On a Delocalised Oxygen Model of Two-Level System Defects in Josephson Junctions

A Thesis Submitted in
Fulfilment of the Requirements
for the Degree of
Doctor of Philosophy (Applied Physics)

by

Timothy Charles DuBois
Bachelor of Science (Applied Sciences) (Honours), RMIT
Bachelor of Science (Physics), RMIT
Certificate of Ancient Languages, Macquarie

School of Applied Sciences
College of Science, Engineering and Health
RMIT University

Primary Supervisor: Associate Professor Jared Cole
Secondary Supervisor: Professor Salvy Russo

May 2015

PHILOSOPHIÆ

NATURALIS

OXYGENIUM

NON LOCALIS

Autore T C. DUBOIS, *RMIT Uni. App. Sci.* Bacalarius Phyſica *Cum Honore*, & Chemica Indiviſible Phyſica Sodali.

I M P R I M A T U R

S. R U S S O, *Chem. Ind. Phys.* P R Æ S E S.

Maius 28. 2015.

B A T M A N I A,

Juſſa *RMIT Univerſitas* ac Typis *Jared Cole*. Proſtat apud plures Bibliopolas. *Anno* MMXV.

Declaration

I hereby declare that the contents of this thesis has been undertaken by me alone except where due acknowledgement has been made. Material herein has has not been submitted previously, in whole or in part, to qualify for any other academic award. The content of this thesis is the result of work which has been carried out subsequent to the official commencement date of my PhD candidature. Any editorial work, paid or unpaid, which was carried out by a third party is acknowledged. Finally, all ethics procedures and guidelines have been followed during all phases of this research.

Timothy DuBois
Melbourne, May 28, 2015

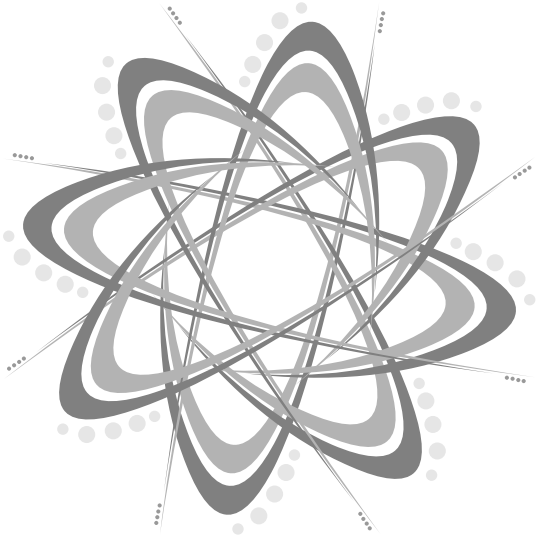

Acknowledgements

This thesis would never have been possible without the ideas, friendship, support and assistance of numerous people. Foremost among them is indubitably Associate Professor Jared Cole. Whilst he may have balked my attempts to investigate the crazy world of AdS/CFT and holographic superconductivity when we first met, this was probably the only time he discouraged the harebrained scheme *du jour*. Always approachable and engaging, it has been a continual joy to work with him—even if photographic evidence suggests I distrust his opinions.

Professor Salvy Russo has been a valued mentor since 2007. If he's not winning karate championships one point away from disqualification, he's writing proofs by intimidation on the white board in front of undergraduates. This take no-prisoners-attitude is reflected in his dedication and meticulousness towards both his research and his students—even with the constant pressures the title of Deputy Dean bestows upon him.

A special thanks to all members past and present of the Chemical and Quantum Physics Group. I owe much to Dr Manolo Per and Dr Akin Budi, both of whom spent a great deal of time tutoring me in the art of numerical modelling during the late undergrad years. Jan Jeske, for putting up with my rants; contributing succinct and useful information each time I ran out of breath. Also for being a wonderful 'working' holiday companion and introducing me to a number of fantastic international colleagues. David Ing and Jesse Vaitkus for actively participating in, if not instigating office shenanigans. Only after offering valuable insight that greatly advances the scientific endeavour, of course. Kelly, Daniel, Jackson, Muhammad, Martin, Nico, Paolo, Chris, George and the non-CQP members of the Engine Room for a great working environment, as well as the constructive and encouraging input over the years. Martin: I hope you uncover more about this elusive defect during your candidature than is printed on these pages.

My supervisors at DSTO: Dr Milan Jamriska and Dr Alex Skvortsov for their patience, understanding and encouragement. Furthermore, their acceptance that I may be off doing something other than my work there—without warning, at the drop of a hat.

Finally, my family for their continual encouragement and for at least attempting to not zone out when I discuss the elegance of `Haskell` or the why the collective odium towards the `jet colormap` is completely warranted.

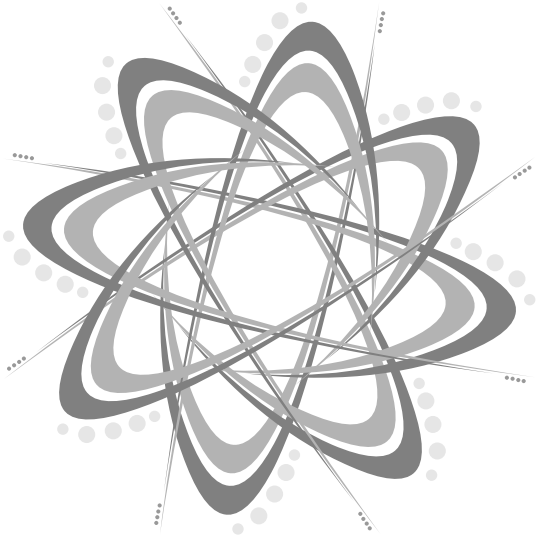

Contents

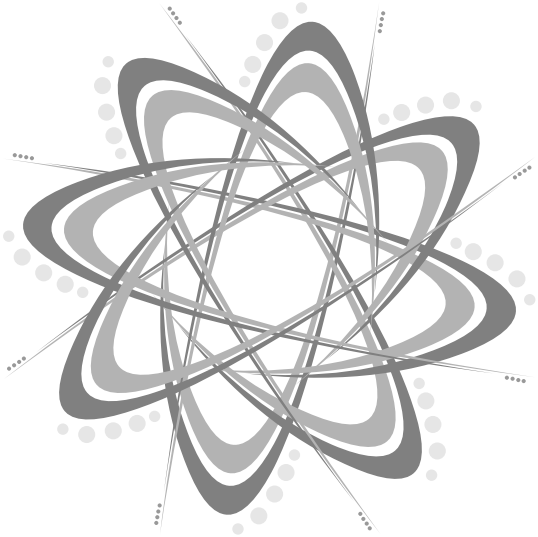

List of Figures

List of Tables

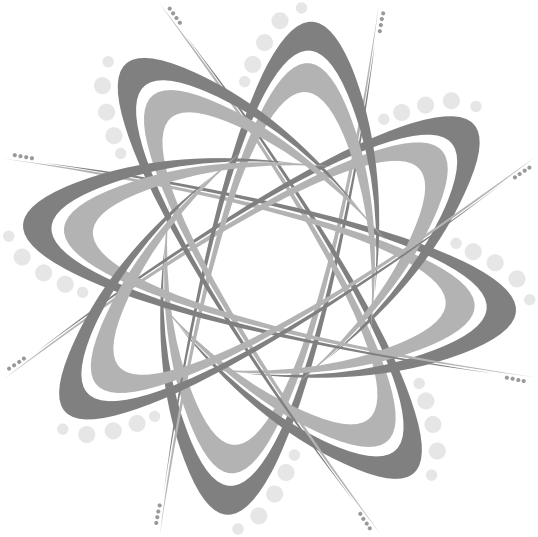

List of Publications

This thesis is based in part on the following published articles:

1. T. C. DuBois, S. P. Russo, and J. H. Cole, "A 3D investigation of delocalised oxygen two-level defects in Josephson junctions", arXiv:1508.05204.

2. T. DuBois, M. Cyster, G. Opletal, S. Russo, and J. Cole, "Constructing *ab initio* models of ultra-thin Al–AlO_x–Al barriers", *Molecular Simulation* **Ian Snook Special Issue**, (2015).

3. T. C. DuBois, S. P. Russo, and J. H. Cole, "Atomic delocalisation as a microscopic origin of two-level defects in Josephson junctions", *New Journal of Physics* **17**, 13 (2015).

4. T. C. DuBois, M. C. Per, S. P. Russo, and J. H. Cole, "Delocalized Oxygen as the Origin of Two-Level Defects in Josephson Junctions", *Physical Review Letters* **110**, 077002 (2013).

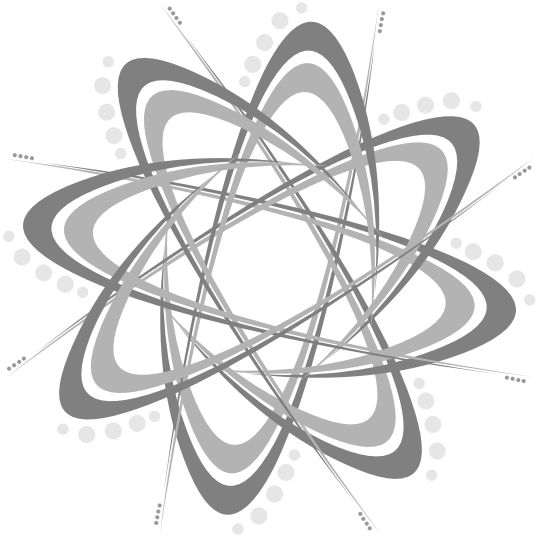

Nomenclature

Abbreviations:

AD	Automatic Differentiation, see section 3.2.
ALD	Atomic Layer Deposition, see subsection 1.6.2.
DC	Direct current.
DFT	Density Functional Theory.
DSTO	Defence Science and Technology Organisation.
EAM	Embedded Atom Model, see section 3.1.
JIT	Just-in-time compilation.
JJ	Josephson junction.
LC	Inductor-Capacitor circuit elements, see section 1.3.
ODE	Ordinary Differential Equation.
PBC	Periodic Boundary Conditions.
PDE	Partial Differential Equation.
SQUID	Superconducting Quantum Interference Device.
STM	Standard Tunneling Model.
TISE	Time Independent Schrödinger Equation.
TLS	Two level system.

Symbols:

Å	Ångström, 1 Å$= 10^{-10}$ m.
$\|\cdot\|$	A pair of atoms displaced equidistantly from an origin, *e.g.* $\|X\| = -X, +X$, see section 4.2.
∇	Del operator. In Cartesian coordinates, $\nabla = \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}\right)$.
$\mathcal{E}_n^{(d)}(x,h)$	Total error for a d^{th} derivative central difference method of order n, see subsection 3.4.3 and sections 3.7 and 3.9.

$\mathcal{F}_\delta$ Cancellation error coefficient for central difference methods, see chapter 3.

$\mathcal{F}_\varepsilon$ Condition error coefficient for central difference methods, see chapter 3.

$\hbar$ Reduced Planck constant, $\hbar \simeq 1.055 \times 10^{-34}$ Js.

$|\cdot\rangle$ Ket vector in Dirac notation. When singular it denotes the wave-function of a system.

$\mathcal{O}$ Big-O notation, signifying the order of a value.

$\wp$ Electronic dipole moment. May also have a directional subscript *e.g.* $\wp_x$ denoting a dipole in the x direction, see equations (1.6) and (4.2).

∂ Partial derivative.

$[f_a|f_b]$ A bracket notation representing Coulomb interaction integrals, see section 3.1.

$\mathcal{TE}_n(x,h)$ Truncation error of order n for central difference methods, see subsection 3.4.1.

Latin Letters:

C_n Approximate truncation coefficient for small step sizes. Expands to $C_n(x,h)$ for all steps, see chapter 3.

E Energy of a system.

e Elementary charge, $e = 1.602176565 \times 10^{19}$ C.

F An electron sitting in a vacancy defect, see subsection 1.6.1.

$\mathcal{F}_n^{(d)}(x,h)$ Finite difference approximation of the d^{th} derivative of $f(x)$ of order n using step size h, see chapter 3.

$f_n^{(d)}(x)$ True d^{th} derivative of $f(x)$ of order n expressed in terms of a finite difference approximation, see chapter 3.

$F_i\left[\rho_i\right]$ Energy required to embed atom i in a local electron density ρ_i. A contribution to the Streitz-Mintmire potential, see (3.6).

f Frequency, predominantly used to describe the $1/f$ noise spectrum, see section 1.4 and subsection 1.5.1.

f_k A short form for the central difference expansions, see sections 3.8 and 3.9.

$G(r)$	Radial distribution function, see (2.2).
H	The Hamiltonian operator.
J_i^0	Self-Coulomb repulsion. Used in the Streitz-Mintmire potential, see (3.2).
L	Strain length, see (6.14).
L_k	Lindblad operator, see (1.4).
m	Mass of a particle.
N_{pairs}	Number of unique atomic pairs in a radial distribution function, see (2.2).
$p(r)$	Average number of atom pairs found at a distance between r and $r + dr$ in a radial distribution function, see (2.2).
$q_{i,j}$	Atomic charge contribution to the Streitz-Mintmire potential, see (3.2).
R_n	High order remainder terms of a Taylor series expansion of $f(x)$ to be truncated when using central difference, see section 3.3.
$\mathbf{r}$	General distance vector. In 1D $\mathbf{r} = x$ for example.
S_{max}	qubit-TLS coupling strength, see figure 1.6 and (1.6).
h	A step size used in the central difference formalism, see chapter 3.
T	Temperature.
$T_1,\ T_2$	Energy (longitudinal) and phase (transverse) relaxation times, see subsection 1.1.1.
V	Total volume contained in a radial distribution function, see (2.2).
$V(\mathbf{r})$	Potential V at point $\mathbf{r}$. Assumed to be the Streitz-Mintmire potential unless otherwise stated, see equations (3.1) and (3.2).
w	Junction width (usually of order $1-2$ nm), see (1.6).
x^*	Any value of x about which a central difference is applied, see chapter 3.

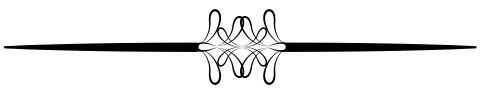

Greek Letters:

δ	Precision of a calculation using floating point numbers. On modern computers using double precision $\delta = 2^{-53}$, see chapter 3.
δ_{ij}	Kronecker Delta.

ε	Magnitude of the most significant digit affected by condition error, see chapter 3.
ζ	An unknown but bounded point used in central difference methods, see chapter 3.
Γ_k	Rate operator, see equations (1.4) and (1.5).
Λ	Unit conversion constant measure distance in Å and energy in μeV rather than SI, see (3.60).
γ	Relative phase correction to the dephasing time, see (1.5).
ρ	Density matrix, see (1.2) to (1.4).
$\Psi_n(x)$	n^{th} state wavefunction in the x direction.
φ_{ij}	Pair potentials (Buckingham and Rydberg) for Streitz-Mintmire, see (3.6).
ξ	ξ metric: a qualitative expression used to interpret eigenspectra, see (4.3).
ω	Angular frequency.

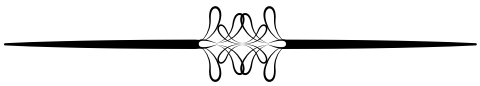

Superscripts:

(d)	The d^{th} derivative used in central difference methods, see chapter 3.

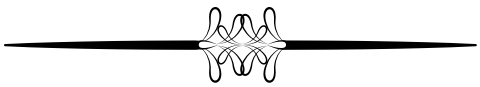

Subscripts:

01	Or some other integer pair. Denotes a difference between two levels. For example the splitting energy of a first excited state E_1 and a ground state E_0 is $E_{01} \equiv E_1 - E_0$, see section 4.1.
EAM	Embeded Atom Model potential contribution, see equations (3.2) and (3.5).
k	Arbitrary summation index.
max	Maximum value of parameter.
n	The n^{th} order expansion of the Taylor series used in central difference methods, see chapter 3.
opt	Optimal value of the parameter.
oxy	Property pertaining to an oxygen atom.

Summary

Constructing circuits which invoke the tunnelling effect of Josephson junctions in the superconducting regime manifest controllable, quantum properties. The development of a quantum computer employing these phenomena is just one of the many advantages to be gained from building such devices.

However, the use of quantum bits for computation is dependent on the ability to operate them in an essentially isolated environment. The phenomenon of decoherence refers to the instability of a quantum state of a system when it interacts with the surrounding environment. Superconducting qubits are sensitive to decoherence mechanisms within the readout leads connecting to the device, and more importantly from the materials from which they are manufactured. Removal or control of these imperfections is required before quantum computers using superconducting circuit architecture can be realised.

One identifiable noise source is the so called 'strongly coupled' two-level system (TLS). Comparable resonance frequencies to the qubit; strong coupling strengths and decoherence times long enough to allow coherent oscillations between the qubit and TLS have been experimentally measured. The premise of this thesis is that positional anharmonicity of oxygen atoms arises within the AlO_x barrier of the Josephson junction solely due to its amorphous construction. This ansatz allows the existence of various spatial configurations throughout the layer, causing unique TLS properties based solely on atomic positions and rotation in relation to the external electric field.

To validate this conjecture, Josephson junction models are constructed using a hybrid *ab initio* and molecular mechanics approach, with various stoichiometry and density properties to reflect experimental observations of the barrier. The resultant atomic positions provide input conditions through a Voronoi classification scheme to a framework describing an oxygen atom that has the capacity to become spatially delocalised as bonds perturb away from a crystalline structure. A direct diagonalisation method is developed for low dimensional descriptions and a Wick-rotated time-dependent Schrödinger equation implementation is used for three dimensional investigations.

Calculated properties are compared to many current experimental strongly coupled TLS measurements in phase qubits, which shows the models' capacity to explain how an oxygen atom can generate a large dipole and appropriate ground to first excited state splitting values expected of a TLS by merely migrating from its preferred lattice position.

“ἀγεωμέτρητος μηδεὶς εἰσίτω.”

—Ἀκαδημία

1 Introduction

Background information and motivation for this thesis. Discussion of current qubit designs, pitfalls and architecture issues. Investigation of current scientific knowledge pertaining to material defects in AlO_x.

The ability to create a device whose macroscopic quantities are quantum variables is advantageous for a myriad of reasons, one in particular is the development of a quantum computer. Constructing circuits which invoke the tunnelling effect of Josephson junctions (JJs) in conjunction with superconducting phenomena manifest controllable quantum properties which can be exploited to this end. However, the use of quantum bits (qubits) for computation is dependent on the ability to operate them in an essentially isolated environment, *i.e.* the system must operate on coherent states. The phenomenon of decoherence refers to the instability of a quantum state of a system when it interacts with the surrounding environment. In general, the stronger the interaction with the environment, the quicker the quantum state decoheres. Removal or control of these imperfections is required before quantum computers using superconducting circuit architecture can be realised.

1.1 A Qubit Primer

A quantum bit differs from its classical counterpart (the bit) in a number of ways. A bit can be represented by either 0 or 1, where the state of the system exists in only one of these positions at any given time, and are usually manifestations of a potential difference inside some macroscopic circuit. Qubits on the other hand are a two-level quantum mechanical system, existing in a Hilbert space or projective Hilbert space. In a given orthonormal basis, the state vectors $|0\rangle$ and $|1\rangle$ represent the two level nature of the system (perhaps corresponding to the spin-up and spin-down states of an electron for example). In addition to the state vectors, a qubit may exist in a superposition of the $|0\rangle$ and $|1\rangle$ basis, which is known as a pure state.

The Bloch sphere [1] is a useful tool to represent qubit states, with the poles interpreted as the two basis vectors, and the points on the surface of the sphere corresponding to the pure states of the system. Figure 1.1 illustrates this for an arbitrary pure state $|\psi\rangle$.

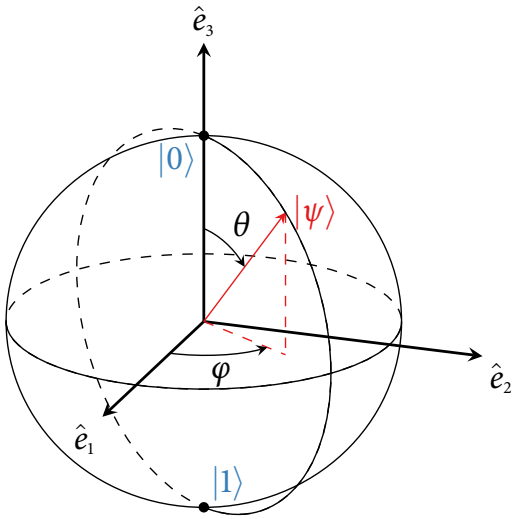

FIGURE 1.1— The Bloch Sphere: a geometrical representation of the pure state space of a two-level quantum system in an arbitrary Hilbert space $\hat{e}_{\{1,2,3\}}$. This example shows a pure state $|\psi\rangle$ calculated using (1.1) with $\varphi = \frac{\pi}{3}$ and $\theta = \frac{2\pi}{9}$.

The physical state of the system is not affected by global phase factors, thus in the complex superposition of the basis vectors, the coefficient of $|0\rangle$ can be set as real and non-negative. This means the pure state can be represented as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\varphi}\sin\left(\frac{\theta}{2}\right)|1\rangle, \tag{1.1}$$

where $0 \leq \theta < \pi$ and $0 \leq \varphi < 2\pi$.

The interior region of the Bloch sphere represents qubit mixed states: statistical ensembles of pure states. As these states cannot be described by ket vectors, a density matrix must be defined

$$\rho = \sum_k p_k |\psi_k\rangle\langle\psi_k|, \tag{1.2}$$

where p_k is the fraction of each pure state ψ_k represented in the mixture.

As a qubit state evolves through time, the equation of motion can be written as

$$i\hbar\frac{\mathrm{d}\rho}{\mathrm{d}t} = [H, \rho] \tag{1.3}$$

which evolves entirely due to the dynamics of a Hamiltonian H, provided the qubit is completely isolated from any environmental perturbations.

1.1.1 Interaction With the Environment

Environmental noise sources are seldom completely removed from qubit operations and in practically all current designs, interaction with this uncontrolled, perturbative source is unavoidable. Noise being an umbrella term for thermal radiation, collisions, lattice vibrations or any other form of energy dissipation phenomena which causes the time evolution of the qubit to become irreversible: taking away information about the state being represented. This (primarily) decreases the quantum coherence of the system and computational operations in turn deteriorate.

Qubit dynamics can therefore no longer be considered as a unitary evolution described by the von Neumann equation (1.3), but instead as an open quantum system. A master equation approach is therefore required to further describe qubit time evolution as it interacts with the environment. For simplicity, we assume that environment perturbations to our qubit come from an arbitrarily large bath of uncorrelated noise sources with a uniform frequency spectrum. Also, interactions between the system and environment are weak with respect to ordinary system dynamics determined by the Hamiltonian H. These conditions describe the most general type of Markovian and time homogeneous master equation, the Lindblad form [2]

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = -\frac{i}{\hbar}[H,\rho] + \sum_k \Gamma_k \left(L_k \rho L_k^\dagger - \frac{1}{2}\left\{ L_k^\dagger L_k, \rho \right\} \right) \tag{1.4}$$

where the Lindblad operators L_k and the rates Γ_k determine how the environment interacts with the system.

It can be shown that a qubit with energy splitting $\hbar\omega$ damped via environment interaction at temperature T looses coherence on two time scales [3]:

$$\frac{1}{T_1} = \Gamma_+ + \Gamma_-, \qquad \frac{1}{T_2} = \gamma + \frac{\Gamma_+ + \Gamma_-}{2}, \qquad \frac{\Gamma_+}{\Gamma_-} = \exp\left(-\frac{\hbar\omega}{k_B T} \right). \tag{1.5}$$

The times T_1 and T_2 are called the *energy (longitudinal)* and *phase (transverse)* relaxation times respectively. T_2 has an additional contribution γ which is a contribution due to random movement of the relative phase caused by perturbations of the environment.

These relaxation processes are the ultimate obstacle to the realisation of a working qubit architecture. Level populations tend to be more robust than the phase (which is easily disturbed by the environment), so generally $T_2 < T_1$. This time equates to a fundamental limit in most experimental systems and is often called the *decoherence* time.

1.2 Qubit Architectures

Physical qubit implementations can nominally be broken up into four categories: ultracold atoms, spin-based systems, quantum optics and superconducting circuits; each with their own strengths and weaknesses.

The following is not an exhaustive list or discussion of qubit conceptions, it is presented to the reader as an overview of the efforts in the wider quantum information science community before focusing on specific devices relevant to this thesis.

A comprehensive roadmap containing many of the advances in quantum computing (unfortunately only up to 2004) can be found via reference 4.

1.2.1 Ion Traps

Electromagnetic fields, light waves or a combination of both can be used to suspend and isolate charged atomic particles (ions) in free space [5]. Ion trapped qubits [6] exploit stable electronic or spin states of the ion to store information, which in turn can be processed and transmitted through applied laser fields and the collective quantised motion (via the Coulomb force) of other ions in the trap.

Two-qubit operations can be invoked by spin-phonon coupling (for example) induced by a laser pulse, and multiple entangled qubits can be coupled through internal states and external motion states in a similar fashion.

Ion traps show significant promise for the future of quantum information processing, as fundamental gating operations such as the C-NOT have also been demonstrated with high fidelity [7]. In fact, entire sets of gating operations can be used to implement algorithms and error correction on a set of coupled qubits [8, 9], and architectures with many coupled qubits are already experimentally accessible [10, 11].

1.2.2 Nitrogen-Vacancy Center of Diamond

The nitrogen vacancy (NV^-) center in diamond [12] is seemingly the most studied crystalline defect in existence. Application areas include ultrasensitive sensing of electric, magnetic and strain fields; biological imaging and sensing; and a room temperature, solid state qubit architecture.

The center exhibits atom-like properties: long-lived spin quantum states and well-defined optical transitions, whilst encased in a robust solid-state device. Both its bound electrons and nearby nuclear spins can be addressed using optical or microwave transitions much like the atomic states in ion traps. On-chip wiring and waveguides govern fast electrical and magnetic control which exploits the solid state nature of the diamond lattice.

The unique properties of the defect and its host lattice give rise to an electron spin lifetime longer than that of a trapped ion (at low temperatures) [13] and coherence times larger than a second at room temperature if nearby nuclear spins are utilised [14]. Ion implantation techniques can assist in stabilising the NV^- sites [15], charge-state initialisation is now understood [16] and individual defects can be addressed [17].

1.2.3 Quantum Dots

Quantum dots are made of semiconductor material and are used to contain and manipulate electrons in microcavity modes, which in turn behave as an *artificial* atom. Manipulation can come about via electronic [18] or optical [19] methods, and the resultant coherent system states can be used as a qubit.

Quantum dots are relatively simple to construct and control with modern semiconductor technology, and their environmental coupling is rather weak providing long coherence times. Even early schemes hinted at the possibility of coupling more than 100 qubits using a single cavity mode and laser fields to mediate coherent interactions between distant quantum dot spins [20], although this has proved difficult in practice.

Recent advances have seen some interesting possibilities in this field, such as reconfigurable architectures [21] and non-local control mechanisms [22].

1.2.4 Superconducting Circuits

Microscopic degrees of freedom (*e.g.* spin) of ions, crystal defects or artificial atoms are well isolated from the environment and therefore operate as good qubits. However, inter-qubit coupling in these systems becomes complex at scale and fast gating operations required for computation is a task that for many architectures introduces decoherence. Hence while the qubit itself is isolated, the quantum computer suffers as a whole.

The superconducting circuit architecture is different from these systems, as it instead uses a collective electrodynamic mode of macroscopic electrical

elements. Therefore, gating is no longer an issue: simple electrical elements (capacitors, inductors *etc.*) can be coupled to the qubits directly. Adaptation of nano-fabrication techniques developed in the semiconductor industry to build scalable integrated quantum devices once operational designs are developed is therefore possible. A tradeoff exists though, as the qubits are no longer easily isolated from the surrounding environment.

The basic features of a superconducting qubit circuit are ultra-low dissipation (from superconductivity), ultra-low noise (from low temperature) and non-linear, non-dissipative elements (specifically in the form of the Josephson junction).

Superconductivity allows electron flow with almost no resistance. Type-I superconductors are primarily used in circuit construction (aluminium is a popular choice) due to their simplicity compared to their Type-II counterparts. Cooling the system to this level is advantageous from a noise argument as well: qubits must be cooled to temperatures where the typical energy $k_B T$ of thermal fluctuations is much less that the energy quantum $\hbar\omega_{01}$ associated with the transition between the $|0\rangle$ and $|1\rangle$ states. This frequency is nominally in the 5–20 GHz range, thus the operating temperature T must be around 20 mK or lower.

There are three basic types of superconducting qubits which form around the Josephson junction. Each design is an attempt to minimise the detrimental effect of flux or charge fluctuations introduced into the circuit via the surrounding environment.

The simplest idea to mitigate charge fluctuations is the Cooper pair box [23, 24], which biases the JJ with a voltage in series with a gate capacitor. Circuits derived from this method are commonly called charge qubits [25, 26] although it's important to note that the name does not imply information is encoded in the charge of the system, but instead refers to the controlling variable of the system. The transmon [27, 28] is a relatively new qubit design that stems from the original charge qubit ideas: shunting the cooper pair box with a large capacitance, whilst shrinking the Josephson junction size in an attempt to squelch decoherence sources.

The second circuit type (the RF-SQUID [29]) can be considered the dual of the Cooper pair box, employing a superconducting transformer in place of the aforementioned gate capacitor. Designs utilizing this method are referred to as Flux qubits [30–32], again this is because a controllable external flux field

is imposed on a superconducting loop across the JJ rather than the idea that information is stored in the magnetic moment of the qubit.

Another method to minimise the effect of charge fluctuations is to bias the JJ with a fixed DC current source. Implementation of this design is referred to as the Phase qubit [33, 34], a circuit diagram of which is depicted in figure 1.2.

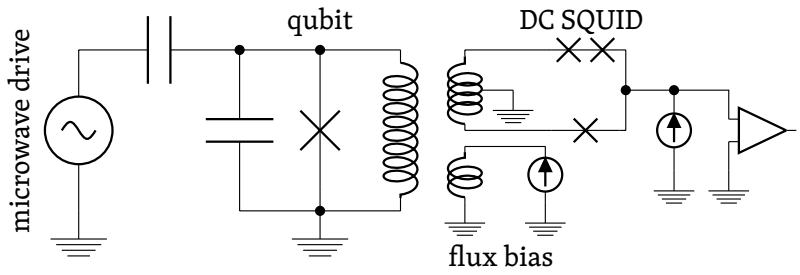

Figure 1.2— Circuit diagram of a phase qubit and readout components. This schematic is representational of current designs, based on improvements made in reference 35.

This design has a highly capable readout mechanism and multi-qubit scaling arrangements are promising. Thus the possibility of a functioning quantum computer using it is encouraging with only one major caveat: the system is mired with decoherence causing defects. In fact, much of the superconducting qubit community has transitioned from active phase qubit experimentation to the transmon because of this phenomena.

On the other hand, this design is useful for studying these decoherence sources (*vide infra* subsection 1.5.1) and is therefore the most relevant architecture to study throughout this thesis.

Architecture differences aside, superconducting qubits share many common positive capabilities—most of which have already been touched upon. Recent research has pushed these systems even further, with controllable reset capability [36], schemes which are approaching active quantum error correction [37] and the exciting prospect of 'autonomously stabilized entanglement', which uses a Bell state as the qubit's steady state and may counteract the detrimental effects of environmental noise [38]. An outlook on the future of superconducting circuits can be found in reference 39.

1.3 Josephson Junctions

The superconducting phase difference across a JJ is the typical degree of freedom of interest in superconducting qubit designs. This phase difference can be treated as a particle moving in a potential landscape defined by the rest of the circuit. A normal LC circuit generates a parabolic potential, which in energy terms is a harmonic oscillator: transitions between neighboring excited states are degenerate. If we wish to use this kind of circuit as a qubit, the transition between state $|0\rangle$ and $|1\rangle$ must be sufficiently different to the higher-lying eigenstates of the system, *i.e.* some degree of anharmonicity must be introduced.

Figure 1.2 depicts an LC circuit in parallel with a barrier component —×—. This symbol represents the Josephson junction, which at low temperature is a strongly non-linear, non-dissipative circuit element*. In fact, a JJ behaves as a pure non-linear inductor in parallel with a capacitor resulting from the parallel plate capacitance inherent in the junctions' construction (see figure 1.3 below). These features give rise to an anharmonic oscillator, a readable energy spectrum and the possibility of a working qubit.

*One reason why the JJ is so ubiquitous in superconducing devices is that for many years it was the *only* non-linear component. Recently the quantum phase slip junction [40–42] (effectively the dual of the JJ) has been showing great promise as well.

This is all possible because of the Josephson effect [43, 44], which is one of the most significant examples of macroscopic quantum phenomena presently known. Current flow between two superconductors separated by a thin insulating layer (*i.e.* a JJ) is possible even in the absence of electrical voltage drop across the junction, making this phenomena a *supercurrent*. Cooper pairs tunneling through the insulating barrier are the cause of this supercurrent, which manifests macroscopically, but depends on the quantum phase of the system.

Josephson junctions may be constructed from any superconducting material with any insulating or non-superconducting metal barrier to invoke a weak link coupling. A popular material choice involves the use of aluminium as the superconducting material, and an amorphous oxide layer as an insulating barrier—an illustration of which is depicted in figure 1.3.

For further discussion of the theory behind, and the engineering prospects of the Josephson junction, see the great review article by Makhlin *et al.* [25].

1.3.1 Formation Process

Shadow evaporation is a common technique used to fabricate systems such as this, where two metallic layers are deposited from different angles with an

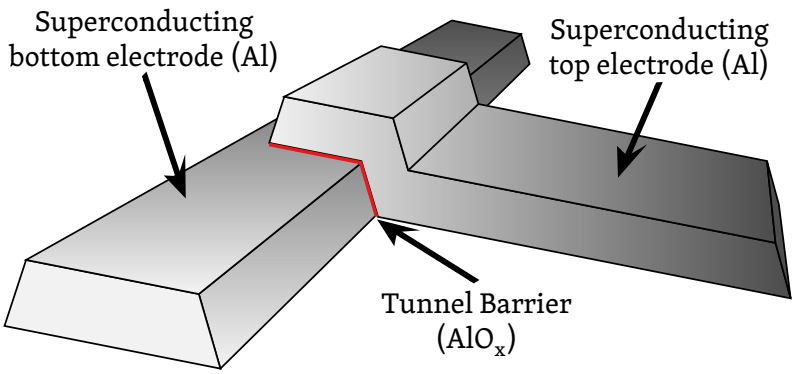

Figure 1.3— Schematic of a Josephson junction constructed with two superconducting thin films of aluminium and an amorphous AlO_x tunnel barrier.

intervening oxidation step. This is usually performed using a Dolan bridge, which obscures part of the substrate during each metal deposition step [45]. It has more recently been shown that junction fabrication can be performed without the requirement of this bridge [46]. Regardless of the process chosen, the oxidation of the aluminium does not result in a set of crystalline monolayers, but a non-uniform amorphous layer varying in stoichiometry [47, 48], density [49] and thickness [50–52] (nominally ~ 2 nm). Although epitaxial growth of aluminium-oxide barriers has been demonstrated [53], this technique is not yet mainstream as it is considerably more difficult than conventional shadow mask evaporation.

1.4 Investigating Decoherence Sources

We have so far considered that our architecture of choice: the superconducting phase qubit (subsection 1.2.4) trades off its relatively simple construction, control and readout capabilities with a qubit design which interacts with its environment in a destructive manner (subsection 1.1.1), causing it to decohere. Identifying the phenomena which cause environmental perturbations (or put simply on our context: electronic noise) is an investigation which spans an entire century of research effort.

The uniform noise assumption of the Lindblad form master equation (1.4) in reality is generally considered to be naïve, as the situation tends to be far more complicated. Whilst a complete decoherence theory discussion is not required for the intricacies of this thesis, it is important to understand a few

more pieces of the puzzle before moving on. Introduction of non-uniform noise, as well as the treatment of spatial correlations between individual noise sources can be discussed through the framework of the Bloch-Redfield master equation [54]. To treat spatial correlations in the noise correctly, one must treat the perturbations explicitly as part of a larger open quantum system. A straightforward, and pertinent example of this treatment can be observed in reference 55. In certain limits, the noise profiles described by Bloch-Redfield can be mapped to a Linbladian form [56, 57], which indicates that the relationship between these two theories is tightly bound.

Extensive experimentation on superconducting qubits has identified that at low frequencies, the spectral density of the noise exhibited in the environment varies approximately as $1/f$ over a sampled frequency range f [58]. In terms of the decoherence times this noise adds a contribution to the dephasing mechanism γ in T_2 (1.5). At higher frequencies, the spectrum is closer to an Ohmic ($\propto \omega$) response, contributing to qubit relaxation (T_1 decay), and it appears these two asymptotes cross at $\omega \sim T$ (where T here is temperature) [59].

These scaling laws are predominantly explained by a phenomenological ensemble of systems with two distinct, almost degenerate sites of equilibrium. A system can easily excite into its second mode, then fluctuate back to the first in a rapid and oscillatory behaviour. This phenomenon can be driven by a small energy transfer, in our case this occurs through the decoherence of a qubit. The literature denotes the name 'two-level system' (TLS), or less frequently: two-level fluctuator, to an individual noise contributor in a bath of similar systems, which constitute a model of environmental perturbations.

1.4.1 Defects in Glasses

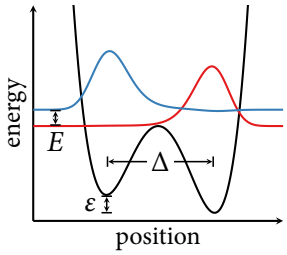

FIGURE 1.4— STM representation of a TLS, a quantum mechanical description by wave functions —— & —— in a double well potential ——. Excitation energies are calculated via $E = \sqrt{\Delta^2 + \varepsilon^2}$.

To investigate the possible origin of these TLS baths, we turn to the materials science literature, as bistable defects in glasses and amorphous solids in general have been understood for some time [60, 61]. Early research on this topic identified defects in a number of imperfect crystalline lattices, assuming that defects formed by individual atoms or small atom clusters tunneling between two almost equivalent lattice positions [61, 62]. This work also saw the genesis of the Standard Tunnelling Model (STM): the conventional phenomenological 1D model which explains the anomalous bulk properties of these amorphous glass systems at low temperature. The two-level system description in the STM requires two parameters: the tunneling energy Δ and asymmetry

energy ε. Both parameters depend on local atomic configuration and lattice strain forces, necessitating a redistribution of charge when parameters shift. TLSs therefore couple to their environment through electric and strain (elastic) dipole moments.

A representation of a quartz-like (SiO_2) tetragonal lattice is shown defect free on the left of figure 1.5. Under some alternative formation process, the crystallinity of the system may be compromised and produce a structure amorphous in nature: depicted to the right. We can use this representation in illustrate examples of the types of TLS defects which can occur in real amorphous solids.

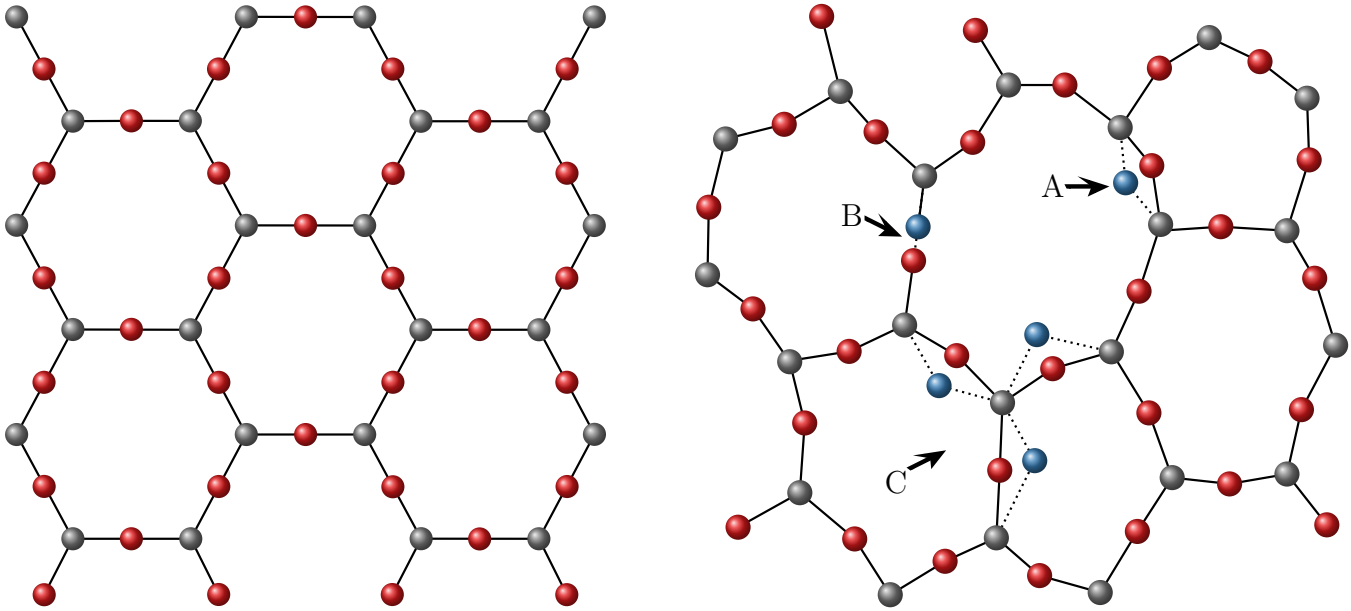

Figure 1.5— Left: Quartz-like crystal lattice, with oxygen ● and silicon (or some other metal such as aluminium) ● atoms aligned in a hexagonal grid. Right: Defected quartz-like lattice strained to an amorphous state. Three possible defect types are identified, where equivalent lattice positions the oxygen atom can occupy are identified as .

