

Constrained Models in Aperiodic Systems

Shobhna Singh

A thesis submitted for the degree of
Doctor of Philosophy.

Cardiff University

September 2025

Abstract

Constrained models play a pivotal role in condensed matter physics by capturing essential features of strongly interacting many-body systems. These models impose local or global restrictions that mirror real physical constraints, enabling the study of complex emergent behaviors. For instance, hard-core dimer models reflect physical systems with strong local constraints due to repulsive interactions. Similarly, loop models, vertex models, and height representations provide constrained frameworks that reveal rich phase structures and critical phenomena. In quantum systems, constraints often emerge naturally from microscopic Hamiltonian and lead to phenomena such as fractionalization and quantum spin liquids. Constrained models act as fertile grounds for theoretical insight, enabling exact solutions, mapping to field theories, and exploring universal behaviors. They also bridge condensed matter with disciplines like statistical mechanics, combinatorics, and computational complexity, offering a unified language to study order, disorder, and transitions in systems ranging from magnets to artificial spin ices and cold atomic gases.

This thesis investigates critical phenomena, strongly correlated models, and optimization problems on aperiodic lattices, with a particular focus on NP-hard problems, dimer models, and their classical and quantum variants. Moving beyond periodic lattices like square and honeycomb structures, this research delves into the less understood area of aperiodic systems, such as quasicrystals, which exhibit long-range order without translational symmetry. These structures challenge traditional paradigms of condensed matter physics and open avenues for discovering new critical behaviors.

I present original contributions that combine analytical methods and numerical algorithms to study constrained systems on quasicrystalline tilings. These include the solution of NP-hard optimization problems like the Hamiltonian cycle on Ammann-Beenker tilings, an energy-based worm algorithm for simulating classical dimers on random graphs and modified Penrose tilings, and the analytical solution of the dimer partition function on the newly discovered aperiodic Spectre tiling. I further extend the study to quantum systems by analytically investigating the quantum dimer model (QDM) on the Spectre tiling and numerically analyzing it on the Ammann-Beenker tiling, revealing distinctive ground-state properties and quantum phase transitions that arise due to aperiodicity and local tiling constraints.

This work builds interdisciplinary bridges between condensed matter physics, graph theory, and computational complexity. It also sheds light on the feasibility and limitations of transferring universal results from periodic systems to aperiodic ones.

Contents

| | |
|--|-------------|
| Abstract | v |
| List of Figures | viii |
| Acknowledgements | xix |
| 1 Introduction | 1 |
| 1.1 Preface | 1 |
| 1.2 Thesis Overview | 3 |
| 1.3 Contribution During the PhD | 4 |
| 1.3.1 Publications on Work Related to the Thesis | 4 |
| 1.3.2 Presentations and Posters on Work Related to the Thesis | 4 |
| 2 Aperiodic Systems: Quasicrystals | 7 |
| 2.1 Ammann-Beenker (AB) Tilings | 8 |
| 2.2 Penrose Tiling | 11 |
| 2.2.1 Trivalent-Penrose Tiling | 11 |
| 2.3 Einstein Problem: New aperiodic Monotilings | 12 |
| 2.3.1 Hat Tiling | 12 |
| 2.3.2 Spectre Tiling | 15 |
| 2.4 Graph Theory Terminology | 18 |
| 3 Dimer and Loop Models | 19 |
| 3.1 Classical dimer model | 20 |
| 3.2 Quantum dimer model | 21 |
| 3.2.1 The General structure of QDM ground state phase diagrams | 22 |
| 3.3 The $O(n)$ Loop Model | 24 |
| 3.4 Height Model Framework | 25 |
| 4 NP-Hard Problems on the AB tiling | 29 |
| 4.1 Ammann-Beenker Tilings - Key Properties | 30 |
| 4.2 Hamiltonian Cycle Problem | 31 |
| 4.2.1 Constructing Fully Packed Loops on AB* | 32 |
| 4.2.2 Constructing Hamiltonian cycles on AB | 32 |
| 4.3 Solution to other NP-Hard problems on AB | 40 |
| 4.3.1 Equal-weight travelling salesperson problem | 40 |
| 4.3.2 Application: ultra - high resolution atomic imaging | 40 |
| 4.3.3 Longest Path Problem | 41 |
| 4.3.4 Application: Adsorption and catalysis | 41 |
| 4.3.5 Three-coloring Problem | 42 |
| 4.3.6 Application: The Potts model and protein folding | 43 |
| 4.3.7 Minimum dominating set problem | 45 |
| 4.3.8 Induced Path Problem | 46 |
| 4.4 Conclusion | 47 |

| | | |
|----------|---|------------|
| 5 | Classical Dimer Model on Random Planar Graphs | 49 |
| 5.1 | Random Graphs | 51 |
| 5.2 | Modified (Trivalent) Penrose Tiling | 54 |
| 5.3 | Numerical methods: Worm algorithm | 54 |
| 5.4 | Results | 55 |
| 5.5 | Conclusion | 61 |
| 6 | Dimer models on the Spectre tiling | 63 |
| 6.1 | Exact Solution to the Classical Dimer Model | 64 |
| 6.2 | Numerical Results | 65 |
| 6.3 | Exact Solution to the Quantum Dimer Model | 67 |
| 6.3.1 | Monomer Deconfinement | 69 |
| 7 | The Quantum dimer model on the AB tiling | 71 |
| 7.1 | Studying QDM on the AB* tiling using MPS-DMRG | 72 |
| 7.1.1 | Methods: MPS-DMRG on AB ladders | 74 |
| 7.1.2 | Results: Ground-State Phase Diagram | 83 |
| 7.1.3 | Phase Transitions | 90 |
| 7.2 | Conclusion | 92 |
| 8 | The Fully packed loop Model on AB tiling | 95 |
| 8.1 | Construction of Fully Packed Loops on AB | 95 |
| 8.2 | Infinite Temperature - Criticality | 96 |
| 8.2.1 | Geometric critical exponents of Loops | 98 |
| 8.3 | Conclusion | 105 |
| 9 | Conclusion and Future Directions | 107 |
| 9.1 | Summary | 107 |
| 9.2 | Outlook | 108 |

List of Figures

| | | |
|------|--|----|
| 2.1 | The Ho-Mn-Zn dodecahedral quasicrystal (left) alongside the 10-fold symmetric diffraction pattern of the first naturally occurring quasicrystal, Icosahedrite (right). (Following Refs. [1, 2]) | 8 |
| 2.2 | Cut-and-Project Method: Generating a 1D Fibonacci quasicrystal by projecting a 2D square lattice \mathbb{Z}^2 (gray circles). A stripe with irrational slope $1/\tau$ (where $\tau = (1 + \sqrt{5})/2$) is chosen, and lattice nodes within the stripe are projected onto a line parallel to it. The resulting points form intervals of two lengths, S (short, orange) and L (long, purple), creating a Fibonacci tiling. This tiling can also be generated by substitution rules: $L \rightarrow SL$, $S \rightarrow L$, producing sequences of lengths $1, 1, 2, 3, 5, \dots$, corresponding to Fibonacci numbers. Tiles S and L have lengths 1 and τ , respectively. | 9 |
| 2.3 | A patch of Ammann-Beenker tiling, created using two types of unit cells: squares and rhombuses. It exhibits quasi-periodicity and local 8-fold rotational symmetry. | 9 |
| 2.4 | The inflation rule σ : Tiles are subdivided into smaller tiles, followed by rescaling (inflation) of all lengths by the silver ratio $\delta_S = 1 + \sqrt{2}$. Red and blue points indicate vertices of single- and double-inflated tiles, respectively. Note that the triangles in the first and second inflations are actually half-square tiles. Also, inflating the square tile breaks its rotational symmetry, so the orientation must be tracked, as indicated by the triangular motif. For clarity, these markings will generally be omitted in subsequent figures. (Following Ref. [3]) | 10 |
| 2.5 | Illustration of discrete scale invariance in AB tilings. The thick black lines with red and blue vertices depict an AB tiling. After applying the inflation rules from Fig. 2.4 twice, a dense underlying tiling emerges, composed of gray and white tiles, which itself forms an AB tiling. | 10 |
| 2.6 | Inflation rules for Penrose tiles, where each tile is subdivided into smaller tiles and rescaled by the silver ratio $\delta_S = 1 + \sqrt{2}$. Red nodes mark the vertices of single-inflated tiles. | 11 |
| 2.7 | A patch of Penrose tiling, created using two types of unit cells: thick and thin rhombuses. It exhibits quasi-periodicity and local 5-fold rotational symmetry. | 12 |
| 2.8 | Trivalent Penrose tiling, depicted with thick black and blue edges, and circled gray nodes. It is derived from the underlying Penrose tiling (light gray lines) using the method described above. Following Ref. [4], I thank Dr. Nick Jones for this idea. | 13 |
| 2.9 | A patch of ‘Hat’ aperiodic monotiling embedded within the periodic hexagonal lattice (solid grey lines). A single hat tile is highlighted in grey, composed of eight smaller kite shapes (dashed grey lines). | 13 |
| 2.10 | A patch of Hat tiling, with reflected (mirror image) Hat tiles shown in dark blue, sparsely distributed throughout the pattern. Non-overlapping clusters of tiles are grouped to form H, T, P, F supertiles by simplifying cluster boundaries (outlined in black). The left panel demonstrates how these supertiles combine to create larger patches of the Hat tiling. (The figure was created using an interactive application, following Ref. [5].) | 14 |
| 2.11 | Inflation rules for the Hat supertiles (H, T, P , and F), illustrating the hierarchical expansion of the Hat tiling. (The figure was created using an interactive application, following Ref. [5].) | 14 |
| 2.12 | The first three iterations of the P supertile using inflation rules. Starting with the pink P supertile, one inflation produces the underlying yellow tiling. A second inflation of the yellow tiling results in the underlying black tiling, and the final iteration yields a dense Hat tiling. This hierarchical growth demonstrates the discrete scale invariance of the Hat tiling. (The figure was created using an interactive application, following Ref. [5].) | 15 |

