinetics. *The Journal of Chemical Physics*, 134(21):214115, June 2011. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.3596751. URL <http://scitation.aip.org/content/aip/journal/jcp/134/21/10.1063/1.3596751>. (page 24.)
- [180] E.M. Stuve, R.J. Madix, and C.R. Brundle. CO oxidation on pd(100): A study of the coadsorption of oxygen and carbon monoxide. *Surface Science*, 146(1):155–178, October 1984. ISSN 0039-6028. doi: 16/0039-6028(84)90235-8. URL <http://www.sciencedirect.com/science/article/pii/0039602884902358>. (page 76.)
- [181] Hugh Stott Taylor. A theory of the catalytic surface. *Proceedings of the Royal Society of London. Series A*, 108(745):105–111, May 1925. ISSN 1364-5021, 1471-2946. doi: 10.1098/rspa.1925.0061. URL <http://rspa.royalsocietypublishing.org/content/108/745/105>. (page 117.)
- [182] Burcin Temel, Hakim Meskine, Karsten Reuter, Matthias Scheffler, and Horia Metiu. Does phenomenological kinetics provide an adequate description of heterogeneous catalytic reactions? *The Journal of Chemical Physics*, 126(20):204711, May 2007. ISSN 0021-9606. doi: 10.1063/1.2741556. URL <http://www.ncbi.nlm.nih.gov/pubmed/17552793>. (pages 23, 24, and 72.)
- [183] L. H. Thomas. The calculation of atomic fields. *Mathematical Proceedings of the Cambridge Philosophical Society*, 23(05):542–548, 1927. doi: 10.1017/S0305004100011683. (page 9.)
- [184] M. Todorova, E. Lundgren, V. Blum, A. Mikkelsen, S. Gray, J. Gustafson, M. Borg, J. Rogal, K. Reuter, J. N. Andersen, and M. Scheffler. The pd(1 0 0)-r27°-o surface oxide revisited. *Surface Science*, 541(1-3):101–112, September 2003. ISSN 0039-6028. doi: 10.1016/S0039-6028(03)00873-2. URL <http://www.sciencedirect.com/science/article/B6TVX-491BHDN-1/2/cd16c6926abf8a34aaf17a9d21cf4023>. (pages 71, 77, 92, and 101.)
- [185] Oleg Trushin, Altaf Karim, Abdelkader Kara, and Talat S. Rahman. Self-learning kinetic monte carlo method: Application to cu(111). *Physical Review B*, 72(11):115401, 2005. doi: 10.1103/