Three unique defect types are depicted, which have been used in the glass community extensively from as early as the 1930s [60].[†] A: the bond length is shortened causing an oxygen to form a dipole perpendicular to the bond axis, B: the bond length is lengthened causing an oxygen to form a dipole parallel to the bond axis, C: a cluster of three oxygens are rotated about a central metal atom. Ensembles of thousands of such defects govern the acoustic, thermal

[†]It is interesting to note that historically, the tunneling of atoms was considered as a phenomena before electron tunneling or α-particle emission was ever discussed.

and dielectric material properties of glassy systems at low temperature. Characterisation of the macroscopic response due to these TLS ensembles via a distribution of defect parameters [63] has seen the STM applied not only to large systems, but more recently to microfabricated circuits.

Early superconducting qubits using tunnel junctions such as the Josephson junction in the presence of strong electric fields were often stymied by even just a few two-level defects [35]. Single electron transistors [64], nanomechanical resonators [65], kinetic inductance single photon detectors and microwave resonators [66] also suffer from noise channels attributed to TLS behaviour; indicating that there may be ultimately many fundamentally different noise sources. Another possible interpretation is that there is only one microscopic origin, which must take into account the chemical composition of all device types.

1.5 Strongly Coupled Two Level Systems

As superconducting circuit designs improve, qubit operating times (before coherence is lost to the environment) increase and new modes of studying the environmental noise spectrum are realised. Developments in charge, flux and phase qubit architectures has enabled the study of so-called 'strongly coupled defects' [67–69], which couple more strongly to the qubit than other environmental processes. These defects have comparable resonance frequencies to the qubit circuit and coupling strengths as well as decoherence times long enough to allow coherent oscillations between the qubit and TLS.

1.5.1 Qubits as Probes of TLSs

Probing environmental noise with qubits has been a useful tool for the development of the $1/f$ and Ohmic spectrum models, as well as solidifying the conjecture that the underlying mechanism is the result of incoherent random TLS transitions [70, 71].

Pioneering work was undertaken in the early 2000's, with new qubit designs insensitive to charge noise (but still susceptible to flux fluctuations), isolating the qubit from dissipation of the bias and measurement leads. Rabi oscillations of the qubit were measured, showing coherent manipulation was possible with a fidelity of 85% [34]. Further experiments identified temperature dependencies of the $1/f$ noise spectrum at low frequencies in the critical

current of SQUIDs [72] and single-electron transistors [73]. Identifying noise sources within the JJ oxide barrier as a major decohernce channel improved superconducing circuit designs further, moving research toward the fabrication junctions with smaller area and better dielectrics [74].

More recently the investigation has focused on a new capability: resolving individual TLSs. Using a phase qubit, Lisenfeld *et al.* developed a coherent single-pulse resonant driving method in which a detuned qubit acts a detection and measurement device to characterise the coherence properties of an individual TLS by applying microwave pulse sequences to the system [69]. Figure 1.6 shows one result of this method: the identification of two individual TLSs coupled to a phase qubit.

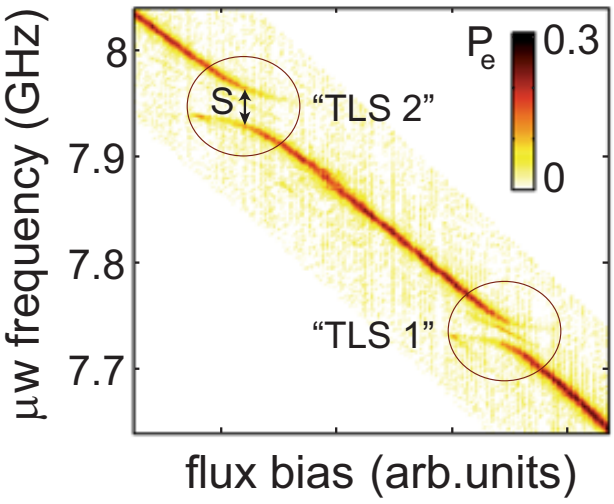

Figure 1.6— Two avoided level crossings indicate the existence of two individual two-level systems extant inside a phase qubit. Avoided level crossings indicate the TLSs are in resonance with the qubit at this frequency, and the probability of detecting an excited qubit state P_e falls dramatically due to the coherent qubit-TLS coupling. (Figure 1 of Lisenfeld *et al.* [69]).

Avoided level crossings indicate the qubit is in resonance with a TLS, with a coupling strength S measured as the magnitude of the crossing (recovered spectroscopically). Whilst this method was not the first to individually resolve a strongly coupled TLS, it was the first to measure TLS coherence times (T_1 and T_2) as a function of temperature. In particular the T_1 time is found to decrease quadratically as a function of temperature, which is a very curious result: the reason for which is still hotly debated.

Further experiments have probed these defects and shown them to be stable, controllable and have relatively long decoherence times [35, 67, 68, 75, 76]. The response of TLS coherence times under qubit deformation (via a localised strain mechanism) has been experimentally investigated [77], and coherent TLS-TLS interactions have also been observed [78].

1.5.2 Phenomenological Models

In an attempt to understand these defects, many phenomenological theories attempting to describe them exist.

The charge dipole [74], perhaps the most straightforward and well understood model, suggests the STM approach can be applied to strongly coupled TLSs, where one would consider a charge hopping from one equilibrium position to another (see figure 1.4) inside the amorphous layer of the JJ, or perhaps at the metal–oxide interface.

The measurable signal of a TLS in a phase qubit in this framework is the resonance of the TLS and qubit splitting energy, E_{01}, with the qubit-TLS coupling S_{max} (identical to S in figure 1.6). This coupling is described by the expression [74]

$$S_{max} = 2\frac{\wp}{w}\sqrt{\frac{e^2}{2C}E_{01}}, \tag{1.6}$$

where the width of the Josephson junction w and its internal capacitance C can be measured or inferred experimentally; e is the charge of an electron and $\wp$ is the dipole moment of the TLS. This expression was identified by Martinis *et al.* in their seminal work discussed above, in which they demonstrate the loss profiles from a range of microwave and qubit measurements can be explained as resonant absorption of two-level defects from within the JJ tunnel barrier [74].

Strongly coupled TLSs are observed in many different qubit designs. Their characteristic energy splitting E_{01} is observed over the full operating range of 1–10 GHz for transmons [79], 4–5 GHz for flux qubits [68] (although recent designs can tune this gap down to the MHz range [80]) and nominally 7–8 GHz for phase qubits [81]. Dipole moment strengths also vary, but are usually on the order of 1 eÅ [75, 81].

Other popular models describe a number of alternative mechanisms which may be contributing to environmental noise.

Direct experimental evidence for surface magnetism in superconducting circuits was observed only recently in 2008, when unpaired surface spins in thin-film SQUIDs were identified [82]. One can consider that these surface spins couple to the magnetic flux of the system as a paramagnetic dipole rather than a charge dipole to the electric field.

De Sousa *et al.* identify individual trapping centers in the Josepson junction's oxide barrier hybridised with electrons in the attached superconducting leads [83]. These centers cause sharp subgap resonances, which may form pairs of Andreev bound states (entangled electron–hole states).

The Kondo effect is the formation of a many-body singlet with conduction electron spins screening out a local trap spin [84], which occurs due to a local trap interacting with a Fermi sea. Electron trapping via this mechanism in shallow subgap states forming at the metal–metal-oxide boundary explains TLS temperature dependencies which other phenomenological models struggle with [85].

A variation in critical current coupled to a qubit results in decoherence. Ku and Yu suggest two noise models describing TLS state dependence of the JJ transparency: a qubit in resonance with a high-frequency TLS which decays via phonon emission and a $1/f$ noise modulation of the critical current [86].

Detailed fitting of experimental data can place limits on these models [81] but the scope of free parameters of each model allows them all to fit experimental results—rendering them presently indistinguishable.

1.5.3 Microscopic Models

Probing individual defects has promoted their bistable nature from hypothesis to observable fact, as well as providing clues to their microscopic origin. The indistinguishability of the phenomenological models makes it very important to construct microscopic models of these systems to increase our understanding of their composition and put further bounds on the TLS phenomenon.

Recent years have seen a number of possible microscopic extensions to phenomenological theories, for example the origin of the local magnetic moments as been proposed as multiple metal-induced gap states arising from localised disorder at the metal–oxide interface [87] or even surface aluminium

ions paramagnetically coupling to ambient molecules [88]. Others suggest polaron dressed electrons [89] or the accidental inclusion of an alien species (primarily hydrogen) in the junction construction process [90–92].

Future fabrication techniques or more robust qubit designs may suppress or diminish the response of such noise sources as has been the case historically [26, 28, 74, 79, 93], although there are still many open questions to be considered. As stated above, it is suggested that a TLS bath is responsible for the weakly coupled, ohmic $1/f$ noise [58]. However, it is unclear if the identified strongly coupled TLSs are from the same origin. Ultimately, several separate microscopic suspects may be identified; although work in this area is not mature enough to speculate further.

1.6 Historical Material Defect Research

Both phenomenological and microscopic models suggest differing TLS origins within the junction or leads, which is still a matter of contention. It is therefore important to change tack for a moment and focus on material properties, considering the identification of a noise source within a system constructed by a material as well studied as aluminium oxide may have already occurred. Condensed matter defects and their properties have been an active area of research for many decades, we should therefore look at this research body in light of the recent findings and advances of the strongly coupled TLS experiments.

1.6.1 Defects in Corundum

Various defects in corundum (the low temperature and pressure phase of aluminium oxide, α–Al_2O_3) have been thoroughly investigated, both experimentally and theoretically for some time. The first *ab initio* corundum papers discussed the partially covalent nature of the aluminium—oxygen bond and the inherent computational complexity of studying this structure [94]. Experimental identification of defects like the F and F^+ electron centres [95] and observations of defect mobility [96] motivated further theoretical study.

Intrinsic defects (such as self-trapped and defect-trapped holes, or oxygen and aluminium vacancies) as well as impurities of transition-metal ions like Co, Fe, Mg, Mn or Ti are all possible [97]. Jacobs and Kotomin also suggest that oxygen vacancies are mobile throughout the lattice, as effective charges

of an oxygen at a saddle point between aluminium atoms is similar to one on a normal lattice site.

Density Functional Theory (DFT) studies of single oxygen vacancies soon followed, using 120 atom supercells of corundum [98] which show the oxygen vacancy introduces a deep and doubly occupied (electronic) defect level (an F center), and was found to be "not so localized". Singly occupied defect levels (the F^+ center) are also discussed by introducing a positive background charge into the calculation.

Further investigations into point defects began to classify systems such as the Schottky-type (two Al^{3+} and three O^{2-} vacancies) and Frenkel-type (*e.g.* one interstitial Al^{3+} and one Al^{3+} vacancy) defects, where the Schottky-type was found to be the dominant class [99]. However, further theoretical studies found oxygen Frenkel defects to have the lower formation energy [100], which precipitated several discussions concerning the appropriateness of empirical potentials fitted to bulk properties in describing defects. DFT studies aligned this discussion with experimental data, presenting formation energies of the classes to be Schottky<Al Frenkel<O Frenkel [101]. In addition, Matsunaga *et al.* calculate formation energies and band structures for aluminium and oxygen, vacancy and interstitial defects; as well as an in-depth discussion on lattice deformation about each defect. Formation energies are dependent on the chemical potential of the oxygen in the local environment, although for a wide range of this value it was found that energies were ranked as $V_{\mathrm{Al}}^{3-} < \mathrm{O}_i^{2-} < V_{\mathrm{O}}^{2+} < \mathrm{Al}_i^{3+}$ (*i.e.* aluminium vacancy, oxygen interstitial, oxygen vacancy, aluminium interstitial). All of which are stable in their ionised states, and additional electronic defects compensate defect charges only at high temperatures [101].

Implanting dopants like Ti^{4+} and Mg^{2+} into corundum to create certain defect classes can be used to investigate this high temperature oxygen diffusion throughout the lattice. If bivalent dopant concentrations dominate, the predominant transport method is through oxygen vacancies. On the other hand, high tetravalent impurity concentration produces oxygen interstitials as the prevalent mechanism for transport [102].

Surface point defects (interstitial and vacancies) have also been compared to their bulk counterparts [103], showing that relative formation costs are similar (*i.e.* the cost of an aluminium versus oxygen vacancy), although the surface defects require much less energy to form and electronic delocalisation is

larger on the surface than in the bulk. This suggests surface defects may play an important role as nucleation centers when bonding to metals (such as aluminium in our case).

1.6.2 Deficiency Defects in Amorphous Aluminium

It has been suggested that an alternative growth process to thermally oxidising aluminum via oxygen diffusion (*i.e.* methods discussed in subsection 1.3.1), such as Atomic Layer Deposition (ALD), may remove defects such as oxygen vacancies from the Josephson junctions' tunnel barrier. ALD works by exposing a substrate to a heat source, then alternating pulses of water and trimethylaluminium. Each pulse is separated by a nitrogen gas flush to assure no gas phase mixing between the pulses. Ligand exchange between the water and trimethylaluminium at the substrate generates growth of the oxide [104]. Recent experiments show promising results in terms of using this process to build Josephson junctions—overcoming the lack of hydroxyl groups on metal surface which assist nucleation [105]. However; it is possible that this process generates oxygen deficient lattices with parameters comparable to an oxygen vacancy in corundum [106].

Although ALD is not presently used in Josephson junction fabrication, the process is actively being studied as a fabrication method to create resistive random access memory. It is well known that oxygen diffusion methods can create low stoichiometries (*i.e.* $x < 1.5$) of amorphous AlO_x [47, 48], whereas this is not the goal for ALD. Hence, research into single oxygen vacancies in ALD grown $AlO_{1.5}$ provide a good intermediary between the discussion of crystalline defects above, and the topic of this thesis (identifying two-level defect mechanisms in amorphous aluminium oxides).

Momida *et al.* investigated single oxygen vacancies denoted as V_O^0: a fully-occupied state, V_O^{1+}: half-filled state, and V_O^{2+}: an empty state by introducing them into a DFT based model of ALD grown $AlO_{1.5}$ [107] (V_O^0 and V_O^{1+} are also denoted as the F and F^+ centers elsewhere in the literature). The defect was placed at a random site in the amorphous lattice and the density of states was calculated. This process was then repeated a number of times to generate an appropriate level of statistics. Relative to the bulk valence band edge, these states appear in the band gap between 0.51–2.51, 1.30–3.58, and 2.39–3.64 eV respectively. Figure 2 in reference 107 is of particular interest: comparing the difference between the stable and localised V_O^0 wavefunction and the V_O^{2+}

wavefunction—which is delocalised throughout the lattice. Furthermore, lattice configurations for each of the defects are also presented, which may be useful for defect classification.

1.7 Chapter Summary and Thesis Outline

Decoherence sources such as the environmental two-level system (including $1/f$, Ohmic and strongly coupled sources) are currently a major hurdle to the realisation of operational superconducting qubits and other Josephson junction based quantum devices. Many properties of the TLS have been probed experimentally, demonstrating relatively long decoherence times whilst showing them to be stable and controllable. However, identifying their true microscopic nature remains an open problem. Controllable superconducting qubit architectures make it possible to study single strongly coupled TLSs, as opposed to weakly coupled ensembles of TLSs that may be responsible for $1/f$ noise. Although the Standard Tunneling Model can be used to explain all existing experimental results, to improve devices we need a detailed microscopic understanding of these defects.

Several microscopic models have now been proposed: surface state interactions, polaron dressed electrons and alien species; all of which are possible decoherence sources and/or explain the origin of strongly coupled TLSs. As stated in subsection 1.5.3: better fabrication methods, higher vacuums and more robust qubit designs have historically suppressed or diminished the response of such noise sources. However, continuing down the path of a perfectly clean but amorphous dielectric may no longer be the optimal direction. Whilst crystalline layers are difficult to construct in this environment, they are possible and show a dramatic decrease (of up to 80%) in visible TLSs over their amorphous counterparts [53].

Alien species such as hydrogen are still introduced in the epitaxial process, which may indicate the origin of some TLS sources, however a decrease of 80% must be explained through the comparison of an epitaxial growth and the non-crystalline nature of the dielectric layer built with standard junction fabrication processes.

1.7.1 Conjecture Concerning TLS Origins

The premise of this thesis is that positional anharmonicity of oxygen atoms arises within the AlO_x barrier of the Josephson junction due to its fundamentally non-crystalline nature. This ansatz allows the existence of many different spatial configurations throughout the layer, causing unique TLS properties based solely on atomic positions and rotation in relation to the external electric field. Additionally, this hypothesis satisfies the discussion above concerning a solution accounting only for the differences in oxide growth processes and possibly explains TLS existence in the devices discussed in subsection 1.4.1, as each device requires a large oxide area in their construction. As such, this is a concrete microscopic ensemble of the classic atomic tunneling model: typically cited as motivation for the STM (which is more general, as the tunneling object need not be a single atom).

As an illustrative example, consider an interstitial oxygen defect in crystalline germanium: the harmonic approximation for atomic positions cannot be applied due to the rotational symmetry of the defect as oxygen delocalises around the Ge–Ge bond axis [108]. A similar effect occurs with oxygen interstitial defects in bulk silicon, forming an anharmonic potential leading to a system with a degenerate ground state, even in a 'perfect' crystal lattice [109]. Figure 1.7 depicts these two scenarios, which produce fundamentally different low-energy-excitation and infrared spectra—showing their characteristic features at different energy scales.

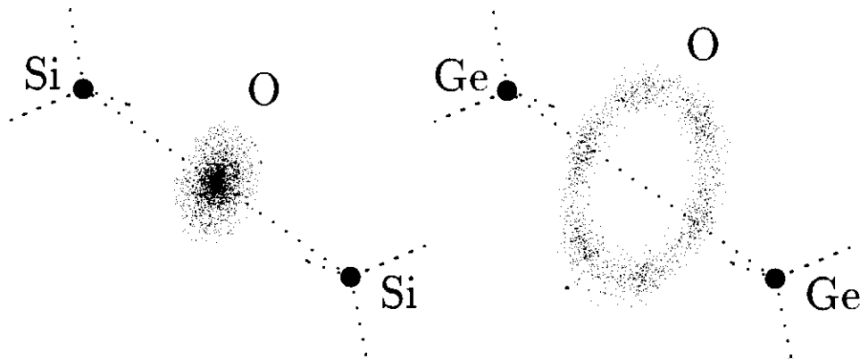

FIGURE 1.7— Quantum delocalisation of oxygen interstitial defects in crystalline silicon (left) and germanium (right), shown perpendicular to the Si–Si and Ge–Ge axis. (Figure 2 of Artacho *et al.* [108]).

If similar potential landscape scenarios exist within amorphous aluminium oxides, this phenomena may additionally explain the strongly coupled

TLS defects residing within the JJ tunnel barrier, considering experiments also observe a range of qubit-TLS coupling strengths (and therefore different TLS energy scales).

1.7.2 Thesis Framework

Strongly coupled TLSs cannot be considered as a simple Lindbladian decoherence channel (1.4), because they are coherent and coupled more strongly to the qubit than other processes. Time scales associated with qubit-TLS coupling (1.6) are shorter than all other decoherence time scales. Therefore we have to consider the TLS as its own quantum system *in isolation*. This thesis asks questions about the properties of the TLS derived from the material properties of the junction it resides in, which has important ramifications for materials science based efforts to reduce the effects of environmental noise.

Chapter 2 presents a methodology for constructing atomic scale computational models of Josephson junctions using a combination of molecular mechanics, empirical and *ab initio* methods. The stability and structure of these barriers as a function of density and stoichiometry are investigated, which are compared to experimentally observed parameters. These realistic values of atomic positions in an amorphous barrier can then be used as input conditions to test any TLS model generated in the proceeding chapters.

A simple single particle Hamiltonian is constructed in chapter 3 to describe the A and B type defects of figure 1.5 with the interpretation that an oxygen atom is indeed the tunnelling source in a defected (amorphous) lattice of aluminium oxide. A model is developed to describe the delocalised oxygen state as it interacts with the localised potential of surrounding aluminium and oxygen atoms in chapter 4, using 1D and 2D constructions to explain some of the fundamental physics of the oxygenic interactions.

This model is then rigourously compared to experimental results in chapter 5, applying at a 2+1D variant to explain presently identified interaction phenomena. A complete 3D model using a different theoretical approach explores a smaller parameter space in chapter 6, additionally explaining the need to use this computationally intensive approach in some instances, but how in general the 1D STM model still adequately describes the relevant physics.

Finally, chapter 7 summarises the model and its relevant results, discusses any and all conclusions that can be deduced from the oxygen delocalisation premise, and outlines the outlook for this model and the ramifications it has on strongly coupled TLS research and the wider quantum information community in general.

“Но природа не справляется с логикой, с нашей человеческою логикой: у ней есть своя, которую мы не понимаем и не признаем до тех пор, пока она нас, как колесом, не переедет.”

—Иван Сергеевич Тургенев

2 Junction Models

Methodology used to construct high precision Josephson junction models of varying stoichiometry and density.

THE JURY IS STILL OUT when it comes to the specific location of strongly coupled TLSs inside Josephson junctions, and as we've seen in subsection 1.5.3, many perchance competing microscopic models exist that attempt to describe their origin. One possible way of distinguishing between these options (including the one posited in subsection 1.7.1) is to develop complete atomistic models of the Josephson junction and study the many spatial configurations of the amorphous tunneling barrier.

From a quantum simulation point of view, this is not a trivial problem. Crystalline symmetries usually invoked to minimise calculations cannot be used to reduce the state space of the oxide barrier of a junction due to its amorphous nature. However, forming such atomistic models using molecular mechanics and *ab initio* methods is the focus of this chapter.

Initially, we discuss the computational complexities of simulating oxide formation directly in sections 2.1 and 2.2, followed by the specific methods that have been implemented in section 2.3. Sections 2.4 to 2.6 examine a number of techniques which assist in validating the resultant models in terms of expected configurational properties and experimental data. Finally a summary in section 2.7 which outlines some technical uses for these models outside the main scope of this thesis.

2.1 Simulating the Junction Formation Process

Josephson junctions may be constructed from any superconducting material with any insulating or non-superconducting metal barrier to invoke a weak link coupling. A popular material choice involves the use of aluminium as the superconducting material, and an amorphous oxide layer as an insulating barrier. Subsection 1.3.1 discusses the shadow evaporation techniques usually undertaken to grow a ~ 2 nm amorphous layer which may vary in stoichiometry, density and thickness.

Simulating oxide layer growth is in general a difficult problem as the time scale of the oxide growth ($\sim$ minutes) is orders of magnitude greater than typ-

ically achievable molecular dynamics timescales (ps–ns). One standard approach is to perform the simulation at elevated temperatures and gas pressures (≥ 1 atm) [110–112]. This accelerates the oxidation process, making the computation feasible on current high performance computing infrastructure. However, it also removes the simulation from the reality of experimental junction formation, where pressures range between 10^{-9} and 10^{-3} atm [113–115]. It remains to be seen whether any fundamental physics is neglected by adopting this approximation.

An alternative approach is to form an amorphous layer via direct melt and quench [116, 117]. This method has the advantage of computational simplicity and speed, however the resulting layers are not necessarily representative of the true physical situation and therefore benchmarking against other methods and experiment is critical. Generating stoichiometry or density gradients across an artificial junction is not something that can be simulated directly using this process, so to investigate the effect of these properties, a number of constant density and stoichiometry models were produced. A more sophisticated method, closely mimicking the oxygen deposition process and examining the effects of layer thickness is the subject of M. Cyster's PhD dissertation.

2.2 Computational Frameworks

To obtain realistic, high precision atomic positions, computational models of the junction were created using a combination of molecular mechanics and DFT.

In the molecular mechanics framework, a group of molecules can be considered as *classical* collections of balls and springs rather than *quantum* collections of electrons and nuclei. It has been shown that under homogeneous compression, solids can be described by an expression of pair–potentials and contributions from non–pair interactions can be expanded in terms of the pair interactions [118]. This effectively includes any non–pair interactions into the pair terms and therefore the kinetic and potential states of the solid can be described on a unified basis. These interactions are mathematically constructed from well known classical mechanics formulae, and are known as *force fields*. The Buckingham potential [119] for example describes the Pauli repulsion energy and van der Waals energy interactions of two unbonded atoms as a function of their distance separation. Three parameters need to be empirically fitted for each atomic species pair the potential wishes to describe.

Density functional theory is a ground state theory which describes an interacting system of fermions via its density rather than its many-body wavefunction

$$\rho(\mathbf{r}_1) = N \int \Psi^*(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N)\Psi(\mathbf{r}_1, \mathbf{r}_2, \ldots, \mathbf{r}_N)\, \mathrm{d}\mathbf{r}_2 \ldots \mathrm{d}\mathbf{r}_N. \tag{2.1}$$

In other words, the ground state properties of a system can be expressed as functionals of the ground state electron density. Practical applications of DFT are based on approximations of a local, Hermitian, energy independent exchange–correlation potential [120], which describes the effects of the Pauli principle and the Coulomb potential beyond a pure electrostatic interaction of the electrons. Possessing the exact exchange–correlation potential for a system requires the many-body problem to be solved exactly, hence the need for an approximation.

The most common approximation is the local density approximation (LDA), which locally substitutes the exchange–correlation energy density of an inhomogeneous system by that of an electron gas evaluated at the local density. While many ground state properties (lattice constants, bulk moduli, *etc.*) are well described in the LDA, the dielectric constant is overestimated by 10–40% compared to experiment. This overestimation stems from the neglect of a polarisation dependent exchange–correlation field. This, and other limitations have forced the invention of other exchange–correlation functionals such as the generalised gradient approximation (GGA). Where LDA fails when the density undergoes rapid changes (such as in molecules), GGA considers the gradient of this density; yielding a vast improvement.

2.3 Model Construction

A 4×4×5 supercell of bulk aluminium (measuring $16.168 \times 16.168 \times 20.183$ Å) representing both the top and bottom slabs was relaxed in the DFT code VASP [121–123] using a projector-augmented wave (PAW) potential [124, 125]. Exchange-correlation interactions were evaluated using the PBE functional [126] (which uses the GGA); a $7 \times 7 \times 7$ Γ centered Monkhorst Pack K point mesh and a plane wave cutoff of 250 eV.

Formation of the amorphous AlO_x layers required a number of preparation steps to accurately represent experimental results. Corundum was used as a

basis for all the constructed junction models as it represents the low temperature and pressure phase of aluminium oxide. Experimental investigations of stoichiometry suggest, in general, an oxygen deficiency with oxide O/Al ratios varying between 0.6 and 1.4 [48], which are highly dependent on the fabrication process. In response to this, we construct models with four stoichiometries: $AlO_{0.8}$, $AlO_{1.0}$, $AlO_{1.25}$ and $AlO_{1.5}$. The oxide density may also be an important formation variable. For simplicity we identify oxide density as a fractional value of the (average) corundum density: 4.05 g/cm^3, and construct junctions with 0.5, 0.625, 0.75, 0.875 & 1.0 density fractions for each stoichiometry listed above. A value of 3.2 g/cm^3 is typical [49] (which corresponds to a density fraction of 0.8).

Using $AlO_{1.25}$ with a density fraction of 0.75 as an example, a $6{\times}6{\times}1$ supercell of corundum was geometry optimised in the software package GULP [127], employing the empirical Streitz-Mintmire potential [128] which can capture the variable oxygen charge states when present in a predominantly metallic environment. This capability is particularly important here, as a Josephson junction has two metal-oxide interfaces. This large superstructure was required due to the trigonal nature of the lattice, as it was then cut down such that the xy plane of the bulk aluminium slab could be covered. A non-periodic slab of corundum measuring $16.168 \times 16.168 \times 11.982$ Å was the result of this process. Oxygen atoms were randomly removed from the corundum lattice until the appropriate stoichiometry of $AlO_{1.25}$ was obtained and the cell was shortened in the z-direction to achieve a 0.75 fractional multiple of the corundum density. These changes add a lot of force onto the structure, so a geometry optimisation (in GULP) was undertaken at this stage to minimise energy contributions. To simulate the oxygen deposition phase and generate the amorphous nature of these layers, the structure was then annealed using NVT molecular dynamics at 3000 K for 3 µs and quenched to 350 K over a 1.5 µs period.

The $AlO_{1.25}$ layer was inserted between two bulk Al supercells described above with 0.5 Å of vacuum space on each side. The junction was further annealed to simulate a metal–metal–oxide interface reconstruction using VASP NVT Molecular Dynamics at 300 K until equilibrium was reached (approximately 250 ionic steps), then geometry optimised using a $2 \times 2 \times 1$ Γ centered Monkhorst Pack K point mesh and a 450 eV plane wave cutoff to obtain the final model, depicted in figure 2.1.

For comparison, junctions were also modelled without the added computational overhead of DFT by solely employing GULP and the Streitz-Mintmire

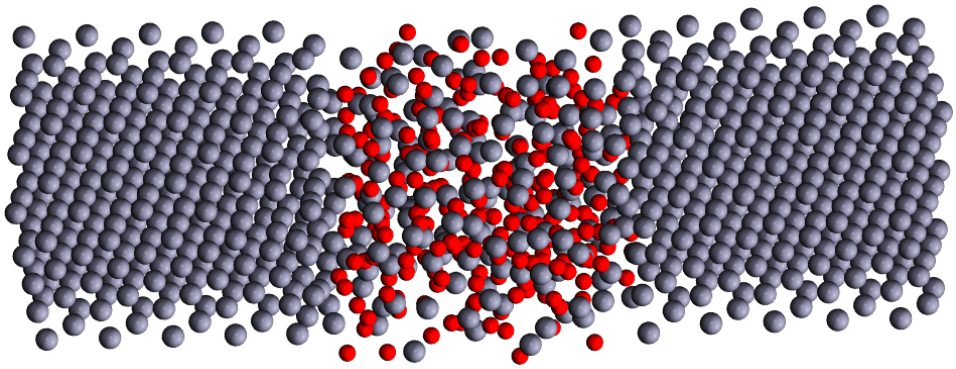

Figure 2.1— Model of a Josephson junction comprised of aluminium ● and oxygen ●. Two superconducting regions composed only of aluminium, separated by an amorphous $AlO_{1.25}$ barrier with a density 0.75 times that of corundum.

potential. The construction process of these models matches the procedure above, but interchanges the *ab initio* optimisations of the oxide layer with an empirical framework.

2.4 Radial Distribution Function $G(r)$

To validate our models against experimental observations, we perform a number of statistical tests to scrutinize the structures. First we must ensure that the oxide layer of the junctions are in fact amorphous in nature. We employ a projected radial distribution function

$$G(r) = \lim_{dr \to 0} \frac{p(r)}{4\pi \left(N_{\text{pairs}}/V\right) r^2 dr} \tag{2.2}$$

where r is the distance between a pair of particles, $p(r)$ is the average number of atom pairs found at a distance between r and $r+dr$, V is the total volume of the system, and N_{pairs} is the number of unique pairs of atoms [129]. This function was calculated for each stoichiometry and density configuration using oxygen as the reference species, and aluminium atoms in the amorphous region along with the superconducting bulk as the projection species. Figure 2.2 depicts the results of this analysis.

A major peak is visible centered around 1.85 Å, which corresponds to contributions from the two Al–O bond distances, 1.852 Å and 1.971 Å of the corundum crystal [130]. For a crystalline $G(r)$ this peak is resolved as to two delta

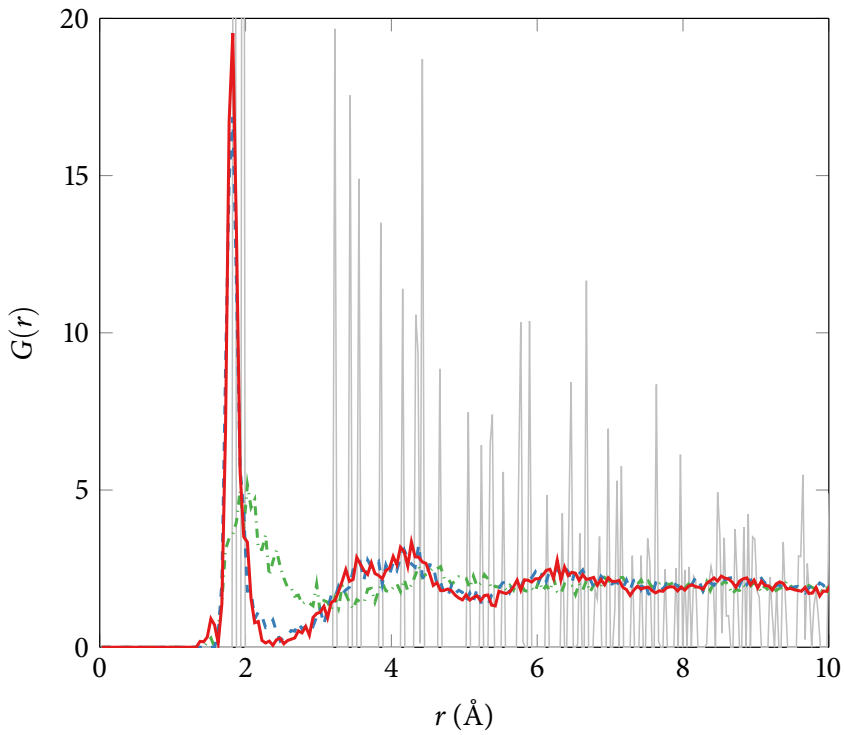

FIGURE 2.2— Evolution of the oxygen projected radial distribution function $G(r)$. Crystalline corundum ——; Metal–oxide interface before - - - and after - - - reconstruction; final optimised geometry ——.

functions (see figure 2.2 and the discussion below). For the amorphous layers here, we see a broadening of the statistics and hence differences in neighbour distances: diverging from a crystalline form. Moving away from this peak to larger distance separations, we see the statistics tending toward a uniform result similar to what a liquid would produce under this analysis. These two features represent an amorphous system quite well, as close range order suggests a connection to the crystalline form whilst long range order no longer agrees with such periodic conditions. It's also significant to note that we don't observe neighbours closer than $\sim$1.5 Å which is a good indication that the models do not have non-physical neighbour forces acting on atoms.

Most importantly, this trend is almost uniform across all the modeled junctions, which indicate the process outlined in section 2.3 is capable of producing amorphous oxides whilst varying other physical parameters of the system. An evolution of the important steps in the procedure is depicted in figure 2.2.

The corundum $G(r)$ —— is a complicated structure due to the 30 atom unit cell of the crystal, however it is clear from this figure where much of the amor-

phous structure originates from. Specifically the 1.852 Å and 1.971 Å Al–O bond distance contributions and the void in the 2–3 Å range. After the melt/quench phase of the procedure ··· the lattice still appears liquid-like. Whilst the quench cycle minimises the possibility of atoms positioning themselves too close to one another due to an excess of kinetic energy, it still appears to exhibit liquid behaviour. This may be a shortcoming of the Streitz-Mintmire potentials ability to capture the relevant physics, however this is rectified after the metal–oxide interface reconstruction is completed - - - using the *ab initio* methods. Finally, the geometry optimisation — yields a smoother 1.85 Å peak and recovers some of the void region around 2 Å.

Whilst optimal $G(r)$ results for both the VASP and GULP simulations are similar (figure 2.3), the GULP simulation actually produces a drastically different final structure. We find under GULP simulation that stoichiometric ratios higher than 1:1 are not stable and oxygen atoms diffuse into the metallic regions until a stoichiometric ratio of at most 1:1 is achieved. As a result of this excess oxygen diffusion, the junction width can increase by up to 30% or more over the course of the simulation. At high densities and stoichiometries (than typical amorphous alumina) some expansion of the oxide region is also seen in the *ab initio* simulations, although this effect is much less pronounced. Higher oxygen mobility in GULP could be attributed to shortcomings of the empirical potential. However, we see little increase in oxide distribution during the optimisation phase—suggesting that the details of the Nosé-Hoover thermostat routine employed during the MD simulation may play a role.

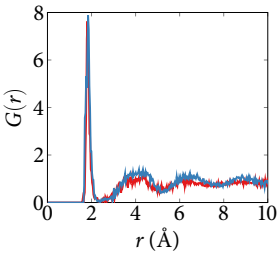

Figure 2.3— Oxygen projected $G(r)$ computed using *ab initio* (VASP) — and empirical (GULP) — methods, showing little statistically significant difference.

2.5 Total Energy and Optimal Conditions

The total energy of a computational model is a good indication of the structures' electronic stability. Due to the stoichiometry changes invoked in the oxygen depleted models, not all structures have the same number of atoms. This gives structures with more atoms (such as $AlO_{1.25}$) additional electronegativity which in turn results in a deeper potential well and a large total energy.

In order to be able to validly compare systems of different stoichiometry, we normalise the total energy of each system by a factor $|F|$. $F = \sum_k \mu_k N_k$ is calculated as the linear combination of the number of atoms of chemical species k (N_k) by the chemical potential of that species (μ_k), where $k = \{\mathrm{Al}, \mathrm{O}\}$. The chemical potential for aluminium, μ_{Al} was obtained by calculating the DFT

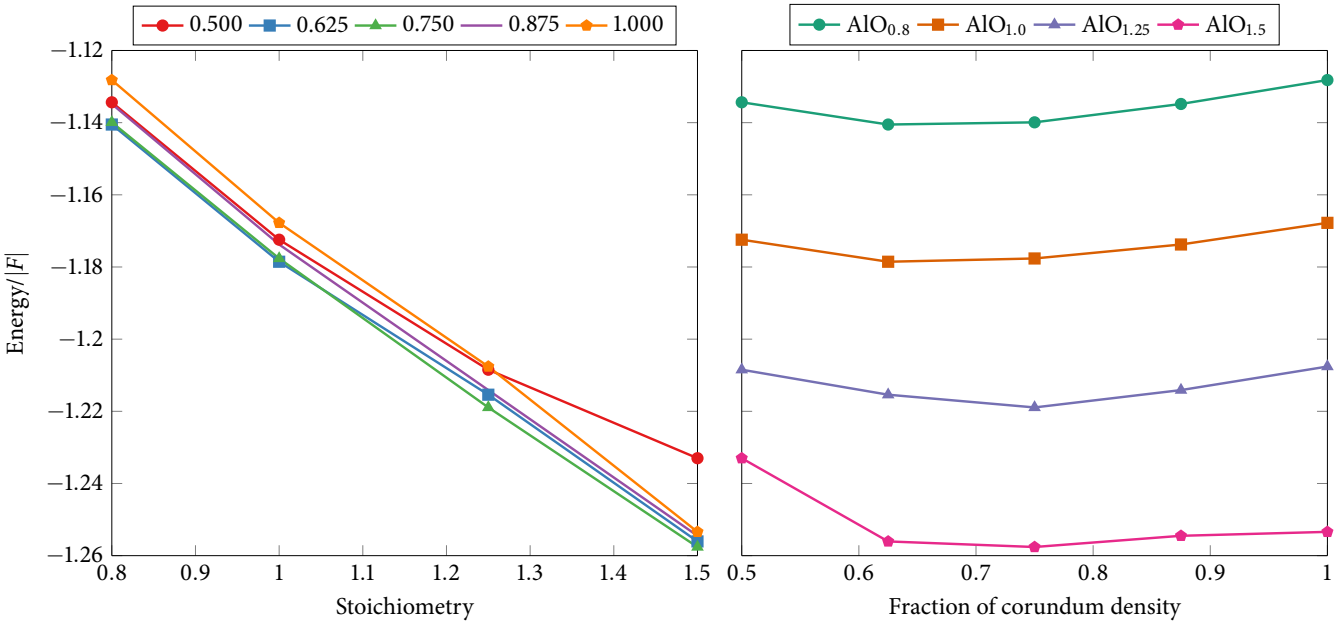

FIGURE 2.4— Normalised total energy for various junction models with stoichiometry (left) and density (right). Although the energy is strongly dependent on stoichiometry, we see a trend to optimal densities of approximately 75% of the density of corundum.

total energy of a 4×4×5 supercell of bulk Al and dividing by the number of atoms in the supercell. Similarly the chemical potential of oxygen was obtained from calculating the DFT total energy of a $2 \times 2 \times 2$ supercell of bulk Al_2O_3 using $\mu_O = (\mu_{Al_2O_3} - 2\mu_{Al})/3$, where $\mu_{Al_2O_3}$ is the total energy of a molecular unit of Al_2O_3. The factor F (essentially the free energy at $T = 0$) effectively allows one to partition the total energy of the system using the chemical potentials of each component species, as a means to compare the energies of systems with differing number of chemical components.