| | | |
|------|--|----|
| 2.13 | (a) A single Spectre tile (above) and a Tile (1,1) with 14 vertices (below). The (black) vertex at the midpoint of a straight edge is implied by edge connections. (b) A patch of Spectre aperiodic monotiling using just rotations and translations of the spectre tile. The tiles shaded in dark green are called the ‘Upper Mystics’. (The figure was created using an interactive application, following Ref. [6].) | 16 |
| 2.14 | Two inflation rules for constructing the Spectre tiling. The rules are based on a single Spectre tile (white) and a two-Spectre combination called a Mystic (dark and light green). The first rule (above) replaces the Spectre with a cluster containing a Mystic and seven reflected Spectres. The second rule (below) replaces a Mystic with a cluster containing a Mystic and six Spectres. (The figure was created using an interactive application, following Ref. [6].) | 16 |
| 2.15 | An iterative process for constructing Spectre and Mystic clusters. In Step 1, a single Spectre tile is defined by four key points (red nodes). These key points guide the placement of additional Spectre tiles and Mystic tiles in Step 2a, which creates a cluster of Spectre and Mystic tiles. In Step 2b, the key points (red nodes) for higher-level clusters are determined, allowing for the continued growth of the tiling. Steps 3a and 3b further expand the structure into superclusters, with Spectre clusters highlighted in red and Mystic clusters in blue. This iterative process demonstrates the hierarchical nature of the Spectre tiling system. (The figure was created using an interactive application, following Ref. [6].) | 17 |
| 3.1 | The general ground state phase diagram of the classical dimer model on regular periodic lattices with aligning interactions features ordered columnar and staggered ground state phases, corresponding to the maximum and minimum number of flippable plaquettes, respectively. These ground states exhibit monomer confinement. On the square lattice, the ground states are fourfold degenerate. At higher temperatures, a critical phase emerges. At infinite temperature ($T = \infty$), the partition function reduces to the unweighted sum over all possible dimer coverings of the square lattice, where the model can be solved exactly. | 20 |
| 3.2 | Generic phase diagram of the quantum dimer model on 2D bipartite lattices. The RK point marks a critical deconfined point. The staggered phase appears immediately to the right of the RK point, while the plaquette phase lies to its left and spans a finite parameter range. For $V/J \rightarrow -\infty$, the columnar phase dominates. The region between the columnar and plaquette phases remains debated. On the honeycomb lattice, the plaquette phase transitions to the columnar phase via a first-order transition around $v/t = -0.2$ [7]. A similar scenario is proposed for the square lattice [8], though some findings hint at a potential ‘mixed’ phase between the plaquette and columnar phases [9]. Representations of all three phases are shown for the square lattice. | 22 |
| 3.3 | A configuration of fully packed loops (thick black lines) on a square lattice (dotted lines). The loops are non-intersecting, self-avoiding, and visit every vertex of the lattice exactly once. | 24 |
| 3.4 | (a) In the square lattice, the sites can be partitioned into two bipartite sets: A (blue) and B (red). (b-c) Two examples of height representations in the fully packed dimer model on a square lattice, constructed using the rules described in Sec. 3.4. The construction begins from the top leftmost plaquette, which is assigned a height zero. These rules apply to any 2D bipartite lattice. | 26 |
| 3.5 | Example of height representations in the fully packed loop model on a square lattice, constructed using the rules described in Sec. 3.4. These rules apply to any 2D bipartite lattice. | 27 |
| 4.1 | The seven vertex configurations and their transformations under inflation (Fig. 4). Each vertex maps to an 8-vertex after at most two inflation steps. (Following Ref.[10]) | 31 |
| 4.2 | Twice-inflations of the two AB prototiles. These represents level one (L_1) of the tiling. Each small tile is at level zero (L_0). The thick black edges visit all but the corner vertices. These edges are called e_0 edges. Since the 8-vertices at L_0 corresponds to the tile corners at L_1 , the union of e_0 constitutes fully packed loops on AB^* (AB without the 8-vertices). Augmenting the red path (switching covered and uncovered edges) places e_0 edges on the two corner vertices while still visiting the original vertices. The red path can be thought of as the twice-inflation of an e_0 edge, and denoted as e_1 . It has the same effect along any tile edge. (Following Ref. [3]) | 33 |

| | | |
|------|--|----|
| 4.3 | A patch of AB tiling showing the eightfold symmetric region W_1 , and W_0 in gray. W_1 is the twice inflation of W_0 , formed from the double-inflated tiles shown in Fig. 4.2. Purple lines show the boundary of these L_1 tiles (and do not belong to the tiling itself). The thick black edges form a fully packed loop set on the AB* patch. (Following Ref. [3]) | 33 |
| 4.4 | The red alternating path in Fig. 4.2 can be thought of as a twice-inflation of an e_0 edge from which loops are constructed at level zero (L_0). It connects nearest neighbors in the L_1 tiling (8-vertices in the L_0 tiling) while following an alternating path in the canonical choice of e_0 placements. I therefore denote this path an e_1 edge. Alternating paths connecting higher-order 8_n -vertices can be constructed by concatenating e_{n+1} edges in the same way that e_1 was formed from a concatenation of e_0 ; levels e_2 and e_3 are shown here. The orientation of e_n is defined according to the arrow. (This figure was made by my supervisor.) | 34 |
| 4.5 | The canonical placements of e_0 edges (black) form e_0 -loops visiting every $8_{n<0}$ vertex (i.e. every vertex which is not an 8-vertex). The red and blue lines form e_1 - and e_2 -loops, respectively. The central 8-vertex is now added onto loops by folding one corner of the level ' n ' inflation (blue star) inwards to break the 8-fold symmetry. Here we have shown the process explained in Sec. 4.2.2 to order $n = 2$, but it can be iterated to any order n . The green outline represents the largest disk that fits within the blue loop, connecting three 8-vertices with green chords. (Following Ref. [3]) | 36 |
| 4.6 | Next, e_1 edges from Fig. 4.4 are placed along the (red) e_1 loops to incorporate 8_0 -vertices into the loops. Subsequently, e_2 edges from Fig. 4.4 are added along the (blue) e_2 loop, with one corner folded, to include the central vertex and 8_1 -vertices into the loops. Note that some blue e_2 edges overlap with the red e_1 edges. This process can be iterated to any desired order. (Following Ref. [3]) | 37 |
| 4.7 | The overlapping e_2 edges and e_1 edges as shown in Fig. 4.6 will augment to leave the L_0 tiling unchanged. Hence we can remove these overlapping e_2 edges and e_1 edges completely. This results in the rewiring of e_1 -loops and the e_2 -loop into a single loop as shown here. Finally augmenting this single loop, consisting of the remaining red and blue edges, will gives the Hamiltonian cycle shown in Fig. 4.8. (Following Ref. [3]) | 38 |
| 4.8 | The Hamiltonian cycle visiting all the vertices of a U_2 region (the AB tiles themselves are omitted for clarity). The image is obtained from Fig. 4.5 by placing e_1 and e_2 edges in alternate orientations along the red and blue loops respectively, as shown in Fig. 4.6, and augmenting. Note the present figure has been rotated through $1/16^{\text{th}}$ of a turn relative to other figures. (Following Ref. [3]) | 39 |
| 4.9 | (a),(b) Tiles used for 3-coloring of the AB tiling. (c) The structure with three (shaded) rhombuses present around every ladder. (Following Ref. [3]) | 42 |
| 4.10 | A small patch of 3-colored AB tiling using the steps as mentioned in Sec 4.3.5. Tile 1 and tile 2 are outlined in gold. The only remaining portions of the tiling are segments of ladders that can be 3-colored consistently based on their surrounding tiles. (Following Ref. [3]) | 43 |
| 4.11 | A fully 3-colored AB tiling, featuring tile 1 and tile 2 (both outlined in gold), with all ladders consistently colored as described in Sec. 4.3.5. (Following Ref. [3]) | 44 |
| 4.12 | A portion of the AB tiling showing the solution to the minimum dominating set problem in red, which, as proved in Sec. 4.3.7, are also vertices of a single inflated AB tiling. (Following Ref. [3]) | 45 |
| 4.13 | FPLs on AB* created using the method in Sec. 4.2.1 also forms induced paths on AB. The total region shown here is W_1 , with W_0 highlighted in grey. The corresponding induced paths are shown in gold. (Following Ref. [3]) | 46 |
| 4.14 | (a) Penrose tiling shown by gray lines. The modified trivalent penrose tiling (black lines) decorated to admit fully packed loops (purple) by construction. (b) Augmenting the alternating cycle surrounding each red cross leads to a cycle (red) which visits every graph vertex on any Penrose tile touching a cross. (This figure was made by my supervisor.) | 48 |
| 5.1 | (a) Honeycomb lattice. (b) Square-octagon lattice. | 50 |
| 5.2 | Modulus heights represented on a 3RPB random graph (generated by plantri). The purple edges form a random perfect dimer matching. (a) Shows the heights obtained using the rules in Sec. 3.4, while (b) demonstrates that taking the modulus-3 of these heights corresponds to a 3-coloring of the graph, with only three possible values $\{0,1,2\}$, each represented by a different color. Note that the exterior tile of the graph has a height of 2. | 52 |