- PhysRevB.72.115401. URL <http://link.aps.org/abstract/PRB/v72/e115401>. Copyright (C) 2009 The American Physical Society; Please report any problems to prola@aps.org. (page 94.)
- [186] M. Tüshaus, W. Berndt, H. Conrad, A. M. Bradshaw, and B. Persson. Understanding the structure of high coverage CO adlayers. *Applied Physics A: Materials Science & Processing*, 51(2):91–98, 1990. ISSN 0947-8396. doi: 10.1007/BF00324270. URL <http://www.springerlink.com/content/hx11r8182813458k/abstract/>. (pages 76, 81, and 86.)
- [187] P. Uvdal, P.-A. Karlsson, C. Nyberg, S. Andersson, and N.V. Richardson. On the structure of dense CO overlayers. *Surface Science*, 202(1-2):167–182, August 1988. ISSN 0039-6028. doi: 16/0039-6028(88)90067-2. URL <http://www.sciencedirect.com/science/article/pii/0039602888900672>. (page 76.)
- [188] R. van Rijn, O. Balmes, A. Resta, D. Wermeille, R. Westerström, J. Gustafson, R. Felici, E. Lundgren, and J. W. M. Frenken. Surface structure and reactivity of pd(100) during CO oxidation near ambient pressures. *Physical Chemistry Chemical Physics*, 13: 13167–13171, 2011. ISSN 1463-9076. doi: 10.1039/c1cp20989b. URL <http://pubs.rsc.org/en/Content/ArticleLanding/2011/CP/c1cp20989b>. (pages 3, 87, and 92.)
- [189] Richard van Rijn, Olivier Balmes, Roberto Felici, Johan Gustafson, Didier Wermeille, Rasmus Westerström, Edvin Lundgren, and Joost W. M. Frenken. Comment on “CO oxidation on pt-group metals from ultrahigh vacuum to near atmospheric pressures. 2. palladium and platinum”. *The Journal of Physical Chemistry C*, 114(14):6875–6876, April 2010. doi: 10.1021/jp911406x. URL <http://dx.doi.org/10.1021/jp911406x>. (pages 4, 71, and 72.)
- [190] R. A. van Santen, A. J. Markvoort, I. A. W. Filot, M. M. Ghouri, and E. J. M. Hensen. Mechanism and microkinetics of the Fischer–Tropsch reaction. *Physical Chemistry Chemical Physics*, 15(40): 17038, 2013. ISSN 1463-9076, 1463-9084. doi: 10.1039/c3cp52506f. URL <http://xlink.rsc.org/?DOI=c3cp52506f>. (page 23.)
- [191] David Vanderbilt. Soft self-consistent pseudopotentials in a generalized eigenvalue formalism. *Physical Review B*, 41(11):7892, April 1990. doi: 10.1103/PhysRevB.41.7892. URL <http://link.aps.org/doi/10.1103/PhysRevB.41.7892>. (page 131.)
- [192] D. G. Vlachos. Multiscale integration hybrid algorithms for homogeneous–heterogeneous reactors. *AIChE Journal*, 43 (11):3031–3041, 1997. ISSN 1547-5905. doi: 10.1002/aic.690431115. URL <http://onlinelibrary.wiley.com/doi/10.1002/aic.690431115/abstract>. (pages 3 and 24.)

- [193] Arthur F. Voter. Introduction to the kinetic monte carlo method. In *Radiation Effects in Solids*, page 1–23. Springer, 2007. URL http://dx.doi.org/10.1007/978-1-4020-5295-8_1. (pages 24, 25, and 30.)
- [194] Arthur F. Voter, Francesco Montalenti, and Timothy C. Germann. Extending the time scale in atomistic simulation of materials. *Annual review of materials research*, 32:321–346, 2002. ISSN 1531-7331. URL <http://cat.inist.fr/?aModele=afficheN&cpsidt=13860930>. (page 26.)
- [195] Bert M. Weckhuysen. Snapshots of a working catalyst: possibilities and limitations of in situ spectroscopy in the field of heterogeneous catalysis. *Chemical Communications*, 2002(2):97–110, January 2002. ISSN 1364-548X. doi: 10.1039/B107686H. URL <http://pubs.rsc.org/en/content/articlelanding/2002/cc/b107686h>. (page 3.)
- [196] C. F. v Weizsäcker. Zur theorie der kernmassen. *Zeitschrift für Physik*, 96(7-8):431–458, July 1935. ISSN 0044-3328. doi: 10.1007/BF01337700. URL <http://link.springer.com/article/10.1007/BF01337700>. (page 10.)
- [197] Wikipedia. List of quantum chemistry and solid-state physics software, October 2013. URL http://en.wikipedia.org/w/index.php?title=List_of_quantum_chemistry_and_solid-state_physics_software&oldid=577618719. Page Version ID: 577618719. (page 24.)
- [198] C. Wu, D.J. Schmidt, C. Wolverton, and W.F. Schneider. Accurate coverage-dependence incorporated into first-principles kinetic models: Catalytic NO oxidation on pt (1 1 1). *Journal of Catalysis*, 286:88–94, February 2012. ISSN 0021-9517. doi: 10.1016/j.jcat.2011.10.020. URL <http://www.sciencedirect.com/science/article/pii/S0021951711003551>. (page 24.)
- [199] Lijun Xu and Graeme Henkelman. Adaptive kinetic monte carlo for first-principles accelerated dynamics. *The Journal of Chemical Physics*, 129(11):114104, 2008. ISSN 00219606. doi: 10.1063/1.2976010. URL http://apps.isiknowledge.com/full_record.do?product=UA&search_mode=GeneralSearch&qid=1&SID=U27H48Bkc5I8D178DCA&page=1&doc=1&colname=WOS. (page 94.)
- [200] J Zetterberg, S Blomberg, J Gustafson, Z W Sun, Z S Li, E Lundgren, and M Aldén. An in situ set up for the detection of CO₂ from catalytic CO oxidation by using planar laser-induced fluorescence. *The Review of scientific instruments*, 83(5):053104, May 2012. ISSN 1089-7623. doi: 10.1063/1.4711130. (page 3.)