It is clear from figure 2.4 that stoichiometry plays a larger role in energy minimisation than density, and that the structures would prefer additional oxygen to minimise internal forces. This suggests that fabrication processes that generate oxygen deficiencies may be inviting the inclusion of alien species or oxygenic site hopping in an attempt to rectify this offset.

Density changes seem to alter the energy contribution marginally. Minimum energies correspond to density fractions between 0.6 and 0.75—slightly lower than typical constructions of 3.2 g/cm³ [49] (an 0.8 density fraction); which may indicate another method of experimentally optimising the junction formation process.

2.6 Coordination Number

Coordination number is a useful metric which allows for some insight into both the crystallinity of the structures being analysed, and their similarity to fabricated junctions. For instance, in the corundum structure every aluminium ion is coordinated with six oxygen ions. In amorphous alumina, the proportion of 6-coordinated aluminium as compared to 4-coordinated aluminium is an experimentally accessible quantity and has been reported on previously [131]. However, in order to establish this ratio it is assumed that there is a bimodal distribution of hexahedral (AlO_6) and tetrahedral (AlO_4) coordination. Ratios of AlO_6:AlO_4 are quoted in a range from 80:20 to 30:70, depending on the method by which the oxide layer was formed [132]. More modern techniques using Nuclear Magnetic Resonance are also able to resolve any AlO_5 coordinations [133].

Figure 2.5 shows the distribution of oxygen coordination about aluminium as a function of density and stoichiometry. These results are calculated using an Al–O bond length cutoff of 2.5 Å, which corresponds to the first minimum after the nearest neighbour peak in the $G(r)$ (see figure 2.2). As one would expect, the coordination number (for Al–O bonding) increases with increasing density or stoichiometry. We also note that there exists a reasonable proportion of 2- and 3-coordinated aluminium atoms, which persists at high density and stoichiometry. In order to compare directly to previous experimental and theoretical work, we compute the ratio of 4-, 5- and 6-coordination for Al–O bonding, matching the stoichiometry of 1.5 and assuming the density fraction closest to experimental values (0.750). The results are presented in table 2.1*. We observe excellent agreement, both before and after the *ab initio* optimisation.

*Coordination results were complied by Martin Cyster.

Table 2.1— Relative proportions of 4-, 5- and 6-coordinated aluminium atoms within the oxide layer for a density of 0.75 and stoichiometry of 1.5.

	4 (%)	5 (%)	6 (%)
VASP (before optimisation)	57	39	4
VASP (after optimisation)	53	43	4
Lee *et al.* [133]; experiment	55 ± 3	42 ± 3	3 ± 2
Momida *et al.* [107]; theory	60.4	29.2	10.4

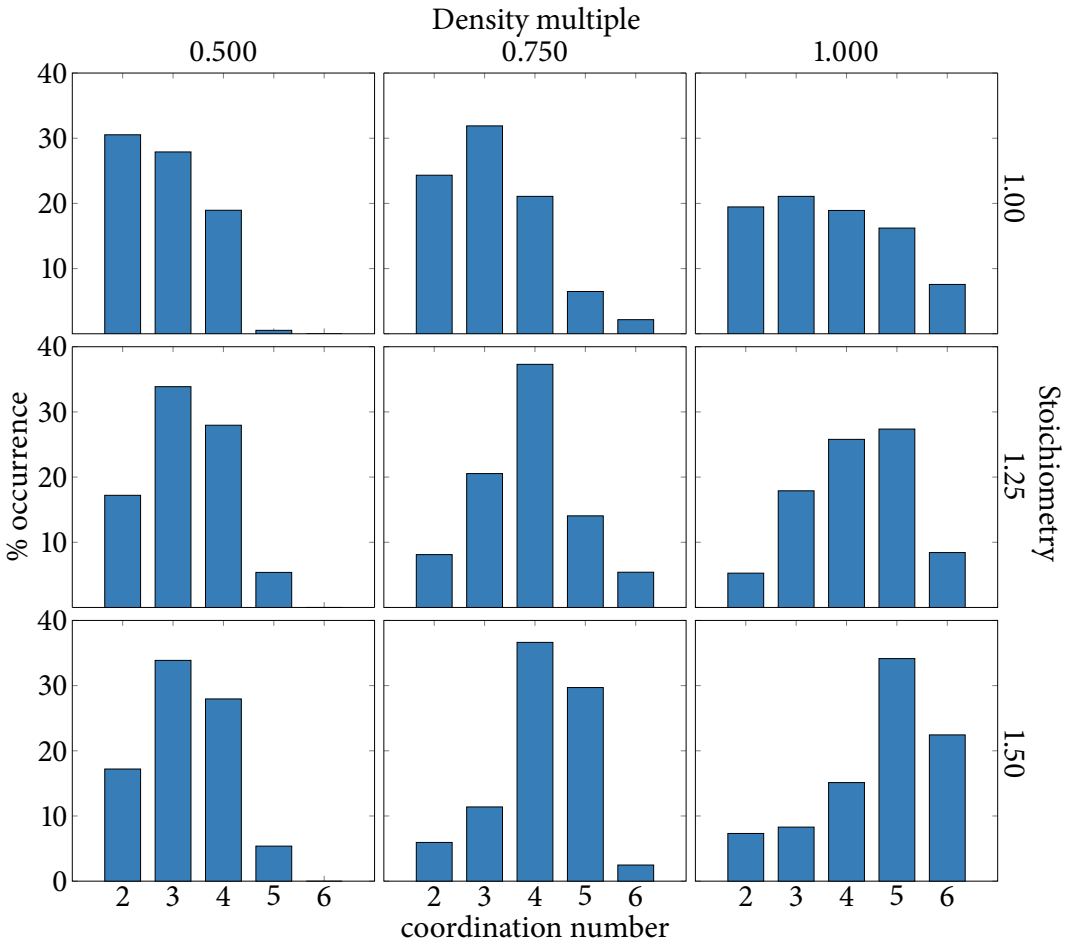

FIGURE 2.5— Distribution of oxygen coordination about aluminium as a function of density and stoichiometry, showing a tendency to higher coordination number with increasing density or stoichiometry.

2.7 Chapter Summary

Precise computational models of Josephson junctions are becoming crucial to efforts to reduce dissipation and loss in superconducting circuits. The limits of computational resources mean that full *ab initio* models are computationally intractable. However, a combination of *ab initio* and empirical models holds promise for developing flexible and efficient simulation approaches. Through comparisons with both previous theoretical analysis and experimental measurements, we have shown that the resulting structures are representative of those fabricated experimentally. The structure of such junctions may now be used as input conditions for the delocalised oxygen models discussed in the subsequent chapters. Additionally, free parameters in existing phe-

nomenological defect models can be determined via information directly obtained from the atomic positions, and other microscopic TLS models may use these structures to investigate their own particular phenomenon.

“Causa latet, vis est notissima.”

—Publius Ovidius Naso

3 Numerical Solution of the TISE

Derivation of an n dimensional formalism which models a delocalised oxygen atom embedded in a surrounding amorphous region of AlO_x.

We are interested in what happens to the oxygen atom as it responds to an external potential exerted via its nearest neighbours, similar to the oxygen interstitial defects discussed in subsection 1.7.1. Particularly if the bonded aluminium atoms are displaced in a manner similar to the defect types identified as A and B in figure 1.5. Type C requires either a complicated many body investigation or some form of pseudo-single body approximation. As a result, assumptions about the defects response to external forces will be compounded which may prove to be drastically misstated. Hence, we do not include this possible defect type in our considerations henceforth, although detailed literature is available in the event that this defect type is later identified as significant [134–137]. Oxygen is considered as the moving atom throughout this thesis due to its mass of 16 amu compared to 27 amu for aluminium. Whilst this mass difference is not huge, the probability of an oxygen atom spatially delocalising over an appropriate distance to yield a dipole strength of order 1 eÅ is greater. In addition to the mass argument, section 5.3 discusses a Bader charge analysis identifying a Bader volume (which approximates electronic charge) of 1.395 ± 0.006 e for oxygen in a junction. Aluminium atoms residing in the amorphous layer have a much smaller volume (on average) of 0.431 ± 0.370 e, which decreases the probability of an aluminium atom exhibiting a large dipole response compared to an oxygen atom.

If we also ignore any time evolution properties of the system for now, an effective single particle Hamiltonian can be derived

$$H = -\frac{\hbar^2}{2m_{oxy}}\nabla^2 + V(\mathbf{r}), \tag{3.1}$$

where m_{oxy} is the mass of an oxygen atom and $V(\mathbf{r})$ is the potential due to the surrounding (mostly amorphous) lattice.

This chapter is primarily concerned with the numerical methods used to solve this Hamiltonian. First, the choice of a suitable potential to describe $V(\mathbf{r})$ is considered in section 3.1. An in-depth analysis of the numerical treatment of ∇^2 follows in sections 3.2 to 3.9, then finally an overview of the construction of the Hamiltonian matrix and unit conversions in section 3.10.

3.1 Potential Configuration

A potential which represents the junction, requires a number of capabilities. It needs to describe interactions between atomic species (in this case, Al–Al, O–O and Al–O interactions), as well as many body interactions which need to be accurate as possible (a requirement to investigate degenerate states). As a complete description of many body interactions does not currently exist, potentials of this type tend to be empirically fitted to experimental data in order to obtain physico-chemical properties of a studied system to high precision. The trade off here is that whilst any given potential may describe some properties of the system well because of certain fitting parameters, other properties may be well out of range as they were not included in the study that constructed the potential. Great care was taken to choose the best potential that accurately represented the junction, in this case the empirical Streitz-Mintmire potential [128], which describes a myriad of aluminium oxide properties over quite a large range of temperatures and pressures with high accuracy when compared against similar potentials. It was also chosen over simpler fixed-charge models [100, 138] due to the complex geometry of the Josephson junction. A variable charge potential such as Streitz-Mintmire can capture the variable oxygen states when present in a predominantly metallic environment through the minimisation of the electrostatic term in the potential (3.2). This capability is particularly important here, as our junction has two metal-oxide interfaces, and our TLS defects may reside close to these boundaries.

The Streitz-Mintmire potential is given by

$$V(\mathbf{r}) = E_{EAM} + \sum_{i}^{N} q_i \chi_i + \frac{1}{2} \sum_{i,j}^{N} q_i q_j V_{ij}, \tag{3.2}$$

where $q_{i,j}$ is the atomic charge, and the terms $\chi_i = \chi_i^0 + \sum_j Z_j([j|f_i] - [f_i|f_j])$ and $V_{ij} = J_i^0 \delta_{ij} + [f_i|f_j]$. J_i^0 is an empirical parameter known as "atomic hardness" or a self-Coulomb repulsion [139], $\delta_{ij} = 1$ when $i = j$ and $\delta_{ij} = 0$ when $i \neq j$, and all summation is calculated for N atoms of the target system. The square bracket notation represents Coulomb interaction integrals between valence charge densities and/or effective core charge densities and take the form [140]:

$$[a|f_b] = \int \frac{f_b(r_b, q_b)}{r_{av}} \, \mathrm{d}V_b \tag{3.3}$$

$$[f_a|f_b] = \iint \frac{f_a(r_a, q_a) f_b(r_b, q_b)}{r_{vv}} \, \mathrm{d}V_a \mathrm{d}V_b \tag{3.4}$$

with $a = i, j; b = i, j; a \neq b$ in (3.2) and $\mathrm{d}V_{a,b}$ are integrating volume units. The value r_{av} is therefore the center distance between atom a and $\mathrm{d}V_b$, and r_{vv} is the center distance between $\mathrm{d}V_a$ and $\mathrm{d}V_b$.

The first term in (3.2): E_{EAM}, does not depend on the partial charges q_i and therefore describes a charge-neutral system, represented here with a quantum mechanical based empirical embedded atom model (EAM) for the Al–Al and Al–O interactions

$$E_{EAM} = \sum_{i}^{N} F_i\left[\rho_i\right] + \sum_{i<j}^{N} \varphi_{ij}(r_{ij}), \tag{3.5}$$

with $F_i\left[\rho_i\right]$ as the energy required to embed atom i in a local electron density ρ_i, and $\varphi_{ij}(r_{ij})$ describing the residual pair-pair interactions by way of Buckingham and Rydberg potentials

$$\varphi_{ij}(r_{ij}) = A\exp\left(-\frac{r_{ij}}{\rho}\right) - B\left[1 + C\left(\frac{r_{ij}}{r_0} - 1\right)\right]\exp\left[-C\left(\frac{r_{ij}}{r_0} - 1\right)\right], \tag{3.6}$$

where r_{ij} is the interatomic (Euclidean) distance between atoms i and j, all other constants are listed in table 3.1. Further formalisms, variable descriptions and parameters can be found in [128, 140], additional implementation is also discussed in [127].

Table 3.1— Empirical constants used in the calculation of the Buckingham and Rydberg pair potentials in the Streitz-Mintmire formalism [127, 128].

Pair	A	ρ	B	C	r_0
Al–Al	4.474755	0.991317	0.159472	5.949144	3.365875
Al–O	62.933909	0.443658	0.094594	9.985407	2.358570
O–O	3322.784218	0.291065	1.865072	16.822405	2.005092

3.2 Numerical Treatment of the Second Derivative

The Streitz-Mintmire potential [128] is an assemblage of many functions that are not completely analytic over all regions of importance. Therefore ∇^2 in (3.1) will also require numerical treatment over a discrete grid of spatial coordinates. A number of methods exist for this problem; most of which introduce errors from approximations and platform limitations. Models exist that effectively remove these errors, but are usually incredibly mathematically complex

and problem specific: for example, a variational model which calculates an optimised three-finite-burn lunar escape trajectory [141].

On the simpler end of the spectrum, finite difference algorithms are useful for boundary value problems (where forward and backward methods are usually applied), and for ordinary and partial differential equations. If a first order ODE (or PDE) of the form $f(x)$ can be evaluated both left and right of x, the central difference method can be used where abscissæ are chosen symmetrically about x, which takes the form

$$f'(x) \approx \frac{f(x+h) - f(x+h)}{2h}, \tag{3.7}$$

where h is some step size. The step size controls the accuracy of the computed derivative, which is unfortunately bound by two factors. If the step size is too small, numerical roundoff errors cause accuracy issues; on the other hand, step sizes which are too large see mathematical truncation errors dominate. Identifying an optimal balance for the step size is also dependent on the specific value of $f'(x)$ being calculated, which may not be appropriate for a similar set of values. This problem, known as 'the step size dilemma' has generated a number of investigations in an attempt to find a middle ground between simple finite difference methods with strong bounding conditions and the highly domain specific models like the variation model described above.

The complex step method (exploiting complex perturbations of the general Taylor series) yields no subtractive cancellation errors (*vide infra* subsection 3.4.2) [142], which exist in the real spaced Taylor series finite difference methods; (3.7) is a simple example of this. However, for many years this statement was only true for the first order derivative—higher orders were shown to have as many cancellation errors as their finite difference counterparts.

This method was therefore excluded as a candidate for the model designed for this thesis, as it would perform on par with standard finite difference methods for the second order derivative we wish to compute. After much of the base code for this model had been constructed, a generalised, arbitrary order derivative complex step algorithm was published [143]. This algorithm may improve the error contribution of the models' final implementation, although this has not been investigated.

A second possible successor to finite difference is automatic differentiation (AD). The premise of this algorithm isn't necessarily mathematical in nature, but capitalises on the fact that computers are methodological reductionists and ultimately only execute simple arithmetic operations no matter how

complicated the actual computation is. Repeatedly applying the chain rule to these operations, AD can compute the derivative to working precision [144]. There a two methods to implement this algorithm; neither is straightforward. One uses special AD preprocessors that analyse each function call, break it up instruction by instruction and generates a new function that computes derivatives. Method two involves operator overloading that can generate code at compile time, which in a JIT accelerated environment like `Matlab` is unfeasible—requiring some hacking of the engine itself.

Both AD and complex step require access to functions at the source code level, meaning calls to third party libraries like `BLAS` and `LAPACK` (which are used in this model's implementation) become increasingly over-complicated.

For these reasons, it was decided to implement our model using a finite central difference method, paying close attention to the inherent error in exchange for relative computational simplicity.

3.3 Understanding the Central Difference Method*

*While most of this and the following section is standard canon information for the finite difference approach; it closely follows sections of Mathur [145], which will be formally introduced in section 3.5.

All finite difference methods involve truncating a Taylor series expansion of a function $f(x)$ about x^* after a certain number of terms

$$f(x) = f(x^*) + f'(x)(x - x^*) + \cdots + R_n. \tag{3.8}$$

The discarded, higher order remainder terms R_n, are considered to contribute a negligible error to the approximation assuming a sufficiently small step size h

$$R_n = \sum_{m=n}^{\infty} \frac{f^{(m)}(x^*)}{m!}(x - x^*)^m. \tag{3.9}$$

It can be shown using the integral calculus derivation of the Taylor series that the d^{th} derivative can be bound over the interval $[x^*, x]$, *i.e.* $a \leq f^{(d)}(x) \leq c$ [146]. Furthermore, $\exists b \in [a, c]$ such that the remainder can be written as

$$R_d = b\frac{(x - x^*)^d}{d!}. \tag{3.10}$$

The Intermediate Value Theorem can be invoked at this stage to posit some point $\zeta \in [x^*, x]$ exists for which $f^{(d)}(\zeta)$ will equal the unknown parameter b. R_d in this form is called the Lagrange remainder. While there is no known

method to determine a value of ζ exactly for a general function, it is useful to express the d^{th} order Taylor series in terms of the Lagrange remainder:

$$f(x) = \sum_{k=0}^{d} \frac{f^{(k)}(x^*)}{k!}(x - x^*)^k + \frac{f^{(d+1)}(\zeta)}{(d+1)!}(x - x^*)^{d+1}, \quad \zeta \in [x^*, x]. \tag{3.11}$$

The most commonly used central difference formula is the first derivative of second order, which can be obtained by applying (3.11) at two abscissæ of length h from a sampling point in $f(x)$ and solving the simultaneous equation for $f'(x)$

$$f(x+h) = f(x) + f'(x)h + \frac{f''(x)}{2}h^2 + \frac{f^{(3)}(\zeta^+)}{2}h^3, \quad \zeta^+ \in [x, x+h] \tag{3.12}$$

$$f(x-h) = f(x) - f'(x)h + \frac{f''(x)}{2}h^2 - \frac{f^{(3)}(\zeta^-)}{2}h^3, \quad \zeta^- \in [x-h, x] \tag{3.13}$$

$$f'(x) = \frac{f(x+h) - f(x-h)}{2h} - \frac{f^{(3)}(\zeta^+) + f^{(3)}(\zeta^-)}{2}\frac{h^2}{3!}. \tag{3.14}$$

The error term in (3.14) contains the average of the third derivative evaluated the two unknown points ζ^+ and ζ^-, which are bounded inside the range $[x - h, x+h]$. Applying the Mean Value Theorem and assuming that $f^{(3)}(x)$ is smooth over the bounded range, a value ζ must exist between ζ^+ and ζ^- which satisfies the average

$$f'(x) = \frac{f(x+h) - f(x-h)}{2h} - \frac{f^{(3)}(\zeta)}{3!}h^2, \quad \zeta \in [x-h, x+h] \tag{3.15}$$

$$\mathcal{F}_2^{(1)}(x,h) = \frac{f(x+h) - f(x-h)}{2h} + \mathcal{O}(h^2). \tag{3.16}$$

Both equations above represent the first derivative of $f(x)$; although they differ in the sense that (3.15) is still the true derivative and (3.16) truncates the error term and represents it as an order of magnitude—which is the essence of the finite difference approximation to the derivative. The notation $\mathcal{F}_n^{(d)}(x,h)$ represents the general form of the approximation of the d^{th} derivative of $f(x)$ of order n using step size h, which can formerly be expressed as:

$$\mathcal{F}_n^{(d)}(x,h) = \frac{\Delta f_n^{(d)}(x,h)}{h^d} + \mathcal{O}(h^n). \tag{3.17}$$

The $\Delta f_n^{(d)}$ term describes the appropriate finite difference expression obtained from the set of (3.11) for particular values of n and d. Frequently used constructions of this form can be found in tables in books such as Mathews and Fink and

web resources like Holoborodko without the need to derive them from first principles, although comprehensive error discussion is uncommon or over-simplified at best [147, 148].

Solving the Schrödinger equation using the Hamiltonian (3.1) requires a second derivative finite difference expression. Truncating the Taylor series at $\mathcal{O}(h^2)$ is the simplest arrangement, which takes the form

$$f''(x) = \frac{f(x-h) - 2f(x) + f(x+h)}{h^2} - \frac{f^{(4)}(\zeta)}{12}h^2, \quad \zeta \in [x-h, x+h] \tag{3.18}$$

$$\mathcal{F}_2^{(2)}(x,h) = \frac{f(x-h) - 2f(x) + f(x+h)}{h^2} + \mathcal{O}(h^2). \tag{3.19}$$

3.4 Contributors to the Step Size Dilemma

Before applying this method, a step size must be chosen; which relies on an accurate description of the competing error contributors.

3.4.1 Truncation Error

The truncation error in (3.19) as $h \to 0$ is $\mathcal{O}(h^2) \to 0$, implying that h should be as small as possible to maximise the accuracy of the calculation. Limits can also be put on the true truncation error as well. Even though ζ is an unknown quantity, as the step size decreases, so does range in which ζ exists. Hence the limit of (3.19) becomes

$$\lim_{h\to 0} \mathcal{O}(h^2) = -\frac{f^{(4)}(x)}{12}h^2. \tag{3.20}$$

This error differs for each formula, depending on the particular values of n and d (for example, the error in 3.15 tends to $-\frac{f^{(3)}(x)}{3!}h^2$). Using (3.17) and (3.19), a general relationship between the true d^{th} derivative and its finite difference approximation

$$f_n^{(d)}(x) = \mathcal{F}_n^{(d)}(x,h) + C(x,h)h^n, \tag{3.21}$$

can be expressed with an undetermined truncation coefficient term $C(x,h)$, which is independent of the derivative parameter d. This coefficient represented in Lagrange form is

$$C(x,h) = a_1 f_n^{(n+d)}(\zeta), \quad \zeta \in [x - a_2 h, x + a_3 h], \tag{3.22}$$

with the set of unknown constants a also being dependent on the finite difference formula in question. As the truncation coefficient is dependent on ζ, its dependence on h is removed as $h \to 0$ (because $\zeta \to x$). Hence, for small values of h, a simplified coefficient can be defined as

$$C_n(x) \equiv a_1 f_n^{(n+d)}(x). \tag{3.23}$$

Simplifying (3.21) with this new coefficient, this relationship now takes the form of a Richardson extrapolation, therefore we can evaluate at two different step sizes h_1 and h_2 (where $h_1 > h_2$)

$$f_n^{(d)}(x) = \mathcal{F}_n^{(d)}(x, h_1) + C_n(x)h_1^n \tag{3.24}$$
$$f_n^{(d)}(x) = \mathcal{F}_n^{(d)}(x, h_2) + C_n(x)h_2^n \tag{3.25}$$

and solving for $C_n(x)$ we find

$$C_n = \frac{\mathcal{F}_n^{(d)}(x, h_2) - \mathcal{F}_n^{(d)}(x, h_1)}{h_1^n - h_2^n}. \tag{3.26}$$

This expression still infers a small step size such that the approximated truncation coefficient stays independent of h (and therefore constant for a valid range of steps). The complete estimate for the truncation error of order n can now be found via (3.21)

$$\mathcal{TE}_n(x, h_1) = C_n h_1^n = \frac{\mathcal{F}_n^{(d)}(x, h_2) - \mathcal{F}_n^{(d)}(x, h_1)}{h_1^n - h_2^n} h_1^n. \tag{3.27}$$

As much of this derivation requires a 'sufficiently small' step size, it's imperative to consider the behaviour of the truncation error at larger values of h as well. If h increases, the approximate truncation coefficient can no longer be used and the magnitude of the true coefficient $C(x, h)$ may become large and unwieldily as there is no restriction of the values over $f^{(n+d)}(x \pm h_{large})$. Additionally, h^n in (3.27) increases, ultimately suggesting that a step size 'too large' will also result in unpredictable truncation error values.

3.4.2 Roundoff Errors

As mentioned in section 3.2, numbers in a computer are represented with a fixed number of binary digits, and operations applied to them have an inherent loss of accuracy. This phenomenon is designated the term roundoff error,

as the extra digits that cannot be held in memory must be discarded—rounded to the nearest tolerance.

Roundoff error has been a hot topic of research since the invention of computers. Notable checks on the accuracy of finite element methods were completed before the moon landings [149] which were paramount to their success. Investigations after a Patriot missile defense system allowed a Scud missile to hit a barracks, killing 28 people; identified the root cause to be numerical in nature. Roundoff via the differencing of floating point numbers introduced errors into the timing register when converting representations [150]. These relative errors caused by subtraction (called cancellation error) tend to decrease as h is increased, thus to minimise this uncertainty h should be as large as possible—contrary to the truncation requirement of a small step size. An upper bound on this error is given by

$$\left|(\alpha-\beta)_{true}-(\alpha-\beta)\right| \leq \delta \max(|\alpha|,|\beta|), \tag{3.28}$$

where α and β are two fixed precision numbers, and δ indicates the precision of the calculation. Usually, modern computer languages operate using standard double precision floating point numbers, thus $\delta = 2^{-53}$.

A second roundoff error type known as condition error creeps in when functions don't use machine precision floating point numbers in their internal routines. For example, a third party function may approximate π to 3.14159, which would artificially round the result of whatever operation was being applied to a precision of 10^{-5}. As with cancellation error, an upper bound can also be formulated for condition error:

$$|f(x)_{true}-f(x)| \leq \varepsilon|f(x)|. \tag{3.29}$$

Here, ε is the magnitude of the most significant digit affected by the condition error. For elementary operations ε should be equal to machine precision, although this cannot be assumed for all functions hence it is a value that should be calculated in general.

Finally, an error which is also important in this instance is named representational error. Not all numbers can be accurately be represented in binary using a fixed number of digits. Many people choose their step sizes in base 10 (*e.g.* 1×10^{-2}, 5×10^{-6} *etc.*), and what's most troubling about this is the fact that no negative power of 10 has an exact binary representation. This error

is small—usually smaller that machine precision in fact, but it has non-trival and cumulative side effects. If one avoids step sizes other than $2^{\mathbb{N}}$ this error can be avoided completely.

With this information in hand, an upper bound to the total roundoff error can now be calculated:

$$|\mathcal{F}_n^{(d)}(x,h)_{true} - \mathcal{F}_n^{(d)}(x,h)| \leq \frac{\varepsilon|\mathcal{F}_\varepsilon| + \delta|\mathcal{F}_\delta|}{h^d}. \tag{3.30}$$

The functions $\mathcal{F}_\varepsilon$ and $\mathcal{F}_\delta$ are derived from $\Delta f_n^{(d)}$ (see 3.17). Forms of (3.19) for instance are

$$\mathcal{F}_\varepsilon = |f(x+h)| + |2f(x)| + |f(x-h)| \tag{3.31}$$

$$\mathcal{F}_\delta = \max(|f(x+h) + f(x-h)|, |2f(x)|). \tag{3.32}$$

3.4.3 Estimation of Total Error

Using the equations for truncation and roundoff errors mentioned above, and assuming a small h such that (3.26) holds, the total error can be bounded by the expression

$$|f^{(d)}(x)_{true} - \mathcal{F}_n^{(d)}(x,h)| \leq \frac{\varepsilon|\mathcal{F}_\varepsilon| + \delta|\mathcal{F}_\delta|}{h^d} + |C_n|h^n. \tag{3.33}$$

Assuming a worse case scenario, the total error can therefore be described as

$$\mathcal{E}_n^{(d)}(x,h) = \frac{\varepsilon|\mathcal{F}_\varepsilon| + \delta|\mathcal{F}_\delta|}{h^d} + |C_n|h^n \tag{3.34}$$

noting the dependence on x comes from the functions $\mathcal{F}_\varepsilon$ and $\mathcal{F}_\delta$, and that ε is still an unknown value.

3.5 Calculating an Optimal Step Size

With all of these caveats in mind, finding an optimal step size h_{opt} is a troublesome undertaking; which is why the complex step and AD methods discussed in section 3.2 were developed.

A simple method proposed in Gill *et al.* suggests minimising an expression similar to (3.34) which trades off the truncation and roundoff errors [147, 151]. However, this method assumes the condition error ε is known and cancellation error is treated in a trivial manner.

Violation of monotonicity is an algorithm which attempts to ignore actually calculating any estimates for truncation and roundoff errors [152]. This method breaks down if (3.26) doesn't hold and therefore is not a general solution. Algorithms specifically designed for forward differences have also been discussed in the literature [153], but are of no benefit for purposes herein.

The method which has been chosen to be implemented in this work is the algorithm from the PhD thesis of **Ravishankar Mathur**, which allows one to calculate an optimal step without *a priori* knowledge of the condition error ε. Additionally, corrections to the step size are introduced to account for the approximate nature of the truncation coefficient C_n, as well as estimations of step size validity and the maximal appropriate step size for a problem are examined [145].

As with Gill *et al.*, Mathur starts by minimising the total error expression

$$\frac{\partial \mathcal{E}}{\partial h} = -d\frac{\varepsilon|\mathcal{F}_\varepsilon| + \delta|\mathcal{F}_\delta|}{h^{d+1}} + n|C_n|h^{n-1} = 0 \tag{3.35}$$

and solving for h to find

$$h_{opt,\mathcal{TE}} = \left[\frac{d}{n}\frac{1}{|C_n|}\left(\varepsilon|\mathcal{F}_\varepsilon| + \delta|\mathcal{F}_\delta|\right)\right]^{1/(n+d)}. \tag{3.36}$$

Note that obtaining correct values of $\mathcal{F}_\varepsilon$ and $\mathcal{F}_\delta$ require the optimal step size $h_{opt,\mathcal{TE}}$, thus an iterative method is required. Once this value is known, (3.36) can be rearranged to find the condition error ε.

The steps of the algorithm which computes these values (as well as discussions of further optimisations) can be found in Chapter 3 of reference 145.

3.6 A Harmonic Approximation to Streitz-Mintmire

A complication arises when attempting to apply this algorithm to ∇^2 in (3.1), which more specifically should be written as $\frac{\mathrm{d}^2\Psi}{\mathrm{d}x^2}$; where $\Psi(x)$ is an eigenstate of H in the time-independent Schrödinger equation $H|\Psi(x)\rangle = E|\Psi(x)\rangle$, defining the spatial coordinate system in this example to be one dimensional for simplicity: $\mathbf{r} \equiv x$. This eigenvalue equation is solved via direct diagonalisation of the Hamiltonian matrix—requiring a discretised treatment of ∇^2 before $\Psi(x)$ can be obtained. The step size algorithm requires an analytical form of the function to which it is applied. Unfortunately there are very few quantum mechanical systems with analytical solutions, and from that small

pool, most are overly simplified constructions that do not exist physically. The functional form of the wavefunction for a particle under the influence of the Streitz-Mintmire potential is not one of these systems. However, certain configurations of the potential approximate to a parabola-like shape (for example see figures 3.1 and 4.4), the quantum harmonic oscillator can therefore serve as an analogue in this limit as it has a simple analytic solution.

The Hamiltonian of a particle in a one dimensional quantum harmonic oscillator can be described as

$$\widehat{H} = \frac{\widehat{p}^2}{2m} + \frac{1}{2}m\omega^2\widehat{x}^2 \tag{3.37}$$

where m is the particle's mass and ω is the angular frequency of the oscillator. Two operators, $\widehat{x} = x$ for position and $\widehat{p} = -i\hbar\frac{\partial}{\partial x}$ for momentum describe the complete Schrödinger equation. After separation of variables, the time independent form becomes

$$-\frac{\hbar^2}{2m}\frac{\partial^2\Psi}{\partial x^2} + \frac{1}{2}m\omega^2x^2\Psi = E\Psi, \tag{3.38}$$

with the family of solutions for the wavefunctions

$$\psi_n(x) = \frac{1}{\sqrt{2^n\, n!}}\left(\frac{m\omega}{\hbar\pi}\right)^{1/4} e^{-\frac{m\omega x^2}{2\hbar}} H_n\left(\sqrt{\frac{m\omega}{\hbar}}x\right), \quad n = 0, 1, 2, \ldots \tag{3.39}$$

$$H_n(x) = (-1)^n e^{x^2}\frac{\mathrm{d}^n}{\mathrm{d}x^n}\left(e^{-x^2}\right) \tag{3.40}$$

$$\therefore \psi_0(x) = \left(\frac{m\omega}{\hbar\pi}\right)^{1/4} e^{-\frac{m\omega x^2}{2\hbar}} \tag{3.41}$$

where $\psi_0(x)$ (3.41) is the ground state wavefunction. Particle mass m in our case is the mass of an oxygen atom, and the angular frequency ω is currently unknown.

Using a small step size, the eigenvalue equation is solved for the Streitz-Mintmire case, and a ground state wavefunction is found. The functional form of which is shown in the left plot of figure 3.1 as —. As the potential form (shown in the right plot) is similar to a parabola, the resultant wavefunction is gaussian-like. Using (3.41), a harmonic form of the ground state wavefunction can now be fitted to the calculated Streitz-Mintmire result through the unknown angular frequency, which is found to be $\omega = 2.241 \times 10^{13}$ rad/s. Both $\psi_0(x)$ and the associated potential of this harmonic approximation are plotted as — to compare with the Streitz-Mintmire result.

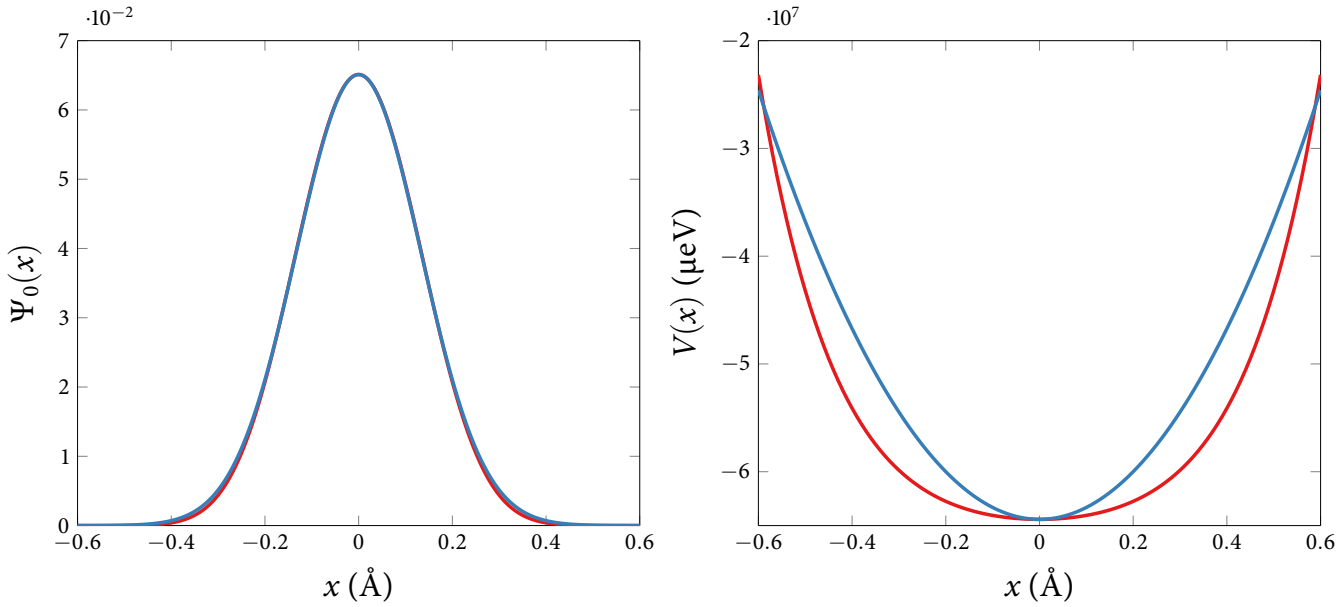

Figure 3.1— Calculated Streitz-Mintmire — and Harmonic approximations — to the ground state wave function $\psi_0(x)$ (left) and potential $V(x)$ (right). The harmonic response is scaled via $\omega = 2.241 \times 10^{13}$ rad/s using (3.41). Units of $V(x)$ in [μeV] with x in [Å] correspond to typical atomic energy and length scales.

The fitted value of ω yields a complete approximation to the Strietz Mintmire ground state wavefunction, which we can now use in conjunction with (3.19) to compute ∇^2 in (3.1); evaluating at $x = 0$ to find the eigenvalue E.

Figure 3.2 displays the absolute value of the truncation error (3.27) over a range of possible step sizes ●. Roundoff errors dominate for small step sizes ($h \lesssim 5 \times 10^{-5}$) with some step sizes resulting in invalid results (*i.e.* $\mathcal{TE}_2 = 0$). These steps are labelled ▲ and are scaled to $2{\times}10^{-8}$ to be displayed on the graph. As the step size increases, truncation error dominates until the region $h \gtrsim 5$, where this error becomes invalid.

The step size optimisation algorithm initially chooses an uncorrected optimal step size ● of $h_{opt,\mathcal{TE}} = 5.775 \times 10^{-5}$ Å. The truncation error (3.27) overestimates the true roundoff error contribution by a proportional amount given by $t^* = (1 + (1/t)^d)/(1 - t^n)$ [145]. The constant $t = 0.65$ is the step size ratio, required by the step size algorithm and is optimal for $d = 2$, $n = 2$. Using this adjustment, the true optimal step is

$$h_{opt,true} = \left(\frac{1}{t^*}\right)^{1/(n+d)} h_{opt,\mathcal{TE}}, \tag{3.42}$$

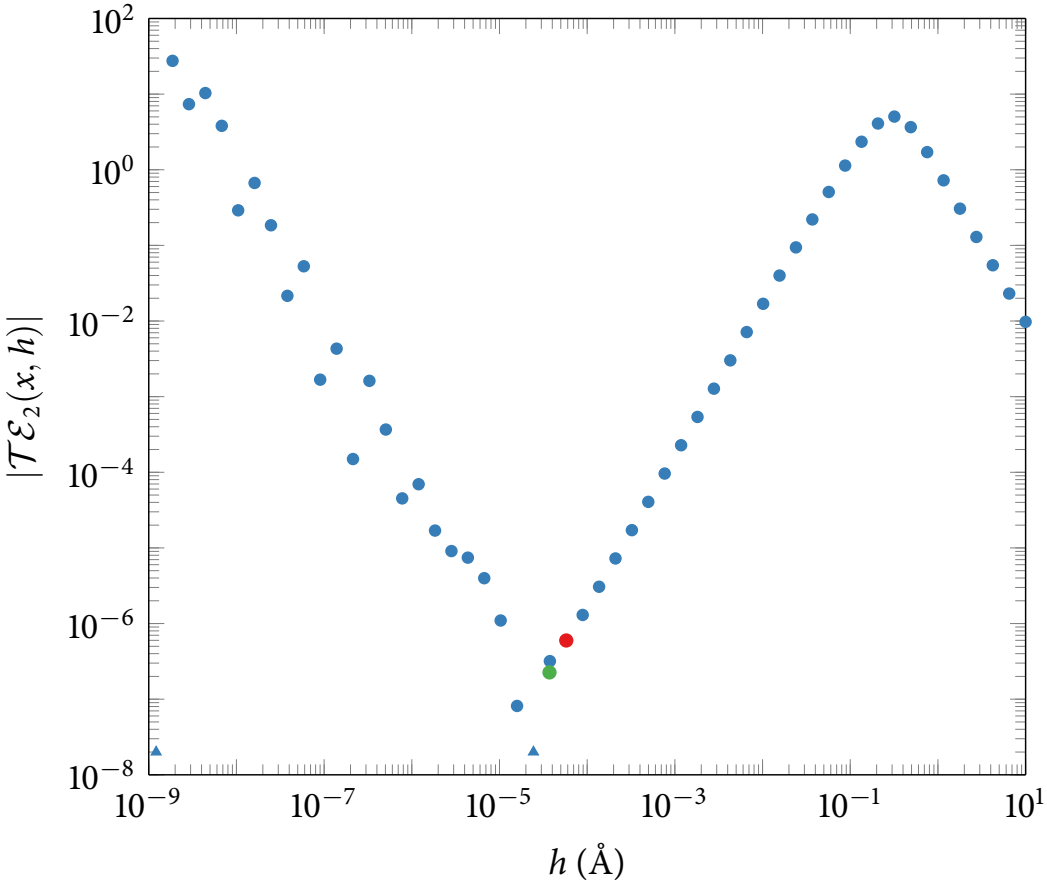

FIGURE 3.2— Step size optimisation of $f_2^{(2)}(x)$ for step sizes $10^{-9} \leq h \leq 10^{1}$ ●. Steps which generate an invalid roundoff error (*i.e.* $\mathcal{TE}_2 = 0$) are translated to 2×10^{-8} for display purposes and labelled ▲. Two optimal step sizes are identified: $h_{opt,\mathcal{TE}}$ ●, found using the optimal step algorithm [145], and $h_{opt,true}$ ●, corrected by (3.42).

labelled as ●. The value of this corrected step is $h_{opt,true} = 3.717\times10^{-5}$ Å, which does not have an accurate power of two representation. As stated in subsection 3.4.2, representation error may be smaller than machine precision; although in particular instances this may have a non-trivial cumulative effect. Hence the optimal step size for the central difference second derivative of second order applied to (3.1) is

$$h_{opt} = 3.717\times10^{-5} \simeq 2^{-15} = 3.052\times10^{-5}\ \text{Å}. \tag{3.43}$$

3.7 Memory Concerns

The complexity of numerical calculations are exacerbated by yet another constraint in the form of finite memory resources. Ignoring technical intricacies, a first order approximation to the required memory footprint just to store the numbers of a 2D plane, 9 Å^2 in area discretised via h_{opt} requires $294,913^2$ double precision floats. On a computer with 64 bit architecture, each double requires 8 bytes of memory, meaning our grid has a footprint of just over 695 Gigabytes.