| | | |
|-----|---|----|
| 5.3 | Schematic of the algorithm for constructing large 3RPB graphs. (a) A closed random walk loop is shown in red, with thick black dots representing its step points. The self-intersection points of the loop are marked by black and green circles. These points form the nodes of a 4-regular graph, with edges given by the loop segments connecting them. The green circled node has a self-loop; removing this node and connecting its adjacent nodes results in a new 4-regular graph, shown in black, with black circle nodes only. (b) A square is drawn around each node of the 4-regular graph, shown in blue, such that each corner of a square lies on the original walk. (c) The resulting 3-regular planar bipartite (3RPB) graph is obtained by removing all self-intersection points inside the squares and their associated edges. Connections between squares are simplified as straight lines by removing intermediate random walk step points. | 53 |
| 5.4 | (a) A height gem generated using our random walk algorithm in Fig. 5.3. (b) Trivalent Penrose: Thick black lines represent the modified Penrose tiling, where all vertices have degree 3, derived from the underlying standard Penrose tiling (thin gray lines). (c) A purely random 3RPB graph generated using the <code>plantri</code> package. Note that all these graphs are represented using the Tutte planar embedding. | 54 |
| 5.5 | Expectation value of the ratio of flippable plaquettes to total plaquettes, $\langle NFP \rangle$, as a function of temperature T for different lattices at $v = -1$. For this case, the system evolves from a disordered state with intermediate flippability at high temperature to a maximally flippable state as $T \rightarrow 0$. The dashed gray line indicates the analytical prediction for the maximum number of flippable plaquettes in the height gem, random, and Penrose lattices as $T \rightarrow 0$ at $v = -1$ | 56 |
| 5.6 | (a) Energy cumulant (EC) as a function of temperature at $v = -1$. (b) Specific heat per site (C_v/N) as a function of temperature at $v = -1$ | 57 |
| 5.7 | Monomer-monomer correlations, $M(r)$, as a function of monomer separation r at $v = -1$. (a) Exact calculations of $M(r)$ at $T = \infty$ using the FKT algorithm, with a logarithmic y -axis, showing a straight-line behavior indicative of exponential decay, $\log(y) = -mx$ where $m > 0$. (b, c, d) Rapid decay of $M(r)$ at finite temperatures (T) for the Penrose, height gem, and square-octagon lattices, respectively, indicating confinement at all $T > 0$ for each lattice type. The separation distance (r) on the x-axis is normalized according to the Tutte planar embedding, which maps the entire lattice within a disk of radius 0.5 centered at (0.5,0.5), ensuring consistent comparison across different graph structures. | 59 |
| 5.8 | Height variance (W^2) as a function of linear system size at finite temperatures with $v = -1$ for the (a) height gem, (b) Penrose, and (c) square-octagon cases. The x-axis is log-scaled and normalized according to the Tutte planar embedding. The results show that W^2 converges to a finite value in the thermodynamic limit, confirming a smooth confining phase. | 60 |
| 6.1 | (a) The dark and light green tiles, collectively called the Mystic (M_0), with the dark green tile referred to as the Upper Mystic (M_0^+). This shows the first inflation of a spectre S_0 and Mystic M_0 . The Upper Mystic can have only two types of perfect matchings, shown in red and blue. Choosing either red or blue matching enforces the same set of purple dimers, thus fixing the purple dimers. The internal edges of S_1 and M_1 tiles allow only two dimer matchings per Mystic. (b) The twice-inflated spectre region, S_2 . Mystic regions are highlighted in pink, with the special gold dimer (arrow) enforcing all green dimers. The two-fold dimer freedom is preserved on the Upper Mystic tiles and the outermost boundary, validating Eq. (6.3). Source:[11] | 65 |
| 6.2 | The free energy per dimer f_N for spectre patches containing N vertices, shown as purple points, corresponds to S_2 through S_6 spectre regions. These data points are calculated numerically using the FKT algorithm. The dashed line represents the analytical result from Eq. 6.5, valid as $N \rightarrow \infty$. The slow convergence is due to the fractal nature of the boundary. | 67 |
| 6.3 | Monomer Deconfinement on Spectre Tiling: If a dimer is removed, creating two monomers (represented by red circles), one monomer will always follow the path of green dimers along the boundary of the S_1/M_1 tiles and move toward the outermost boundary. The other monomer, however, will become trapped on an upper mystic dark green tile. These possible paths for the two red monomers are shown in yellow—one reaching the outer boundary and the other confined to an upper mystic tile. Similarly any pair of monomers can be separated to infinite distances at zero energy cost, as their movement does not require crossing any M_0^+ tiles during separation. (Following Ref. [11]) | 68 |

| | | |
|------|---|----|
| 6.4 | Possible Escape Routes for a Monomer: All possible paths (gold) by which a monomer (black), created on any upper mystic tile, can escape to the boundary of the S_1/M_1 tiles — and ultimately to the outermost boundary — by following the green dimer paths shown in Fig. 6.3 [11]. (This figure was made by my supervisor.) | 69 |
| 7.1 | A patch of AB* Tiling, which is an AB tiling with all 8-vertices and their connected edges removed. The ground states for the classical dimer model on AB*, with Hamiltonian given in Eq. 3.1 (or defined as aligning interactions in Sec.V of Ref. [10]), will never hold a dimer on gray edges. Consequently, for both the classical and quantum dimer models, the AB* tiling can be decomposed into independent loops and ladders, represented by thick purple and black edges, respectively. The purple edges form 8-star loops, which allow only two trivial dimer configurations. The only remaining portion of the tiling with nontrivial phases consists of ladders made up of yellow and blue tiles, outlined here with thick black edges. | 74 |
| 7.2 | These three components are the fundamental building blocks that make up ladders of any size in an AB tiling. The black circles show the nodes included in each of the components. For instance, the largest ladder in Fig. 7.1 can be expressed in terms of these components as $[..sAssBssAssBssAssBssAssBs..]$, with periodic boundary. | 75 |
| 7.3 | All possible basis states for dimer configurations on each of the constituent ladder components. Green and purple edges represent rung and leg dimers, respectively. The reason behind assigning specific edges on rhombus tiles as rungs or legs will be explained later in the text. The matrix representation of the basis states for component ‘s’ is given by $h_1 = (10)$ and $h_2 = (01)$. Similarly, each basis state in $\{r_i^A; i = 1, 2, 3, 4\}$ and $\{r_i^B; i = 1, 2, 3, 4, 5\}$ is represented as a column vector of size 4×1 and 5×1 , respectively, with the $(i - 1)$ -th entry equal to 1 and all other entries set to 0. | 76 |
| 7.4 | Illustration of the step-by-step process used to construct an exact Matrix Product State (MPS) representation of any quantum state, utilizing a sequence of singular value decompositions. | 77 |
| 7.5 | Schematic representation of L - tensors at the boundaries and within the bulk of MPS chains: the left diagram illustrates $L_{1,a_1}^{\sigma_1}$, the row vector at the left boundary; the right diagram shows $L_{a_{N-1},1}^{\sigma_N}$, the column vector at the right boundary. The central diagram depicts $L_{a_{l-1},a_l}^{\sigma_l}$ for the bulk. The σ_n are the physical legs, while the a_n are the auxiliary legs of the MPS tensors. | 77 |
| 7.6 | Fusion rules and dimer constraints: The <i>top row</i> illustrates how the left environment is constructed during a left-to-right MPS sweep. Incoming labels $\{0, 1\}$ at each site (black circles) encode the left subsystem state and fuse with the onsite basis—depending on whether the site is of type s , A , or B —to produce outgoing labels. These define the updated subsystem state and are passed recursively to the next site. A similar process builds the right environment. The <i>middle row</i> shows all valid fusion rules for any ladder, detailing allowed combinations of incoming labels and onsite basis states. The <i>bottom row</i> presents their shorthand notation. | 78 |
| 7.7 | The thick and thin black edges represent the assigned rungs and legs, respectively, in the three fundamental ladder components. The rationale behind this assignment is that edges forming a standard square-ladder equivalent are designated as rungs, while all remaining edges are classified as legs. | 79 |
| 7.8 | (a) Onsite matrix product operators W_n with physical indices σ, σ^* and horizontal matrix indices. (b) A chain of MPOs acting on an MPS, with contracted horizontal indices and vertical contractions applying the MPO to the MPS. | 80 |
| 7.9 | Tensor network representation of Eq. 7.12, which is then optimized to determine the ground state and its corresponding energy. The left side illustrates the expectation value $\langle \Psi H \Psi \rangle$, while the right side depicts the squared norm $\langle \Psi \Psi \rangle$. (Following Ref. [12]) | 82 |
| 7.10 | Tensor network representation corresponding to the equation $Hv - \lambda v = 0$, derived by differentiating the equation in Fig. (7.9) to optimize a single site $v = L_{a_{l-1},a_l}^{\sigma_l}$, which is circled in both the left and right networks. This formulation leads to a generalized eigenvalue problem, where the Hamiltonian H is represented by the tensor network enclosed in the blue box. The networks in the two orange boxes depict the left and right environments of the current site l , labeled as LE and RE, respectively. Meanwhile, the tensor network within the purple boxes contracts to yield the identity. | 82 |

| | | |
|------|--|----|
| 7.11 | Tensor network for calculating the onsite expectation value of any operator O^i . This is also used to calculate the onsite probability density for different dimer basis states. The tensors in green and blue represent the left and right environments, respectively, while the red diamond-shaped tensors are the Schmidt matrices on any bond. | 84 |
| 7.12 | Ground state phase diagram of the QDM on the AB ladder, as defined by Eq. 7.10, plotted as a function of the parameters $\{v_r, v_l\}$ with $v_{lr} = 0.01$. The diagram displays five distinct phases described in the text, along with their phase boundaries. The thick black line denotes the isotropic line where $v_r = v_l$. Solid (pink, green) lines indicate first-order phase transitions, while dashed (pink, green, blue) lines correspond to second-order transitions. The orange and black circles mark the values of v_r at which various observables are evaluated to study the nature of the transitions in Sec. 7.1.3. The purple circle highlights the region where all five phases converge, suggesting a transitional regime. | 84 |
| 7.13 | Probability density profiles for the five distinct phases on the AB ladder Λ_4 with 432 sites. A small segment of the ladder is shown for illustration, along with its site components $\{s, A, B\}$. Blue points indicate the probability of either a rung dimer on s sites or a leg-rung dimer basis on $\{A, B\}$ sites. Orange points represent the probability of no rung dimer on s sites or an all-leg dimer basis on $\{A, B\}$ sites. The legend in (a) applies to all subfigures. | 85 |
| 7.14 | Probability density profiles for the five distinct phases on the AB ladder Λ_5 with 1296 sites. A small segment of the ladder is shown for illustration, along with its site components $\{s, A, B\}$. Blue points indicate the probability of either a rung dimer on s sites or a leg-rung dimer basis on $\{A, B\}$ sites. Orange points represent the probability of no rung dimer on s sites or an all-leg dimer basis on $\{A, B\}$ sites. The legend in (a) applies to all subfigures. | 86 |
| 7.15 | Pictorial representation of phase I on the Λ_1 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. Sites ‘ s ’ always host a rung dimer (green). Sites ‘ A ’ are in an equal superposition of $\{ r_1^A\rangle, r_3^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$. Sites ‘ B ’ are in a superposition of $\{ r_2^B\rangle, r_3^B\rangle, r_4^B\rangle\}$ with coefficients depending on v_r : $(d, e, f) = (0, 1/\sqrt{2}, 1/\sqrt{2})$ for $v_r \ll 0$, $(d, e, f) = (1/\sqrt{3}, 1/\sqrt{3}, 1/\sqrt{3})$ for $v_r < 0$, $(d, e, f) = (1/\sqrt{2}, 1/2, 1/2)$ for $v_r > 0$, and $(d, e, f) = (1, 0, 0)$ for $v_r \gg 0$ | 87 |
| 7.16 | Pictorial representation of phase II on the Λ_1 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges indicate the absence of a dimer. Sites ‘ s ’ never host a dimer. In any AB ladder, the system exhibits one of two configurations: (i) ‘ A ’ and ‘ B ’ sites in the pure states $ r_2^A\rangle$ and $ r_1^B\rangle$, or (ii) ‘ A ’ and ‘ B ’ sites in the pure states $ r_4^A\rangle$ and $ r_5^B\rangle$ | 88 |
| 7.17 | Pictorial representation of phase III on the Λ_2 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. For ladder segments of the form $(sAsAsA\dots)$, the ‘ A ’ sites alternate between equal superpositions of $\{ r_3^A\rangle, r_4^A\rangle\}$ and $\{ r_1^A\rangle, r_2^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$. The ‘ s ’ sites alternate between hosting a rung dimer and being empty, depending on the neighboring ‘ A ’ state, in order to satisfy the hard-core dimer constraint. For ladder segments of the form $(AssB\dots)$, the ‘ ss ’ sites do not host rung dimers but only leg dimers. The ‘ A ’ sites are in an equal superposition of $\{ r_1^A\rangle, r_3^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$. The ‘ B ’ sites are in a superposition of $\{ r_2^B\rangle, r_3^B\rangle, r_4^B\rangle\}$ with coefficients $(d, e, f) = (1/\sqrt{2}, 1/2, 1/2)$ | 88 |
| 7.18 | Pictorial representation of phase IV on the Λ_2 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. The states on ‘ A ’ and ‘ B ’ sites are identical to those in phase II. In this phase, alternating ‘ ss ’ components on square plaquettes host two rung dimers, depending on their neighbouring configurations. | 89 |
| 7.19 | Pictorial representation of phase V on the Λ_2 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. The ‘ A ’ sites are in an equal superposition of $\{ r_1^A\rangle, r_3^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$, while the ‘ B ’ sites are in a superposition of $\{ r_2^B\rangle, r_3^B\rangle, r_4^B\rangle\}$ with coefficients $(d, e, f) = (1/\sqrt{2}, 1/2, 1/2)$. In this phase, ‘ ss ’ components on square plaquettes alternately host either two rung dimers or two leg dimers. | 89 |