- [201] C. J. Zhang and P. Hu. CO oxidation on pd(100) and pd(111): A comparative study of reaction pathways and reactivity at low and medium coverages. *Journal of the American Chemical Society*, 123(6):1166–1172, February 2001. doi: 10.1021/ja002432f. URL <http://dx.doi.org/10.1021/ja002432f>. (page 77.)
- [202] Changjun Zhang and Angelos Michaelides. Quantum nuclear effects on the location of hydrogen above and below the palladium (100) surface. *Surface Science*, 605(7–8):689–694, April 2011. ISSN 0039-6028. doi: 10.1016/j.susc.2011.01.004. URL <http://www.sciencedirect.com/science/article/pii/S0039602811000070>. (page 133.)
- [203] Yongsheng Zhang, Jutta Rogal, and Karsten Reuter. Density-functional theory investigation of oxygen adsorption at pd(11N) vicinal surfaces (n=3,5,7) : Influence of neighboring steps. *Physical Review B*, 74(12):125414, 2006. doi: 10.1103/PhysRevB.74.125414. URL <http://link.aps.org/doi/10.1103/PhysRevB.74.125414>. (page 38.)
- [204] Yongsheng Zhang, Volker Blum, and Karsten Reuter. Accuracy of first-principles lateral interactions: Oxygen at pd(100). *Physical Review B*, 75(23):235406, June 2007. doi: 10.1103/PhysRevB.75.235406. URL <http://link.aps.org/doi/10.1103/PhysRevB.75.235406>. (pages 21, 72, 76, 77, and 96.)
- [205] V.P. Zhdanov and B. Kasemo. Kinetic phase transitions in simple reactions on solid surfaces. *Surface Science Reports*, 20(3):113–189, April 1994. ISSN 0167-5729. doi: 10.1016/0167-5729(94)90009-4. URL <http://www.sciencedirect.com/science/article/pii/0167572994900094>. (page 90.)
- [206] G. Zheng and E. I. Altman. The reactivity of surface oxygen phases on pd(100) toward reduction by CO. *The Journal of Physical Chemistry B*, 106(5):1048–1057, February 2002. ISSN 1520-6106. doi: 10.1021/jp013395x. URL http://apps.isiknowledge.com/full_record.do?product=UA&search_mode=GeneralSearch&qid=19&SID=N11ILgijLnFDGk18Mim&page=1&doc=8&colname=WOS. (pages 3 and 71.)
- [207] G. Zheng and E.I. Altman. The oxidation mechanism of pd(1 0 0). *Surface Science*, 504:253–270, April 2002. ISSN 0039-6028. doi: 10.1016/S0039-6028(02)01104-4. URL <http://www.sciencedirect.com/science/article/pii/S0039602802011044>. (page 81.)
- [208] Robert M. Ziff, Erdagon Gulari, and Yoav Barshad. Kinetic phase transitions in an irreversible surface-reaction model. *Physical Review Letters*, 56(24):2553, June 1986. doi: 10.1103/

PhysRevLett.56.2553. URL <http://link.aps.org/doi/10.1103/PhysRevLett.56.2553>. (pages 46, 47, and 90.)