Computational resources of that magnitude are infeasible when a simple solution can both minimise memory requirements and increase the accuracy of the calculation simultaneously. Recall the total error (3.34), which depends on h via

$$\mathcal{E}_n^{(d)}(x,h) \propto \frac{1}{h^d} + h^n. \tag{3.44}$$

For our purposes, d is fixed at 2, and n was also chosen as 2, but only due to the fact that this generates the simplest second derivative central difference formula. Using order of magnitude arguments, a step size $h = 1\times10^{-6}$ with an order parameter $n = 2$ contributes $\mathcal{O}(10^{-12})$ to the total error from the h^n term. If the order is increased to $n = 6$, a much larger step size $h = 1\times10^{-2}$ yields the same $\mathcal{O}(10^{-12})$ contribution.

A back of the envelope calculation for a step size that large over the same range as above only requires 901^2 doubles, which equates to a much more feasible memory requirement of 6.5 Megabytes.

3.8 Second Derivative of Sixth Order

Using the formalism outlined in section 3.3, an expression for the second derivative of sixth order can be calculated. Starting with the Taylor series of the Lagrange remainder (3.11) using six abscissæ

$$\begin{aligned} f(x \pm h) = {} & \\ & f(x) \pm f'(x)h + \frac{f''(x)h^2}{2!} \pm \frac{f^{(3)}(x)h^3}{3!} + \frac{f^{(4)}(x)h^4}{4!} \pm \frac{f^{(5)}(x)h^5}{5!} \\ & + \frac{f^{(6)}(x)h^6}{6!} \pm \frac{f^{(7)}(x)h^7}{7!} + \frac{f^{(8)}(\zeta^{\pm})h^8}{8!}, \\ & \qquad \zeta^+ \in [x, x+h], \qquad \zeta^- \in [x-h, x] \end{aligned} \tag{3.45}$$

$$f(x \pm 2h) = f(x) \pm 2f'(x)h + \frac{4f''(x)h^2}{2!} \pm \frac{8f^{(3)}(x)h^3}{3!} + \frac{16f^{(4)}(x)h^4}{4!} \pm \frac{32f^{(5)}(x)h^5}{5!} + \frac{64f^{(6)}(x)h^6}{6!} \pm \frac{128f^{(7)}(x)h^7}{7!} + \frac{256f^{(8)}(\zeta^{\pm})h^8}{8!}, \quad \zeta^+ \in [x, x+2h], \quad \zeta^- \in [x-2h, x] \tag{3.46}$$

$$f(x \pm 3h) = f(x) \pm 3f'(x)h + \frac{9f''(x)h^2}{2!} \pm \frac{27f^{(3)}(x)h^3}{3!} + \frac{81f^{(4)}(x)h^4}{4!} \pm \frac{243f^{(5)}(x)h^5}{5!} + \frac{729f^{(6)}(x)h^6}{6!} \pm \frac{2187f^{(7)}(x)h^7}{7!} + \frac{6561f^{(8)}(\zeta^{\pm})h^8}{8!}. \quad \zeta^+ \in [x, x+3h], \quad \zeta^- \in [x-3h, x] \tag{3.47}$$

Followed by removing the odd degree terms

$$f(x+h) + f(x-h) = 2f(x) + f''(x)h^2 + \frac{2f^{(4)}(x)h^4}{4!} + \frac{2f^{(6)}(x)h^6}{6!} + \frac{2f^{(8)}(\zeta^{\pm})h^8}{8!} \tag{3.48}$$

$$f(x+2h) + f(x-2h) = 2f(x) + 4f''(x)h^2 + \frac{32f^{(4)}(x)h^4}{4!} + \frac{128f^{(6)}(x)h^6}{6!} + \frac{512f^{(8)}(\zeta^{\pm})h^8}{8!} \tag{3.49}$$

$$f(x+3h) + f(x-3h) = 2f(x) + 9f''(x)h^2 + \frac{162f^{(4)}(x)h^4}{4!} + \frac{1458f^{(6)}(x)h^6}{6!} + \frac{13122f^{(8)}(\zeta^{\pm})h^8}{8!} \tag{3.50}$$

leaving a set of equations with fourth and sixth degree terms which need be eliminated in order to arrive at an equation similar to (3.18). The equation $72\times$(3.49)$-2\times$(3.50)$-270\times$(3.48) eliminates these terms, arriving at

$$f''(x) = \frac{2f_{-3} - 27f_{-2} + 270f_{-1} - 490f_0 + 270f_1 - 27f_2 + 2f_3}{180h^2} - \frac{f^{(8)}(\zeta)h^6}{560}, \tag{3.51}$$

$$\mathcal{F}_2^{(6)}(x,h) = \frac{2f_{-3} - 27f_{-2} + 270f_{-1} - 490f_0 + 270f_1 - 27f_2 + 2f_3}{180h^2} + \mathcal{O}(h^6), \tag{3.52}$$

where $\zeta \in [x-3h, x+3h]$ and both the true and finite difference expressions are displayed in a simplified form. The notation can be read as

$$f_k = f(x_k), \quad x_k = x^* + kh, \quad \{k \in \mathbb{Z} | -(N-1)/2 \leq k \leq (N-1)/2\} \tag{3.53}$$

for example $f_{-3} \equiv f(x-3h)$. The form of (3.52) adheres to (3.17), where

$$\Delta f_6^{(2)} = \frac{1}{180} 2f_{-3} - 27f_{-2} + 270f_{-1} - 490f_0 + 270f_1 - 27f_2 + 2f_3. \tag{3.54}$$

Applying this equation is no different to (3.19), and the only additional overhead is the requirement of six abscissæ rather than two.

An optimal step size for ∇^2 in (3.1) can be found using (3.52) and the harmonic approximation of $\Psi_0(x)$ (3.41). As with the calculation in section 3.6 using (3.19), we will evaluate at $x = 0$ and apply the optimal step size algorithm [145]; the results of which are shown in figure 3.3.

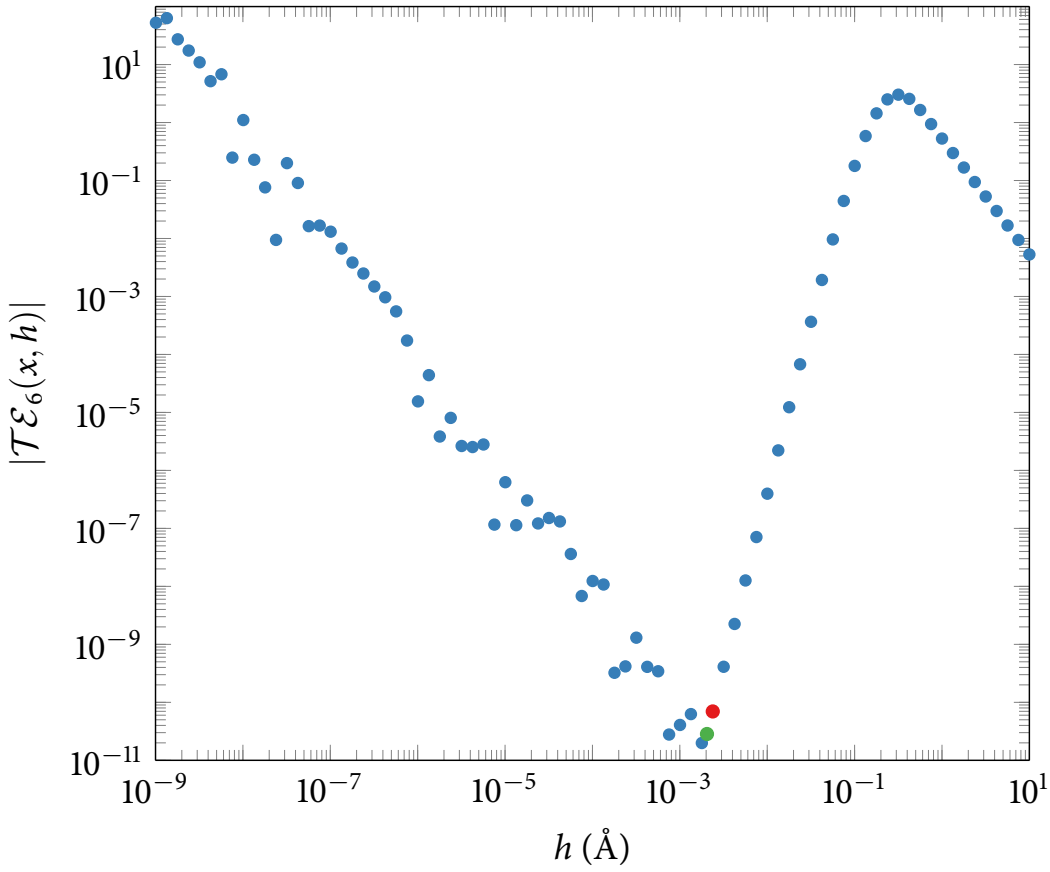

Figure 3.3— Step size optimisation of $f_6^{(2)}(x)$ for step sizes $10^{-9} \leq h \leq 10^{1}$ ●. Two optimal step sizes are identified: $h_{opt,\mathcal{TE}}$ ●, found using the optimal step algorithm [145], and $h_{opt,true}$ ●, corrected by (3.42).

In comparison to the results of the second order method (see figure 3.2), which was calculated over the same step size range, it's clear that the estimations in section 3.7 hold. The absolute truncation error $|\mathcal{TE}_6|$ minimum is in fact four orders smaller, and the optimal step sizes two orders larger.

The uncorrected step ● was found to be $h_{opt,\mathcal{TE}} = 2.381 \times 10^{-3}$ Å, and corrected using (3.42) (with $t = 0.75$, optimised for $d = 2$, $n = 6$) to $h_{opt,true} = 2.044 \times 10^{-3}$ Å, again labelled as ●.

Moving this value to a power of two representation, the optimal step size for the central difference second derivative of sixth order applied to (3.1) is

$$h_{opt} = 2.044 \times 10^{-3} \simeq 2^{-9} = 1.953 \times 10^{-3} \text{ Å}. \tag{3.55}$$

A step of this size for a 9 Å^2 grid requires $4,609^2$ doubles with a memory footprint of around 170 Megabytes.

3.9 Condition and Total Error Calculation

The harmonic approximation was useful for finding the optimal step of the Streitz-Mintmire method only because the change in x is equivalent over the calculated range. Actual error values on the other hand can not be considered in the same manner. The condition error of the harmonic ground state wave function (3.41) is dependent on the accuracy of the input variables and the exponential function, which in `Matlab`, is computed by the built-in function `exp`. On the other hand, the Streitz-Mintmire potential is a custom coded implementation built for the purpose of this thesis, and calls `Matlab`, `C++` and `LAPACK` (implemented in `Fortran`) routines to calculate a potential value for a given x. Then, to obtain $\Psi_0(x)$, the Hamiltonian matrix is diagonalised through the `eigs` function, also calling `LAPACK` through a `C++` wrapper.

Many of these steps may contribute to the condition error, which can be calculated by rearranging the optimal step size formula (3.36), now that $h_{opt} = 2^{-9}$ is known

$$\varepsilon = \frac{1}{|\mathcal{F}_\varepsilon|}\left(\frac{n}{d}|C_n|h_{opt}^{n+d} - \delta|\mathcal{F}_\delta|\right). \tag{3.56}$$

C_n (3.26) is calculated during the optimal step algorithm [145] when the steps h_1 and h_2 are found to be in the valid truncation error range. The absolute value of which is $|C_6| = 3.888 \times 10^5$ for the sixth order central difference. Functional forms of the condition error $\mathcal{F}_\varepsilon$ and cancellation error $\mathcal{F}_\delta$ coefficients are generated from (3.54)

$$\mathcal{F}_\varepsilon = \frac{1}{180}2|f_{-3}| + 27|f_{-2}| + 270|f_{-1}| + 490|f_0| + 270|f_1| + 27|f_2| + 2|f_3| \tag{3.57}$$

$$\mathcal{F}_\delta = \frac{1}{180}\max(|27f_{-2} + 490f_0 + 27f_2|, |2f_{-3} + 270f_{-1} + 270f_1 + 2f_3|). \tag{3.58}$$

Applying this formula to the harmonic approximation, the condition error is found to be $\varepsilon = 7.958 \times 10^{-19}$. This result is beneath machine precision ($\delta = 2^{-53} \simeq 1.110 \times 10^{-16}$), which is expected considering exp is a built-in function. The total error (3.34) is now a trivial undertaking, using the above values $\mathcal{E}_6^{(2)}(0, h_{opt}) = 8.634 \times 10^{-11}$.

Calculating these values for the Streitz-Mintmire case however is a much more daunting task. As the wavefunctions' form is unknown before the eigenvalue problem is solved, computing the required variables: C_n, $\mathcal{F}_\varepsilon$ and $\mathcal{F}_\delta$ is not possible.

However, estimates of the condition error can be made. The sparse matrix eigenvalue solver eigs has a documented precision of 2^{-53} [154], in other words: machine precision δ. One of the pivotal advances Mathur accomplishes is the ability to obtain an optimal step without *a priori* knowledge of ε [145]. Turned on its head: with a known step size for a function one can estimate the condition error. Therefore, applying (3.52) to the Streitz-Mintmire potential to calculate $\frac{\mathrm{d}^2 V}{\mathrm{d}x^2}$ can obtain an optimal step, then a condition error for $V(x)$. This process yields a value of $\varepsilon = 3.054 \times 10^{-19}$, again below machine precision.

Whilst these values cannot quantify the Streitz-Mintmire total error, they generate at least some confidence that the values are in the same order of magnitude as depicted in figure 3.3. One further test of stability that can be used in this instance is a convergence test of the energy E. Solving the eigenvalue equation over the range of step sizes used previously, figure 3.4 displays how the energy fluctuates. As expected, steps with high roundoff error contributions cause fluctuations in the calculated value of E, and invalid truncation errors also generate large deviations from the acceptable value at $h_{opt} = 2^{-9}$.

3.9.1 Acceptable Maximum Step Size

The 170 Megabyte memory footprint calculated for a 9 Å^2 plane using the sixth order second derivative in section 3.8 is completely acceptable if one plane was all that was required. Below in section 3.10 and chapter 4, the need for calculating at least 6 excited state wavefunctions using direct matrix diagonalisation is presented. This, along with other computational overheads sees the memory requirement balloon and the problem again becomes intractable. For instance, the 1D configuration used to compute energies for figure 3.4 lies within $x \in [-0.6, 0.6]$ Å. Although a step of $h = 1 \times 10^{-8}$ Å needs just under 1 Gigabyte

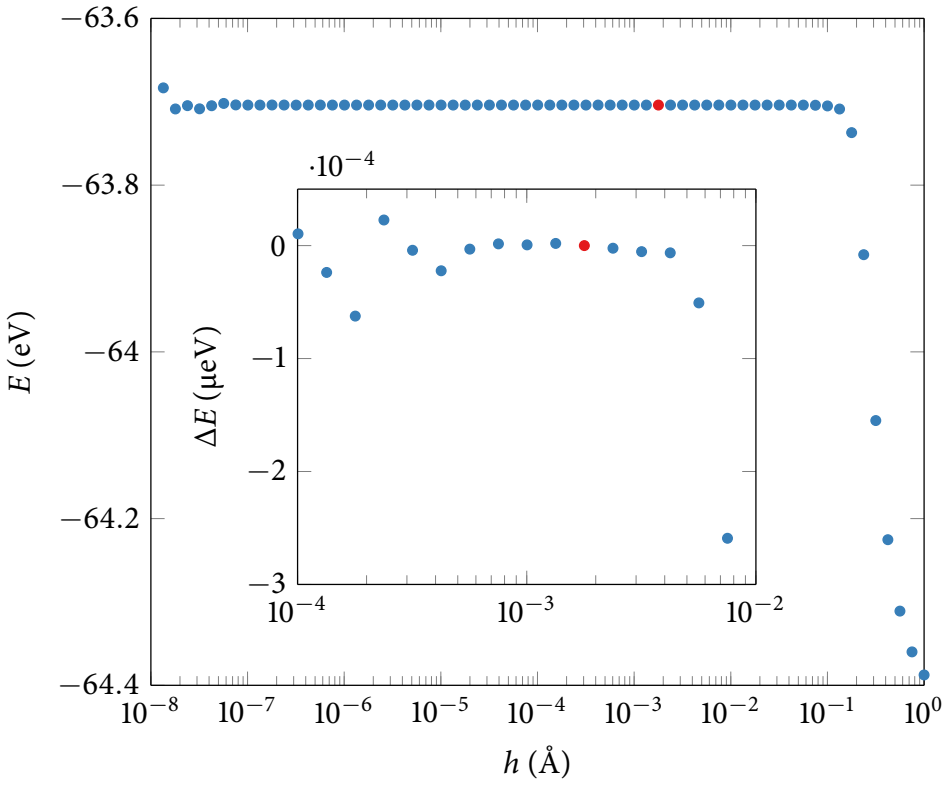

FIGURE 3.4— Energy Convergence for step sizes $10^{-8} \leq h \leq 10^{0}$ •, with h_{opt} labelled as •. The invalid truncation range from figure 3.3 is visible in the sense that large step sizes $h \gtrsim 2 \times 10^{-1}$ contribute sizable error. The inset shows two decades close to h_{opt}, where the energy scale is normalised to $E - E(h_{opt})$. Left of h_{opt} sees roundoff error contributions, and right depicts truncation error.

of memory to store this line, the total calculation cannot be completed on a machine with 32 GB of RAM and 128 GB of swap space.

Following the argument in section 3.7, an instinctive solution to this issue would be to increase the order of the central difference method to $n = 8$ or perhaps even $n = 10$. This would increase the optimal step size to $\mathcal{O}(10^{-1})$ Å, and the total error contribution would decrease further. Juxtaposing figures 3.2 and 3.3, two problems with this solution are elucidated. Whilst the optimal step size becomes larger as n increases, this also increases the gradient which obtains C_n and defines the region of valid truncation. Also, the region of invalid truncation error remains constant for the system in question (in this case $\Psi_0(x)$). If the order is increased to the point at which the optimal step size is almost at the lower bound of the invalid truncation range, C_n will become a large contribution to error. Section 3.2 also discusses the complication that an

optimal step size may not be valid for all points in a system. Pushing the optimal step to be too large may encounter non-optimal position errors, where one (optimal) point has an error of 10^{-11} associated with it, and the next (non-optimal) point contributes 10^{1} from invalid truncation.

It is shown in Mathur that the optimal step size is valid over the range $x \in [x^* - a_2 h_{max}, x^* + a_3 h_{max}]$ if $f^{(n+d)}(x)$ can be shown to be constant over the same range. This is in turn proven by showing $f^{(n+d)}(\zeta)$ is constant over $[h_{opt}, h_{max}]$. The derivation is related to (3.22), where a_2 and a_3 originate. The value h_{max} is simply the point at which the truncation error becomes valid as step size decreases. Thus $h_{max} = 2^{-3} \simeq 0.125$ Å for the sixth order second derivative from figure 3.3 so long as $f^{(8)}(\zeta)$ is constant over $[2^{-9}, 2^{-3}]$, which is indeed the case if one applies the regression algorithm from reference 145. As a result, the estimated truncation error calculated for h_{opt} is considered to be consistent over the range $x \in [x^* - 3 \cdot 0.125, x^* + 3 \cdot 0.125]$ Å using the sixth order method. Put another way: the calculated value of $h_{opt} = 2^{-9}$ is only the optimal value at $x = 0 \pm 0.375$ Å. Require values anywhere else on the grid and a new h_{opt} should be calculated if you're a purist.

Unfortunately, compromises need to be made in practice. So, keeping the above analysis in mind, the following concessions will be made throughout this thesis:

- h_{opt} will be considered optimal over a larger range than 0 ± 0.375 Å to avoid iterative and overtly complicated treatments of the eigenvalue equation. In most cases this range will be within $x \in [-3.3, 3.3]$ Å, although some calculations may require $x \in [-4.5, 4.5]$ Å. Care has been taken on the choice of box size such that the wavefunction generally tends to zero as x increases, meaning values close to the centre of the box are the most important regardless.

- Step sizes optimised for a range in x are assumed valid in all spatial dimensions r.

- Truncation error at h_{max} contributes a consistent $\leq 1 \times 10^{-3}$ μeV difference in energy to the value calculated at h_{opt} over tested configurations. Large calculations that require only relative energy values, such as the phase maps in sections 4.5, 5.1 and 5.2 have been calculated with step sizes close to h_{max}. This allows the systems to fit in memory and/or not take the age

of the universe to finish computing. However, specific energy values stated herein will still be calculated using h_{opt} unless otherwise stated.

3.10 Hamiltonian Construction

A complete description of the Hamiltonian (3.1) can now be implemented to solve the eigenvalue equation

$$\left[-\frac{\hbar^2\Lambda}{2m_{oxy}}\nabla^2 + V(\mathbf{r})\right]|\Psi(\mathbf{r})\rangle = E|\Psi(\mathbf{r})\rangle, \tag{3.59}$$

the implementation of which requires the construction of a sparse matrix representation of H, which can be fed into the `eigs` function of `Matlab`; computing the wavefunction Ψ and energy E of the system.

In one dimension, this matrix is straightforward to construct (outlined below), although higher dimensions require special treatment (see sections 4.1 and 6.1 for specifics).

Starting with the prefactor of ∇^2, m_{oxy} is of course the mass of an oxygen atom: 2.657×10^{-26} kg, and the unit conversion constant defined as

$$\Lambda \equiv \frac{(1\times10^{10})^2\cdot 1\times10^{6}}{e} = 6.243\times10^{46} \tag{3.60}$$

describes the system in terms of distance in Å and energy in μeV. This prefactor is then multiplied by -490, becoming the f_0 contribution from (3.52). Each subsequent position value ($\mathbf{r}=x$) also requires this contribution as well, which are all placed down the diagonal of a sparse matrix once the corresponding $V(x)$ has been added. Off diagonals are also required to fulfil the remaining f_k values in (3.52), where the k^{th} off diagonal holds the f_k coefficient and the ∇^2 prefactor.

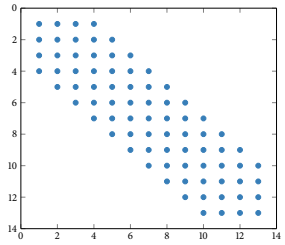

FIGURE 3.5— Sparsity of the Hamiltonian matrix H for a simple one dimensional system with 13 discrete points in x.

The resulting matrix can then be diagonalised, obtaining eigenvectors and eigenenergies for the system in question. Assuming that errors are of the same order of those calculated for the harmonic approximation in section 3.9, differences between energy levels can be acquired with precision better than 10 kHz (4×10^{-5} μeV) using this method. In contrast, the energy scale for JJ defects observed in experiments is $\lesssim$ 10 GHz (40 μeV) [67–69]. This energy splitting, while large for qubit experiments, is very small when compared to typical electronic structure calculations. For comparison, the resolution of *ab*

initio methods such as DFT has a lower bound of order 120 GHz (5 meV): above the superconducting gap and therefore too imprecise to be considered useful for JJ defect studies other than providing atomic positions as input.

In subsequent chapters we use this approach to investigate what affect changes in the local potential landscape impose on an oxygen atom by systematically introducing (initially) aluminium atoms at Al–O bond distances and perturbing their positions. With an understanding of the interplay between atomic positions and the energy spectrum of the oxygen, we slowly increase the complexity of the potential environment until the oxygen is embedded in a cluster characterised by the JJ models constructed in chapter 2.

“ᠮᠣᠩᠭᠣᠯ ᠬᠡᠯᠡ ᠮᠢᠨᠦ
ᠲᠡᠭᠦᠰ ᠬᠦᠮᠦᠨ-”

4 Delocalised Oxygen TLSs In Low Dimensions

Investigating the effects of local potential perturbations on an oxygen atom. Calculation of candidate TLS properties such as E_{01} energy splitting and dipole moments.

Equipped with a framework describing the behaviour of an oxygen atom delocalising in the presence of an external potential, we are now in a position to investigate which atomic configurations may demonstrate TLS characteristics. Whilst other properties have been measured, we will focus on obtaining respectable values for E_{01} and dipole strength, assuming our model defect is embedded inside a fictitious phase qubit. Expected values are approximately $E_{01} = 8$ GHz and $\wp = 1.2\ e$Å [81]. This will enable us to directly compare measured qubit-TLS couplings via (1.6) to our calculated splitting energies and dipole values.

Starting with the model outlined in chapter 3, section 4.1 extends these methods to treat oxygen delocalisation in two dimensions, and introduces an expression to calculate the effective dipole of the oxygen relative to an external field. Sections 4.2 to 5.1 begin with a minimal example of this model and slowly add complexity, so interactions can be examined and understood in a systematic way. Section 4.2 considers a defect comprising of an Al–O–Al chain in one dimension, perturbed from a crystalline lattice, simulating defects A and B in figure 1.5. In reality, an oxygen atom in the amorphous layer of a Josephson junction will be surrounded by atoms in all three dimensions (which can be expressed using a cluster of atoms simulated in chapter 2). Moving towards a model representation of this, section 4.3 extends the model to two dimensions, with four aluminium atoms confining an oxygen in a plane.

4.1 2D Formalism and Dipole Expression

To extend the Hamiltonian H (3.1) into a two dimensional matrix representation, one must arrange data in a way which is understood by the `eigs` function, which expects a two dimensional sparse matrix.

If $\mathbf{r} = (x, y)$ in (3.59), we can treat each dimension separately and follow the method outlined in section 4.1 for x then y. With two sparse matricies, the

Kronecker product can be applied to generate a block matrix ready for diagonalisation.

$$H = H_x \otimes \mathbb{I}_y + \mathbb{I}_x \otimes H_y. \tag{4.1}$$

The dipole element is computed using numerical integration of the ground- and first-excited states (ψ_0, ψ_1), where

$$\wp_x = \iint \psi_0^*(x,y) x \psi_1(x,y)\, \mathrm{d}x\mathrm{d}y \tag{4.2}$$

is an example of the dipole in the x direction.

Relative differences in energy levels (*i.e.* energy splittings) are an important measure of the model. We therefore define a convention where $E_{ij} = E_j - E_i$, such that the ground state (E_0) to first excited state (E_1) energy splitting is defined as E_{01}. For comparison with experimental results, energy is expressed in frequency units throughout this discussion.

4.2 TLS Defects as Perturbed Bond Angles in a Lattice

Consider the two simplest cases in figure 1.5: defect type A; where the aluminium–oxygen bond distance is shortened, forcing the oxygen to occupy two off axis positions, and defect type B; where the opposite occurs: the aluminium–oxygen bond distance is lengthened, allowing two preferred oxygen positions on axis.

These two defect types can be modelled by solving our system Hamiltonian (3.1) for three atoms: an oxygen with two aluminium atoms at a lattice coordinate apart, displaced toward and away from the oxygen. For example, corundum has an Al–O bond distance of $\sim$1.85 Å. If we define the oxygen position to be at an origin, the aluminium atoms can be considered as pairs ($x = -X,\ +X$ = $\{\pm 1.85$ Å$\}$) lying on a cardinal axes; which are identified henceforth as $|X|$ or simply the *defect pair*. Displacing $|X|$ equidistantly from this origin (*i.e.* moving away from optimal crystalline configuration) will yield either an A or B type defect, depending on the direction of displacement.

An eigenspectrum of the six lowest energy levels of this system over a continuum of values in $|X|$ is depicted in figure 4.1. Each energy is measured relative to the ground state, which shows two particular regions where E_{01} (green, solid line) is degenerate (labelled sections A and B: both associated with the

respective defect type). There exists a third (an)harmonic region (section C), which reaches a harmonic state at a separation distance of $|X| \sim 1.85$ Å: the optimal corundum Al–O bond distance. At this distance the spatial harmonic approximation holds and the oxygen can be considered to be localised.

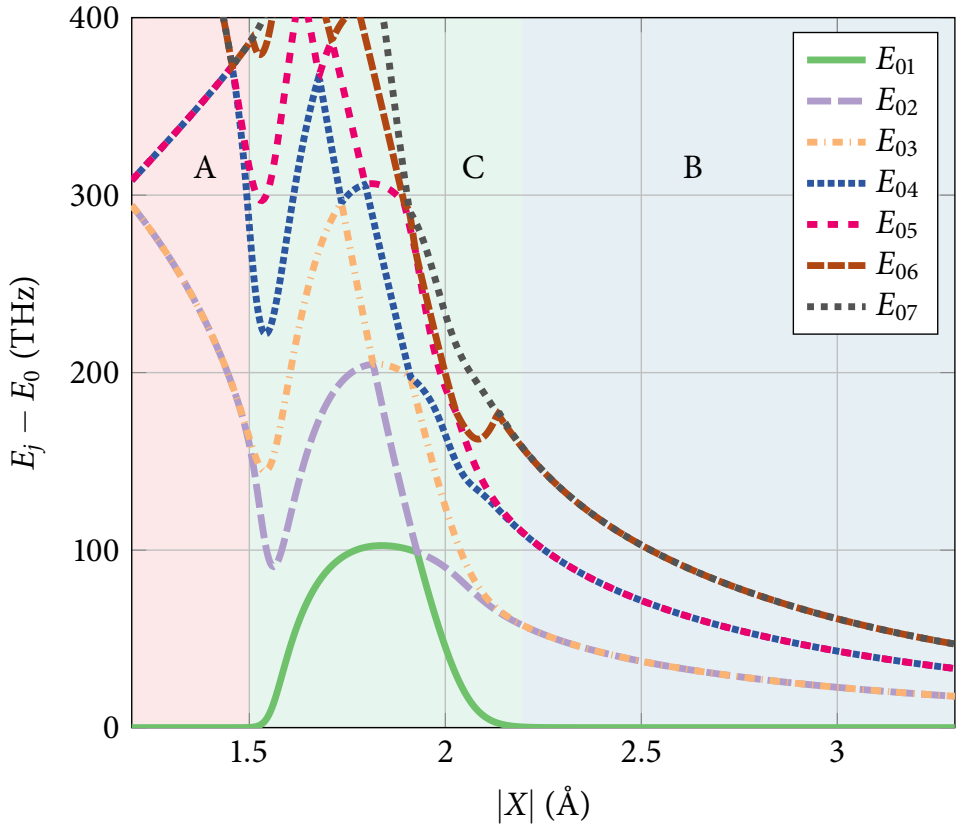

Figure 4.1— Eigenspectrum of a three atom system Al–O–Al, over a varying distance separation. Each excited state has been normalised with the ground state, which shows two regions with a degeneracy at the lowest level. Section A is indicative of the A type defect (figure 4.2), section B of the B type (figure 4.3). An (an)harmonic crossover point is also extant, labelled as section C, which is approximately centered about the optimal Al–O bond distance of corundum (1.85 Å).

To investigate the potential landscape and the resultant oxygen wavefunction in sections A and B of figure 4.1, we choose two separation distances: $|X| = 1.5$ Å (figure 4.2), which lies in the A type defect section, and $|X| = 2.2$ Å (figure 4.3), which exists in the B type defect section. Although the Al–O–Al chain is arranged in a line, we consider delocalisation of the oxygen in two dimensions so that both defect types can be identified on a continuum separation in $|X|$ (as figure 4.1 depicts) rather than using two separate coordinate systems.

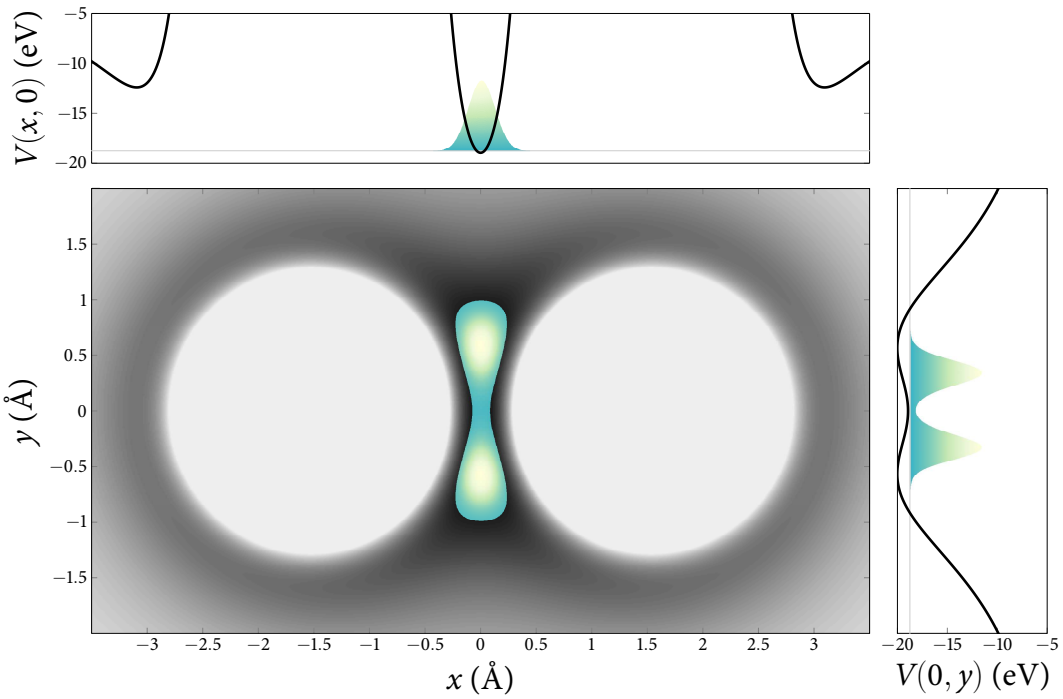

FIGURE 4.2— A top down view of the potential seen by an oxygen atom with two aluminium atoms closer than an appropriate lattice distance (truncated at 0 eV [white], potential minimum [black]). The ground state eigenmode of the oxygen can be seen in the center (green/yellow). This configuration has a separation distance of $|X| = 1.5$ Å and is representative of an A type defect (see figure 1.5). 1D potential values and projected wavefunctions, which are anchored to the ground state energy — are plotted in the outsets to better indicate the depth of the potential well.

Both figures show a top down view of the potential exerted on the oxygen by the two aluminium atoms (gray), whose positions lie at the centre of the circles. We truncate potential energy values above 0 eV for clarity, where energy is measured relative to the zero point set by the Streitz-Mintmire potential's electronegativity correction. The ground state wavefunction (green/yellow) indicates the spatial probability of the oxygen atom in these configurations. It can be seen in figure 4.3, where the pair is far apart, two local minima exist in the form of rings around the base of each aluminium position.

Equivalences between the A and B type defects illustrated in figure 1.5 are clearly visible: figure 4.2 showing a shortened bond length, causing an oxygen dipole perpendicular to the bond axis, and figure 4.3 depicting a lengthened bond and an oxygen dipole parallel to the bond axis.

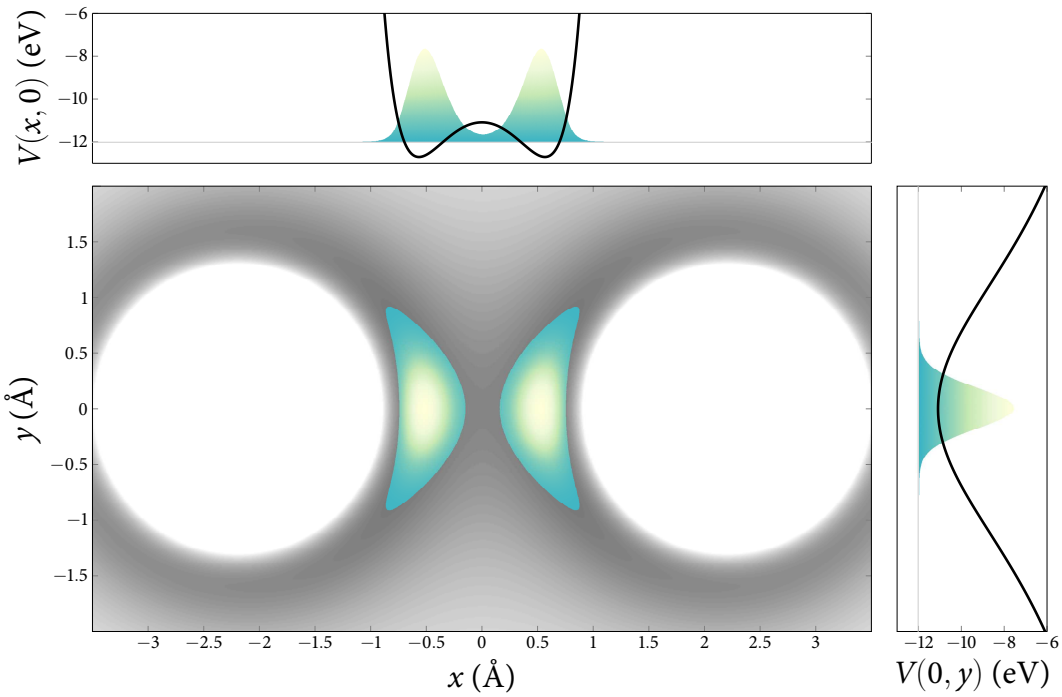

FIGURE 4.3— A top down view of the potential seen by an oxygen atom with two aluminium atoms further apart than an appropriate lattice distance (truncated at 0 eV [white], potential minimum [gray]). The ground state eigenmode of the oxygen can be seen in the center (green/yellow). This configuration has a separation distance of $|X| = 2.2$ Å and is representative of a B type defect (see figure 1.5). 1D potential values and projected wavefunctions, which are anchored to the ground state energy — are plotted in the outsets to better indicate the depth of the potential well.

The functional form of the 1D potential (outset axes of figures 4.2 and 4.3, —; either $V(x,0)$ or $V(0,y)$) in the direction of the defect is a double well. Also depicted in these outsets are projected wavefunctions, which have been scaled for visual purposes, but anchored at the ground state energy — . This representation is equivalent to the standard two level physics depiction of the TLS, and as such, potential offsets ε and tunneling matrix elements Δ (see subsection 1.4.1) of our model may be estimated in this limit.

4.3 TLS Defect Confined in Two Dimensions

The ideal case discussed in section 4.2 ignores many real world complications, in particular any potential constraints from nearest neighbour atoms that undoubtedly surround the TLS. In an attempt to add complexity to the model

gradually, we start with two additional aluminium atoms on the same plane as figures 4.2 and 4.3, confining the defect in the $|Y|$ direction (*i.e.* $y = -Y, +Y$).

Throughout the discussion we use the term *defect pair* to refer to the pair of aluminium atoms which provide the dominant contribution to the potential felt be the oxygen. The term *confining pair* in contrast is used to refer to additional neighbouring aluminium atoms which also apply major influence to perturb the symmetries of the potential.

4.4 Classifying Eigenspectrum Dynamics

As the values of $|X|$ or $|Y|$ are altered, a complex interplay between the excited states of the model can result. Simple two-level degeneracy and harmonic states are no longer the only possibilities. To interpret what is occurring in a certain domain, we define a metric using the ground and four lowest excited state energies

$$\xi = \frac{E_1 - E_0}{E_2 - E_0} + \frac{E_3 - E_0}{E_4 - E_0}. \tag{4.3}$$

This metric ranges from 0 to 2 and can give a qualitative understanding of the eigenspectrum of the defect. To understand why this is useful, we need first to consider the new symmetries which are possible in a 2D potential.

To begin we generate a phase space diagram where ξ is plotted as a function of the distance to the confining aluminium atoms $(|X|, |Y|)$. Each phase diagram is split into at least four domains, where the properties of these domains can be explained through the interplay of potential configuration and dipole alignment (discussed in section 5.2). Focusing for now on the influence of potential shape, the 2D potential can be approximated as two 1D potentials: one projected in the x direction and the other along y. A simple harmonic potential in 2D is quadratic in both x and y. However, for TLS physics there are two relevant configurations: a set of two double wells (tetra-well) or a set of a double and harmonic well (hemi-tetra-well); which are both illustrated in figure 4.4. It is clear from the outset potential projections of figures 4.2 and 4.3 that both A and B type defects reside in hemi-tetra-wells.