| | | |
|------|--|-----|
| 7.20 | First-Order Phase Transitions: These transitions are characterized by the variation of the order parameter (OP) and the ground state energy per site (E_N/N) as functions of v_l , near the boundaries between different phases, for different fixed values of v_r . Subfigures (a,c) show the transition from Phase II to Phase III, while (b,d) correspond to the transition from Phase IV to Phase I. In all cases, the OP exhibits a discontinuous jump from 0 to 1, and the ground state energy per site (E_N/N) shows a distinct kink. This discontinuous behavior of both the order parameter and the energy confirms the first-order nature of these transitions. Vertical dashed lines indicate the location of the transitions. | 91 |
| 7.21 | Second-Order Phase Transitions: These transitions are characterized by the variation of the order parameter (OP) and the ground state energy per site (E_N/N) as a function of v_l , near the boundaries between different phases, for different fixed values of v_r . Subfigures (a,b) show the transition from Phase III to Phase V, (c,d) correspond to the transition from Phase II to Phase IV, and (e,f) illustrate the transition from Phase V to Phase I. In all cases, the OP increases continuously from 0 to 1, and the ground state energy per site (E_N/N) varies smoothly without any kink. The continuous behavior of both the order parameter and the energy confirms the second-order nature of these transitions. | 93 |
| 8.1 | An example of a fully packed loop configuration (shown by purple edges) on the U_1 AB tiling (shown by thin gray edges), resulting from multiple flip move sweeps. The corresponding discrete height field is also illustrated, with each plaquette labeled by its respective height, derived from the steps in Sec. 3.4. Loops with blue nodes have clockwise orientation, while those with green nodes have anti-clockwise orientation. | 96 |
| 8.2 | The Rokhsar-Kivelson (RK) flip, which shifts dimers (highlighted in purple) from one set of alternating edges to the complementary set within a plaquette. | 97 |
| 8.3 | Plot of $[-\log(\langle h(\mathbf{p}) ^2 \rangle)]$ vs $[\log(\mathbf{p})]$: (a) for the square lattice FPL model, where the slope of the linear fit (gray line) deviates from 2, confirming the smooth phase at $T = \infty$. (b) for the AB tiling FPL model, where the slope of the linear fit (gray line) is very close to 2, confirming rough phase or criticality at $T = \infty$ | 98 |
| 8.4 | Colormap of the net height variance (W^2) in real space (Eq. 3.7) for each plaquette, derived from an ensemble of loop configurations (at infinite T) for the U_2 AB tiling. Darker shades represent lower height variance, while lighter shades correspond to higher height variance. This spatial pattern indicates a rough phase, as the variance increases with system size when examining concentric regions from the center to the boundary. | 99 |
| 8.5 | Plot of W^2 versus $\ln(L)$, where L is the linear system size defined as the radius of the concentric circles in the U_2 region shown in (a). The linear fit demonstrates logarithmic divergence as predicted by Eq. 3.8, confirming the rough phase for the AB tiling at $T = \infty$ | 100 |
| 8.6 | Log-log plot of area density $P(s)$ of loops versus loop length (s) for a square lattice of size 64×64 . The linear fit (gray line) indicates a power-law decay with a slope of $\tau \sim 2.14$ | 101 |
| 8.7 | The average loop size $s(R)$ plotted against its radius R on a log-log scale shows power-law behavior with weak periodic corrections. When $s(R)R^{-d_f}$ is plotted against R/L , the data for different system sizes collapse onto a finite-size scaling form (Eq. 8.6), with a fractal dimension $d_f = 1.568$, revealing weak log-periodic modulations to the power-law behavior. | 103 |
| 8.8 | $P(s)$, the area density of loops plotted against their size s on a log-log scale, exhibits power-law behavior with periodic corrections. When $P(s)s^\tau$ is plotted against sL^{-D_f} , the data for different system sizes collapse onto a finite-size scaling form, as per Eq. 8.11, with $\tau = 2.282$ and $D_f = 1.56$, revealing log-periodic modulations to the power-law behavior. | 104 |
| 8.9 | Two-point correlation function $C(r)$ plotted on a log-log scale against the separation between two nodes r in real space for the U_1 and U_2 AB tiling regions. The linear fit gives the critical exponent x and reveals the power-law decay. | 105 |

Acknowledgements

First and foremost, I would like to express my deepest gratitude to my supervisor, Felix Flicker, for all the support you have provided over the past three years. Thank you for your wisdom, patience, and unwavering confidence in me throughout this PhD journey. I am especially thankful for the regular and insightful discussions on the many topics covered in this work, for your guidance in shaping the direction of the research, and for encouraging me to explore new areas, which ultimately contributed to the broad scope of this thesis. I am also grateful for the many opportunities you offered—from travelling to new places and attending conferences, to engaging in meaningful scientific discussions with researchers around various universities.

I would like to extend my sincere thanks to Sean Giblin for always being there whenever I needed help at the university, and for your consistent support and encouragement. In the second half of my PhD, I was fortunate to have many friendly conversations with Miguel Ángel and Lizzy Reith—about both research and life in general. Thank you for your willingness to help whenever needed.

I am grateful to Dr. Natalia Chepiga for her invaluable guidance on tensor networks and the DMRG algorithm—topics that were entirely new to me. Her technical expertise and help were essential for this project, and I truly appreciate her patience and availability. Without her, I would likely still be stuck in the early stages of the coding process.

While the acknowledgements page of a thesis usually focuses on academic contributions, this journey became much happier thanks to the support of amazing people beyond academia. To all the friends I made in Cardiff: having come such a long way from home, I gradually began to feel part of a new family—Prachi, Vikalp, Pranay, Vageesh, Sanyam, Ananya—thank you for being fellow travellers on this PhD journey. I am incredibly grateful for all the conversations we shared, whether about science or not, and for the memories of table tennis games, evening walks, dinner parties, board games, movie nights, and some fun trips.

A special thank you to my dear friend Sridevi—since IISER days, you’ve stuck with me through all the chaos. From late-night studies to daily PhD video calls, your support and humour made everything feel a bit more like home. Our yearly trips were the best bonus—and thank goodness for the same time zone!

To my mother and father: thank you for believing in me and allowing me to pursue what I like. Papa, I remember you telling me to “always think big.” I try my best to live by those words, even if my mind occasionally holds me back. Mummy, you taught me to stay humble and to believe in the power of hard work above all else. I hope I’ve done justice to that wisdom. To my brother, Bhaiyu (Anirudh): I love you, and I’m so thankful for our many conversations. They brought just the right dose of chaos and joy to keep me going. Finally, I want to sincerely thank the Department of Physics at Cardiff University and the administrative staff for their invaluable support throughout the course. I would also like to thank the EPSRC for funding my PhD and supporting this research journey.

Introduction

1.1 Preface

Condensed matter physics explores the diverse collective behaviors that emerge when large numbers of particles, such as electrons, atoms, or molecules, interact within materials. A central focus within this field is the study of strongly correlated systems, where interactions between particles lead to complex phenomena that cannot be easily explained by considering individual components alone. These models have long fascinated physicists and have proven to be very useful tools for tackling real-world problems. In addition to their mathematical appeal, constrained models serve as convenient toy models for understanding various physical phenomena, bridging the gap between different branches of physics and fostering interdisciplinary links.

Understanding these systems is crucial for the development of advanced materials and technologies. Among many strongly correlated systems, the classical and quantum dimer models have been extensively studied. The classical dimer model examines arrangements of dimers — pairs of adjacent atoms or molecules — on lattices, subject to the constraint that each lattice site is part of only one dimer. It has been instrumental in understanding phenomena such as adsorption on surfaces and the statistical mechanics of close-packed structures [13, 14]. It has significant applications across various scientific fields. In physics, it has been instrumental in solving the Ising model on planar lattices, a major result in statistical physics [15, 16]. In computational biology, the dimer model has been utilized to study protein folding mechanisms and interactions, contributing to a deeper understanding of biological processes [17, 18].

A natural progression from classical dimers is the study of their quantum counterparts, notably the Rokhsar-Kivelson (RK) quantum dimer model [19, 20]. This model incorporates quantum fluctuations, allowing dimers to resonate between different configurations. Introduced to describe high-temperature superconductors, the RK model has provided insights into resonating valence bond (RVB) states and quantum spin liquids [21]. Both classical and quantum dimer models impose a "hard-core" constraint, ensuring that dimers do not overlap.

The investigation of critical phenomena in statistical physics provides profound insights into the behavior of systems undergoing phase transitions [22]. To analyze low-energy excitations and critical properties, physicists often employ field-theoretical approaches, notably the renormalization group (RG) approach [23]. By mapping lattice models onto continuous field theories, they can utilize robust analytical tools to explore emergent phenomena. Notably, conformal field theory (CFT) has been instrumental in describing systems at criticality, characterized by scale invariance [24]. CFT offers a framework to classify universality classes of phase transitions

and has been effectively applied to various two-dimensional lattice models, yielding exact results for critical exponents and correlation functions. These universal behaviors are observed across diverse physical systems, where macroscopic properties near critical points exhibit common features despite microscopic differences.