ACRONYMS

ASE	Atomic Simulation Environment. 59
BFGS	Broyden-Fletcher-Goldfarb-Shanno. 58, 128, 134
CI-NEB	Climbing Image Nudged Elastic Band. 55, 59
DFT	Density Functional Theory. 7, 131
FIRE	Fast Inertial Relaxation Engine. 57–66
fNEB	freezing Nudged Elastic Band. 57–63, 65–67
FS	Freezing String. 56, 57, 67
HRXPS	High Resolution X-ray Photoelectron Spectroscopy. 4
kMC	kinetic Monte Carlo. 3, 4, 53
LAMMPS	Large-scale Atomic/Molecular Massively Parallel Simulator. 59
LBFGS	Limited Memory Broyden-Fletcher-Goldfarb- Shanno. 58–60, 63–66
MEP	Minimum Energy Path. 56
NEB	Nudged Elastic Band. 55–63, 65–67
PW	plane-wave. 14
SCF	Self-Consistent Field. 13, 14, 55, 59, 64–66
STM	Scanning Tunneling Microscope. 4

ACKNOWLEDGEMENTS

The completion of this thesis would not have been possible without the support and input of many people. I would like to thank my parents Ursula and Karl Friedrich for always supporting me and my education and giving me so much freedom in pursuing my interests. I would like to thank my supervisor Prof. Karsten Reuter for initiating this project, trusting in my skills, and giving me sufficient time to work on it. I also would like to thank Prof. Matthias Scheffler from the Fritz-Haber Institut in Berlin for initiating and accelerating the project by generously providing computational resources.

I am furthermore indebted to Antonia Aravena for being always a good listener and friend, Hakeem Abrar Ahmad for having a good time at the ICTAC-14 conference, Daniel Berger for being a fair colleague and conference room mate, Thomas Becker for an enlightening introduction to project management, Christopher Bronner for proof-reading the introduction, Sam Chill for initiating an open database on transition state searches, Juscele Duraes for being my girl friend and best friend in all the years, Frank Essenberg for valuable discussions, Ingebord Weidl and Christian Guggenberger from the Rechenzentrum for administering the AIMS cluster smoothly and providing for a stellar user experience, Steffen Hartung for having a fun time around kinetic Monte Carlo on batteries thanks for the good discussion (the discussions during his stay made me rewrite the kmos back end from scratch which was crucial for getting the model done), Graeme Henkelman for discussing NEB parameters and his forum and generous support on NEB, Sergey Levchenko for providing precious feedback and ideas on the software development part and for crucial encouragement, Sebastian Matera for a lot of clarifying discussion on statistical physics and numerical problems many good discussion on multi-scaling late-night encouragements during bug hunts, Reinhard Maurer for support in the submission process, Jörg Meyer for being a true-blue mentor and friend and fruitful discussions and setting a great example for a vigorous scientist, Randall Meyer for support and fruitful discussions on chemical engineering, Erik McNellis for giving valuable hints on scientific practice and software development, Hakim Meskine for giving inspiration and guidance, Claudia Mangold for having a lot of patience of discussing my confusions over subtleties of rate constants and for printing the poster for the FBR, Robert C. Nordvall for all the good times, Horia Metiu for introducing me to kMC and kinetic modeling, Arwed Niestroj for stimulating discussions and valuable advice, Harald Oberhofer for always having an open door to life's persistent questions, Udo

Pohl for stimulating discussions and great advice, Lorenz Rognoni for holding up the spirit of Berlin left and right of the Isar river, Marek Rudnicki for a lot of stimulating discussions and introducing me to mindfulness, Andriy Samsonyuk for being a super fair partner in and around technical support, Aditya Savara for valuable feedback around `kmos` and advice, Florian Schweinberger for very intriguing discussions around information management, new ideas and a whole lot a motivation, Christoph Scheurer for always having a watching eye on file servers and network security and being an angel for the whole group and my ideas, Christoph Schober for picking up my responsibility in the group so heroically and by way of this allowing me to finish this thesis, Daniel Strobusch for being a great discussion partner around python, numerics, workflows, and parallelization, Tongyu Wang for being a good discussion partner and office mate, Maxime Van den Bossche for being a brave first user of `kmos` and stimulating discussions around kinetic Monte Carlo, Greg Wilson for teaching me software development, David Weikersdorfer for stimulating discussions about geometric algorithms and project management, the group members of the chair of Theoretical Chemistry, and the countless unnamed heroes contributing to free and open source software.