The ξ metric is capable of identifying the tetra- $(\xi = 0)$ and hemi-tetra- $(\xi = 1/2)$ domains, as well as harmonic $(\xi = 5/4)$ and unique ground state (rotationally symmetric, Mexican hat-like) $(\xi = 2)$ regimes and finally the location of bifurcations or transitions $(\xi = 1)$. Figure 4.5 shows the corresponding

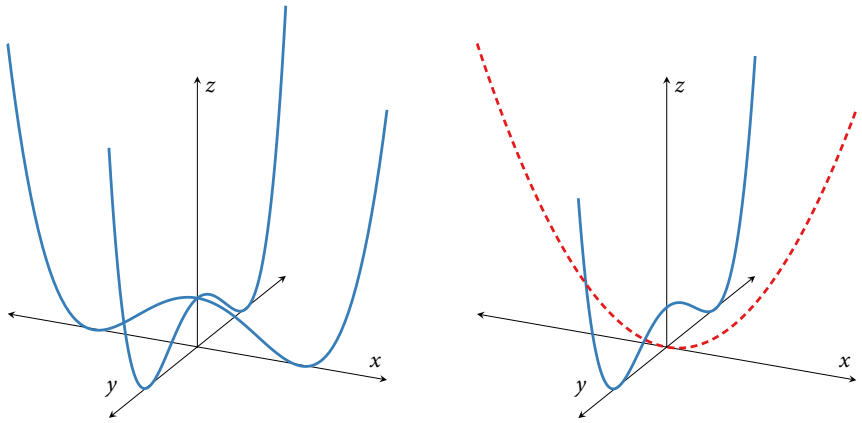

Figure 4.4— 1D double wells —— and harmonic wells - - - can be used to represent simple projections of a 2D potential onto the x and y axes. Left: two projected double wells is an example of a tetra-well. Right: a combination of one double well and a harmonic well reflects the hemi-tetra- case.

layouts of each interplay. It is worth noting that $\xi = 3/2$ can also be considered harmonic for the lowest three levels.

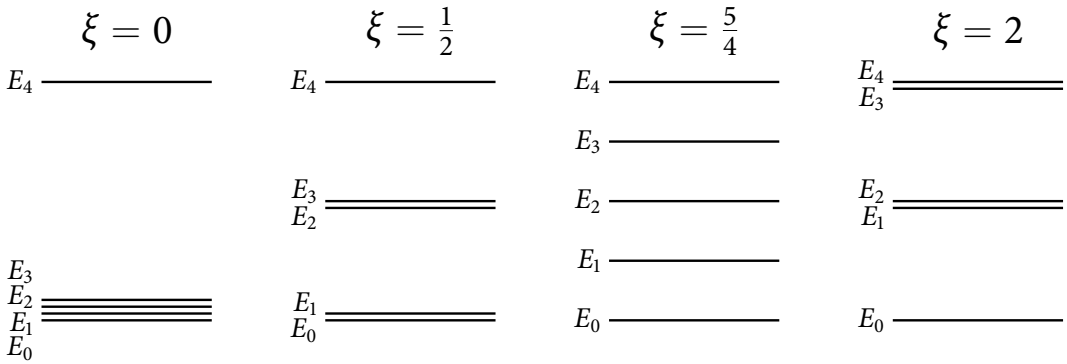

Figure 4.5— Energy level representation of the lowest five eigenenergies of a candidate defect and their associated ξ value. $\xi = 0$: the tetra-well domain, $\xi = 1/2$: the hemi-tetra-well domain, $\xi = 5/4$: the harmonic region and $\xi = 2$: the unique ground state region. The bifurcation/transition region, $\xi = 1$ is not clearly defined thus an example is not shown here.

4.5 Visualising Phase Space and Identifying TLSs

Using this ξ metric and varying the values of both $|X|$ and $|Y|$ generates figure 4.6, a phase map of the interplay of low energy states of the oxygen atom confined in two dimensions by aluminium atoms. The x and y axes show the

$|X|$ and $|Y|$ pair separation distances respectively over a range of $1.85-4$ Å and the phase colour indicates which ξ region a particular configuration exists in.

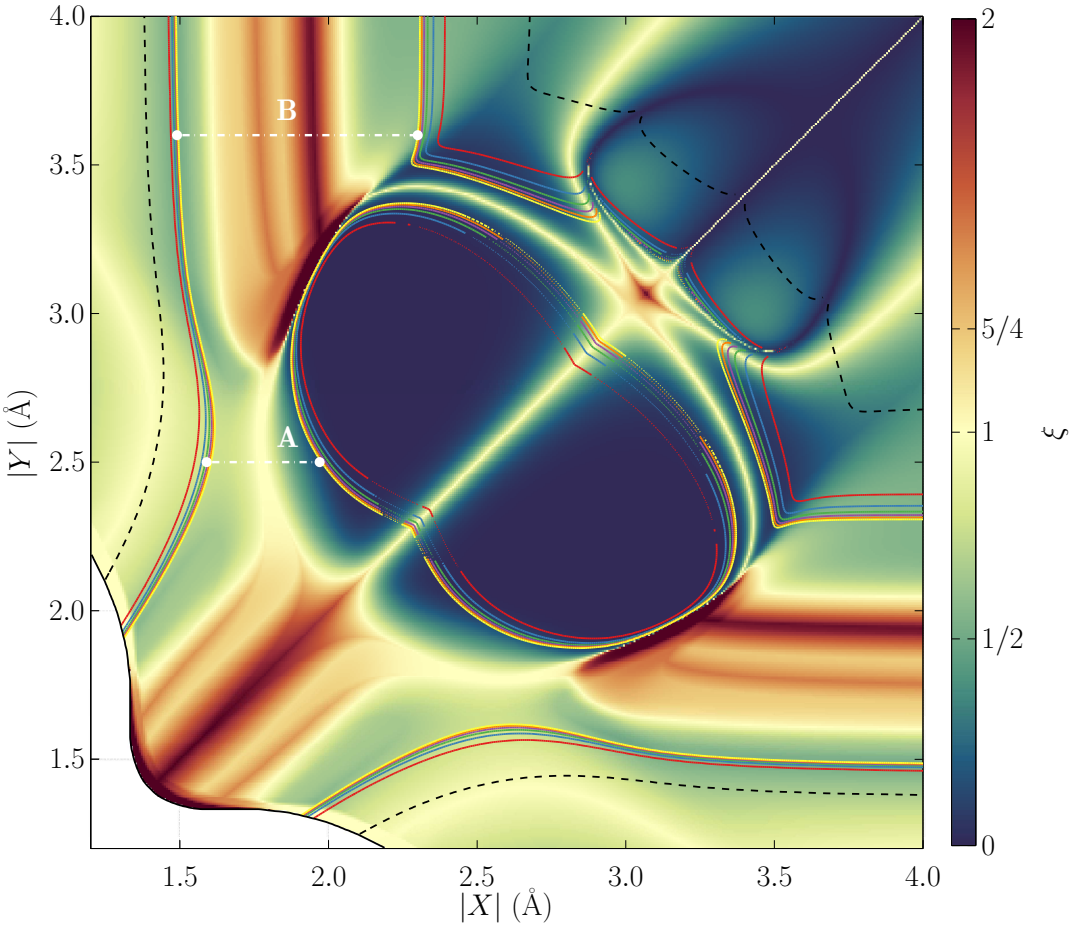

FIGURE 4.6— Map of the ξ metric of the delocalized oxygen 2D model. The $|X|$ and $|Y|$ axes represent aluminium pair position separations. - - - represents a minimum resolvable energy splitting of 10 kHz. The white (blank) section indicates where the aluminium atoms are so close, the oxygen confinement region no longer exists. Overlayed contour lines corresponding to $E_{01} = 0.5-10$ GHz (red to yellow) are comparable to existing experimental qubit results. Cases where quad-degeneracy exists are denoted as dotted rather than solid contours. Two traces (white, dash-dotted lines) are also depicted. The first, at $|Y| = 2.5$ Å (labelled A) is plotted in figure 4.7. In contrast, the trace at $|Y| = 3.6$ Å (labelled B) yields an equivalent eigenspectrum to the 1D case in figure 4.1.

Regions that are deemed to be unimportant when searching for TLS behaviour are those where $\xi \geq 1$, as these correspond to eigenspectra which do not possess a doubly degenerate ground state.

As discussed in section 4.2, in the one dimensional confinement case, both A and B type defects exist in a hemi-tetra-well ($\xi = 1/2$). This is also the

case for the two dimensional landscape—contour lines corresponding to $E_{01} = \{0.5, 2, 4, 6, 8, 10\}$ GHz (red–yellow) on figure 4.6 show the configuration properties that result in TLS E_{01} splittings in qubit architectures described in subsection 1.5.2. For example: anywhere an orange, 8 GHz line exists on this phase map, the $|X|, |Y|$ coordinates of the model generate a cluster configuration that yields the same E_{01} observed in phase qubit experiments [81].

Many of these lines lie completely within a $\xi = 1/2$ (green) region. Consider the contour set on the left of figure 4.6 where $|X| \simeq 1.5$ Å: Section 4.2 states this distance is indicative of an A type defect. The addition of the $|Y|$ aluminium pair does little to perturb the potential at larger distances ($|Y| \gtrsim 3$ Å) and causes only slight deformations at smaller distances ($2 \lesssim |Y| \lesssim 3$ Å). As the phase space is symmetric about $|X| = |Y|$, A type defects exist at the bottom of the plot where $|Y| \simeq 1.5$ Å as well.

B type defects also exist in a $\xi = 1/2$ region, when $|X|$ or $|Y| \simeq 2.4$ Å. The orthogonal pair separation distance is at least 1 Å larger than the defect pair in these configurations and therefore have no bearing on the potential seen by the oxygen. Complications arise when the orthogonal pair is closer and begins to confine the defect. To explain this response, a better understanding of the domains of the map is required.

4.6 Analysis of Phase Space Domains

The case of $\xi = 0$ is particularly interesting in this scenario: a tetra-well region, causing a quad degeneracy in the ground state. In this region, the characteristics of a B type defect are strongly modified. To understand why, we first explain the large rounded $\xi = 0$ domains in the centre of figure 4.6's phase space.

Tetra-well domains exist when the confinement potential acting on the oxygen consists of two double wells (see left plot in figure 4.4). This phenomena emerges when both the defect pair and the confining pair of atoms are close enough to interact. Consider an A type defect with a defect pair in $|X|$ and a confining pair at a constant value of $|Y| = 2.5$ Å. If $|X|$ is varied from an initial value of 1.6 Å to 2 Å, this extension spans phase space from one point where $E_{01} = 8$ GHz to another. This separation is traced on figure 4.6 (white dash-dotted line labelled A), where the defect visibly moves from a hemi-tetra-regime ($\xi = 1/2$) and goes through a transition region before reaching the

tetra-well regime ($\xi = 0$). Figure 4.7 depicts the eigenspectrum of this trace. This is in contrast to larger confining pair values such as $|Y| = 3.6$ Å (also traced on figure 4.6: white dash-dotted line labelled B) where the eiginspectrum is largely unchanged from the one dimensional case presented in figure 4.1.

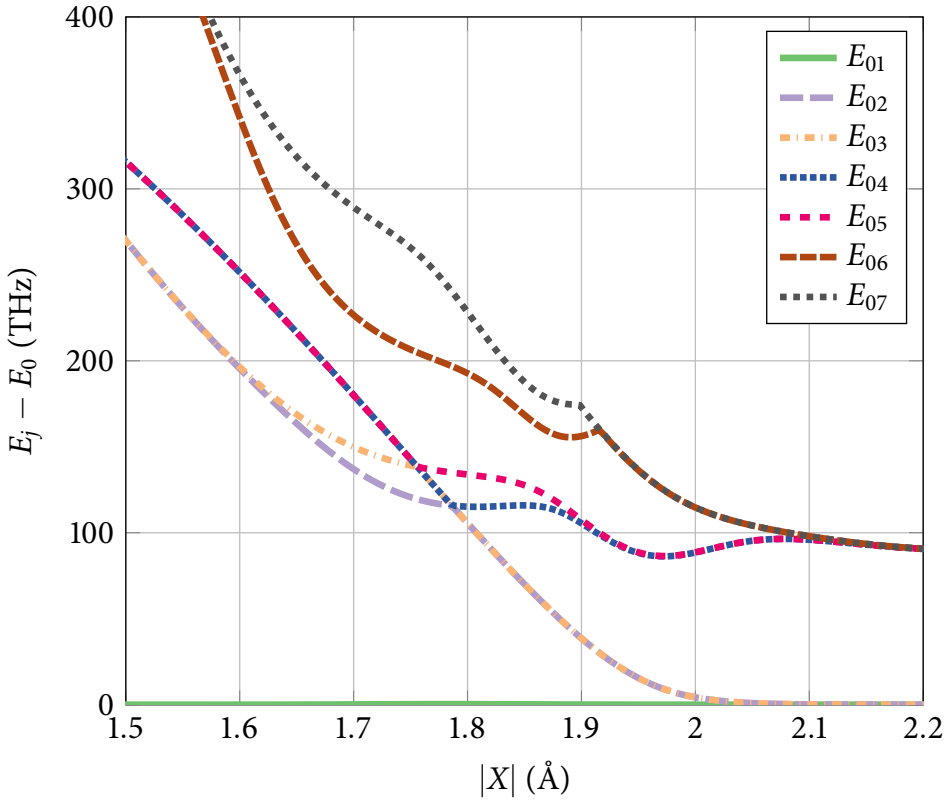

FIGURE 4.7— Eigenspectrum of a 2D TLS: an oxygen atom caged with four aluminum atoms. Each excited state has been normalised with the ground state. The confinment pair $|Y|$ is held at a constant distance separation of 2.5 Å and the $|X|$ range shows how higher eigenvalues behave as the confinement atoms force the defect into a tetra-well regime. Figure 4.6 depicts this data in terms of ξ (white dash-dotted trace labelled A).

As the $|X|$ separation distance is increased, the effect on E_{01} is negligible on the scale of the higher energy levels (although differs greatly on the TLS splitting energy scale). However, the degenerate pair E_{23} (a degenerate second and third excited state pair) rapidly shifts from a level much higher than ground to degenerate at ground. Complete quad-degeneracy is not always extant in the tera-well regime, configurations of $|X|$ and $|Y|$ in these domains usually have two degenerate pairs which are still observable, akin to the hemi-tetra-

domains, with the difference *between* the pairs approaching the difference *of* the pairs: $E_{01} = E_{23} \approx E_{12}$ ($1.95 \lesssim |X| \leq 2$ Å in figure 4.7 for example). In other words, when the ground to first excited and second to third excited state differences are equal, the first to second excited state difference trends from a much larger value to one effectively equivalent to the aforementioned pairs in this region.

When considering the influence of higher lying excited states, it's important to keep the fundamental energy scales of the problem in mind. In typical qubit experiments, the superconducting properties of the device put a rigorous upper limit on TLS frequencies of interest. At frequencies greater than approximately 100 GHz, there is enough energy to dissociate Cooper-pairs and therefore it can be viewed as an operational upper bound for Josephson junction devices (typical operating frequencies are however device specific, and are much lower in practice). For much of the tetra-well domain $E_{12} \geq 100$ GHz and consequently can effectively be ignored, the system can be considered as a two-level defect even in this quad degeneracy domain.

The B type defects' sudden behaviour change as $|X| \rightarrow |Y|$ is also caused by this response (see figure 4.6). Confinement pairs start interacting with the defect, causing E_{23} to approach the value of E_{01} (again, $1.95 \lesssim |X| \leq 2$ Å in figure 4.7 depicts this phenomena). The ξ metric does not clearly differentiate between two degenerate pairs that are marginally separated and two pairs that are actually degenerate.

If however, $E_{12} < 100$ GHz, higher lying energy states must still be considered and the model exhibits true quad degenerate behaviour. Regions in which this occurs are denoted in figure 4.6 as dotted contours, which over the entirety of phase space are extremely rare—suggesting a reason why quad-level systems are yet to be experimentally observed.

The final domain yet to be discussed on figure 4.6 is the upper right hand corner where both $|X|$ and $|Y|$ are large. This region is tetra-well dominated but can be considered as a region where the TLS model breaks down. Each of the four potential minima exist localised about the four confining aluminium atoms and as such should not be considered as TLS candidates.

4.7 Chapter Summary

The number of possible atomic configurations in an amorphous structure is immense. Even within an ultra-thin barrier such as the JJ oxide layer, the state

space is still vast. Before investigating even a small subset of these possibilities, this chapter shows that analysis of local variations at the nearest neighbour level provides an abundance of useful information an oxygens properties.

Understanding how the eigenspectrum of an oxygen atom responds in the presence of an external potential in this comprehensive manner allows us, in chapter 5, to compare calculated TLS properties from this model to measurements obtained from experiments.

“子曰，視其所以，觀其所由，察其所安。人焉廋哉，人焉廋哉。”

——孔夫子

5 Phase Space and Experimental Comparisons

Identifying possible TLS candidates by comparing model results to experimental coupling variables such as the ground to first excited state energy splitting and dipole moment.

To understand the properties of a delocalised oxygen, we have considered confining aluminium atoms in both a line and in a plane. In reality, aluminium atoms will surround the oxygen in all three dimensions.

This chapter expands on the 2D model discussed in section 4.3, adding an additional aluminium pair in section 5.1, which extends oxygen confinement into three dimensions (with six aluminum atoms), whilst still projecting oxygen delocalisation on a plane. Although in general an oxygen can delocalise in three dimensions, for this investigation we focus on an effective 2+1D model, minimising both computational and descriptive complexity while still modelling the relevant behaviour of the system. Chapter 6 will consider possible ambiguities to this approach when juxtaposing it with a 3D implementation. The following sections then apply the 2+1D model and compare delocalised oxygen responses to experimental TLS data. Section 5.3 discusses qubit–TLS coupling and section 5.4 observes the effect of mechanical strain.

5.1 TLS Defect Confined in Three Dimensions

Two more confining atoms are now added into the system in the z direction labelled as $|Z|$. Interactions with these atoms in the third dimension are now considered, although the model is still two dimensional (*i.e.* 2+1D); thus oxygen continues to be confined to the xy plane. An illustration of this cluster configuration is displayed in figure 5.1 and representations of the cage potential and oxygen wavefunctions are shown in figure 5.2 for A and figure 5.3 for B type defects.

The selection of a fixed $|Z|$ distance changes the phase landscape in a manner that can be qualitatively extrapolated between two arbitrary values even a few angstroms apart. Values of $|Z| = 2.75$ Å (figure 5.4) and $|Z| = 2.25$ Å (figure 5.5) were chosen to analyse in detail.

TLS behaviour can be observed on both maps and each value of $|Z|$ has been selected based on model parameters. Oxygen confinement occurs for $|Z|$ values lower than 2.25 Å (*i.e.* $E_{01} \gg 100$ GHz). $|Z|$ values larger than 2.75 Å show

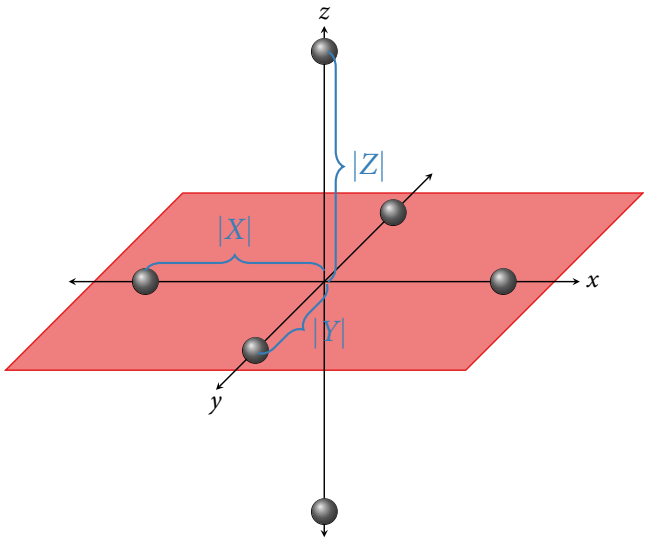

FIGURE 5.1— Illustration of an idealised cluster producing a void in 3D. Six cage aluminium atoms ● sit in pairs on each cardinal axis, equidistant from the origin. Separation distances $|X|$ and $|Y|$ are labeled for reference (see text). The plane ▭ at $z = 0$ is a representation of the 2D delocalisation of an oxygen atom.

similar phase behaviour to that of figure 4.6, which is completely unbound in z. Large $|Z|$ separation distances also decrease the validity of the 2+1D model, in addition: the radial distribution analysis in section 2.4 suggests large separation distances for nearest neighbour atoms have a low probability of occurrence.

As the pair separation distance $|Z|$ decreases, the tetra-well ($\xi = 0$) regimes diminish in size and no longer exhibit TLS behaviour. This suggests that quad-degenerate defects, whilst quite rare in phase space already, are extremely rare in reality. For one to exist in a junction, the amorphous layer would have to be disordered in such a way that an oxygen atom's nearest neighbour atom pair exists at a distance greater that ~ 3 Å along one orthogonal basis vector.

Whilst this configuration of six cage aluminium atoms on cardinal axes around a central oxygen is still an idealised system, figure 5.5 is confined in all three spatial dimensions with pragmatic distances and is therefore considered to be the most 'realistic' representation of the TLS phase space for the 2+1D model. TLS candidates lie well within hemi-tetra- ($\xi = 1/2$) domains, are not mired by higher energy level complexities (*i.e.* quad-degeneracies) and can be identified as A type and B type defects separated by an (an)harmonic boundary.

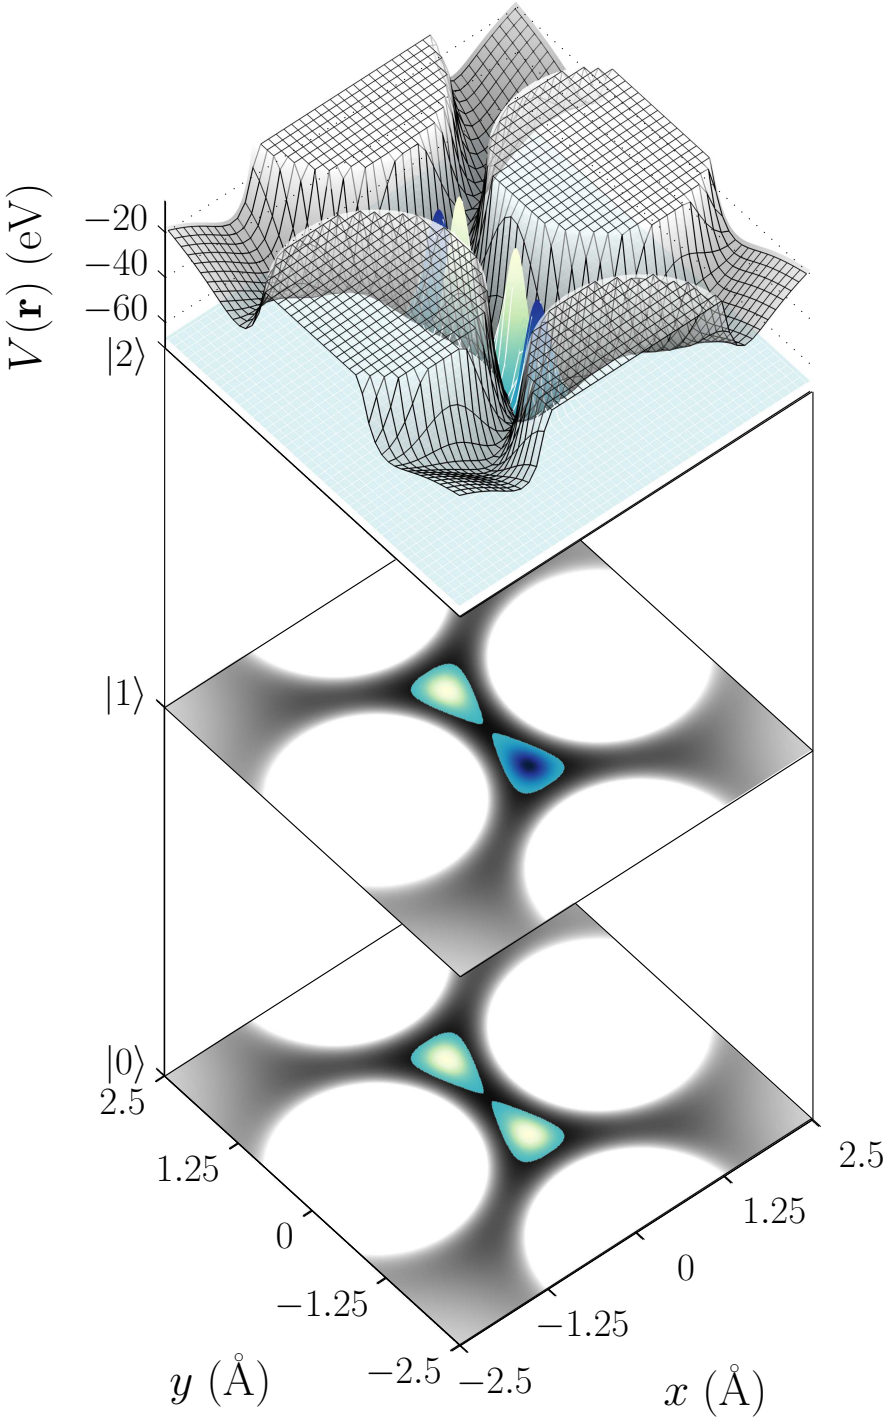

Figure 5.2— Cage potential and the lowest three eigenfunctions of a cluster with the values $|X| = 1.53$, $|Y| = 2.52$, $|Z| = 2.25$ Å. The top image in the stack is presented with the apparent 'depth' of the potential well on the z-axis and the second excited state $|2\rangle$ scaled accordingly. Ground $|0\rangle$ and first excited $|1\rangle$ states are displayed in a projected representation underneath. This cluster configuration is representative of an A type defect (see figure 4.2); with a ground to first excited state splitting of $E_{01} = 8.1$ GHz and a ground to second excited state splitting of $E_{02} = 202.2$ THz.

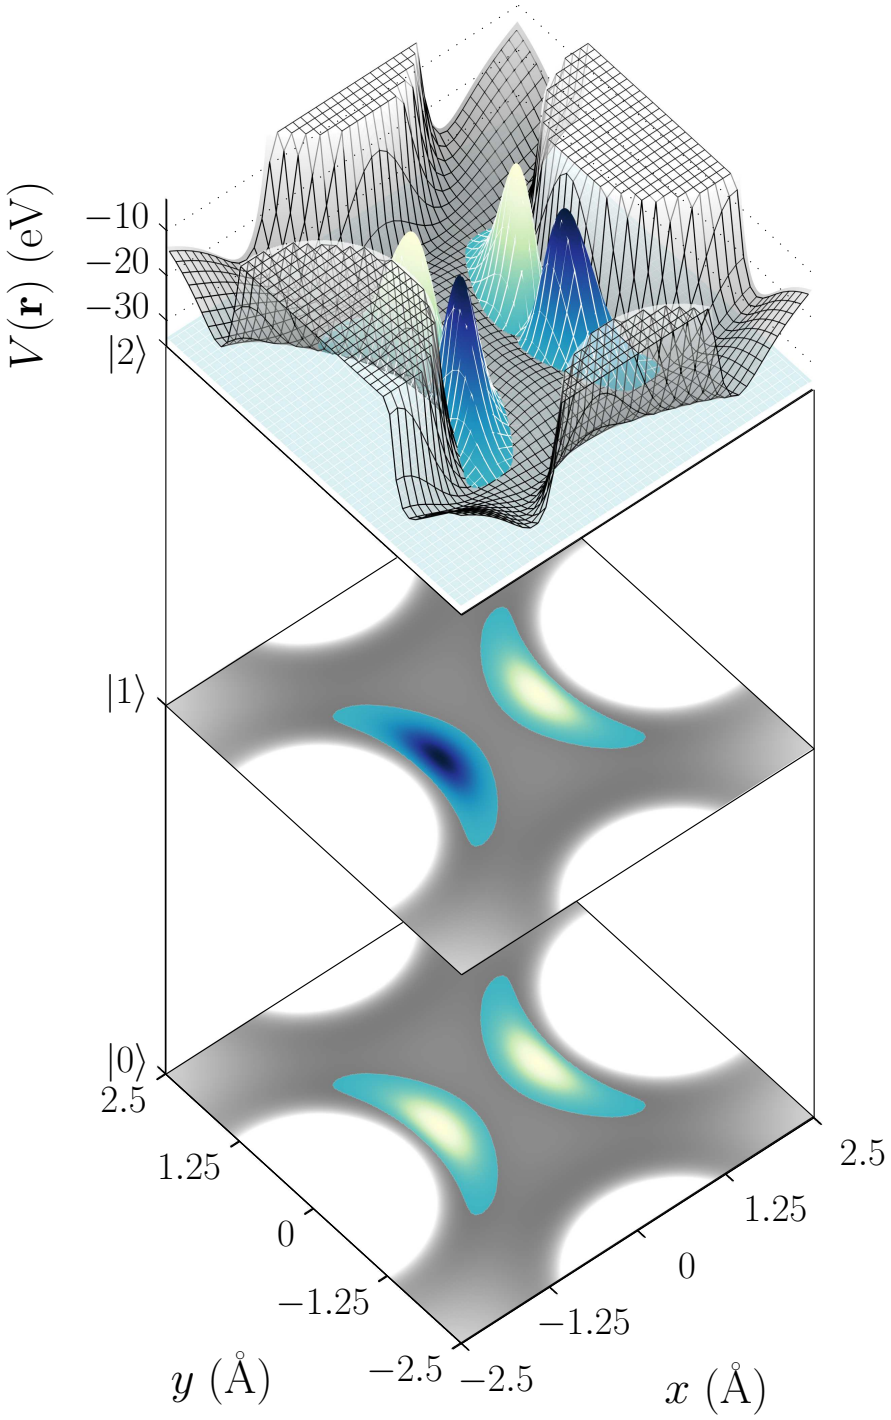

FIGURE 5.3— Cage potential and the lowest three eigenfunctions of a cluster with the values $|X| = 2.38$, $|Y| = 3.19$, $|Z| = 2.75$ Å. The top image in the stack is presented with the apparent 'depth' of the potential well on the z-axis and the second excited state $|2\rangle$ scaled accordingly. Ground $|0\rangle$ and first excited $|1\rangle$ states are displayed in a projected representation underneath. This cluster configuration is representative of a B type defect (see Figure 4.3); with a ground to first excited state splitting of $E_{01} = 8.4$ GHz and a ground to second excited state splitting of $E_{02} = 36.3$ THz.

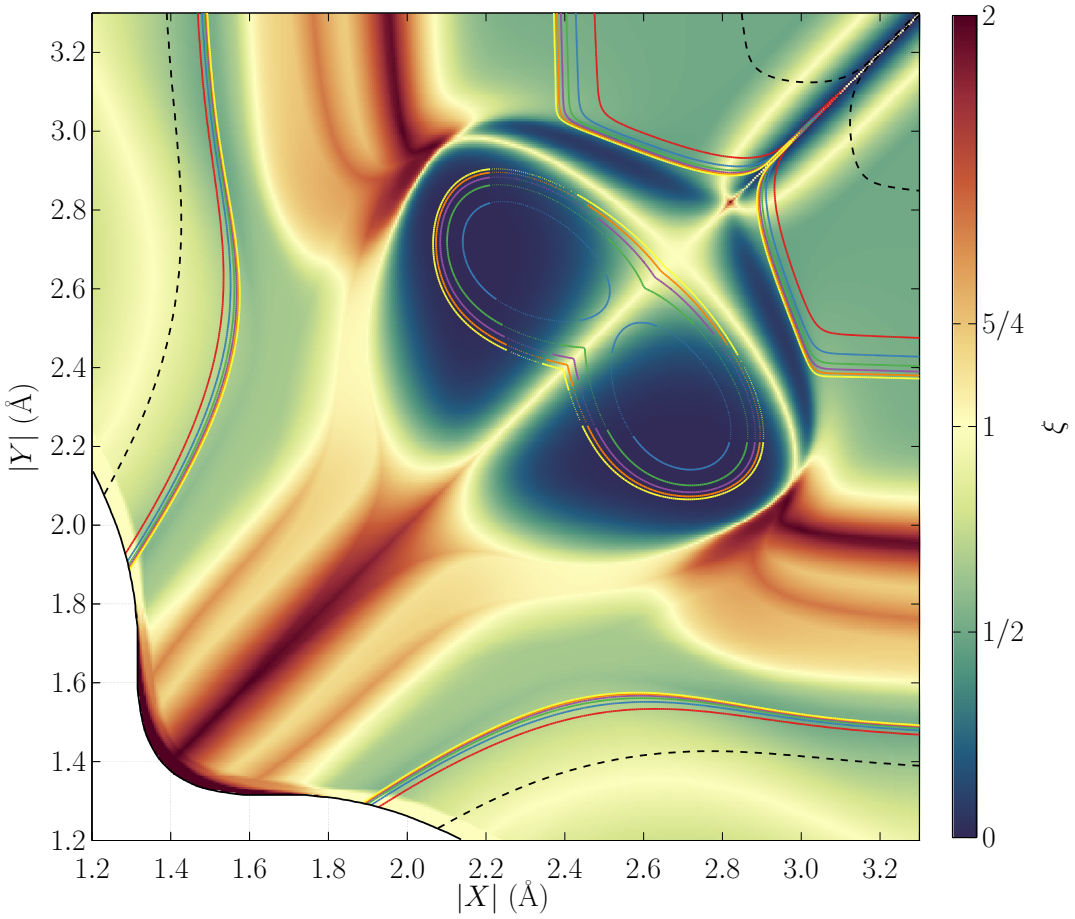

Figure 5.4— Map of the ξ metric of the delocalized oxygen model confined in two dimensions. The $|X|$ and $|Y|$ axes represent aluminium pair positions with $|Z| = 2.75$ Å. **- - -** represents a minimum resolvable energy splitting of 10 kHz. The white (blank) section indicates where the aluminium atoms are so close, the oxygen confinement region no longer exists. Overlayed contour lines corresponding to $E_{01} = 0.5-10$ GHz (red to yellow) are comparable to existing experimental qubit results. Cases where quad-degeneracy exists are denoted as dotted rather than solid contours.

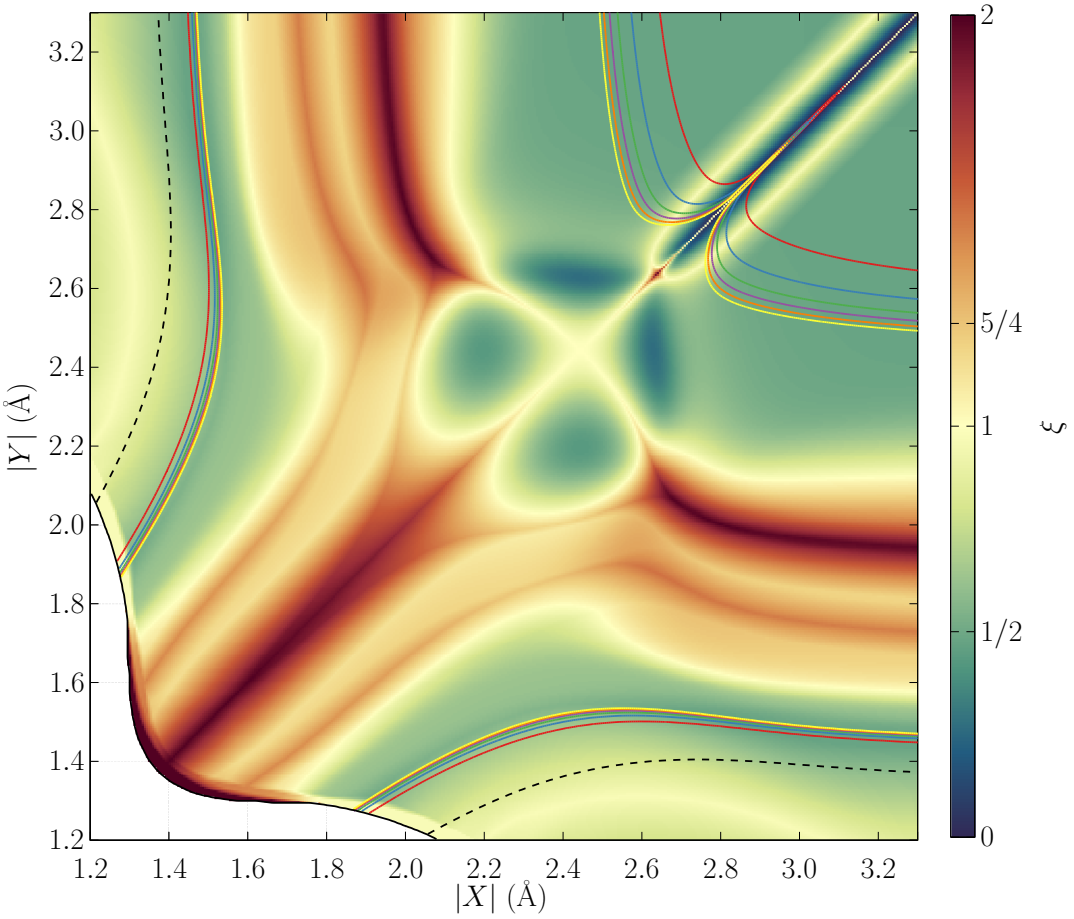

FIGURE 5.5— Map of the ξ metric of the delocalized oxygen model confined in two dimensions. The $|X|$ and $|Y|$ axes represent aluminium pair positions with $|Z| = 2.25$ Å. **- - -** represents a minimum resolvable energy splitting of 10 kHz. The white (blank) section indicates where the aluminium atoms are so close, the oxygen confinement region no longer exists. Overlayed contour lines corresponding to $E_{01} = 0.5-10$ GHz (red to yellow) are comparable to existing experimental qubit results.

Phase space is also dominated on this map with $\xi = 5/4$ (harmonic) and $\xi = 2$ (unique ground state) domains where the oxygen atom can be considered under the spatial harmonic approximation (*i.e.* localised). This is significant, as TLS observations in experiments are not statistically dense. We expect few defects in the junction compared to the number of atoms extant.

5.2 Charge Dipoles

As stated in subsection 1.5.2, another important experimentally measurable property of the TLS is its strong electric dipole moment. Using the same $|X|, |Y|$ and $|Z|$ parameters from the phase maps in the previous section, the dipole elements in the x direction $\wp_x$, and the y direction $\wp_y$ can be calculated via (4.2). Figure 5.6 shows the same phase space as figure 5.4, where $|Z| = 2.75$ Å, and figure 5.7 matches figure 5.5, where $|Z| = 2.25$ Å. The colourmap for both figures now represent dipole strengths of each cluster configuration rather than energy level splittings (represented through the ξ metric). These computed dipole moments correspond well to observed values, assuming $\mathcal{O}$(nm) junction widths [74, 81].

The dipole moments are presented as $\left(|\wp_y| - |\wp_x|\right)/e$ rather than separate plots because the dipole elements are discontinuous at the bifurcation points (*i.e.* when $|\wp_x| > 0, |\wp_y| = 0$ and *vice versa*). Comparing these plots to the phase maps in section 5.1, it is apparent that only the tetra- and hemi-tetra- domains ($\xi < 1$) exhibit a dipole response—which is appropriate for our model as E_{01} splittings representative of a TLS only appear in these regions. Localised oxygen atoms ($\xi > 1$) are also expected to not elicit dipole behaviour. With this information, the domain boundaries and bifurcations on each phase map can now be fully explained.

Two variables alter the landscape: dipole and potential. Clusters with tight z confinement (those without tetra-well regions such as $|Z| = 2.25$ Å) have four unique regions where a TLS may reside. The dipole direction dominates two of these domains: an A type region when the confining pair is collinear to the dipole, and a B type region when the confining pair is perpendicular. A symmetry bifurcation (at $|X| = |Y|$) separates the dipole domains into four regions which can therefore be identified in terms of dipole moment ($|\wp_x|$ or $|\wp_y|$) and defect type (A or B).