Traditionally, these models have been analyzed on periodic lattices, where translational symmetry simplifies their theoretical treatment. Studies on square, kagome, and honeycomb lattices have revealed rich phase diagrams and critical behaviors, enhancing our understanding of classical and quantum phase transitions [25, 26, 20, 27, 28, 29]. For instance, the Heisenberg model on these lattices has been extensively studied to understand magnetic properties and phase transitions [30, 31]. Additionally, the Hofstadter butterfly provides insights into energy spectra of electrons on a two-dimensional square lattice under a perpendicular magnetic field, showcasing complex fractal structures and topological phase transitions [32].

While periodic lattices have traditionally underpinned condensed matter studies, aperiodic structures like quasicrystals present new frontiers. Not all real-world physical systems or phenomena can be accurately modeled using periodic lattices alone, as many lack translational symmetries and contain defects. Discovered in 1982 by Dan Shechtman, quasicrystals exhibit long-range order without periodicity, challenging traditional crystallographic notions [33]. Their unique structural properties, such as non-repeating patterns and unconventional symmetries, have profound implications for electronic, photonic, and mechanical behaviors [34, 35, 36, 37].

In this thesis, I aim to extend studies of critical phenomena, phase transitions, universality classes, and optimization problems to aperiodic systems such as quasicrystals, employing various numerical and analytical techniques. This extension will help determine whether the widely accepted results for periodic lattices hold true for aperiodic tilings. Additionally, I am investigating strongly correlated models on quasicrystalline lattices for several important reasons: 1). The lack of periodicity in quasicrystals can lead to novel phases of matter and electronic states that differ fundamentally from those in periodic systems, potentially hosting new quantum phases [38, 39, 40], 2). Quasicrystals exhibit extended symmetries, such as five-fold rotational symmetry, which are forbidden in periodic crystals. These unique symmetries can influence the nature of excitations and correlations in the system, and 3). Experimental realizations and advances in material synthesis have made it possible to create quasicrystalline structures in both hard and soft matter systems, providing platforms to test theoretical predictions [41, 42, 43, 44]. Other studies have revealed the intriguing properties of quasicrystalline systems, for instance, the exact dimer ground state of a quantum antiferromagnet on a two-dimensional quasicrystal has been identified, highlighting the potential for novel magnetic behaviors in these structures [45].

Interdisciplinary Connections and Computational Complexity: The study of constrained physical models on lattices intersects with various disciplines, notably computational complexity [46]. Certain problems in condensed matter physics, such as determining the ground state energy of specific Hamiltonians, are computationally challenging. For instance, the numerical sign problem, which arises in quantum Monte Carlo simulations of fermionic systems, has been shown to be NP-hard [47]. A further notable result is the proof in Ref. [48] that the general spectral gap problem is undecidable: there exists no algorithm that can decide whether a given translationally invariant quantum many-body Hamiltonian on a two-dimensional lattice is gapped or gapless. Their construction relies on aperiodic tilings combined with Hamiltonian complexity techniques, offering another striking example of the role of aperiodic structures in physics and linking naturally to the complexity-theory. This interplay between physics and computational complexity not only deepens our understanding of material behaviors but also informs algorithm development in computer science, leading to more efficient methods for tackling complex optimization problems.

Optimization problems on graphs, encountered in mathematics and computer science, exemplify significant

interdisciplinary connections [49]. These problems, such as the Hamiltonian cycle problem, traveling salesperson problem, 3-coloring problem, and minimum dominating set problem, involve finding optimal solutions from a vast set of possibilities while adhering to stringent constraints. They find applications in logistics, machine learning, telecommunications networks, and other fields. The study of hard optimization problems intersects with physics through analogies with physical systems such as polymers and loop/spin models.

Exploring strongly correlated models and hard optimization problems on aperiodic lattices, such as quasicrystals, offers a rich avenue for discovering new physical phenomena, mathematical results, and deepening our understanding of complex systems. By integrating concepts from condensed matter physics, field theory, and computational complexity, this research can lead to advances in both fundamental science and practical applications. Quasicrystals, characterized by their non-repeating patterns and unique symmetries, challenge traditional notions of periodicity and order. Investigating these structures can lead to advances materials science, particularly in understanding how aperiodicity influences electronic, magnetic, and mechanical properties. Moreover, the mathematical complexity inherent in aperiodic tilings intersects with computational challenges, providing insights into optimization problems and computational complexity theory. This interdisciplinary approach not only enhances our understanding of quasicrystalline materials but also contributes to the development of novel algorithms and mathematical frameworks applicable across various scientific domains. I will now delve into a summary of the main contributions of this thesis.

1.2 Thesis Overview

This thesis comprises nine chapters, each contributing to the exploration of physical and mathematical phenomena in aperiodic systems. In Chapter 2, I provide a comprehensive background on aperiodic systems, focusing on quasicrystalline tilings and their properties. This foundation supports the subsequent discussions throughout the thesis. Additionally, I introduce general graph terminologies that are utilized consistently in later sections.

In Chapter 3, I present the theoretical framework and summarize already known key results of the strongly correlated models examined in this thesis. This includes classical and quantum dimer models, as well as the $O(n)$ loop model. I also discuss the construction of field theory frameworks in terms of height models relevant to these dimer models.

Chapters 4,5,6,7,8 will be based on my original published and upcoming work. In Chapter 4, I analytically study the NP-hard Hamiltonian cycle problem on the Ammann-Beenker (AB) tilings. Leveraging properties of AB tilings, I derive exact solutions to various non-trivial optimization problems. I also explore how mathematical solutions to these complex problems can apply to physical applications.

In Chapter 5, I focus on the classical dimer model on random graphs and a modified version of the Penrose tiling, analyzed using a numerical energy-based worm algorithm. I demonstrate the absence of a phase transition between high-temperature disordered and low-temperature crystalline phases, contrasting with behaviors observed on periodic lattices. Further, I explore height model results and monomer-monomer correlations to support these findings.

In Chapter 6, I determine the exact solution of the classical dimer model on the recently discovered aperiodic monotiling known as the Spectre tiling. This work extends the understanding of dimer models on complex aperiodic structures. Numerical results are provided to validate the analytical findings. Shifting focus to quantum systems, I also present an analytical solution for the quantum dimer model (QDM) on the Spectre

tiling.

In Chapter 7, I numerically investigate the ground state phase diagram of the QDM model on an AB* tiling (AB tiling with all 8_0 vertices removed) by adapting the Density Matrix Renormalization Group (DMRG) algorithm to incorporate hard-core dimer constraints, enhancing its efficiency. The results include ground-state phase diagrams and insights into various quantum phase transitions.

In Chapter 8, I explore the $O(n)$ loop model on the AB tiling, detailing the construction of fully packed loops and their sampling. I analyze the critical behavior of this model, determine the geometric critical exponents of loops at infinite temperature, and compare them with known results on square lattices.

Chapter 9 summarizes the findings presented throughout the thesis, drawing conclusions from the diverse results obtained. I also propose potential directions for future research in this field.

Portions of this thesis were proof read by online tools such as Grammarly and ChatGPT. I only took assistance for proofreading, identifying grammatical errors, correcting sentence structure, and suggesting appropriate synonyms.

1.3 Contribution During the PhD

The author of this thesis has contributed to the following publications and presentations for the completion of this PhD.

1.3.1 Publications on Work Related to the Thesis

- Michael Saccone, Arjen Van den Berg, Edward Harding, Shobhna Singh, Sean R. Giblin, Felix Flicker, and Sam Ladak, “Exploring the Phase Diagram of 3D Artificial Spin-Ice,” *Commun. Phys.* 6, 217 (2023).

- Shobhna Singh and Felix Flicker, “Exact Solution to the Quantum and Classical Dimer Models on the Spectre Aperiodic Monotiling,” *Phys. Rev. B* 109, L220303 (2024) [11].

The author of this thesis produced all figures, contributed to some parts of the main text, and generated all numerical data underlying the findings. This work is detailed in Chapter 6 and the first part of Chapter 7.

- Shobhna Singh, Jerome Lloyd, and Felix Flicker, “Hamiltonian Cycles on Ammann-Beenker Tilings,” *Phys. Rev. X* 14, 031005 (2024) [3].

The author of this thesis produced all figures except Figures 4, 7, 10, and 15, contributed to parts of the introduction and Section III, and wrote Sections IV, Appendix A, B.1, and B.3. This work is presented in Chapter 4. This research is discussed in Chapter 5. This study is elaborated upon in Chapter 7.

Accompanying Codes: All the chapters required significant computational work. The necessary code developed and used for this research is available upon request.

1.3.2 Presentations and Posters on Work Related to the Thesis

- **Invited Talk:** “Hamiltonian Cycles on AB tilings”, at 11th International Conference on Aperiodic Crystals, Caen, France. (June 2024)

- **Invited Talk:** “Exact Solution to the Quantum and Classical Dimer Models on the Spectre Aperiodic Monotiling”, at British Applied Mathematics Colloquium (BAMC), Newcastle, UK. (April 2024)

- **Talk:** “2D MPS-DMRG: Quantum dimer and loop models on a two-dimensional quasicrystal”, at APS March Meeting, Minneapolis, USA. (March 2024)
- **Invited Talk:** “The $O(n)$ loop model on Ammann-Beenker Tilings”, at The Grimm Network Meeting, Open University, Milton Keynes, UK. (July 2023)
- **Poster Presentation:** Condensed Matter and Quantum Materials Conference (CMQM), University of Birmingham. (June 2023)
- **Conferences / Schools attended:**
 - Delft Many-Body Workshop Series, Tensor networks for constrained systems, TU Delft, Netherlands. (October 2023)
 - European Tensor Network 2023 School - Tensor Network based approaches of Quantum Many-Body Systems, Abingdon, UK. Sept 2023
 - 3rd Plenary meeting of International Quantum Tensor Network, TU Munich, Germany. July 2023
 - The HatFest : Celebrating the discovery of an Aperiodic Monotile, The Mathematical Institute, Oxford University, UK. July 2023

Aperiodic Systems: Quasicrystals

Quasicrystals are structures characterized by an aperiodic arrangement of atoms that still exhibit order [50, 51]. In contrast, conventional crystals have a single unit cell that repeats periodically to fill space without gaps. For quasicrystals, the space is filled using two or more distinct unit cells arranged aperiodically, but without gaps [52]. Despite the lack of periodicity, quasicrystals maintain long-range order, meaning the positions of atoms remain predictable, even though they lack translational symmetry. This long-range order gives rise to symmetric Bragg diffraction patterns [53, 51]. The aperiodicity of quasicrystals is reflected in their diffraction patterns, which exhibit unusual rotational symmetries. Unlike the two-, three-, four-, or six-fold rotational symmetries allowed by the crystallographic restriction theorem, quasicrystals can display five-, seven-, eight-, ten-, or twelve-fold symmetries. Importantly, quasicrystals can possess only local centers of 5-fold (or similar) symmetry, and at most one global center of rotational symmetry.