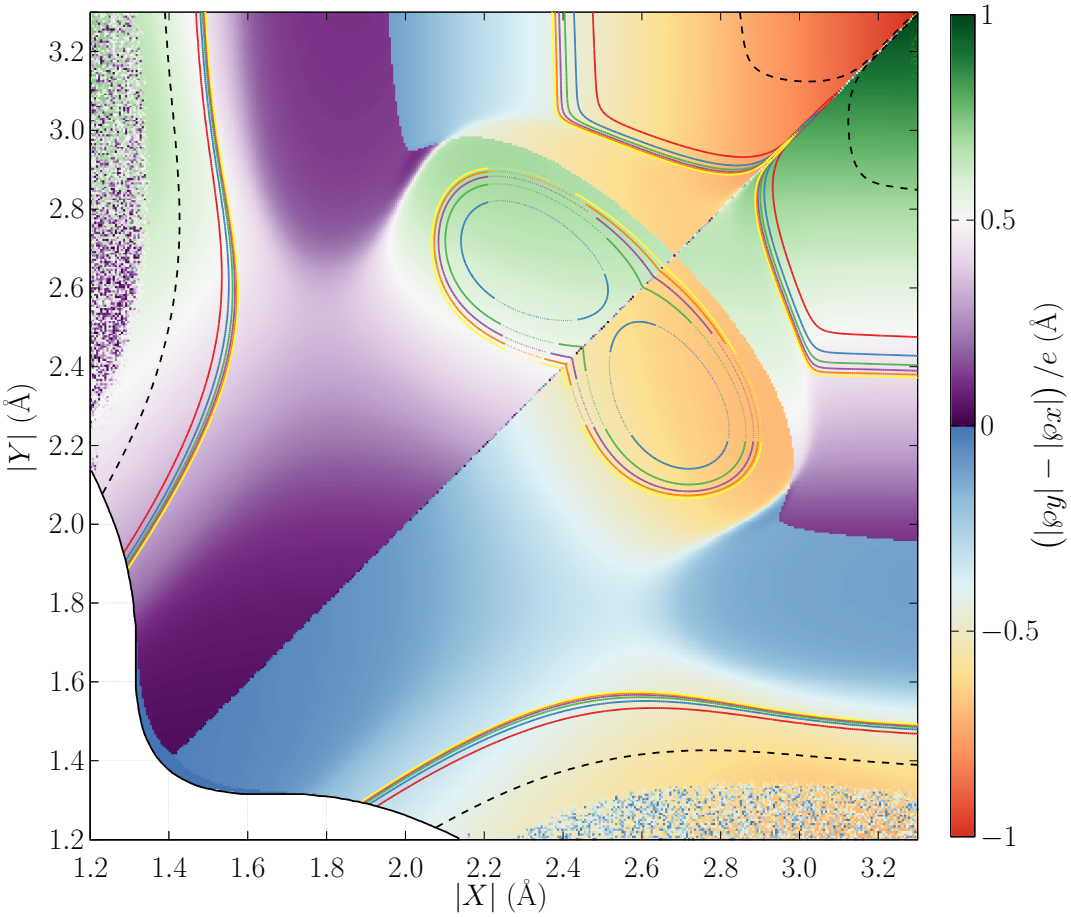

FIGURE 5.6— The difference between the absolute dipole moment (in x- and y-directions) over the same range for $|Z| = 2.75$ Å. We see either $|\wp_x|$ (red) or $|\wp_y|$ (green) dominated behaviour in the tetra- and hemi-tetra- domains but none in other regions.

Clusters without as much z confinement (such as $|Z| = 2.75$ Å) exhibit tetra-well behaviour: generating two additional regions. Tetra-well domains, as discussed in section 4.6, are caused when the confining pair of aluminium atoms start interacting with the defect pair (and hence the oxygen as well). If we consider the A type, $|\wp_y|$ domain in figure 5.6, it is clear that the dominant dipole direction remains constant as $|X|$ separation is in increased and the tetra-well domain is entered. The same $|X|, |Y|$ parameters on figure 5.7 cross a bifurcation line, changing dipole direction and the model indicates B type defect properties. Increased confinement in z induces a deeper potential well in the xy plane and removes any major landscape changing contributions from the confining atom pair—*effectively reducing the model back to a 1D description*. A

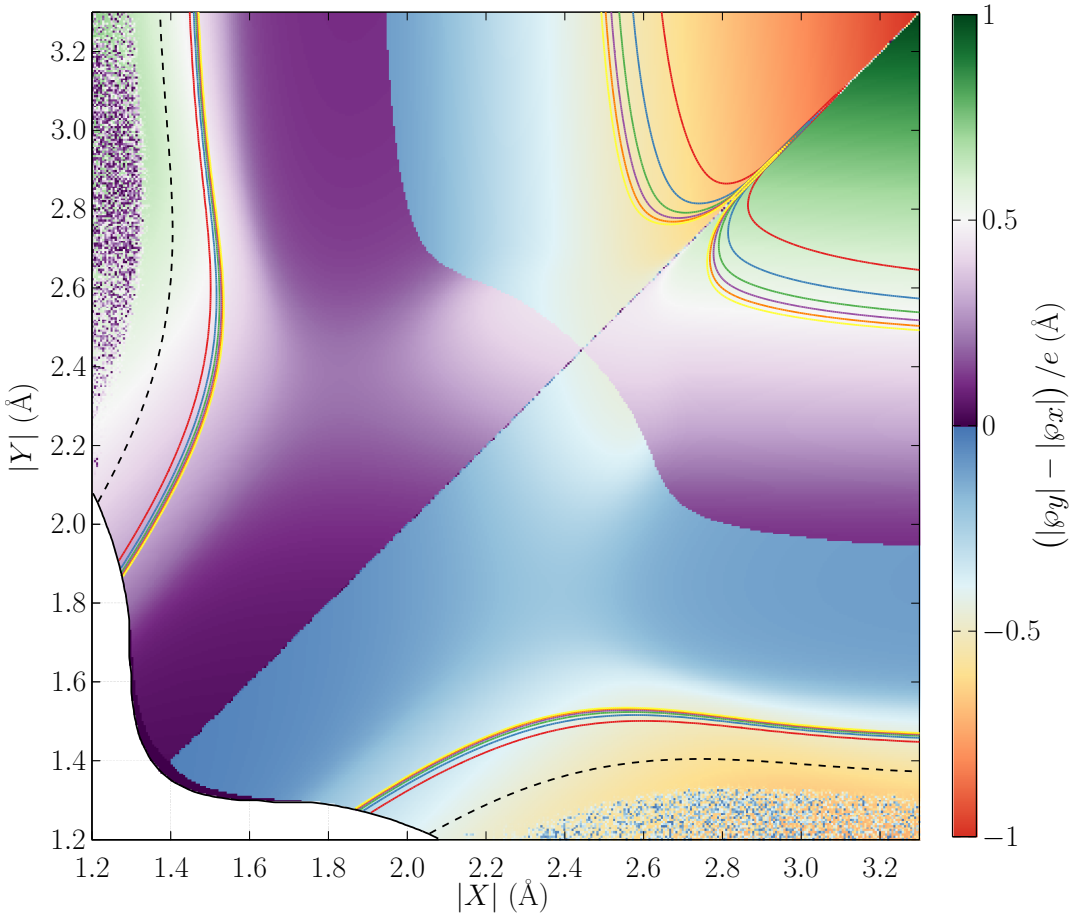

Figure 5.7— The difference between the absolute dipole moment (in x- and y-directions) over the same range for $|Z| = 2.25$ Å. We see either $|\wp_x|$ (red) or $|\wp_y|$ (green) dominated behaviour in the tetra- and hemi-tetra- domains but none in other regions.

cluster with a comparatively shallow potential that generates tetra-well domains does so when conditions are advantageous for an A type confinement pair to become a B type defect pair (a dipole direction change is not required for this to occur). A trace like the $|Y| = 2.5$ Å (labelled A) on figure 4.6 (and an associated eigenspectrum response similar to figure 4.7) is an example of this behaviour. If we do not consider the small portion of quad-degenerate defects in this domain; the measurable properties (*i.e.* E_{01} and dipole strength) tetra-well, B type, $|\wp_y|$ TLSs are identical to B type, $|\wp_y|$ TLSs that reside in a hemi-tetra-well domain.

5.3 Qubit Coupling

To compare our TLS model directly to experiments, we have assumed that our JJ lies within a phase qubit, although the model applies equally for any device comprised of amorphous junctions. The measurable signal of a TLS in a phase qubit is the resonance of the TLS and qubit splitting energy, E_{01}, with the qubit-TLS coupling. For the phase qubit [74], the qubit-TLS coupling S_{max} (1.6) is a function of E_{01} and $\wp$ [155], the effective dipole moment due to an electric field applied in the direction of delocalization. Throughout this discussion we assume a junction width $w = 2$ nm and capacitance $C = 850$ fF.

The dipole magnitudes in figures 5.6 and 5.7 are calculated against the electron charge for simplicity, although as we are discussing an oxygen atom, the effective charge on the atom may in fact be larger. Using our Josephson junction DFT models from chapter 2, we partition the charge density associated with atoms across the lattice into Bader volumes [156]. The charge enclosed within each Bader volume is an adequate approximation to the total electronic charge of an atom. An average value of 1.395 ± 0.006 e is found for oxygen atoms in a junction comprised of $AlO_{1.25}$ at a density of 0.8 times that of corundum. We can use this value such that

$$\widetilde{\wp}_x = 1.4e\wp_x \tag{5.1}$$

(for a dipole in the x direction) to gain a better estimate of S_{max}.

In figure 5.8 we plot contour lines representing constant values of E_{01} corresponding to the purview of qubit resonant frequencies observed experimentally for constant values of $|Z| = 2.25$ Å. The S_{max} (1.6) response to these frequencies is plotted as a function of $|X|$, in which we see maximum coupling strengths which correspond exceptionally well with experimental observations [68, 75, 81]. The most comprehensive of these studies (reference 75) measures S_{max} values of 3–45 MHz.

Whilst the S_{max} response is neither smooth nor singular over the investigated phase space, the value range is notably small. Figure 5.9 shows the range of S_{max} couplings for all $E_{01} = 8$ GHz configurations calculated with confinements in $|Z|$ from 2.25 Å to unbound and $|X|$ (hence $|Y|$ as well from symmetry arguments) from 1.2 to 4 Å. The entire range of S_{max} values is only 60 MHz wide, which suggests an explanation as to why large couplings (of order 500 MHz for example) have not been seen experimentally.

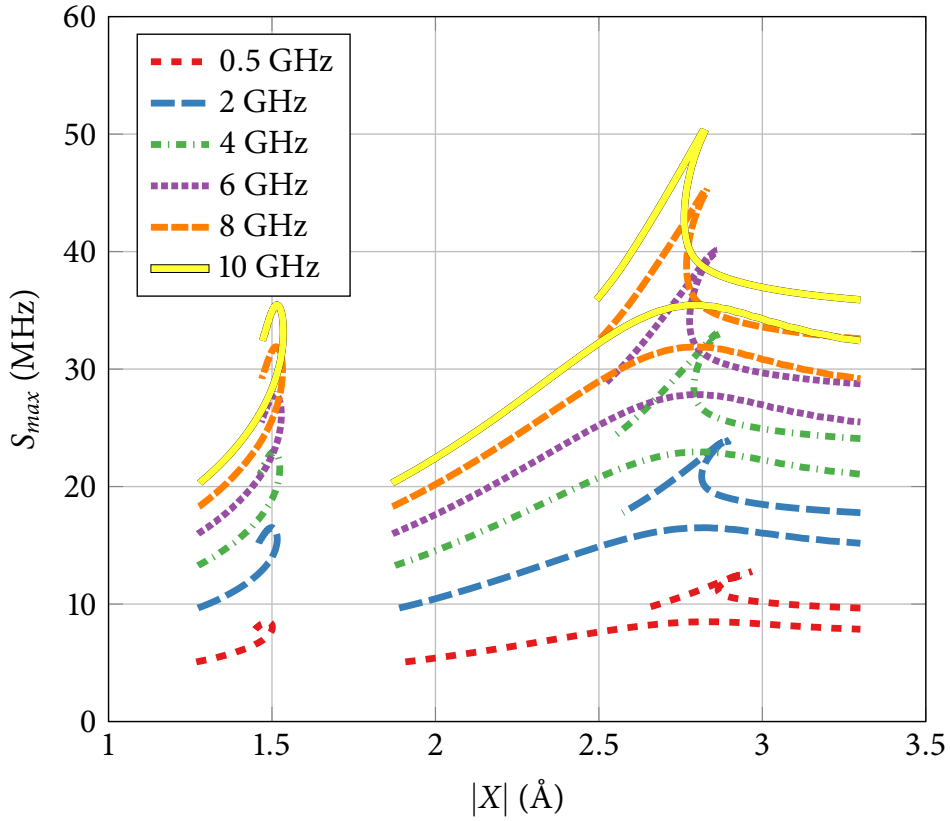

Figure 5.8— Coupling strength to a fictitious phase-qubit S_{max} as a function of $|X|$ in the domains where both dipoles $|\widetilde{\wp}_{x,y}|$ are dominant (see 1.6) for a set of constant E_{01} splitting frequencies and $|Z| = 2.25$ Å.

5.4 Strain Response

Unusually long coherence times of strongly coupled defects are a key observation of TLS-qubit experiments [67, 157]. As our model assumes a charge-neutral defect*, coherence times are linked to the dipole element (for charge noise) and the strain response (for phonons). Mechanical deformation of a phase-qubit has recently been observed directly [77], which we can mimic in our 2+1D model by introducing a series of deformations onto the cluster, and subsequently measure the E_{01} response.

*More specifically: the defect is globally charge neutral, whilst the oxygen possess a net charge.

Deformation types that were tested are depicted in figure 5.10a. All deformations were tested in each of the four hemi-tetra- regions discussed in section 5.1, not only in the x-direction as shown, but also in y. Of the tested deformations we find the response of one (translation of $|X|$, highlighted in figure 5.10a) to be 10^5 times stronger than the others. Such a deformation corresponds to a translation of both aluminium atoms in the same direction and

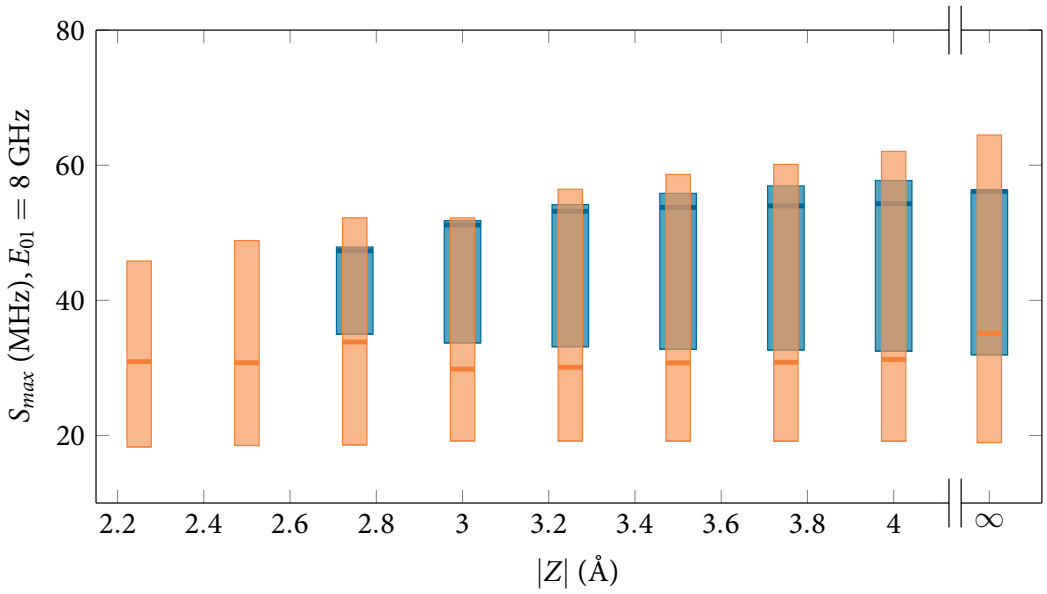

FIGURE 5.9— S_{max} response range for a constant E_{01} splitting frequency of 8 GHz as a function of $|Z|$ separation. Each box on the figure represents the minimum to maximum S_{max} values in both $|\widetilde{\wp}_x|$ and $|\widetilde{\wp}_y|$ dominant domains over the phase space of $|X|$ distance separations. The box at $|Z| = 2.25$ Å represents all $E_{01} = 8$ GHz values on figure 5.8 for example. Median values are plotted as solid lines inside each box. Orange boxes represent degenerate regions and blue boxes indicate quad-degenerate regions.

relative to the oxygen, along the axis of the dominant dipole. This suggests an explanation for the long TLS coherence times, as a delocalized oxygen is only sensitive to a small subset of available phonon modes (as well as coupling to charge noise only through its electric dipole).

Figure 5.10b shows this response for the optical mode (at positions corresponding to phase qubit-TLS coupling strengths ~ 8 GHz where $|Z| = 2.5$ Å), displaying a characteristically hyperbolic response. This is typical of a two-level system in the STM framework: E_{01} is hyperbolically dependent on the asymmetry energy ε (see subsection 1.4.1) [61, 62, 77]. This compares well with the observed strain response in reference 77. Fitting the data in the inset of figure 5.10b using the harmonic approximation (assuming an optical mode), we obtain characteristic frequencies in the range $0.3 - 1.0$ THz. This is well above the energy of typical qubit experiments, again limiting the coupling of the defect to phonons; although it is still possible that slow noise fluctuations couple to the defect via a similar deformation (*vide infra*). Finally, figure 5.10c shows the linear strain gradient plotted along the $E_{01}/h = 8$ GHz contour for

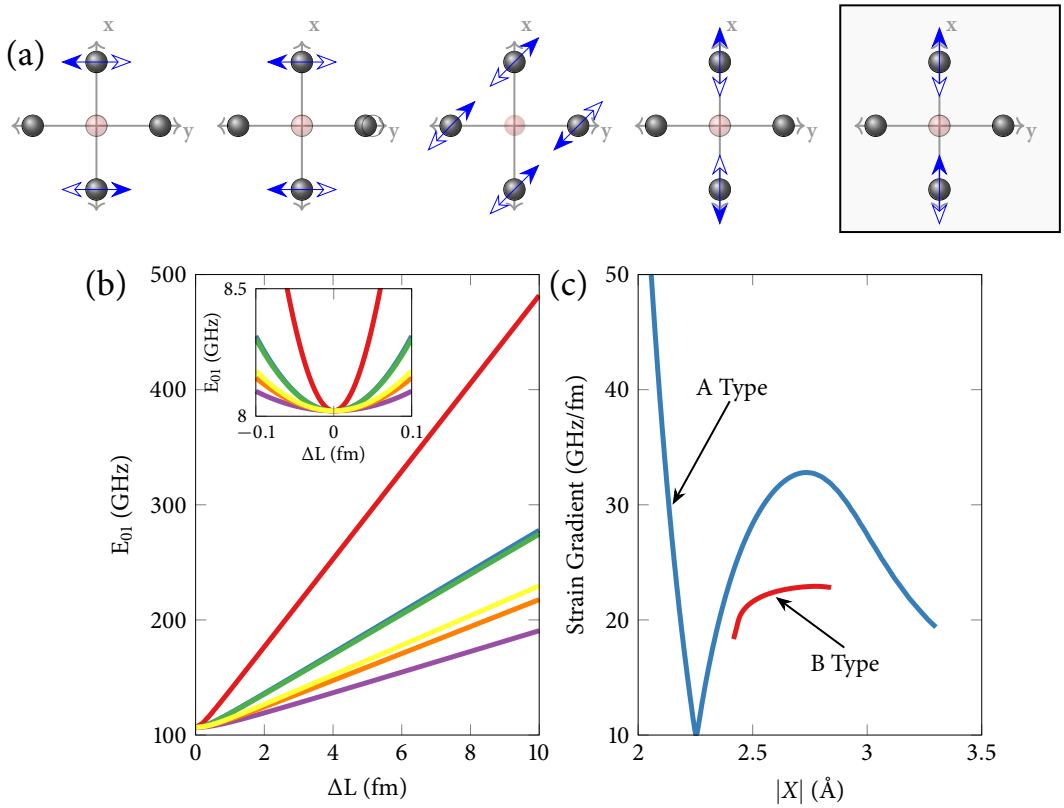

Figure 5.10— (a) Depiction of a number of deformations which were applied to the aluminium atoms in the x-y plane. The response of one of the strain types is orders of magnitude larger than the others (highlighted), which is indicative of an optical phonon mode, generating extremely large frequency splittings for picometer deformations. A deformation range of 10 fm was applied at points along 8 GHz contours (each line colour representing a different starting point, corresponding to phase qubit-TLS coupling strengths), with $|Z| = 2.5$ Å, yielding responses that are both hyperbolic and symmetric. (b) Over the range 2–10 fm the response is linear with a strain gradient, shown in (c) for both the A and B type defects in $|\wp_x|$ domains.

the tetra- and hemi-tetra- regions in the $|\wp_x| \neq 0, |\wp_y| = 0$ domains (when $|Z| = 2.5$ Å)—which may be an important property for a future model describing TLS-phonon coupling.

Whilst the large, hyperbolic response of the $|X|$ translation seems to be observed in experiments, it is not the only strain mode possible. Most deformation types in figure 5.10a do not exhibit any substantial change to defect properties, although two types (depicted again in figure 5.11a) show an active response over a length change ΔL of 1 pm. Four clusters are chosen that lie on the 8 GHz contour when $|Z| = 2.25$ Å, indicated as (b), (c), (d) and (e) in figure 5.11a. As can be seen in figure 5.7, each of these cases are within a $|\wp_x|$ dominant domain.

Clusters (b) and (c) are both identified as A type defects and are insensitive to $|X|$ dilation and $|Y|$ translation. Dilation of $|Y|$ for these defects is a different way of saying 'extending the defect pair separation'—which has been described in the above sections. The intensity of the response however is noticeably larger as the confinement distance ($|X|$) is increased.

B type defects, labelled (d) and (e), respond discordantly depending on their configuration. Dilation in $|Y|$, whilst a sizable strain source for A type configurations, initiates no response from the (e) configuration at all. B type defects located at this position in phase space have their defect pair in $|X|$, and point (e) specifically is confined with $|Y| = 3.1$ Å. As discussed in the sections above, this separation distance is close to being practically unbound. Point (d) on the other hand has a tighter separation distance and begins to confine the system. Dilation in $|X|$ responds in the opposite manner effectively: expanding defect separation distance whilst keeping the confinement distance static ultimately confines the wavefunction. The final response that the B type defects respond to is $|Y|$ translation. As point (e) is essentially unbound in $|Y|$ we do not expect any response from a strain of this magnitude. Point (d) on the other hand is interacting with its tighter confinement distance, with the TLS dipole collinear to the $|X|$ direction. Translation of $|Y|$ forces local minima of the potential landscape from points on the x axis to locations off axis, yet still within the minima rings about the aluminium atoms. In other words, the wavefunction is slightly rotated around the defect axis.

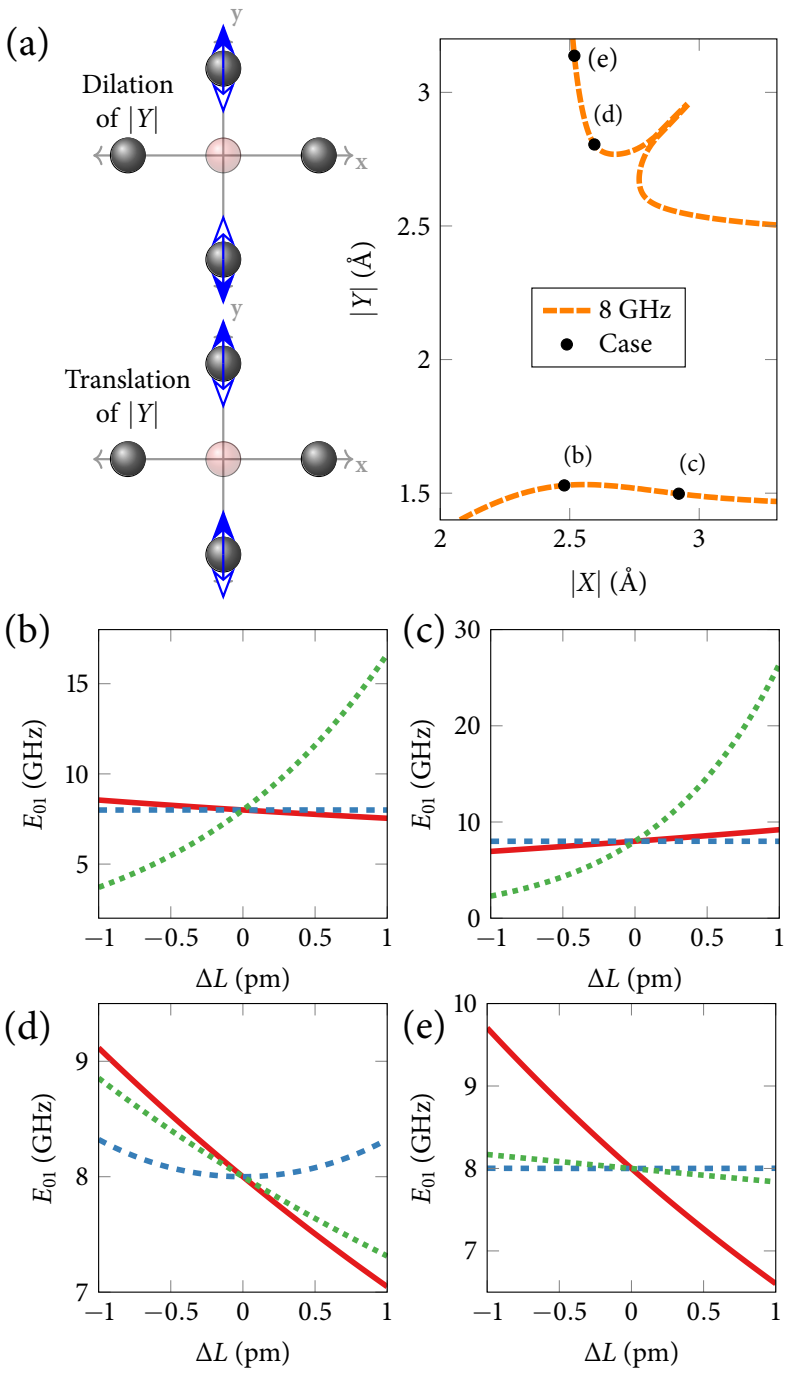

Figure 5.11— Strain response of four candidate TLS clusters. (a) Left: depictions of two deformation types that show active responses. A dilation: where an $|Y|$ (shown) or $|X|$ pair is stretched from their original position and a translation: where an $|Y|$ (shown) or $|X|$ pair is moved from their original position. Right: locations along the 8 GHz contour from the $|Z| = 2.25$ Å phase maps. Each case is deformed and translated in both x and y directions by ± 1 pm and plotted as (b), (c), (d) and (e). Dilation of $|X|$ ——, Translation of $|Y|$ - - -, Dilation of $|Y|$ ····. Note that all four cases lie in the $|\wp_x|$ dominant domain (see figure 5.7). Translation of $|X|$ is not shown on any graph as its' E_{01} response is of $\sim 10^4$ GHz over (ΔL) for all cases.

5.5 Summary of the Low Dimensional TLS Model

Despite the extreme idealisation of this model, it allows the prediction of experimentally measured properties of strongly coupled TLSs with atomic positions as the only input parameters and therefore shows as a proof of concept that these defects can arise in AlO_x without any alien species present.

Even when considering delocalisation in only two dimensions, a range of different behaviour can be seen. The existence of effective two-state systems can come about through different potential shapes, which in turn arise due to various atomic position configurations. To understand these different configurations, we consider both the symmetries of the eigenspectrum and the effective charge dipole of the defect.

We find that two-level systems with equivalent properties to those seen in experiments can be formed from atomic configurations which are entirely consistent with the material properties of amorphous aluminium oxide barriers. Most interestingly we find that the expected qubit-TLS coupling strength and the TLS strain response correspond very well to that observed experimentally.

A complication that occurs with phenomenological models that attempt to describe TLS behaviour is their free parameter count is high enough to describe all facets of the observed experimental parameters without being distinguishable from other, non-complimentary models [81]. Whilst this delocalised oxygen model only uses one type of parameter as input (atomic positions), it still requires an x, y, z coordinate set for up to 6 cage atoms. Realistic values for the atomic positions obtained by the junction model in chapter 2 are still yet to be applied to the model, which may require more than 6 atoms and thus even more free parameters. A method to account for this parameter runaway is discussed later in the next chapter.

„… ein Loch ist ja irgendetwas. Aber dort ist nichts."

—Blubb das Irrlicht

6 TLS Identification In 3D

Extending the search for TLS candidates to three spatial dimensions. An alternative method for solving the Schrödinger equation is introduced. TLS properties inside a model JJ are investigated.

Memory limitations were the greatest concern with the implementation of the direct diagonalisation approach discussed in the previous chapter. The 2+1D framework may capture all relevant physics and therefore the need to extend the system to a complete 3D may be moot. However, the only way to verify this statement is to compare a 3D representation of the TLS model with its 2+1D counterpart.

In this chapter we focus on a delocalised oxygen TLS model in three spatial dimensions. Section 6.1 discusses the pitfalls and concerns of extending the current 2+1D to 3D at low resolution such that the solution fits into memory. As this method can not be trusted to yield accurate results, we introduce an alternative method in section 6.2 which can be executed on highly distributed clusters over many CPU cores.

Sections 6.3 to 6.7 discuss the capabilities and results that can be obtained with this method. Starting in section 6.3 a comparison of the Al–O–Al chain discussed in section 4.2 is examined, followed by descriptions of the wavefunctions exhibited by various atomic configurations, some of which can be identified as A and B type defects in sections 6.4 and 6.5. Section 6.6 generates clusters of atoms surrounding an oxygen vacancy defect in Al_2O_3, then using the same method section 6.7 calculates TLS properties in junction models constructed in chapter 2. Finally, a small summary of the alternate method is given in section 6.8.

6.1 Extending Direct Diagonalisation To 3D

Constructing the Hamiltonian H (3.1) in three dimensions is a quick extension to the two dimension block matrix solution (4.1)

$$H = \left[H_x \otimes (\mathbb{I}_y \otimes \mathbb{I}_z)\right] + \left[(\mathbb{I}_x \otimes H_y) \otimes \mathbb{I}_z\right] + \left[(\mathbb{I}_x \otimes \mathbb{I}_y) \otimes H_z\right]. \tag{6.1}$$

The triviality of this extension stops there though. As previously established in subsection 3.9.1, the maximum acceptable step size one can apply to the

model is $h_{max} = 2^{-3} \simeq 0.125$ Å. Larger step values will introduce invalid truncation error into the solution and cannot be treated as valid solutions. The same machine used to calculate the values required for figure 3.4 and the associated discussion is not capable of fitting a 3D calculation into its memory and swap space with a step size any smaller than $h \approx 0.7$ Å.

Additionally, many excited state wavefunctions exhibit sign changes on scales much less than 0.7 Å, provoking another systemic error: the selection of an incorrect excited state value which does not have such a rapid sign change in the region.

6.2 An Alternative Method

Sudiarta and Geldart introduced a method which exploits a Wick-rotated time-dependent Schrödinger equation to solve for time-independent solutions in three dimensions [158].

Equation (3.59) ignores any time evolution of our TLS model, although the general case (6.2) includes this description. Wick rotations [159] are primarily used to find solutions to Minkowski space problems by an Euclidean space mapping. Here we can use the same method to transfer our Schrödinger equation to an imaginary time basis $t = i\tau$ (6.3)

$$i\hbar\frac{\partial}{\partial t}\Psi(\mathbf{r},t) = H\Psi(\mathbf{r},t) \tag{6.2}$$

$$\Rightarrow -\hbar\frac{\partial}{\partial \tau}\Psi(\mathbf{r},\tau) = H\Psi(\mathbf{r},\tau), \tag{6.3}$$

which yields a general solution to the wavefunctions

$$\Psi(\mathbf{r},\tau) = \sum_{k=0}^{\infty} a_k \psi_k(\mathbf{r}) e^{-E_k \tau}. \tag{6.4}$$

Here, a_k are coefficients based on initial conditions of the system where $k = 0$ indicates the ground state, $k = 1$ the first excited state *etc.* and E_k is the corresponding eigenenergy. As $E_0 < E_1 < E_2 < \ldots$, evolving (6.4) to large imaginary time will provide a good approximation to the ground state influenced by the time-independent potential $V(\mathbf{r})$, where $\mathbf{r} = (x, y, z)$ in this context.

This solution is obtained by numerically approximating the spatial derivatives with finite differences. The wavefunction is averaged in imaginary time to improve stability [160] and for precision, we discretise the space using the 7-point central difference method discussed comprehensively in chapter 3.

To obtain a function which evolves the system through imaginary time, we start by assuming our grid is uniform in all directions such that $a \equiv \Delta x = \Delta y = \Delta z$, and define a difference vector

$$\mathbf{D} = \frac{1}{a^2}\left[2, -27, 270, -490, 270, -27, 2\right] \tag{6.5}$$

in line with the coefficients of (3.52). In conjunction with a matrix-valued wavefunction field

$$\widehat{\psi} = \begin{bmatrix} \psi(x-3,y,z,\tau) & \psi(x-2,y,z,\tau) & \cdots & \psi(x+3,y,z,\tau) \\ \psi(x,y-3,z,\tau) & \psi(x,y-2,z,\tau) & \cdots & \psi(x,y+3,z,\tau) \\ \psi(x,y,z-3,\tau) & \psi(x,y,z-2,\tau) & \cdots & \psi(x,y,z+3,\tau) \end{bmatrix}, \tag{6.6}$$

an approximate solution to the right hand side of (6.3) can be obtained:

$$H\Psi(\mathbf{r},\tau) \approx$$
$$\frac{1}{360a^2 m_{oxy}} \sum_i^3 \left(\mathbf{D}\cdot\widehat{\psi}\right)_i - \frac{1}{2}V(x,y,z)\left[\psi(x,y,z,\tau) + \psi(x,y,z,\tau+\Delta\tau)\right]. \tag{6.7}$$

Equation (6.3) can now be rearranged using (3.52) and (6.7) to produce an expression to step $\psi(x,y,z,\tau)$ through imaginary time

$$\psi(x,y,z,\tau+\Delta\tau) = A\psi(x,y,z,\tau) + \frac{B\Delta\tau}{2m_{oxy}} \sum_i^3 \left(\mathbf{D}\cdot\widehat{\psi}\right)_i \tag{6.8}$$

$$A = \frac{1-\frac{\Delta\tau}{2}V(x,y,z)}{1+\frac{\Delta\tau}{2}V(x,y,z)}, \quad B = \frac{1}{1+\frac{\Delta\tau}{2}V(x,y,z)}. \tag{6.9}$$

6.2.1 Gram-Schmidt Procedure

One important property of a hermitian operator (in our case H), is orthogonality. That is: distinct eigenvalues belong to a systems' eigenfunctions which are orthogonal. The proof of this theorem neglects degenerate states, although it can be shown that any linear combination of eigenfunctions which share the same eigenvalue are indeed eigenfunctions of the system.

Using the Gram-Schmidt orthogonalisation procedure [161, 162], orthogonal eigenfunctions within a degenerate subspace can be constructed; *chosen* to be orthogonal to the systems' basis.

Consider a non-orthonormal basis of vectors $|a_1\rangle, |a_1\rangle, \ldots |a_n\rangle$. A new, orthonormal basis for these vectors can by found be first normalising the initial vector

$$|a_1'\rangle = \frac{|a_1\rangle}{\||a_1\|}, \tag{6.10}$$

then finding the projection of the second vector along the first and subtracting it. This vector is orthogonal to $|a_1'\rangle$ and normalising it obtains

$$|a_2'\rangle = \frac{|a_2\rangle - |a_1'\rangle\langle a_1' \,|\, a_2\rangle}{\big\||a_2\rangle - |a_1'\rangle\langle a_1' \,|\, a_2\rangle\big\|}. \tag{6.11}$$

Similarly, the third vector requires the same treatment, although projections along both $|a_1'\rangle$ and $|a_2'\rangle$ are needed

$$|a_3'\rangle = \frac{|a_3\rangle - |a_1'\rangle\langle a_1' \,|\, a_3\rangle - |a_2'\rangle\langle a_2' \,|\, a_3\rangle}{\big\||a_3\rangle - |a_1'\rangle\langle a_1' \,|\, a_3\rangle - |a_2'\rangle\langle a_2' \,|\, a_3\rangle\big\|}. \tag{6.12}$$

This process is repeated until the orthonormal vector $|a_n'\rangle$ is calculated.

Applying this procedure to our particular case, consider a converged, orthonormal ground state $|\psi_0\rangle$ and an initial, non-orthonormal guess for the first excited state $|\psi_1'\rangle$. Subtracting the projection of the excited state along the ground state from the initial guess yields a vector orthogonal to the ground state

$$|\widetilde{\psi_1'}\rangle = |\psi_1'\rangle - |\psi_0\rangle\langle\psi_0 \,|\, \psi_1'\rangle. \tag{6.13}$$

This resultant vector can then be normalised and evolved through the same process as the ground state (*i.e.* the evolution of 6.4 to large imaginary times) until convergence is achieved. Similarly, the second excited state requires the same treatment, although projections along both $|\psi_0\rangle$ and $|\psi_1\rangle$ are needed.

6.2.2 Implementation

Further theoretical discussion and a link to a rudimentary code base which implements the method can be found in reference 163. This software discretises over a 3-point central difference solution and does not iteratively solve higher order excited state wavefunctions, just merely approximates the first excited state.

The results obtained within this chapter were calculated using a highly modified fork of the original software specifically designed for this thesis. An active repository of this code can be found via reference 164 and may be packaged for general release in the future.

6.3 Representation of an Al–O–Al Chain in 3D

With this new approach, we have the capability to investigate oxygen delocalisation in 3D, using realistic atomic positions within an amorphous layer of AlO_x. However this generates an extensive state space, and whilst a highly distributed computing cluster can be invoked to share the load, phase maps similar to those in chapter 5 are still impractical with the added time burden required for accurate calculations ($\sim$ 200 hours for 6 states in one spatial configuration).

Hence we start this investigation using a simple three atom system comprising of an oxygen atom and a confining aluminium pair $|X|$, representing a perturbed crystalline system as section 4.2 did with the direct diagonalisation approach. An eigenspectrum is generated by increasing or decreasing the pair separation distance from the corrundum lattice distance 1.85 Å using the ground and five lowest excited energy levels of this system, depicted in figure 6.1 (for comparison with the 2D model, see figure 4.1).

In this figure an (an)harmonic division about $|X| \sim 1.85$ Å separates two regions with dissimilar properties. In the region where $|X| > 1.95$ Å, a degenerate E_{01} is observed and as the separation distance has increased, can be identified as a possible B type defect (see figure 1.5). The second region ($|X| < 1.55$ Å) has a tightly bound eigenspectrum compared to the rest of the map although this yields no degeneracy.

In three dimensions, the potential minima manifests as a sphere around the location of an atom. As a consequence the three atom Al–O–Al chain produces a unique, rotationally symmetric ground state which is reminiscent of an oxygen interstitial defect in crystalline germanium (figure 1.7) and depicted in figure 6.2a&b. Comparatively, this region in the 2D case indicated the presence of an A type defect, as the 2D potential minima ring would be projected onto a plane, collapsing a degree of freedom.

Two other observations concerning the 2D to 3D transition is the dramatically reduced total energy values as wavefunctions are no longer artificially confined, and the higher excited states in the B type region now exhibit a quad degeneracy rather than two split doubly degenerate pairs—again because the states are not artificially confined in the extra dimension.

This simplistic three atom description can now be built upon to understand the interactions of a delocalised oxygen in a three dimensional space. A confining pair is introduced in an additional dimension to the configuration in

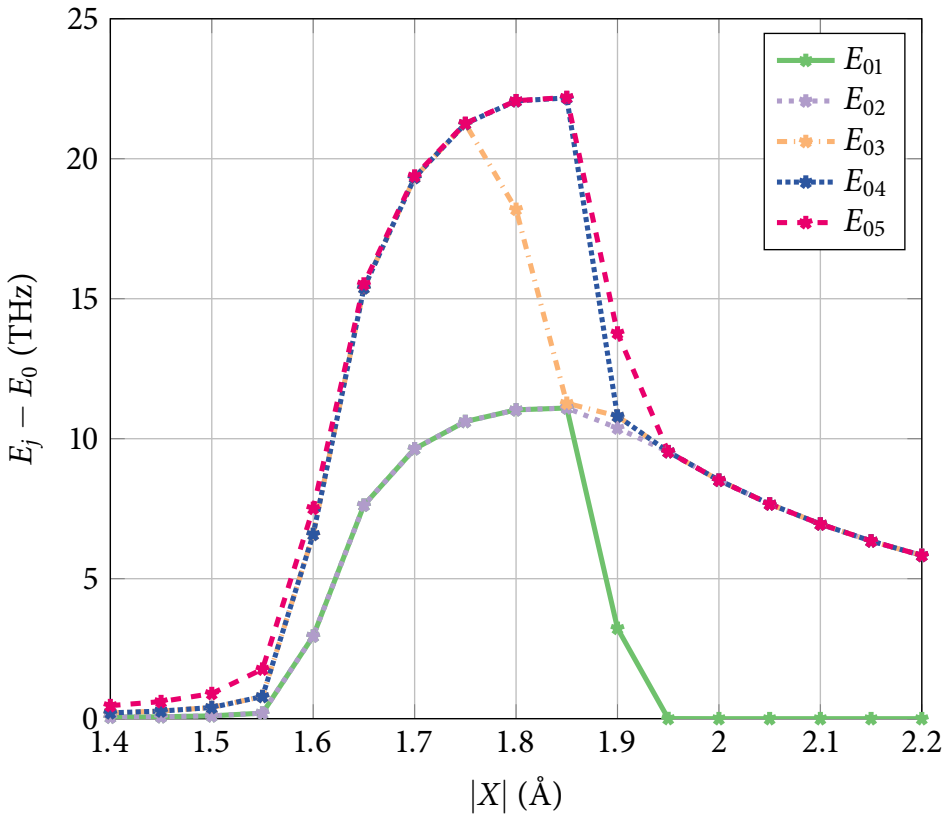

FIGURE 6.1— Eigenspectrum of a three atom system Al–O–Al, over a varying distance separation. Each excited state has been normalised with the ground state, which shows a clear E_{01} degeneracy at large separation distance—indicative of a B type defect. An intermediate region exists as the separation distance decreases, which is approximately centered about the optimal Al–O bond distance of corundum (1.85 Å). At small separation distances, the eigenspectrum is distributed over comparatively small energy levels, but is not degenerate (see figure 6.2).

figure 6.2a&b, and the unique ground state is lifted to a degenerate E_{01} pair; observing an A type TLS defect. Figure 6.2c&d illustrate this effect with a confining pair $|Y| = 3.5$ Å. This distance is an arbitrary choice, as the shift occurs even at the cut-off limits imposed by the potential portion of the model [128] (see section 3.1). In other words, any additional potential contribution which does not share the axial symmetry of the Al–O–Al chain will lift the degeneracy and result in a TLS.