The discovery of these unique diffraction patterns, now recognized as hallmarks of quasicrystals, was first made in 1982 by the materials scientist Dan Shechtman [2]. While studying an aluminum-manganese alloy, he observed a ten-fold electron diffraction pattern. This groundbreaking finding earned him the Nobel Prize in Chemistry in 2011. Subsequently, in 2018, Luca Bindi and Paul Steinhardt discovered the first naturally occurring quasicrystal, a 10-fold symmetric mineral named icosahedrite, with the composition $\text{Al}_{63}\text{Cu}_{24}\text{Fe}_{13}$ [54]. Since then, hundreds of quasicrystals have been reported, commonly in aluminum alloys, as well as other compositions such as Cd–Yb, Ti–Zr–Ni, Zn–Mg–Ho, Zn–Mg–Sc, and In–Ag–Yb [55].

Quasicrystals possess a range of potential applications. For example, their hard and brittle nature makes them ideal for low-friction, high-strength surface coatings [56]. They are also excellent candidates for hydrogen storage materials [35, 57]. Shechtman proposed innovative applications that leverage their low coefficient of friction and hardness, such as embedding particles in plastic to produce strong, durable, low-friction gears. Furthermore, the low thermal conductivity of some quasicrystals makes them suitable for heat-insulating coatings [58]. Remarkably, quasicrystals can exhibit smooth and flat surfaces despite their irregular atomic structure, as two incommensurate surfaces tend to slide with minimal friction, further enhancing their utility [59].

Recent experiments have revealed tunable quasicrystal geometries in systems like twisted trilayer graphene [43] and optical lattices [44]. These discoveries have expanded the scope of quasicrystal research, showcasing their ability to host fascinating physical phenomena, including exotic criticality [60], charge ordering, and topological effects [61]. The lack of periodicity often leads to novel and unexpected behaviors.

From a mathematical perspective, crystals are modeled using periodic lattices, such as square or honeycomb




Figure 2.1. The Ho-Mn-Zn dodecahedral quasicrystal (left) alongside the 10-fold symmetric diffraction pattern of the first naturally occurring quasicrystal, Icosahedrite (right). (Following Refs. [1, 2])

lattices, which are constructed by repeating a single unit cell. However, quasicrystals are modeled using aperiodic tilings or lattices [10, 53, 62]. These tilings mostly involve two or more distinct tiles that cover an infinite space aperiodically using a set of substitution rules. They can also be derived by projecting higher-dimensional periodic lattices onto lower dimensions [52]. For example, 1D quasiperiodic Fibonacci tiling can be constructed from a 2D square lattice using a cut-and-project method, as illustrated in Fig. 2.2.

Some scholars have suggested that the earliest representations of quasicrystalline patterns can be traced back to early Islamic art and architecture, such as the Gunbad-i-Kabud tomb tower and the Darb-e Imam shrine [63]. Interestingly, the oldest known anthropogenic quasicrystals were created during the Trinity nuclear bomb test, which produced icosahedral quasicrystals [64].

2.1 Ammann-Beenker (AB) Tilings

One of the most well-known 2D aperiodic tilings is the Ammann-Beenker (AB) tiling, which consists of two basic prototiles: a square and a rhombus with an acute angle of $\pi/4$, both having unit edge lengths. These tilings exhibit an eight-fold symmetric diffraction pattern and lack the translational symmetry characteristic of periodic lattices, making it challenging to derive mathematical results in the thermodynamic limit. Instead of translational symmetry, AB tilings display discrete scale symmetry, which is a fundamental aspect of their unique properties [10, 53, 62].

The discrete scale symmetry of AB tilings allows them to be constructed iteratively through inflation rules shown in Fig. 2.4. Starting from any ‘legal’ patch of tiles, these rules enable the creation of larger patches by subdividing each tile into smaller square and rhombic tiles [10]. The edges of these tiles are then rescaled (or ‘inflated’) by a factor equal to the silver ratio, $\delta_S = 1 + \sqrt{2}$, restoring the tiles to unit edge lengths. This hierarchical construction is the basis for the discrete scale invariance observed in AB tilings. This is also demonstrated in Fig. 2.5. Starting with an AB tiling formed by thick black edges and red/blue vertices, applying the inflation rules twice results in an underlying dense (with a higher number of vertices) AB tiling composed of gray and white tiles. This property ensures that, despite the lack of periodicity, the positions of vertices at arbitrarily large distances remain predictable, thereby maintaining long-range order.

The edges and vertices of the AB tiling form a bipartite graph, where the vertices can be divided into two subsets such that edges only connect vertices from different subsets. With each application of the inflation rules,




Figure 2.2. Cut-and-Project Method: Generating a 1D Fibonacci quasicrystal by projecting a 2D square lattice \mathbb{Z}^2 (gray circles). A stripe with irrational slope $1/\tau$ (where $\tau = (1 + \sqrt{5})/2$) is chosen, and lattice nodes within the stripe are projected onto a line parallel to it. The resulting points form intervals of two lengths, S (short, orange) and L (long, purple), creating a Fibonacci tiling. This tiling can also be generated by substitution rules: $L \rightarrow SL$, $S \rightarrow L$, producing sequences of lengths $1, 1, 2, 3, 5, \dots$, corresponding to Fibonacci numbers. Tiles S and L have lengths 1 and τ , respectively.




Figure 2.3. A patch of Ammann-Beenker tiling, created using two types of unit cells: squares and rhombuses. It exhibits quasi-periodicity and local 8-fold rotational symmetry.




Figure 2.4. The inflation rule σ : Tiles are subdivided into smaller tiles, followed by rescaling (inflation) of all lengths by the silver ratio $\delta_S = 1 + \sqrt{2}$. Red and blue points indicate vertices of single- and double-inflated tiles, respectively. Note that the triangles in the first and second inflations are actually half-square tiles. Also, inflating the square tile breaks its rotational symmetry, so the orientation must be tracked, as indicated by the triangular motif. For clarity, these markings will generally be omitted in subsequent figures. (Following Ref. [3])




Figure 2.5. Illustration of discrete scale invariance in AB tilings. The thick black lines with red and blue vertices depict an AB tiling. After applying the inflation rules from Fig. 2.4 twice, a dense underlying tiling emerges, composed of gray and white tiles, which itself forms an AB tiling.




Figure 2.6. Inflation rules for Penrose tiles, where each tile is subdivided into smaller tiles and rescaled by the silver ratio $\delta_S = 1 + \sqrt{2}$. Red nodes mark the vertices of single-inflated tiles.

the number of tiles increases exponentially, and the infinite tiling is obtained in the limit:

$$T = \lim_{n \rightarrow \infty} \sigma^n(T_0), \quad (2.1)$$

where T represents the infinite tiling, σ denotes the inflation rule, and T_0 is the initial patch. Here, ‘tiling’ refers to the infinite structure, while ‘patch’ denotes any finite, connected set of tiles. For simplicity, the edges of the graph are considered to have unit length.

2.2 Penrose Tiling

Another notable aperiodic tiling is the Penrose tiling, named after mathematician and physicist Roger Penrose, who discovered it in the 1970s [4, 65]. This two-dimensional aperiodic tiling is composed of two non-overlapping prototiles: a thick rhombus with angles of 72° , 108° , 72° , and 108° , and a thin rhombus with angles of 36° , 144° , 36° , and 144° (see Fig. 2.7). Penrose tilings exhibit a fivefold symmetric diffraction pattern, a hallmark of their aperiodic nature.

Like the Ammann-Beenker tiling, the Penrose tiling possesses a self-similar structure, allowing it to be generated through a set of inflation rules. These rules, illustrated in Fig. 2.6, enable the iterative construction of larger patches, ultimately converging to the thermodynamic limit. This hierarchical construction provides the Penrose tiling with discrete scale invariance, which underpins its long-range order [4].

2.2.1 Trivalent-Penrose Tiling

We observe that while the Penrose and AB tilings share similar properties, such as their hierarchical structures, a surprising distinction exists: the AB tiling admits perfect matchings [10], whereas the Penrose tiling does not [4]. A perfect matching of a graph is a set of independent edges such that no two edges share a common vertex, and together they cover every vertex of the graph. As will be demonstrated in later chapters, the existence of perfect matchings is crucial for studying dimer models on these tilings. To address this limitation, I introduce a modified version of the standard Penrose tiling that admits perfect matchings. Furthermore, this modification ensures that all vertices have the same degree of 3, a property that proves advantageous when defining a height field theory on such aperiodic tilings. I refer to this modified tiling as the ‘trivalent Penrose tiling.’

To construct this modified tiling, begin by drawing rectangles inside each tile of the Penrose tiling, as shown by the thick blue lines in Fig. 2.8, with the original Penrose tiling shown in thin gray lines. Next, connect




Figure 2.7. A patch of Penrose tiling, created using two types of unit cells: thick and thin rhombuses. It exhibits quasi-periodicity and local 5-fold rotational symmetry.

the nearest vertices of each pair of adjacent rectangles. The resulting graph, with thick black and blue edges and vertices circled in gray, represents our desired tiling, where all vertices have degree 3. It is evident that this modified graph also admits perfect matching, achieved by placing dimers on the opposite edges of each blue rectangle.

2.3 Einstein Problem: New aperiodic Monotilings

As discussed in previous sections, it was long believed that at least two different tiles were required to create an aperiodic tiling. A long-standing open problem in mathematics, known as the Einstein problem, asked whether there exists a single tile capable of covering an infinite plane aperiodically, without overlaps or gaps. In 2023, a group of researchers solved this problem by introducing a family of tiles called the ‘Hat’ [5] and the ‘Spectres’ [6]. These tiles, along with their rotations and reflections, can tile the plane aperiodically.

2.3.1 Hat Tiling

A single ‘Hat’ tile is an asymmetric structure, depicted in shaded gray in Fig. 2.9. Surprisingly, this tile is embedded within the periodic hexagonal lattice, outlined by light gray edges. When combined with its mirror image, the Hat tile forms an aperiodic tiling, represented by thick black edges in the same figure [5].