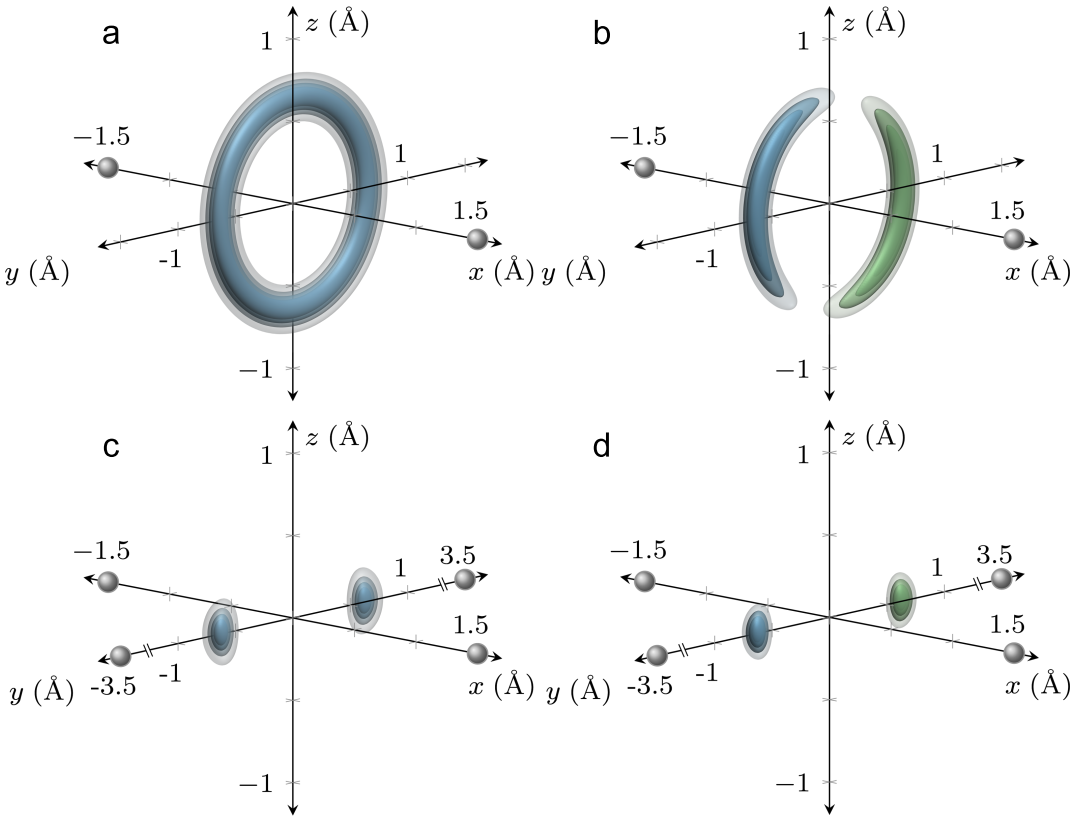

Figure 6.2— Ground (left) and first excited (right) state wavefunctions of an Al–O–Al chain with $|X| = 1.5$ Å(aluminium atoms are depicted as ●). Top axes show a unique ground state, the E_{12} pair is degenerate in this case. Two more aluminium atoms are introduced at $|Y| = 3.5$ Å on the bottom axes, which causes degeneracy in the E_{01} pair.

6.4 Oxygenic Orbitals

A complete picture of the ground and excited state wavefunctions in 3D is now possible, which allows us to further investigate the properties of the TLS and illustrate the importance of crystalline dielectrics in future Josephson junction devices.

A confined harmonic state can exist with many atomic configurations, and as with the low dimensional model, setting variables symmetrically is one of these cases. Figure 6.3 depicts the four lowest energy wavefunctions of an oxygen with six confining aluminium atoms: $|X| = |Y| = |Z| = 1.6$ Å. Here the oxygen atom has no additional local minima to occupy, identifying this configuration as spatially localised, the harmonic approximation holds and the system is not considered as a TLS candidate.

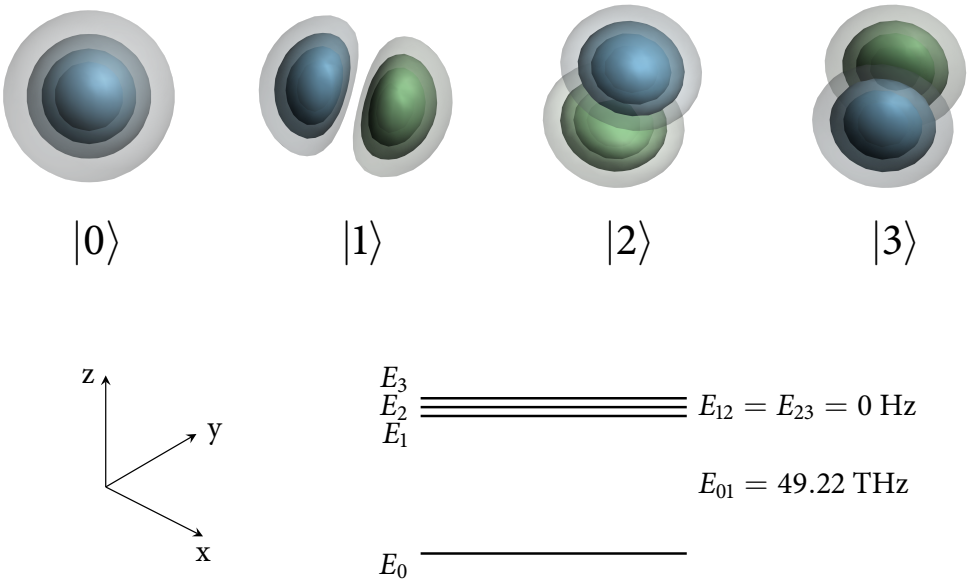

FIGURE 6.3— Wavefunctions of an oxygen confined by six equidistant aluminium atoms at $|X| = |Y| = |Z| = 1.6$ Å. This configuration exhibits a unique ground state (as shown in the energy level diagram) and hence is considered to be spatially localised.

From this localised case, we extend the $|Y|$ confinement out to 2.1 Å. Using the figure 6.1 eigenspectrum we can predict this configuration to be a B type defect; where the $|Y|$ separation distance is increased, and an oxygen dipole forms parallel to the y axis. This is indeed the case as shown in figure 6.4.

A particularly complex phenomenon emerged from the low dimensional model: quad degenerate ground states, where the energy difference between two degenerate pairs E_{01} and E_{23} approach the difference of the pairs themselves, *i.e.* $E_{01} = E_{23} \approx E_{12}$ (see section 4.5 and subsequent sections). Figure 6.5 shows a configuration expressing this behaviour.

The discussion in section 5.1 indicated that this form of degeneracy has a low probability of occurrence compared to A or B type defects. Additionally, if $E_{12} \geq 100$ GHz the higher states (E_2 and E_3) can be ignored completely, as discussed previously and hence can be viewed as an operational upper bound for Josephson junction devices. Splitting energies in 3D are much smaller than the 2+1D case, which suggests that the probability of these defects being experimentally visible may be higher than predicted with the low dimensional model based solely on this measure. However, a comprehensive analysis of

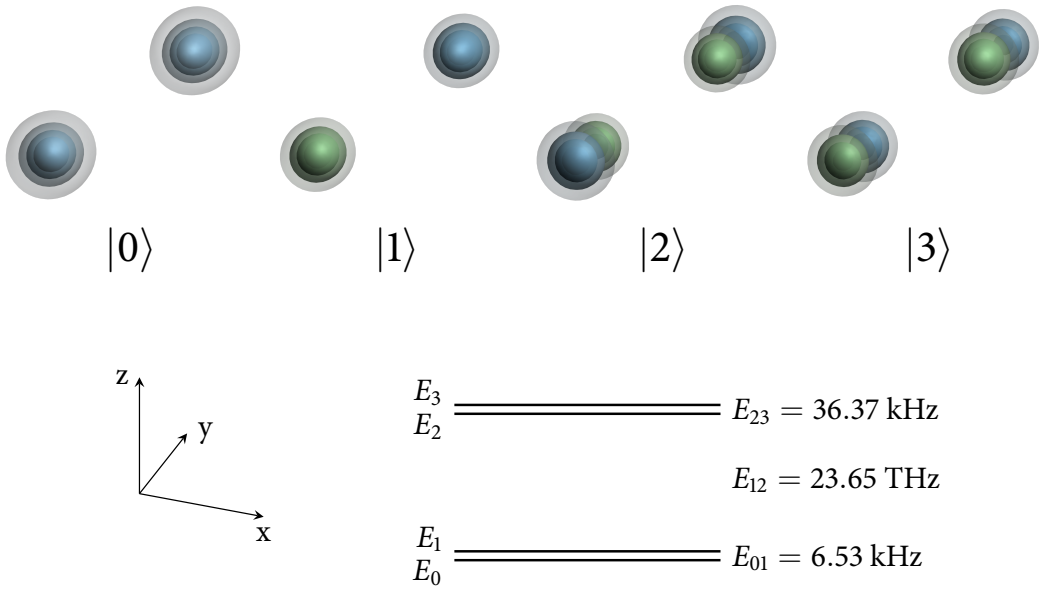

Figure 6.4— Wavefunctions of an oxygen confined by six aluminium atoms, where the symmetry of figure 6.3 has been broken. $|X| = |Z| = 1.6, |Y| = 2.1$ Å, which manifests as a B type defect.

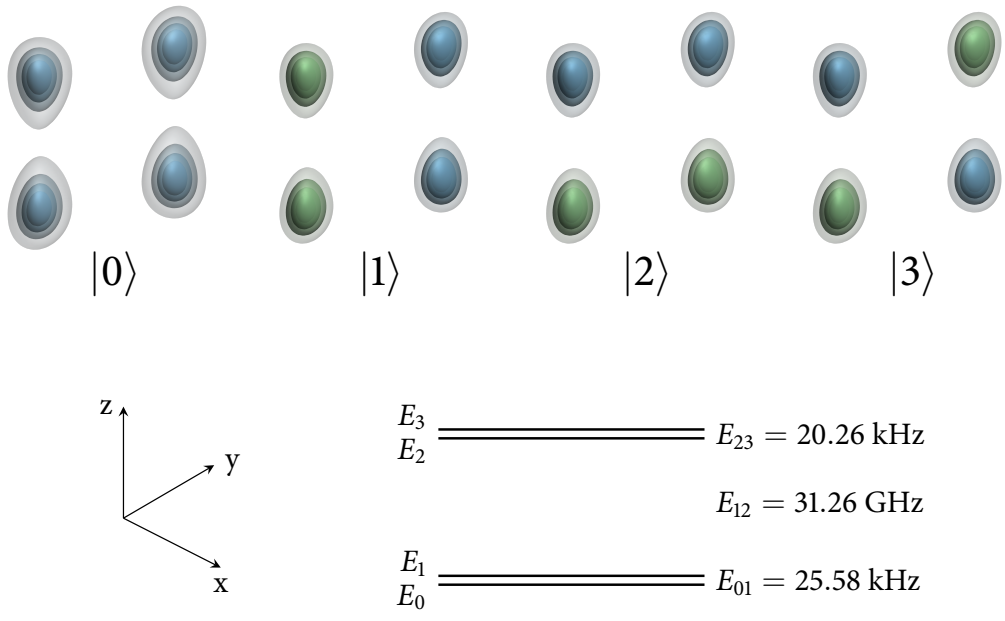

Figure 6.5— Wavefunctions of an oxygen confined by six aluminium atoms, where the confinement configuration is studiously pathological. $|X| = 1.6, |Y| = 2.1, |Z| = 2.0$ Å, yielding a quad degeneracy in the ground state due to the E_{12} splitting existing below the superconducting gap ($\sim$100 GHz, see text).

the three dimensional configuration space would need to be undertaken before real estimates of this behaviour can be given.

Another interpretation of this phenomena has to do with the energy resolution of the spectroscopy discussed in subsection 1.5.1. The minimum resolvable energy using this process is approximately 1 MHz, which may indicate that S_{max} anti-crossing measurements are manifestations of qubit-TLS couplings at the E_{12} level if a configuration such as figure 6.5 is scrutinised—E_{01} and E_{23} splitting levels would be hidden in their degenerate subspace and manifest as an effective two-level defect with $\widetilde{E}_{01} = 31.26$ GHz.

6.5 A Type Defects and Dipole Considerations

Whilst adding additional confinement pairs causes the unique, rotationally symmetric ground state of $|X| < 1.55$ Å (figure 6.2a&b) to become degenerate; the third dimension yields a complication in the dipole measurement for the A type region.

Consider a system with parameters $|X| = 1.53, |Y| = 1.97, |Z| = 1.95$ Å, illustrated in figure 6.6. This system exhibits TLS-like behaviour, with $E_{01} = 40.63$ MHz, and a dipole strength in y of 0.26 eÅ: perpendicular to the confining x axis as expected. The leftmost axis of figure 6.6 shows a 2D projection of the first excited state, illustrating no major differences in the response of the 3D and 2+1D models (refer to figure 5.2 for comparison).

However, a small change in the y confinement alters the system in a non-trivial manner. Moving $|Y|$ from 1.97 to 1.9 Å (for example) crosses a bifurcation in state space. As $|Z|$ is now the least confining pair, the dominant dipole direction flips to z as shown in figure 6.7.

Ultimately this does not add any complexity to the model: minimal confinement in y generates an A type defect with a dipole in y, perpendicular to x; minimal confinement in z generates an A type defect with a dipole in z, perpendicular to x. As strongly coupled defects are experimentally identified via avoided level crossings in qubit spectroscopic diagrams [69], the dipole moment of the defect couples to the electric field across the junction [74]. Therefore, to be identified as a TLS, the defect must be aligned to this field. It follows then, that the dipole directions for an A type defect in this model can be considered equivalent. Additionally, as with B type and any other possible TLS defect, both A type alignments will not appear on a spectroscopic scan as an

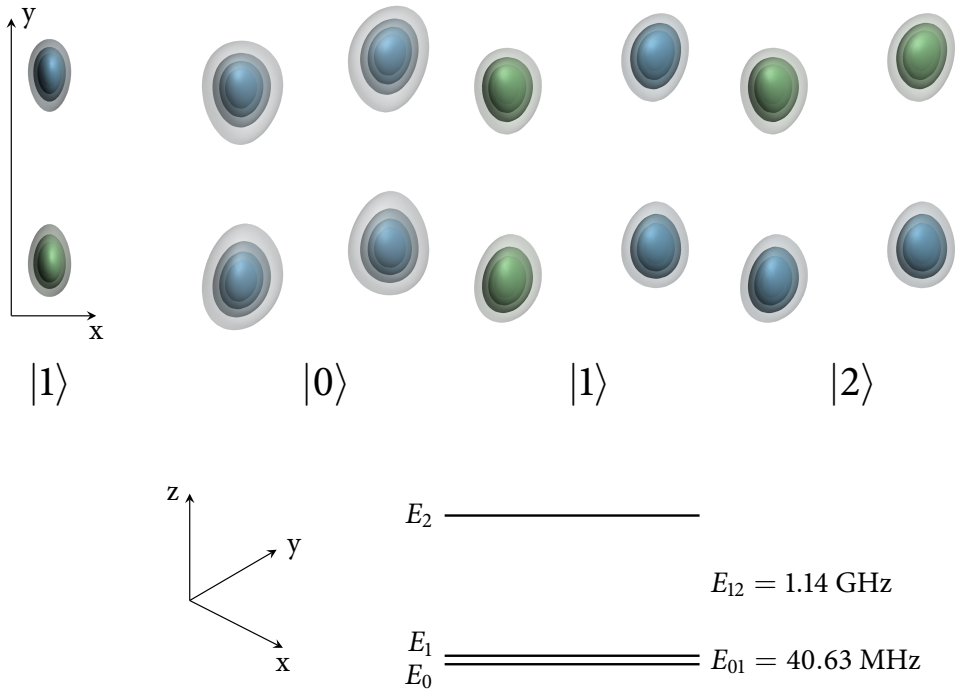

Figure 6.6— Wavefunctions of an A type defect with confinement $|X| = 1.53, |Y| = 1.97, |Z| = 1.95$ Å which yields dominant dipole in the y direction with a strength of $0.26\ e$Å. Leftmost axis shows an xy projection of the first excited state to compare with 2+1D results (figure 5.2).

avoided level crossing if their dipole does not have a component in the direction of the external field.

6.6 Defects in Corundum

Results such as figures 6.1 and 6.3 identify configurations which exhibit harmonic eigenspectrums and localised oxygenic positions. This behaviour is expected when an oxygen atom is confined appropriately and has no reason to tunnel. One way to check this assumption (and consequently further validate the TLS model) is to calculate the potential landscape seen by a single oxygen atom from bulk corundum (α–Al_2O_3). The TLS model, delocalising an oxygen across all potential space about its minimum energy position, should yield a localised, harmonic wavefunction positioned at the location where the oxygen was removed (*i.e.* the lattice position associated with the vacancy).

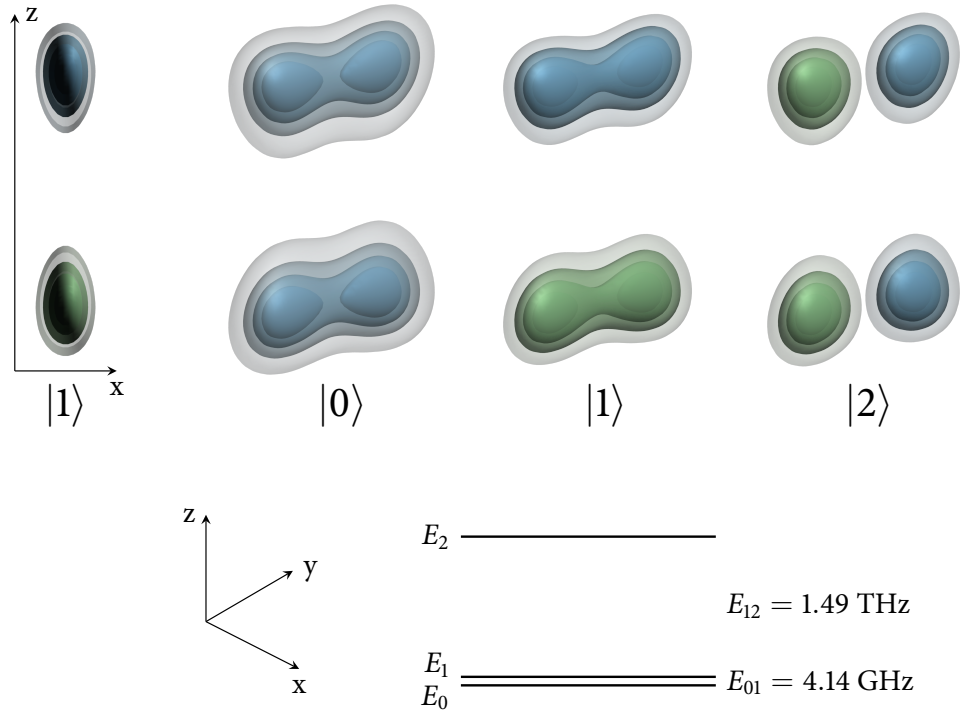

FIGURE 6.7— Wavefunctions of an A type defect with confinement $|X| = 1.53, |Y| = 1.9, |Z| = 1.95$ Å which yields dominant dipole in the z direction with a strength of 0.21 eÅ. Leftmost axis shows an xz projection of the first excited state to compare with 2+1D results (figure 5.2).

6.6.1 Cluster Identification With Voronoi

Both empirical and *ab initio* models of crystals primarily rely upon the simplification that periodic boundary conditions (PBC) impose. A unit cell representation containing just a few atoms and their coordinates are connected through the lattice to a mirror of themselves in every direction due to perfect tiling of the cell and the symmetries of the crystal. This tiling therefore represents an infinite crystal with only a few atomic interactions. Crystals with defects, require more atoms in their unit cell (usually referred to as a 'super cell' in the literature) to hide any mirror effects the defects may observe from close PBCs. More importantly, amorphous materials have no periodicity and therefore are difficult to describe with PBCs as small box sizes introduce extreme regularity to the model's lattice.

As a consequence, the TLS model does not include any PBC capability and if we wish to investigate a vacancy defect in corundum an approximation to the local potential must be made. We must choose how many of the surrounding

aluminium and oxygen atoms to include in the calculation of the local potential. We do this by applying a Voronoi tessellation [165] to the lattice coordinates of each system using a Euclidean distance metric. A polytope representing the convex hull encompassing an origin position symbolises the oxygen vacancy. Atoms situated within polytopes sharing an edge with the origin polytope are selected as the model cluster. These atoms constitute a first order screening of the surrounding lattice potential, representing an acceptable approximation to higher order screenings (*e.g.* also including second nearest neighbours, which add appreciable computational intensity).

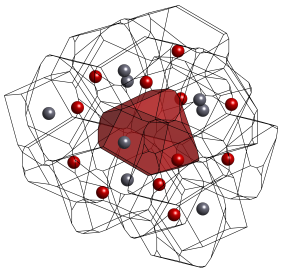

Figure 6.8— Voronoi cell indicating an oxygen vacancy in Al_2O_3 ■. Nearest neighbour aluminium ● and oxygen ● atoms are identified via their encapsulating polytopes.

For this section and section 6.7 below, we use Streitz Mintmire potentials generated by GULP [127], which also account for monopole–monopole and self energy interactions to counteract any net charge generated by selecting an arbitrary cluster of atoms from a lattice with periodic boundaries.

The resulting cluster of atoms found by the Voronoi process is used to calculate the wavefunctions and energies of an oxygen atom located inside the potential landscape generated by the corundum lattice. As predicted, the result is a localised oxygen atom locked into the lattice position of the vacancy. Figure 6.9 shows ground and first excited state wavefunctions, which behave as the harmonic system calculated previously in figure 6.3.

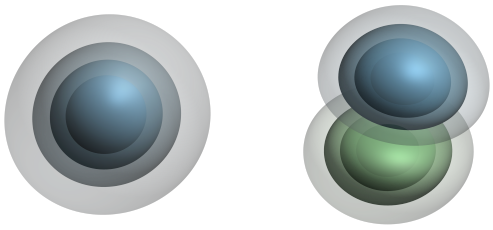

Figure 6.9— Ground (left) and first excited (right) state wavefunctions of an oxygen atom introduced to a vacancy defect in Al_2O_3. $E_{01} = 13.48$ THz for this cluster configuration. The energy states of the oxygen are harmonic and the atom can be considered as localised in its position in the surrounding lattice.

6.6.2 Straining the Local Lattice

In an attempt to put bounds on the amount of amorphousness required for a near-crystalline structure to exhibit TLS properties, we apply a localised strain tensor on the defect cluster and observe its effects.

As a trigonal crystal, corundum has six elastic constants [166], meaning a force exerted on the structure will act on many different crystal planes in a non-trivial manner. The materials science complications of this phenomena are well beyond the scope of this investigation, and instead we approximate a local strain by the Poisson effect [167] using the values of the adiabatic bulk and shear moduli to obtain a Poisson's ratio $\nu = 0.234$. This value is calculated as an equivalent isotropic polycrystalline aggregate to Sapphire (a mineral variety of corundum) [168].

The relationship

$$\left(1 + \frac{\Delta L}{L}\right)^{-\nu} = 1 - \frac{\Delta L'}{L} \tag{6.14}$$

can be used to strain our atom cluster for example, with a length increase of ΔL in the x direction, and a length decrease of $\Delta L'$ in the y and z directions. Figure 6.10 shows the response of two strain ratios $\frac{\Delta L'}{L} = \mp 0.019$. The left axis represents a compressive force along the x axis, and the right is a tensile force also along x. These values attempt to simulate a local strain from the localised Al–O separation distance of 1.85 Å in the direction of the A type and B type regions respectively (see figure 6.1).

Comparing the E_{01} values of the strained results to the original crystalline system in figure 6.9, we see an increase in splitting energy. This response is opposite to the simple straining of figure 6.1 where the E_{01} energy was a maximum at the crystalline Al–O separation distance. This suggests that local Poissonian strain does not give rise to immediate TLS behaviour—at least along the sampled axis x. Recall section 5.4: there exists a preferred strain axis which responds rapidly in comparison to other directions, although this axis is difficult to acquire directly from a completely harmonic case. It is therefore pertinent to investigate the amorphous configurations in more detail.

6.7 Defects in Josephson Junctions

The Vononoi classification scheme in subsection 6.6.1 finally allows us to investigate the myriad atomic configurations generated by the junction models in chapter 2.

Using the structure with an $AlO_{1.25}$ amorphous barrier with a density 0.75 that of corundum (figure 2.1) we classify clusters for all oxygen atoms extant

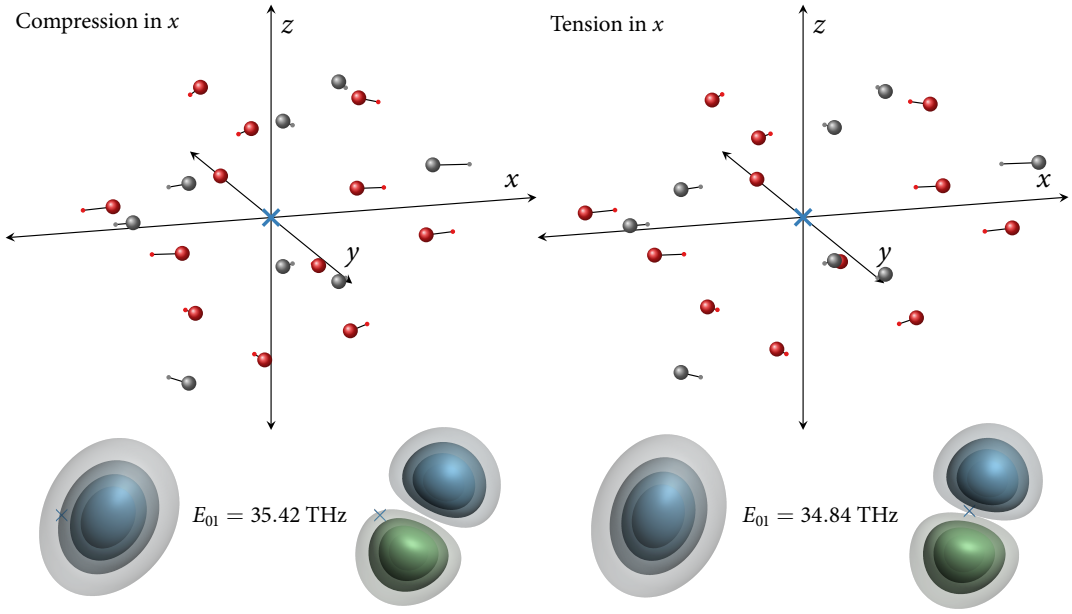

Figure 6.10— A corundum crystal cluster under compression (left) and tension (right) along the x axis. Aluminium ● and oxygen ● markers depict their strained positions and their pre-strained coordinates are shown as small dots. Below, the ground and first excited state wavefunctions are no longer spherical in nature and the compressed clusters minimum has visibly moved from the origin position X.

in the tunnel barrier. The volume of each Voronoi cell can be calculated (as the cell is essentially a convex hull of some vertices describing a Delaunay tessellation) and we choose the five cells with the largest volume to investigate. It is posited that a large void space around an oxygen may allow it to delocalise over a large region and manifest as a B type defect (however, this is pure speculation). These clusters and their properties are displayed in figure 6.11. Another conjecture is that an asymmetric cell (obtained by some form of sheer in the local cluster configuration) yields TLS behaviour, thus a final cluster is investigated which has the most asymmetry—identified using the length parameter of the spherocylinder formalism outlined in reference 169 (indicated as ■ in figure 6.11).

As with the cluster results of the crystalline and Poisson strain investigations of section 6.6 above, the few amorphous clusters studied failed to identify any TLS behaviour. With such a small sample size it is difficult to draw any direct conclusions from this result, although it is noteworthy that the asymmetric case ■ exhibits a smaller E_{01} energy than the clusters with large Voronoi

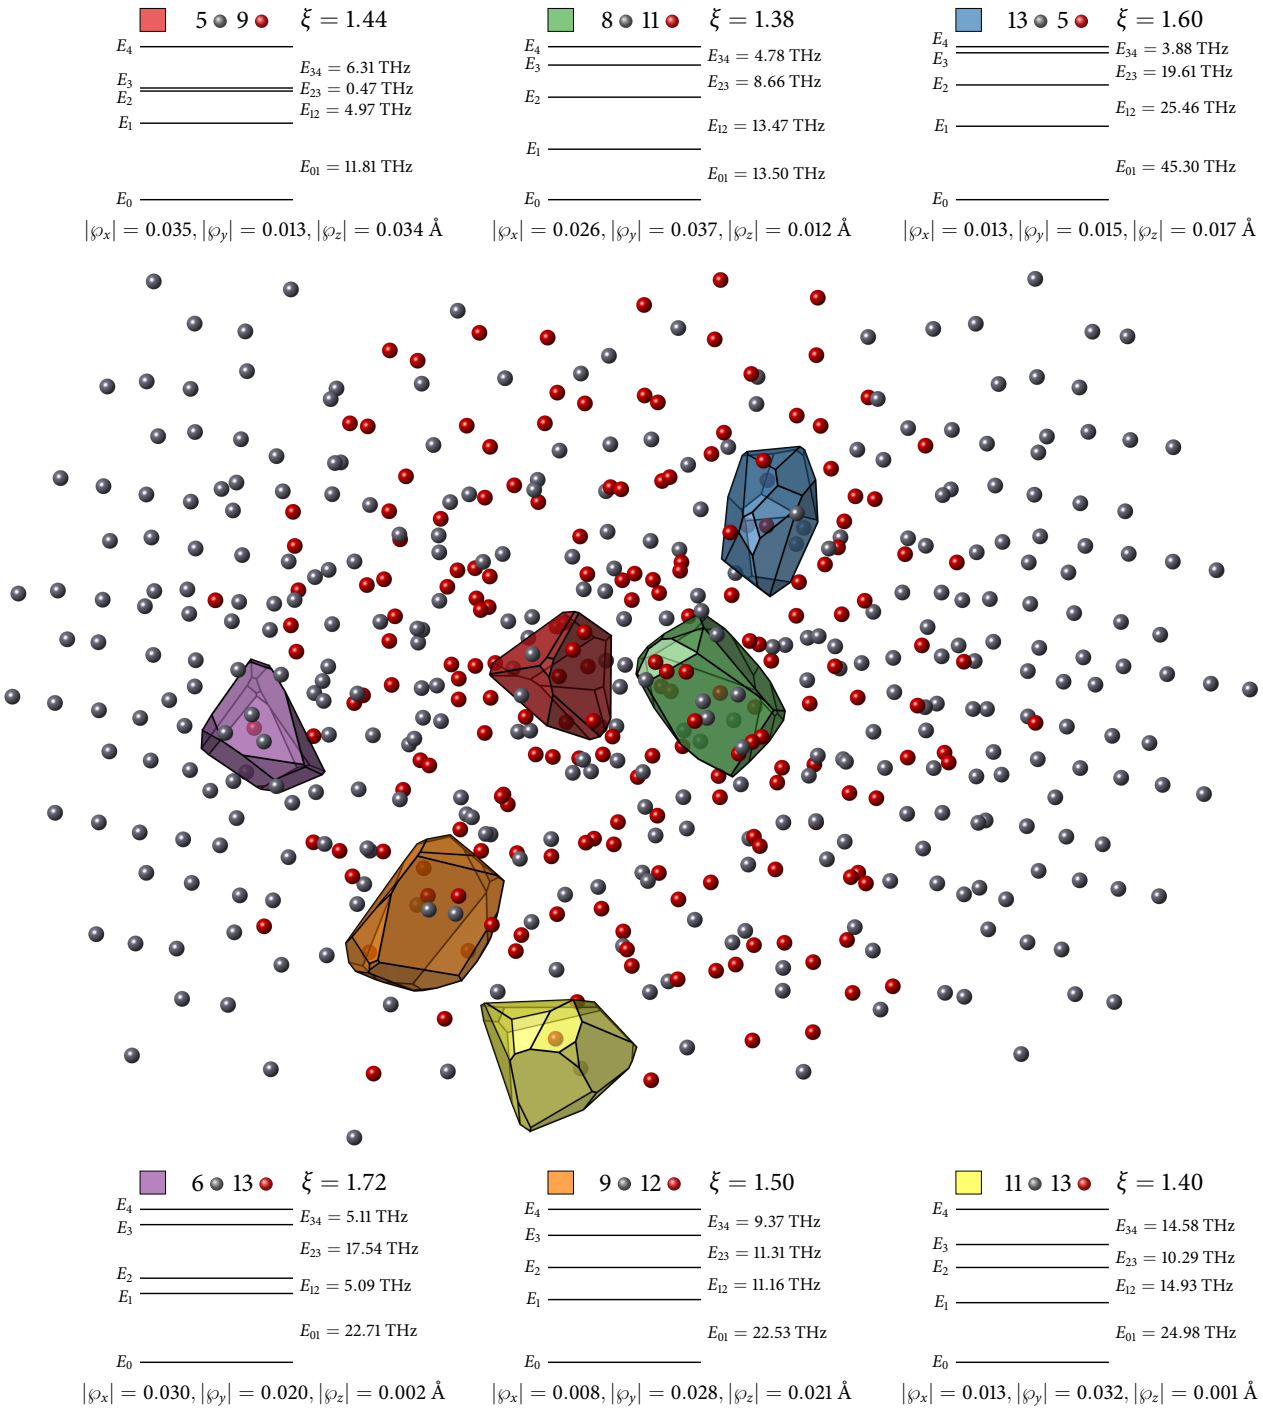

FIGURE 6.11— Properties of various oxygen atoms in a Josephson junction. Aluminium ● and oxygen ● atoms in the center depict the amorphous tunnel barrier, with crystalline aluminium at either end representing the superconducting sections of the junction.

cells as well as the crystalline cluster. A comprehensive exploration of this and JJ models with different stoichiometries and densities may yield clusters with further asymmetry, which may in turn identify a TLS-like cluster.

Other methods of identification such as extending the Euclidian distance Voronoi classification discussed here to the Voronoi S-networks considered in reference 169, or considering not just single body interactions, but a multiplied identification approach similar to Paz *et al.* may assist in TLS classification [170]. Once candidates are determined, the methods outlined herein can classify each system and ascertain their TLS properties or confirm their simple atomic nature.

6.8 Chapter Summary

Using the direct diagonalisation method in three dimensions is seemingly intractable with current memory limitations, although this does not exclude investigations in the 3D domain. The Wick-rotated time-dependent Schrödinger equation method has proved to be a valid tool to study this space, albeit with computational time disadvantages.

Whilst the 2+1D model allows us to investigate phase space in a much more flexible way, one of the major variations in the models is the final energy values, as the 3D model is no longer collapsed into fewer dimensions. This result suggests that if one is attempting to identify exact values for a delocalised oxygen TLS, they should be utilizing the 3D method. Contrarily, if trends are required, the 2+1D model provides the more efficient solution. Whilst splitting energies and dipole results may not completely reflect the atomic positions of the surrounding clusters in this case, it is evident from discussions like those in section 6.5 that there are 2D to 3D equivalencies. For instance, the E_{01} splitting has the same behaviour when comparing figures 4.1 and 6.1, albeit with an energy response an order of magnitude lower in the latter. The direct comparisons of the 2+1D model to experiment in sections 5.3 and 5.4 should therefore not deviate from their trends when using the 3D method, however the positions of the atoms in each model will be dissimilar.

Using a Voronoi classification scheme to identify atom clusters which describes the local potential environment of an oxygen atom is shown to be a powerful method, describing a lattice position in crystalline corundum as well

as the properties of any selected oxygen atom in a model junction. The free parameter concern stated in section 5.5 is somewhat mitigated through this process, as clusters with as many atoms can be described with one Voronoi cell (a cluster of 24 atoms is described in figure 6.11 for example).

The affect of local strain on a crystalline cluster observed behaviour which does not align with the simple strain investigations undertaken with the idealised 6 atom model, suggesting that the deformation axis chosen is not the preferred strain axis identified in section 5.4 and that localised strain may not lead to TLS behaviour directly.

A small number of TLS candidates inside a JJ model constructed via the methods outlined in chapter 2, identified through Voronoi (and asymmetry) classification were studied but did not yield any TLS behaviour. Alternate classification systems were discussed which may be of future use to this research.

“Enfermé dans le navire, d’où on n’échappe pas, le fou est confiné à la rivière aux mille bras, à la mer aux mille chemins, à cette grande incertitude extérieure à tout. Il est prisonnier au milieu de la plus libre, de la plus ouverte des routes: solidement enchaîné à l’infini carrefour. Il est le Passager par excellence, c’est-à-dire le prisonnier du passage. Et la terre sur laquelle il abordera, on ne la connaît pas, tout comme on ne sait pas, quand il prend pied, de quelle terre il vient. Il n’a sa vérité et sa patrie que dans cette étendue inféconde entre deux terres qui ne peuvent lui appartenir.”

—Michel Foucault

7 Conclusions

Thesis summary, relevant conclusions and general outlook.

The current state of strongly coupled TLS models represents an embarrassment of riches: copious seemingly valid ideas in both phenomenological and microscopic flavours. It's possible that the phenomenological indistinguishability problem discussed in subsection 1.5.2 may be circumvented with the recent emergence of the microscopic descriptions. Furthermore, the ultimate reality of the situation may be that many of the microscopic models are not competing, but each represent a certain percentage of the total noise that presently mires Josephson junction devices. However, several of these models fail to account for the drastic reduction in TLS counts in epitaxial junctions [53] examined in section 1.7. This thesis has attempted to identify how an 80% reduction in visible TLSs can be observed solely by altering the JJ tunnel barrier architecture from amorphous to crystalline in nature.

Initially, we constructed atomic Josepson junction models using a hybrid *ab initio* and molecular mechanics approach in chapter 2 with various stoichiometry and density properties to reflect experimental observations of the barrier's makeup. The structure of the amorphous layer has been historically difficult to capture computationally, although our analysis shows excellent agreement with experimentally obtained coordination values. These junctions not only assist the analysis in this thesis, but also offer a tool for other microscopic models. For example, hydrogen could be introduced into the structure to investigate the dangling bond hypothesis.

Starting from the charge dipole phenomenological model, chapter 3 constructs the framework required to describe an oxygen atom which has the capacity to become spatially delocalised as bonds perturb away from a crystalline structure and become amorphous. Care was taken in choosing an empirical potential capable of describing both metallic and insulator regions due to the varying charge states an oxygen atom observes when present in a predominantly metallic environment (such as the metal–oxide boundary of a JJ). Treatment of errors introduced via the central difference method were also studied in detail, as the models precision (resolving splitting energies around 10 kHz) is of paramount importance. The resolution of *ab initio* methods such as DFT for comparison, has a lower bound of order 120 GHz—well above the

scale of crystal defect energies and therefore too imprecise to be considered useful for JJ study other than providing atomic positions as input.

Chapter 4 probes the capabilities of the model in low dimensions (1D and 2D), subsequently chapter 5 expands the dimensionality of the model to 2+1D and compares its results to current experimental strongly coupled TLS measurements in phase qubits. The model shows its capacity to explain how an oxygen atom can generate a large dipole and appropriate ground to first excited state splitting values expected of a TLS by merely migrating from its preferred lattice position. A prediction concerning a quad-degenerate rather than a doubly-degenerate ground state in very rare circumstances permits a possible validation method for experiments to scrutinise. Possible explanations of the relatively small range of observed qubit-TLS coupling strengths and behaviour when in contact with an external strain field are also presented.

To examine a complete three dimensional representation of the TLS model, chapter 6 outlines the pitfalls that memory limitations present us when using direct diagonalisaton approach in 3D. It introduces an alternate method exploiting a Wick-rotated time-dependent Schrödinger equation to obtain time-independent ground state wavefunctions, using Gram-Schmidt orthogonalisation to extract excited states. Equivalences between the low dimensional and 3D approaches were examined and found that the 3D solution is optimal when accurate values of TLS properties are required and the 2+1D solution to be more efficient when trends over large parameter ranges are needed. A Voronoi classification scheme is employed, which mitigates free parameter concerns and can be used to successfully describe the potential landscape in the vicinity of a defect region or oxygen atom residing in a JJ model. Properties of such atoms are investigated, as is the influence of external strain on a crystal of Al_2O_3.

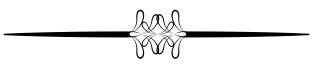

There are multiple ways in which the work contained in this thesis can be expanded upon to gain a better understanding of the strongly coupled TLS and perhaps verify the possible connection to the $1/f$ noise phenomenon.

Active research into a procedure which mimics the formation process of the Josephson junction tunnel barrier is already underway by M. Cyster as stated in chapter 2. It is unclear if the melt and quench cycle discussed in this

thesis ignores any relevant physics that may hold the key to environmental noise in these systems. Nucleation sites near the metal surface may be an example of such a phenomenon.

A complete investigation into TLS-phonon coupling may shed more light on the strain experiments [77] and extend the calculable properties of the delocalised oxygen TLS model to include dephasing and decoherence times. Similarly, a TLS-TLS interaction model such as the one proposed by Faoro and Ioffe [171] could be used to investigate the coherent interactions observed in reference 78. Further study focusing on the mapping between model dimensions would also be a fruitful endeavour. Whilst the STM continues to produce relevant information for TLS research in a concise and analytic manner, realistic values for the tunneling energy Δ and asymmetry energy ε could be obtained from the delocalised oxygen model rather than assuming some arbitrary amount if a consistent transformation of variables from the 3D implementation could be projected down to 1D.

Experimentally, it is suggested that the community moves away from ever cleaner fabrication processes and focuses its attempts on implementing more accessible methods of obtaining epitaxial junctions similar to Oh *et al.* [53]. An alternate route may be to investigate the substitution of oxygen to a heavier non metal in the insulating layer: lowering the tunneling probability of the atom regardless of the amorphous structure it resides in.

Or, perhaps a more pragmatic approach is to continue the active research into transmons with their small surface area and minimal TLS interactions. Perhaps also, the quantum phase slip junction may completely eradicate the need for this investigation all together. Then it seems that the road to take boils down to a personal philosophy: do you prefer to run from the unknown or tackle it head on?

Bibliography

1. F. Bloch, "Nuclear Induction", *Physical Review* **70**, 460–474 (1946). Cited on page 23.

2. G. Lindblad, "On the generators of quantum dynamical semigroups", *Communications in Mathematical Physics* **48**, 119–130 (1976). Cited on page 25.

3. S. Stenholm and K.-A. Suominen, *Quantum Approach to Informatics*. Wiley, 2005. Cited on page 25.

4. Los Alamos National Security LLC, "Quantum Information Science and Technology Roadmap", 2004. http://qist.lanl.gov/. Cited on page 26.

5. W. Paul, "Electromagnetic traps for charged and neutral particles", *Reviews of Modern Physics* **62**, 531–540 (1990). Cited on page 26.

6. J. Cirac and P. Zoller, "Quantum Computations with Cold Trapped Ions", *Physical Review Letters* **74**, 4091–4094 (1995). Cited on page 26.

7. F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, "Realization of the Cirac-Zoller controlled-NOT quantum gate.", *Nature* **422**, 408–411 (2003). Cited on page 26.

8. P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt, "A quantum information processor with trapped ions", *New Journal of Physics* **15**, 123012 (2013). Cited on page 26.

9. B. P. Lanyon, P. Jurcevic, M. Zwerger, C. Hempel, E. A. Martinez, W. Dür, H. J. Briegel, R. Blatt, and C. F. Roos, "Measurement-Based Quantum Computation with Trapped Ions", *Physical Review Letters* **111**, 210501 (2013). Cited on page 26.

10. T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. a. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, "14-Qubit Entanglement: Creation and Coherence", *Physical Review Letters* **106**, 130506 (2011). Cited on page 26.

11. D. Nigg, M. Müller, E. A. Martinez, P. Schindler, M. Hennrich, T. Monz, M. A. Martin-Delgado, and R. Blatt, "Quantum computations on a topologically encoded qubit.", *Science* **345**, 302–305 (2014). Cited on page 26.

12. A. Gruber, "Scanning Confocal Optical Microscopy and Magnetic Resonance on Single Defect Centers", *Science* **276**, 2012–2014 (1997). Cited on page 26.

13. A. Jarmola, V. M. Acosta, K. Jensen, S. Chemerisov, and D. Budker, "Temperature- and magnetic-field-dependent longitudinal spin relaxation in nitrogen-vacancy ensembles in diamond", *Physical Review Letters* **108**, 1–5 (2012). Cited on page 27.

14. P. C. Maurer, G. Kucsko, C. Latta, L. Jiang, N. Y. Yao, S. D. Bennett, F. Pastawski, D. Hunger, N. Chisholm, M. Markham, D. J. Twitchen, J. I. Cirac, and M. D. Lukin, "Room-Temperature Quantum Bit Memory Exceeding One Second", *Science* **336**, 1283–1286 (2012). Cited on page 27.

15. T. Yamamoto, S. Onoda, T. Ohshima, T. Teraji, K. Watanabe, S. Koizumi, T. Umeda, L. P. McGuinness, C. Müller, B. Naydenov, F. Dolde, H. Fedder, J. Honert, M. L. Markham, D. J. Twitchen, J. Wrachtrup, F. Jelezko, and J. Isoya, "Isotopic identification of engineered nitrogen-vacancy spin qubits in ultrapure diamond", *Physical Review B* **90**, 081117 (2014). Cited on page 27.

16. Y. Doi, T. Makino, H. Kato, D. Takeuchi, M. Ogura, H. Okushi, H. Morishita, T. Tashima, S. Miwa, S. Yamasaki, P. Neumann, J. Wrachtrup, Y. Suzuki, and N. Mizuochi, "Deterministic Electrical Charge-State Initialization of Single Nitrogen-Vacancy Center in Diamond", *Physical Review X* **4**, 011057 (2014). Cited on page 27.

17. M. V. Hauf, P. Simon, N. Aslam, M. Pfender, P. Neumann, S. Pezzagna, J. Meijer, J. Wrachtrup, M. Stutzmann, F. Reinhard, and J. A. Garrido, "Addressing Single Nitrogen-Vacancy Centers in Diamond with Transparent in-Plane Gate Structures", *Nano Letters* **14**, 2359–2364 (2014). Cited on page 27.

18. L. L. Chang, L. Esaki, and R. Tsu, "Resonant tunneling in semiconductor double barriers", *Applied Physics Letters* **24**, 593–595 (1974). Cited on page 27.

19. R. Dingle, W. Wiegmann, and C. Henry, "Quantum States of Confined Carriers in Very Thin $Al_xGa_{1-x}As-GaAs-Al_xGa_{1-x}As$ Heterostructures", *Physical Review Letters* **33**, 827–830 (1974). Cited on page 27.

20. A. Imamoğlu, D. D. Awschalom, G. Burkard, D. P. DiVincenzo, D. Loss, M. Sherwin, and A. Small, "Quantum Information Processing Using Quantum Dot Spins and Cavity QED", *Physical Review Letters* **83**, 4204–4207 (1999). Cited on page 27.

21. D. M. Zajac, T. M. Hazard, X. Mi, K. Wang, and J. R. Petta, "A Reconfigurable Gate Architecture for Si/SiGe Quantum Dots", arXiv:1502.01624. Cited on page 27.

22. P. Stano, J. Klinovaja, F. R. Braakman, L. M. K. Vandersypen, and D. Loss, "Fast long-distance control of spin qubits by photon-assisted cotunneling", *Physical Review B* **92**, 075302 (2015). Cited on page 27.

23. Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, "Coherent control of macroscopic quantum states in a single-Cooper-pair box", *Nature* **398**, 4 (1999). Cited on page 28.

24. V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret, "Quantum Coherence with a Single Cooper Pair", *Physica Scripta* **T76**, 165 (1998). Cited on page 28.

25. Y. Makhlin, G. Schön, and A. Shnirman, "Quantum-state engineering with Josephson-junction devices", *Reviews of Modern Physics* **73**, 357–400 (2001). Cited on pages 28 & 30.

26. D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. H. Devoret, "Manipulating the quantum state of an electrical circuit", *Science* **296**, 886–9 (2002). Cited on pages 28 & 38.

27. A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, "Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics", *Nature* **431**, 162–167 (2004). Cited on page 28.

28. J. Schreier, A. Houck, J. Koch, D. Schuster, B. Johnson, J. Chow, J. Gambetta, J. Majer, L. Frunzio, M. Devoret, S. Girvin, and R. Schoelkopf, "Suppressing charge noise decoherence in superconducting charge qubits", *Physical Review B* **77**, 180502 (2008). Cited on pages 28 & 38.

29. A. Barone and G. Paternò, *Physics and applications of the Josephson effect*. Wiley, New York, 1982. Cited on page 28.

30. J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal, and S. Lloyd, "Josephson Persistent-Current Qubit", *Science* **285**, 1036–1039 (1999). Cited on page 28.

31. C. H. van der Wal, A. C. ter Haar, F. K. Wilhelm, R. N. Schouten, C. J. P. M. Harmans, T. P. Orlando, S. Lloyd, and J. E. Mooij, "Quantum superposition of macroscopic persistent-current states", *Science* **290**, 773–777 (2000). Cited on page 28.

32. I. Chiorescu, Y. Nakamura, C. J. P. M. Harmans, and J. E. Mooij, "Coherent quantum dynamics of a superconducting flux qubit", *Science* **299**, 1869–1871 (2003). Cited on page 28.

33. J. Clarke, A. N. Cleland, M. H. Devoret, D. Esteve, and J. M. Martinis, "Quantum mechanics of a macroscopic variable: the phase difference of a josephson junction", *Science* **239**, 992–7 (1988). Cited on page 29.

34. J. Martinis, S. Nam, J. Aumentado, and C. Urbina, "Rabi Oscillations in a Large Josephson-Junction Qubit", *Physical Review Letters* **89**, 117901 (2002). Cited on pages 29 & 34.

35. R. Simmonds, K. Lang, D. Hite, S. Nam, D. Pappas, and J. Martinis, "Decoherence in Josephson Phase Qubits from Junction Resonators", *Physical Review Letters* **93**, 077003 (2004). Cited on pages 29, 34 & 36.

36. K. Geerlings, Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, "Demonstrating a Driven Reset Protocol for a Superconducting Qubit", *Physical Review Letters* **110**, 120501 (2013). Cited on page 29.

37. J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, "State preservation by repetitive error detection in a superconducting quantum circuit", *Nature* **519**, 66–69 (2015). Cited on page 29.

38. S. Shankar, M. Hatridge, Z. Leghtas, K. M. Sliwa, A. Narla, U. Vool, S. M. Girvin, L. Frunzio, M. Mirrahimi, and M. H. Devoret, "Autonomously stabilized entanglement between two superconducting quantum bits", *Nature* **504**, 419–422 (2013). Cited on page 29.

39. M. H. Devoret and R. J. Schoelkopf, "Superconducting Circuits for Quantum Information: An Outlook", *Science* **339**, 1169–1174 (2013). Cited on page 29.

40. A. Bezryadin, C. N. Lau, and M. Tinkham, "Quantum suppression of superconductivity in ultrathin nanowires", *Nature* **404**, 971–974 (2000). Cited on page 30.

41. J. E. Mooij and Y. V. Nazarov, "Superconducting nanowires as quantum phase-slip junctions", *Nature Physics* **2**, 169–172 (2006). Cited on page 30.

42. O. V. Astafiev, L. B. Ioffe, S. Kafanov, Y. A. Pashkin, K. Y. Arutyunov, D. Shahar, O. Cohen, and J. S. Tsai, "Coherent quantum phase slip", *Nature* **484**, 355–358 (2012). Cited on page 30.

43. B. Josephson, "Possible new effects in superconductive tunnelling", *Physics Letters* **1**, 251–253 (1962). Cited on page 30.

44. B. Josephson, "Supercurrents through barriers", *Advances in Physics* **14**, 419–451 (1965). Cited on page 30.

45. G. J. Dolan, "Offset masks for lift-off photoprocessing", *Applied Physics Letters* **31**, 337–339 (1977). Cited on page 31.

46. F. Lecocq, I. M. Pop, Z. Peng, I. Matei, T. Crozes, T. Fournier, C. Naud, W. Guichard, and O. Buisson, "Junction fabrication by shadow evaporation without a suspended bridge", *Nanotechnology* **22**, 315302 (2011). Cited on page 31.

47. B. G. Park, J. Y. Bae, and T. D. Lee, "Growth characteristics of Al oxide formed by ozone in magnetic tunnel junctions", *Journal of Applied Physics* **91**, 8789 (2002). Cited on pages 31 & 40.

48. E. Tan, P. Mather, A. Perrella, J. Read, and R. Buhrman, "Oxygen stoichiometry and instability in aluminum oxide tunnel barrier layers", *Physical Review B* **71**, 161401 (2005). Cited on pages 31, 40 & 48.

49. J. C. Barbour, R. G. Copeland, R. G. Dunn, N. Missert, L. P. Montes, K.-A. Son, and J. P. Sullivan, "The Electrical Properties of Native and Deposited Thin Aluminum Oxide Layers on Aluminum: Hydration Effects", in *The Electrochemical Society Meeting*. Sandia National Laboratories, Nov., 1998. http://www.osti.gov/scitech/servlets/purl/1916. Cited on pages 31, 48 & 52.

50. K. Gloos, P. J. Koppinen, and J. P. Pekola, "Properties of native ultrathin aluminium oxide tunnel barriers", *Journal of Physics: Condensed Matter* **15**, 1733–1746 (2003). Cited on page 31.

51. T. Aref, A. Averin, S. van Dijken, A. Ferring, M. Koberidze, V. F. Maisi, H. Nguyen, R. M. Nieminen, J. P. Pekola, and L. D. Yao, "Characterization of aluminum oxide tunnel barriers by combining transport measurements and transmission electron microscope imaging", *Journal of Applied Physics* **116**, 4 (2014). Cited on page 31.

52. L. J. Zeng, P. Krantz, S. Nik, P. Delsing, and E. Olsson, "The atomic details of the interfacial interaction between the bottom electrode of Al/AlO_x/Al Josephson junctions and HF-treated Si substrates", *Journal of Applied Physics* **117**, 163915 (2015). Cited on page 31.

53. S. Oh, K. Cicak, J. S. Kline, M. A. Sillanpää, K. D. Osborn, J. D. Whittaker, R. W. Simmonds, and D. P. Pappas, "Elimination of two level fluctuators in superconducting quantum bits by an epitaxial tunnel barrier", *Physical Review B* **74**, 100502 (2006). Cited on pages 31, 41, 133 & 135.

54. H. J. Carmichael, *Statistical Methods in Quantum Optics 1*. Springer, Berlin, 1999. Cited on page 32.

55. N. Vogt, J. H. Cole, M. Marthaler, and G. Schön, "Influence of two-level fluctuators on adiabatic passage techniques", *Physical Review B* **85**, 174515 (2012). Cited on page 32.

56. J. Jeske and J. H. Cole, "Derivation of Markovian master equations for spatially correlated decoherence", *Physical Review A* **87**, 052138 (2013). Cited on page 32.

57. J. Jeske, D. Ing, M. B. Plenio, S. F. Huelga, and J. H. Cole, "Bloch-Redfield equations for modeling light-harvesting complexes", *The Journal of Chemical Physics* **142**, 064104 (2015). Cited on page 32.

58. P. Dutta and P. Horn, "Low-frequency fluctuations in solids: $1/f$ noise", *Reviews of Modern Physics* **53**, 497–516 (1981). Cited on pages 32 & 38.

59. O. Astafiev, Y. A. Pashkin, Y. Nakamura, T. Yamamoto, and J. S. Tsai, "Quantum noise in the josephson charge qubit", *Physical Review Letters* **93**, 267007 (2004). Cited on page 32.

60. W. H. Zachariasen, "The Atomic Arrangement In Glass", *Journal of the American Chemical Society* **54**, 3841–3851 (1932). Cited on pages 32 & 33.

61. P. W. Anderson, B. I. Halperin, and C. M. Varma, "Anomalous low-temperature thermal properties of glasses and spin glasses", *Philosophical Magazine* **25**, 1–9 (1972). Cited on pages 32 & 106.

62. W. A. Phillips, "Tunneling states in amorphous solids", *Journal of Low Temperature Physics* **7**, 351–360 (1972). Cited on pages 32 & 106.

63. C. Enss and S. Hunklinger, *Low-Temperature Physics*. Springer-Verlag, Berlin/Heidelberg, 2005. Cited on page 34.

64. G. Zimmerli, T. M. Eiles, R. L. Kautz, and J. M. Martinis, "Noise in the Coulomb blockade electrometer", *Applied Physics Letters* **61**, 237 (1992). Cited on page 34.

65. K.-H. Ahn and P. Mohanty, "Quantum Friction of Micromechanical Resonators at Low Temperatures", *Physical Review Letters* **90**, 085504 (2003). Cited on page 34.

66. J. Gao, J. Zmuidzinas, B. a. Mazin, H. G. LeDuc, and P. K. Day, "Noise properties of superconducting coplanar waveguide microwave resonators", *Applied Physics Letters* **90**, 102507 (2007). Cited on page 34.

67. M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, E. Lucero, A. O'Connell, H. Wang, A. N. Cleland, and J. M. Martinis, "Process tomography of quantum memory in a Josephson-phase qubit coupled to a two-level state", *Nature Physics* **4**, 523–526 (2008). Cited on pages 34, 36, 78 & 105.

68. A. Lupaşcu, P. Bertet, E. F. C. Driessen, C. J. P. M. Harmans, and J. E. Mooij, "One- and two-photon spectroscopy of a flux qubit coupled to a microscopic defect", *Physical Review B* **80**, 172506 (2009). Cited on pages 34, 36, 78 & 104.

69. J. Lisenfeld, C. Müller, J. H. Cole, P. Bushev, A. Lukashenko, A. Shnirman, and A. V. Ustinov, "Measuring the Temperature Dependence of Individual Two-Level Systems by Direct Coherent Control", *Physical Review Letters* **105**, 230504 (2010). Cited on pages 34, 35, 78 & 122.

70. A. Shnirman, G. Schön, I. Martin, and Y. Makhlin, "Low- and High-Frequency Noise from Coherent Two-Level Systems", *Physical Review Letters* **94**, 127002 (2005). Cited on page 34.

71. A. Pourkabirian, M. V. Gustafsson, G. Johansson, J. Clarke, and P. Delsing, "Nonequilibrium Probing of Two-Level Charge Fluctuators Using the Step Response of a Single-Electron Transistor", *Physical Review Letters* **113**, 256801 (2014). Cited on page 34.

72. F. C. Wellstood, C. Urbina, and J. Clarke, "Flicker ($1/f$) noise in the critical current of Josephson junctions at 0.09-4.2 K", *Applied Physics Letters* **85**, 5296–5298 (2004). Cited on page 35.

73. M. V. Gustafsson, A. Pourkabirian, G. Johansson, J. Clarke, and P. Delsing, "Thermal properties of charge noise sources", *Physical Review B* **88**, 245410 (2013). Cited on page 35.

74. J. Martinis, K. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. Osborn, K. Cicak, S. Oh, D. Pappas, R. Simmonds, and C. Yu, "Decoherence in Josephson Qubits from Dielectric Loss", *Physical Review Letters* **95**, 210503 (2005). Cited on pages 35, 36, 38, 101, 104 & 122.

75. Y. Shalibo, Y. Rofe, D. Shwa, F. Zeides, M. Neeley, J. M. Martinis, and N. Katz, "Lifetime and Coherence of Two-Level Defects in a Josephson Junction", *Physical Review Letters* **105**, 177001 (2010). Cited on pages 36 & 104.

76. D. Gunnarsson, J.-M. Pirkkalainen, J. Li, G. S. Paraoanu, P. Hakonen, M. Sillanpää, and M. Prunnila, "Dielectric losses in multi-layer Josephson junction qubits", *Superconductor Science and Technology* **26**, 085010 (2013). Cited on page 36.

77. G. J. Grabovskij, T. Peichl, J. Lisenfeld, G. Weiss, and A. V. Ustinov, "Strain tuning of individual atomic tunneling systems detected by a superconducting qubit", *Science* **338**, 232–4 (2012). Cited on pages 36, 105, 106 & 135.

78. J. Lisenfeld, G. J. Grabovskij, C. Müller, J. H. Cole, G. Weiss, and A. V. Ustinov, "Observation of directly interacting coherent two-level systems in an amorphous material", *Nature Communications* **6**, 6182 (2015). Cited on pages 36 & 135.

79. J. Koch, T. Yu, J. Gambetta, A. Houck, D. Schuster, J. Majer, A. Blais, M. Devoret, S. Girvin, and R. Schoelkopf, "Charge-insensitive qubit design derived from the Cooper pair box", *Physical Review A* **76**, 042319 (2007). Cited on pages 36 & 38.

80. M. J. Schwarz, J. Goetz, Z. Jiang, T. Niemczyk, F. Deppe, A. Marx, and R. Gross, "Gradiometric flux qubits with a tunable gap", *New Journal of Physics* **15**, 045001 (2013). Cited on page 36.

81. J. H. Cole, C. Müller, P. Bushev, G. J. Grabovskij, J. Lisenfeld, A. Lukashenko, A. V. Ustinov, and A. Shnirman, "Quantitative evaluation of defect-models in superconducting phase qubits", *Applied Physics Letters* **97**, 3 (2010). Cited on pages 36, 37, 81, 89, 101, 104 & 110.

82. S. Sendelbach, D. Hover, A. Kittel, M. Mück, J. Martinis, and R. McDermott, "Magnetism in SQUIDs at Millikelvin Temperatures", *Physical Review Letters* **100**, 227006 (2008). Cited on page 37.

83. R. De Sousa, K. B. Whaley, T. Hecht, J. Von Delft, and F. K. Wilhelm, "Microscopic model of critical current noise in Josephson-junction qubits: Subgap resonances and Andreev bound states", *Physical Review B* **80**, 094515 (2009). Cited on page 37.

84. K. G. Wilson, "The renormalization group: Critical phenomena and the Kondo problem", *Reviews of Modern Physics* **47**, 773–840 (1975). Cited on page 37.

85. L. Faoro and L. Ioffe, "Microscopic origin of critical current fluctuations in large, small, and ultra-small area Josephson junctions", *Physical Review B* **75**, 132505 (2007). Cited on page 37.

86. L.-C. Ku and C. Yu, "Decoherence of a Josephson qubit due to coupling to two-level systems", *Physical Review B* **72**, 024526 (2005). Cited on page 37.

87. S. Choi, D.-H. Lee, S. Louie, and J. Clarke, "Localization of Metal-Induced Gap States at the Metal-Insulator Interface: Origin of Flux Noise in SQUIDs and Superconducting Qubits", *Physical Review Letters* **103**, 197001 (2009). Cited on page 37.

88. D. Lee, J. DuBois, and V. Lordi, "Identification of the Local Sources of Paramagnetic Noise in Superconducting Qubit Devices Fabricated on α-Al_2O_3 Substrates Using Density-Functional Calculations", *Physical Review Letters* **112**, 017001 (2014). Cited on page 38.

89. K. Agarwal, I. Martin, M. Lukin, and E. Demler, "Polaronic model of two-level systems in amorphous solids", *Physical Review B* **87**, 144201 (2013). Cited on page 38.

90. J. R. Jameson, D. Ngo, C. Benko, J. McVittie, Y. Nishi, and B. Young, "Dielectric relaxation study of hydrogen exposure as a source of two-level systems in Al_2O_3", *Journal of Non-Crystalline Solids* **357**, 2148–2151 (2011). Cited on page 38.

91. A. M. Holder, K. D. Osborn, C. J. Lobb, and C. B. Musgrave, "Bulk and Surface Tunneling Hydrogen Defects in Alumina", *Physical Review Letters* **111**, 065901 (2013). Cited on page 38.

92. L. Gordon, H. Abu-Farsakh, A. Janotti, and C. G. Van de Walle, "Hydrogen bonds in Al_2O_3 as dissipative two-level systems in superconducting qubits", *Scientific Reports* **4**, 7590 (2014). Cited on page 38.

93. A. Houck, J. Schreier, B. Johnson, J. Chow, J. Koch, J. Gambetta, D. Schuster, L. Frunzio, M. Devoret, S. Girvin, and R. Schoelkopf, "Controlling the Spontaneous Emission of a Superconducting Transmon Qubit", *Physical Review Letters* **101**, 080502 (2008). Cited on page 38.

94. M. Causà, R. Dovesi, C. Roetti, E. Kotomin, and V. Saunders, "A periodic ab initio Hartree-Fock calculation on corundum", *Chemical Physics Letters* **140**, 120–123 (1987). Cited on page 38.

95. E. Kotomin, I. Tale, P. Butlers, and P. Kulis, "Kinetics of non-steady state diffusion-controlled tunnelling recombination of defects in insulating crystals", *Journal of Physics: Condensed Matter* **1**, 6777–6785 (1989). Cited on page 38.

96. P. Kulis, Z. Rachko, M. Springis, I. Tale, and J. Jansons, "Defect assisted intrinsic luminescence in Al_2O_3 crystals", *Radiation Effects and Defects in Solids* **119-121**, 963–968 (1991). Cited on page 38.

97. P. W. M. Jacobs and E. A. Kotomin, "Modeling of Point Defects in Corundum Crystals", *Journal of the American Ceramic Society* **77**, 2505–2508 (1994). Cited on page 38.

98. Y.-N. Xu, Z.-Q. Gu, X.-F. Zhong, and W. Ching, "Ab initio calculations for the neutral and charged O vacancy in sapphire", *Physical Review B* **56**, 7277–7284 (1997). Cited on page 39.

99. S. K. Mohapatra and F. A. Kroger, "The Dominant Type of Atomic Disorder in α-Al_2O_3", *Journal of the American Ceramic Society* **61**, 106–109 (1978). Cited on page 39.

100. C. Catlow, R. James, W. Mackrodt, and R. Stewart, "Defect energetics in α-Al_2O_3 and rutile TiO_2", *Physical Review B* **25**, 1006–1026 (1982). Cited on pages 39 & 58.

101. K. Matsunaga, T. Tanaka, T. Yamamoto, and Y. Ikuhara, "First-principles calculations of intrinsic defects in Al_2O_3", *Physical Review B* **68**, 085110 (2003). Cited on page 39.

102. A. Heuer and K. Lagerlof, "Oxygen self-diffusion in corundum (α-Al_2O_3): A conundrum", *Philosophical Magazine Letters* **79**, 619–627 (1999). Cited on page 39.

103. J. Carrasco, J. Gomes, and F. Illas, "Theoretical study of bulk and surface oxygen and aluminum vacancies in α-Al_2O_3", *Physical Review B* **69**, 064116 (2004). Cited on page 39.

104. S. M. George, "Atomic layer deposition: an overview", *Chemical reviews* **110**, 111–31 (2010). Cited on page 40.

105. A. J. Elliot, G. Malek, L. Wille, R. Lu, S. Han, J. Z. Wu, J. Talvacchio, and R. M. Lewis, "Probing the Nucleation of Al_2O_3 in Atomic Layer Deposition on Aluminum for Ultrathin Tunneling Barriers in Josephson Junctions", *IEEE Transactions on Applied Superconductivity* **23**, 1101405 (2013). Cited on page 40.

106. T. V. Perevalov, O. E. Tereshenko, V. A. Gritsenko, V. A. Pustovarov, A. P. Yelisseyev, C. Park, J. H. Han, and C. Lee, "Oxygen deficiency defects in amorphous Al_2O_3", *Journal of Applied Physics* **108**, 013501 (2010). Cited on page 40.

107. H. Momida, S. Nigo, G. Kido, and T. Ohno, "Effect of vacancy-type oxygen deficiency on electronic structure in amorphous alumina", *Applied Physics Letters* **98**, 042102 (2011). Cited on pages 40 & 53.

108. E. Artacho, F. Ynduráin, B. Pajot, R. Ramírez, C. Herrero, L. Khirunenko, K. Itoh, and E. Haller, "Interstitial oxygen in germanium and silicon", *Physical Review B* **56**, 3820–3833 (1997). Cited on page 42.

109. E. Artacho, A. Lizón-Nordström, and F. Ynduráin, "Geometry and quantum delocalization of interstitial oxygen in silicon", *Physical Review B* **51**, 7862–7865 (1995). Cited on page 42.

110. T. Campbell, R. Kalia, A. Nakano, P. Vashishta, S. Ogata, and S. Rodgers, "Dynamics of Oxidation of Aluminum Nanoclusters using Variable Charge Molecular-Dynamics Simulations on Parallel Computers", *Physical Review Letters* **82**, 4866–4869 (1999). Cited on page 46.

111. X. W. Zhou and H. N. G. Wadley, "Atomistic simulation of AlO_x magnetic tunnel junction growth", *Physical Review B* **71**, 1–8 (2005). Cited on page 46.

112. A. Hasnaoui, O. Politano, J. M. Salazar, G. Aral, R. K. Kalia, A. Nakano, and P. Vashishta, "Molecular dynamics simulations of the nano-scale room-temperature oxidation of aluminum single crystals", *Surface Science* **579**, 47–57 (2005). Cited on page 46.

113. S. Morohashi and S. Hasuo, "Experimental investigations and analysis for high-quality Nb/Al—AlO_x/Nb Josephson junctions", *Journal of Applied Physics* **61**, 4835 (1987). Cited on page 46.

114. H. Kohlstedt, G. Hallmanns, I. Nevirkovets, D. Guggi, and C. Heiden, "Preparation and properties of Nb/Al—AlO_x/Nb multilayers", *IEEE Transactions on Appiled Superconductivity* **3**, 2197–2200 (1993). Cited on page 46.

115. L. P. H. Jeurgens, W. G. Sloof, F. D. Tichelaar, and E. J. Mittemeijer, "Structure and morphology of aluminium-oxide films formed by thermal oxidation of aluminium", *Thin Solid Films* **418**, 89–101 (2002). Cited on page 46.

116. P. Vashishta, R. K. Kalia, A. Nakano, and J. P. Rino, "Interaction potentials for alumina and molecular dynamics simulations of amorphous and liquid alumina", *Journal of Applied Physics* **103**, 083504 (2008). Cited on page 46.

117. H. W. Sheng, E. Ma, and M. J. Kramer, "Relating Dynamic Properties to Atomic Structure in Metallic Glasses", *Journal of The Minerals, Metals & Materials Society* **64**, 856–881 (2012). Cited on page 46.

118. V. V. Ogorodnikov and Y. I. Rogovoi, "Buckingham-potential parameters and relations between them for solids", *Powder Metallurgy and Metal Ceramics* **33**, 327–333 (1995). Cited on page 46.

119. R. A. Buckingham, "The Classical Equation of State of Gaseous Helium, Neon and Argon", *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **168**, 264–283 (1938). Cited on page 46.

120. W. Kohn and L. J. Sham, "Self-Consistent Equations Including Exchange and Correlation Effects", *Physical Review* **140**, A1133–A1138 (1965). Cited on page 47.

121. G. Kresse and J. Hafner, "Norm-conserving and ultrasoft pseudopotentials for first-row and transition elements", *Journal of Physics: Condensed Matter* **6**, 8245–8257 (1994). Cited on page 47.

122. G. Kresse and J. Furthmüller, "Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set", *Computational Materials Science* **6**, 15–50 (1996). Cited on page 47.

123. G. Kresse, "Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set", *Physical Review B* **54**, 11169–11186 (1996). Cited on page 47.

124. G. Kresse, "From ultrasoft pseudopotentials to the projector augmented-wave method", *Physical Review B* **59**, 1758–1775 (1999). Cited on page 47.

125. P. E. Blöchl, "Projector augmented-wave method", *Physical Review B* **50**, 17953–17979 (1994). Cited on page 47.

126. J. P. Perdew, K. Burke, and M. Ernzerhof, "Generalized Gradient Approximation Made Simple", *Physical Review Letters* **77**, 3865–3868 (1996). Cited on page 47.

127. J. D. Gale and A. L. Rohl, "The General Utility Lattice Program (GULP)", *Molecular Simulation* **29**, 291–341 (2003). Cited on pages 48, 59 & 125.

128. F. Streitz and J. Mintmire, "Electrostatic potentials for metal-oxide surfaces and interfaces", *Physical Review B* **50**, 11996–12003 (1994). Cited on pages 48, 58, 59 & 118.

129. B. G. Levine, J. E. Stone, and A. Kohlmeyer, "Fast Analysis of Molecular Dynamics Trajectories with Graphics Processing Units-Radial Distribution Function Histogramming.", *Journal of Computational Physics* **230**, 3556–3569 (2011). Cited on page 49.

130. N. Ishizawa, T. Miyata, I. Minato, F. Marumo, and S. Iwai, "A structural investigation of α-Al_2O_3 at 2170 K", *Acta Crystallographica Section B* **36**, 228–230 (1980). Cited on page 49.

131. S. M. El-mashri, R. G. Jones, and A. J. Forty, "An Electron-Yield EXAFS Study of Anodic-Oxide and Hydrated-Oxide Films on Pure Aluminum", *Philosophical Magazine A* **48**, 665–683 (1983). Cited on page 53.

132. A. J. Bourdillon, S. M. El-mashri, and A. J. Forty, "Application of TEM extended electron energy loss fine structure to the study of aluminium oxide films", *Philosophical Magazine A* **49**, 341–352 (1984). Cited on page 53.

133. S. K. Lee, S. B. Lee, S. Y. Park, Y. S. Yi, and C. W. Ahn, "Structure of amorphous aluminum oxide", *Physical Review Letters* **103**, 095501 (2009). Cited on page 53.

134. U. Buchenau, N. Nücker, and A. Dianoux, "Neutron Scattering Study of the Low-Frequency Vibrations in Vitreous Silica", *Physical Review Letters* **53**, 2316–2319 (1984). Cited on page 57.

135. A. Heuer, "Microscopic view of the low-temperature anomalies in glasses", in *Tunneling systems in amorphous and crystalline solids*, P. Esquinazi, ed., ch. 8. Springer, 1998. Cited on page 57.

136. K. O. Trachenko, M. T. Dove, M. J. Harris, and V. Heine, "Dynamics of silica glass: two-level tunnelling states and low-energy floppy modes", *Journal of Physics: Condensed Matter* **12**, 8041–8064 (2000). Cited on page 57.

137. J. Reinisch and A. Heuer, "What Is Moving in Silica at 1 K? A Computer Study of the Low-Temperature Anomalies", *Physical Review Letters* **95**, 155502 (2005). Cited on page 57.

138. G. Dienes, D. Welch, C. Fischer, R. Hatcher, O. Lazareth, and M. Samberg, "Shell-model calculation of some point-defect properties in α-Al_2O_3", *Physical Review B* **11**, 3060–3070 (1975). Cited on page 58.

139. A. K. Rappe and W. A. Goddard, "Charge equilibration for molecular dynamics simulations", *The Journal of Physical Chemistry* **95**, 3358–3363 (1991). Cited on page 58.

140. X. Zhou, H. Wadley, J.-S. Filhol, and M. Neurock, "Modified charge transfer–embedded atom method potential for metal/metal oxide systems", *Physical Review B* **69**, 035402 (2004). Cited on pages 58 & 59.

141. C. A. Ocampo and R. Mathur, "Variational Model for Optimization of Finite-Burn Escape Trajectories Using a Direct Method", *Journal of Guidance, Control, and Dynamics* **35**, 598–608 (2012). Cited on page 60.

142. W. Squire and G. Trapp, "Using Complex Variables to Estimate Derivatives of Real Functions", *SIAM Review* **40**, 110–112 (1998). Cited on page 60.

143. G. Lantoine, R. P. Russell, and T. Dargent, "Using Multicomplex Variables for Automatic Computation of High-Order Derivatives", *ACM Transactions on Mathematical Software* **38**, 1–21 (2012). Cited on page 60.

144. G. Kedem, "Automatic Differentiation of Computer Programs", *ACM Transactions on Mathematical Software* **6**, 150–165 (1980). Cited on page 61.

145. R. Mathur, *An Analytical Approach to Computing Step Sizes for Finite-Difference Derivatives*. PhD thesis, University of Texas, 2012. http://hdl.handle.net/2152/ETD-UT-2012-05-5275. Cited on pages 61, 67, 69, 70, 73, 74, 75 & 77.

146. M. D. Greenberg, *Foundations of Applied Mathematics*. Prentice Hall, 1978. Cited on page 61.

147. J. H. Mathews and K. D. Fink, *Numerical Methods Using Matlab*. Pearson, 4th ed., 2004. Cited on pages 62, 63 & 66.

148. P. Holoborodko, "Central Differences", 2009. http://www.holoborodko.com/pavel/numerical-methods/numerical-derivative/central-differences/. Cited on page 63.

149. N. J. Cyrus and R. E. Fulton, "Accuracy study of finite difference methods", Tech. Rep. NASA-TN-D-4372, NASA, 1968. http://hdl.handle.net/2060/19680007599. Cited on page 65.

150. R. Skeel, "Roundoff Error Cripples Patriot Missile", *SIAM News* **25**, 11 (1992). Cited on page 65.

151. P. Gill, W. Murray, and M. Wright, *Practical Optimisation*. Emerald, 1982. Cited on pages 66 & 67.

152. R. S. Stepleman and N. D. Winarsky, "Adaptive Numerical Differentiation", *Mathematics of Computation* **33**, 1257 (1979). Cited on page 67.

153. R. R. Barton, "Computing Forward Difference Derivatives In Engineering Optimisation", *Engineering Optimization* **20**, 205–224 (1992). Cited on page 67.

154. Mathworks, "eigs", 2014. http://www.mathworks.com.au/help/matlab/ref/eigs.html. Cited on page 75.

155. A. Kofman, Q. Zhang, J. Martinis, and A. Korotkov, "Theoretical analysis of measurement crosstalk for coupled Josephson phase qubits", *Physical Review B* **75**, 014524 (2007). Cited on page 104.

156. W. Tang, E. Sanville, and G. Henkelman, "A grid-based Bader analysis algorithm without lattice bias", *Journal of physics. Condensed matter : an Institute of Physics journal* **21**, 084204 (2009). Cited on page 104.

157. J. Lisenfeld, C. Müller, J. H. Cole, P. Bushev, A. Lukashenko, A. Shnirman, and A. V. Ustinov, "Rabi spectroscopy of a qubit-fluctuator system", *Physical Review B* **81**, 100511 (2010). Cited on page 105.

158. I. W. Sudiarta and D. J. W. Geldart, "The finite difference time domain method for computing the single-particle density matrix", *Journal of Physics A: Mathematical and Theoretical* **42**, 285002 (2009). Cited on page 114.

159. G. C. Wick, "Properties of Bethe-Salpeter wave functions", *Physical Review* **96**, 1124–1134 (1954). Cited on page 114.

160. A. Taflove and K. R. Umashankar, "The Finite-Difference Time-Domain Method for Numerical Modeling of Electromagnetic Wave Interactions", *Electromagnetics* **10**, 105–126 (1990). Cited on page 114.

161. J. Gram, "Ueber die Entwickelung reeller Functionen in Reihen mittelst der Methode der kleinsten Quadrate", *Journal für die reine und angewandte Mathematik* **94**, 41–73 (1883). Cited on page 115.

162. E. Schmidt, "Zur Theorie der linearen und nichtlinearen Integralgleichungen", *Mathematische Annalen* **63**, 433–476 (1907). Cited on page 115.

163. M. Strickland and D. Yager-Elorriaga, "A parallel algorithm for solving the 3d Schrödinger equation", *Journal of Computational Physics* **229**, 6015–6026 (2010). Cited on page 116.

164. T. C. DuBois, "Wafer", 2015. http://wafer.tcqp.science. Cited on page 116.

165. G. Voronoi, "Nouvelles applications des paramètres continus à la théorie des formes quadratiques. Premier mémoire. Sur quelques propriétés des formes quadratiques positives parfaites", 1908. Cited on page 125.

166. J. Bass, "Elasticity of minerals, glasses and melts", in *Mineral physics and crystallography: a handbook of physical constants*, T. Ahrens, ed., pp. 45–63. American Geophysical Union, Washington DC, 1995. Cited on page 126.

167. S. D. Poisson, "Mémoire sur l'équilibre et le mouvement des corps élastiques", in *Mémoires de l'Académie des Sciences*, p. 8. Académie Des Sciences, Paris, 1829. Cited on page 126.

168. H. Gercek, "Poisson's ratio values for rocks", *International Journal of Rock Mechanics and Mining Sciences* **44**, 1–13 (2007). Cited on page 126.

169. A. V. Anikeenko, M. G. Alinchenko, V. P. Voloshin, N. N. Medvedev, M. L. Gavrilova, and P. Jedlovszky, "Implementation of the Voronoi-Delaunay Method for Analysis of Intermolecular Voids", *Computational Science and Its Applications – ICCSA 2004* **3045**, 217–226 (2004). Cited on pages 127 & 129.

170. A. P. Paz, I. V. Lebedeva, I. V. Tokatly, and A. Rubio, "Identification of structural motifs as tunneling two-level systems in amorphous alumina at low temperatures", *Physical Review B* **90**, 224202 (2014). Cited on page 129.

171. L. Faoro and L. B. Ioffe, "Interacting tunneling model for two-level systems in amorphous materials and its predictions for their dephasing and noise in superconducting microresonators", *Physical Review B* **91**, 014201 (2015). Cited on page 135.