A larger finite patch of Hat tiling is displayed in Fig. 2.10, where the reflected Hat tiles (mirror images) are shown in dark blue, sparsely and evenly distributed among the unreflected tiles. Within this tiling, non-overlapping groups of tiles, called clusters, can be identified. These clusters, shown on the right in Fig. 2.10, are used to construct the tiling. By simplifying the boundaries of these clusters, larger composite tiles, or ‘supertiles’, can be formed. The supertiles — denoted as H, T, P, F — are outlined by thick black boundaries in the figure. The left panel of Fig. 2.10 illustrates how these supertiles combine to form a large patch of Hat tiling.

The Hat tiling exhibits a hierarchical structure, similar to other aperiodic tilings. This hierarchical property arises from the fact that supertiles can be decomposed using a defined set of inflation rules, as illustrated in Fig. 2.11. By rescaling the edges of these supertiles, the tiling can be expanded indefinitely to cover any desired plane.




Figure 2.8. Trivalent Penrose tiling, depicted with thick black and blue edges, and circled gray nodes. It is derived from the underlying Penrose tiling (light gray lines) using the method described above. Following Ref. [4], I thank Dr. Nick Jones for this idea.




Figure 2.9. A patch of 'Hat' aperiodic monotiling embedded within the periodic hexagonal lattice (solid grey lines). A single hat tile is highlighted in grey, composed of eight smaller kite shapes (dashed grey lines).




Figure 2.10. A patch of Hat tiling, with reflected (mirror image) Hat tiles shown in dark blue, sparsely distributed throughout the pattern. Non-overlapping clusters of tiles are grouped to form H, T, P, F supertiles by simplifying cluster boundaries (outlined in black). The left panel demonstrates how these supertiles combine to create larger patches of the Hat tiling. (The figure was created using an [interactive application](#), following Ref. [5].)




Figure 2.11. Inflation rules for the Hat supertiles (H, T, P , and F), illustrating the hierarchical expansion of the Hat tiling. (The figure was created using an [interactive application](#), following Ref. [5].)




Figure 2.12. The first three iterations of the P supertile using inflation rules. Starting with the pink P supertile, one inflation produces the underlying yellow tiling. A second inflation of the yellow tiling results in the underlying black tiling, and the final iteration yields a dense Hat tiling. This hierarchical growth demonstrates the discrete scale invariance of the Hat tiling. (The figure was created using an [interactive application](#), following Ref. [5].)

To illustrate the hierarchical nature of Hat tiling further, Fig. 2.12 demonstrates how a large patch is constructed by applying the inflation rules iteratively to an H supertile. The first three iterations are shown in black, pink, and yellow, respectively, emphasizing the discrete scale invariance inherent in Hat tiling.

2.3.2 Spectre Tiling

As mentioned above, the hat tiling requires the use of its mirror image to form an aperiodic tiling. This reliance on reflected and unreflected tiles effectively introduces two distinct types of tiles. To address this limitation, the same team of researchers developed a modified tile known as the ‘Spectre’ tile (shown in Fig. 2.13(a)) with curved edges. Since we are only concerned with graph connectivity, we can straighten the curved edges of the Spectre tiles (also shown in Fig. 2.13(a)), resulting in what is termed ‘Tile (1,1)’ in Ref. [6]. This tile admits only aperiodic tilings even when reflections are not permitted.

The inflation rules for constructing the Spectre tiling are shown in Fig. 2.14. These rules involve two basic building blocks: a single Spectre tile and a structure of two Spectre tiles (depicted in two shades of green), collectively referred to as the Mystic. We refer to the dark green tile as the Upper Mystic, which has the unique property of being the only tile with two vertices of degree 4. This fact will be very useful in later chapters. A single Spectre tile is replaced by a cluster consisting of one Mystic and seven Spectres (all reflected). Similarly, a Mystic is replaced by a cluster containing another Mystic and six Spectres (all reflected). Notably, these reflections preserve the chirality of the tiling at each level of inflation. Consequently, any Spectre tiling can be decomposed into non-overlapping congruent clusters derived from these inflation rules.

While the inflation rules describe how individual Spectre or Mystic tiles are replaced by clusters, they do not specify the spatial arrangement of these clusters. A second set of rules, shown in Fig. 2.15, governs the placement and alignment of these clusters. Together with the inflation rules (Fig. 2.14), these assembly rules provide a complete framework for constructing Spectre tilings of arbitrary size. Below is a concise summary of




Figure 2.13. (a) A single Spectre tile (above) and a Tile (1,1) with 14 vertices (below). The (black) vertex at the midpoint of a straight edge is implied by edge connections. (b) A patch of Spectre aperiodic monotiling using just rotations and translations of the spectre tile. The tiles shaded in dark green are called the ‘Upper Mystics’. (The figure was created using an [interactive application](#), following Ref. [6].)




Figure 2.14. Two inflation rules for constructing the Spectre tiling. The rules are based on a single Spectre tile (white) and a two-Spectre combination called a Mystic (dark and light green). The first rule (above) replaces the Spectre with a cluster containing a Mystic and seven reflected Spectres. The second rule (below) replaces a Mystic with a cluster containing a Mystic and six Spectres. (The figure was created using an [interactive application](#), following Ref. [6].)




Figure 2.15. An iterative process for constructing Spectre and Mystic clusters. In Step 1, a single Spectre tile is defined by four key points (red nodes). These key points guide the placement of additional Spectre tiles and Mystic tiles in Step 2a, which creates a cluster of Spectre and Mystic tiles. In Step 2b, the key points (red nodes) for higher-level clusters are determined, allowing for the continued growth of the tiling. Steps 3a and 3b further expand the structure into superclusters, with Spectre clusters highlighted in red and Mystic clusters in blue. This iterative process demonstrates the hierarchical nature of the Spectre tiling system. (The figure was created using an [interactive application](#), following Ref. [6].)

the construction process, with full details available in Ref. [6]:

- Step 1: Identify four key vertices of the Spectre tile and connect them to form a quadrilateral (for visualization purposes; not required for construction).
- Step 2a: Build a Spectre cluster by:
 - Placing a Spectre at the desired location and orientation.
 - Attaching neighboring tiles by aligning shared key points (each tile shares two key points with its neighbors).
 - Rotating new tiles to maintain the correct relative orientation.
 - Translating tiles to align their key points accurately.
- Step 2b: Identify the key points of the completed Spectre cluster and define the Mystic cluster by removing one even-numbered Spectre.
- Step 3a: Assemble Spectre and Mystic clusters into a Spectre supercluster.
- Step 3b: Extract the key points of the supercluster from its constituent clusters and define the Mystic supercluster by removing one cluster.

- Iteration: Repeat the process iteratively to construct larger patches and scale up to the desired size.
- Observation: The handedness of the tiles alternates with each successive level of construction.

2.4 Graph Theory Terminology

In subsequent chapters, I will frequently use several graph theory terminologies [66]. To provide clarity and a point of reference, I will define these terms here.

A *graph* G consists of a set of vertices V , connected by a set of edges E . The *cardinality* of a graph G refers to the total number of vertices in V and is denoted by $|G|$. The *degree* $d(n)$ of a node $n \in V$ is the total number of edges connected to that node.

In this thesis, we will focus on graphs derived from various aperiodic tilings discussed earlier, using their respective vertex (V) and edge (E) sets. For example, if we refer to $G = AB$, it indicates a graph formed from the infinite AB tiling. When a graph is constructed from finite patches of a tiling, this will be made clear from the context. All graphs considered in this work will be undirected, meaning that edges have no inherent direction, and both directions of any edge are treated equivalently.

Bipartite Graphs : A graph is called *bipartite* if its vertex set V can be divided into two non-overlapping subsets (V_A, V_B) such that all edges in E connect a vertex from V_A to one in V_B . We will refer to the vertices in these subsets as having two distinct bipartite ‘charges’. All the aperiodic tilings discussed earlier — AB, Penrose, Hat, and Spectre — are bipartite.

Paths, Loops, and Matchings : A *path* is a sequence of edges connecting a series of distinct vertices. When the start and end vertices of a path coincide, it forms a *loop*. A *matching* (M) in a graph G is a subset of edges such that no two edges share a common vertex. A *perfect matching* is a matching in which every vertex in V is covered by an edge in M .

For a graph G and a given matching M : An *alternating path* is a path in which every other edge belongs to M . An *augmenting path* is an alternating path where the first and last edges are not part of M . *Augmenting* an alternating path involves switching the inclusion of its edges in M , ensuring a new matching is formed.

Dimer and Loop Models

Dimer models form a cornerstone in the study of statistical mechanics, graph theory, and condensed matter physics. Emerging in the early 20th century, the concept of dimers — a pair of particles or entities connected by a bond — was first systematically explored in mathematical contexts, such as the enumeration of perfect dimer matchings on lattices [67, 27]. In physics, dimer models gained prominence through the pioneering work of Fowler and Rushbrooke (1937), who investigated adsorption phenomena on surfaces [14], and Kasteleyn (1961), who introduced elegant combinatorial techniques to compute the exact partition function of dimers on planar graphs [67, 68]. Dimer models have since been applied to a wide range of physical phenomena, including zero modes in electronic tight-binding models [69, 70] and magnetism, where dimers are used to construct analytical solutions to problems like the Ising model [71].

The mathematical formulation of dimer models, particularly Kasteleyn’s method and Fisher’s mapping to spin systems [67, 27], firmly established their relevance to statistical mechanics. These models exhibit fascinating properties, such as critical behavior and long-range correlations, which are crucial to understanding phase transitions and disordered systems [25, 26]. The classical dimer model, in which dimers occupy edges of a lattice under the hard constraint that each vertex belongs to exactly one dimer, serves as an elegant and minimal framework for probing these phenomena.

Quantum dimer models, introduced by Rokhsar and Kivelson (1988), extend this classical framework into the quantum domain [19, 20, 72, 73]. These models explore quantum superpositions of dimer coverings, incorporating quantum dynamics via off-diagonal terms in the Hamiltonian. It also enforces a hard constraint: no two quantum dimers can share a vertex. Originally introduced to capture the physics of resonating valence bond (RVB) states in lattice spin systems, a concept first proposed by American physicist P.W. Anderson and Indian theoretical physicist Ganapathy Baskaran in 1987 to explain high-temperature superconductivity [21]. Quantum dimer models have since become a fundamental tool for investigating the low-energy physics of frustrated magnetic systems. They are instrumental in studying exotic phases of matter, such as quantum spin liquids [74], and provide profound insights into quantum criticality, nature of quantum phase transitions and the interplay between geometry, frustration, and entanglement.

The intersection of dimer models with quantum many-body physics has led to significant advances in understanding topological phases, fractionalization, and quantum entanglement entropy [20]. These models have also found experimental realization in systems such as cold atoms and Rydberg arrays [75]. Additionally, their connection to gauge theories underscores the universality of dimer models, bridging condensed matter physics




Figure 3.1. The general ground state phase diagram of the classical dimer model on regular periodic lattices with aligning interactions features ordered columnar and staggered ground state phases, corresponding to the maximum and minimum number of flippable plaquettes, respectively. These ground states exhibit monomer confinement. On the square lattice, the ground states are fourfold degenerate. At higher temperatures, a critical phase emerges. At infinite temperature ($T = \infty$), the partition function reduces to the unweighted sum over all possible dimer coverings of the square lattice, where the model can be solved exactly.

and high-energy field theories. Recently, they have also been realized experimentally in programmable quantum simulators, opening up new opportunities for exploring the model in controlled environments [76, 77].

In this chapter, I describe the theoretical foundations of dimer models, spanning from their classical origins to their quantum generalizations.

3.1 Classical dimer model

The interacting perfectly packed classical dimer model is a fundamental model in condensed matter physics. Originally developed to study the efficient packing of short linear molecules on crystal surfaces, this model has since evolved into a powerful framework for investigating unusual collective behaviors and phase transitions. The partition function for the model with aligning interactions is defined as [26]:

$$Z = \sum_c \exp\left(-\frac{E_c}{T}\right),$$

where

$$E_c = v \cdot N_{fp} \tag{3.1}$$

Here, N_{fp} denotes the total number of flippable plaquettes in a given dimer configuration c . A plaquette in the context of lattice models refers to the smallest closed loop or minimal face bounded by edges on a lattice. A plaquette is deemed flippable if it contains dimers on every alternate edge. The parameter v determines the nature of interactions between nearest-neighbor dimers, influencing the system's phase behavior.

The classical dimer model has been extensively studied on periodic lattices [26, 25, 28, 78, 79], leading to a well-established ground state phase diagram, as shown in Figure 3.1. At $T = 0$, where thermal fluctuations are absent, two ordered crystalline phases emerge: the columnar phase and the staggered phase.

- **Columnar Phase ($v < 0$):** The system minimizes energy by maximizing the number of flippable plaquettes. On a square lattice, this configuration is achieved by arranging dimers in alternating columns.

- Staggered Phase ($v > 0$): The system minimizes energy by reducing the number of flippable plaquettes to the minimum. On a square lattice, this corresponds to a staggered arrangement of dimers.

These phases are confined, meaning that breaking a dimer into two monomers and moving them apart demands an increasingly large amount of energy as the distance between the monomers (r) grows. In a confined state, the monomer-monomer correlation function $M(r)$ decays exponentially to zero with r . At $T = 0$, deconfinement occurs only at a critical interaction value v_c , where monomers can separate infinitely with a finite energy cost. In a deconfined state, $M(r)$ approaches a nonzero constant as $r \rightarrow \infty$.

At finite temperature, however, thermal fluctuations can drive the system into a critical (algebraically confined) phase, where $M(r)$ still vanishes as $r \rightarrow \infty$, but only as a power law rather than exponentially. On the square lattice, for $v = -1$, the system undergoes a Kosterlitz-Thouless (KT) transition from a high-temperature critical phase to a low-temperature crystalline phase [26].

At infinite temperature ($T = \infty$), the partition function reduces to an unweighted sum over all possible dimer coverings of the lattice, a scenario that has been exactly solved for periodic lattices [67, 27]. At this point, the system exhibits critical behavior, with correlation functions decaying algebraically with distance [80]. For example: Dimer-dimer correlation functions decay as $1/r^2$. Monomer-monomer correlation functions decay as $1/\sqrt{r}$ at large distances.

3.2 Quantum dimer model

The Hamiltonian for the Rokhsar-Kivelson quantum dimer model (RK-QDM) on a square lattice is:

$$H_{\text{QDM}} = \sum_{\text{plaquettes}} -J \left(\hat{P}_f + \hat{P}_f^\dagger \right) + v_r \left(\hat{P}_v \right) + v_l \left(\hat{P}_h \right), \quad (3.2)$$

where

$$\hat{P}_f = |\begin{array}{|c|c|} \hline \blacksquare & \blacksquare \\ \hline \end{array}\rangle \langle \begin{array}{|c|c|} \hline \blacksquare & \blacksquare \\ \hline \end{array}|, \quad \hat{P}_v = |\begin{array}{|c|c|} \hline \blacksquare & \blacksquare \\ \hline \end{array}\rangle \langle \begin{array}{|c|c|} \hline \blacksquare & \blacksquare \\ \hline \end{array}|, \quad \hat{P}_h = |\begin{array}{|c|c|} \hline \blacksquare & \blacksquare \\ \hline \end{array}\rangle \langle \begin{array}{|c|c|} \hline \blacksquare & \blacksquare \\ \hline \end{array}|.$$

Here $J > 0$ is the coupling constant for the kinetic term, which describes the flipping of dimers between alternating configurations on a plaquette. v_r and v_l are the coupling constants for the potential terms that counts the total number of flippable plaquettes with vertical (rung) and horizontal (leg) dimers, respectively. The operator $\hat{P}_f / \hat{P}_v / \hat{P}_h$ represents the projection operator for the respective dimer configurations on a plaquette.

In this formulation, quantum dimers represent singlet bonds of two nearest-neighbor spins on a lattice. In Eq. 3.2, a quantum dimer is denoted by a thick black bond. The hard-core constraint is analogous to that in the classical dimer model: no two dimers are allowed to overlap on the same vertex, ensuring that each spin participates in exactly one singlet bond. Hence, the Hilbert space consists of all valid dimer coverings of the lattice. The RK-QDM describes the quantum dynamics of a phase dominated by singlet bonds. The kinetic term in the Hamiltonian models a ‘resonance’ between two different dimer configurations on a plaquette, analogous to the resonance between the two dimerizations in a benzene ring. This resonance is at the heart of the concept of resonating valence bond (RVB) states. The potential term, on the other hand, energetically favors or penalizes certain plaquette configurations, thereby tuning the balance between static order and resonance dynamics.

Together, these ingredients provide a microscopic realization of RVB states, where the ground state is not a single dimer covering but a coherent quantum superposition of many such coverings. This framework allows one to explore exotic phases of matter, including quantum spin liquids, whose emergence is intimately tied to




Figure 3.2. Generic phase diagram of the quantum dimer model on 2D bipartite lattices. The RK point marks a critical deconfined point. The staggered phase appears immediately to the right of the RK point, while the plaquette phase lies to its left and spans a finite parameter range. For $V/J \rightarrow -\infty$, the columnar phase dominates. The region between the columnar and plaquette phases remains debated. On the honeycomb lattice, the plaquette phase transitions to the columnar phase via a first-order transition around $v/t = -0.2$ [7]. A similar scenario is proposed for the square lattice [8], though some findings hint at a potential ‘mixed’ phase between the plaquette and columnar phases [9]. Representations of all three phases are shown for the square lattice.

the interplay between local constraints and resonant dynamics.

3.2.1 The General structure of QDM ground state phase diagrams

In Eq. 3.2, we can combine the two different coupling constants v_l and v_r into a single variable v , which counts the total number of flippable plaquettes with either horizontal or vertical dimers. By using the ratio v/J , we can describe the ground state of the quantum dimer model (QDM) using a single parameter. This simplifies the phase diagram of the model, which is now governed by this ratio.

For 2D bipartite lattices, the generic phase diagram for the QDM is known, such as square and honeycomb lattices [20]. Fig. 3.2 shows this generic phase diagram. I will further discuss in detail the different phases that are known and explore the key features in the phase diagram.

It is important to note that for the quantum dimer model, the primary focus is on the low-temperature limit ($T \sim 0$), where quantum effects dominate and the system is typically in its ground state.

Rokhsar-Kivelson (RK) Point : This is a special point in the phase diagram where $v = J$, and at this point, the ground state wavefunction $|\psi\rangle$ of the quantum dimer model (Eq. 3.2) is a quantum superposition of all possible perfect dimer matchings:

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_c |c\rangle \quad (3.3)$$

where $|c\rangle$ denotes a perfect dimer configuration, and the sum runs over a topological sector. Each sector contains N distinct dimer coverings that can be connected by local resonance moves, and is therefore closed under repeated action of the kinetic term in Eq. 3.2. In other words, we can construct such an equal-amplitude superposition in every topological sector. All these ground states will be degenerate solutions of Eq. 3.2 at $v = J$.

For 2D bipartite lattices, the RK point is a special critical point that separates different crystalline phases of the system. Another fascinating feature of the RK point is the behavior of monomers. When a quantum dimer is broken, a pair of monomers is created. At the RK point, these monomers are deconfined, meaning that they can be separated to infinite distances at a finite energy cost.

In the limit $v/J \rightarrow -\infty$, the quantum dimer system minimizes its energy by converging to a state with the maximum number of flippable plaquettes. Consequently, a columnar state emerges, resembling the configuration observed in the classical dimer model. Columnar correlations persist up to a critical value $(v/J)_c < 1$. Thus, the region $-\infty < v/J < (v/J)_c$ is referred to as the *columnar phase*.

In the opposite limit, $v/J \rightarrow \infty$, the quantum dimer system minimizes its energy by converging to a state with the minimum number of flippable plaquettes. If a lattice admits dimer configurations with no flippable plaquettes, such states will naturally be favored. This leads to the so-called staggered state, analogous to the ground-state configuration in the classical dimer model. Importantly, such staggered states are highly degenerate, as there exist many distinct dimer coverings with no flippable plaquettes, all of which are exact zero-energy ground states. An example of a staggered state on the square lattice is shown in Fig. 3.2. These states persist in the region just beyond the RK point, $1 < v/J < \infty$, which is therefore referred to as the *staggered phase*.

The nature of the phases between the columnar phase and the RK point $((v/J)_c \leq v/J \leq 1)$ depends on the underlying lattice structure. For higher-dimensional bipartite lattices, the phases in this region remain largely unknown. However, for 2D bipartite lattices, a *plaquette phase* has been observed in this region for both square [8, 81] and honeycomb lattices [7].

In Fig. 3.2, all phases can be interpreted in a mean-field sense: the thick lines represent bonds with a high probability of hosting a dimer, while the thin lines indicate bonds where this probability is low. In the plaquette phase, dimers are predominantly found on all edges of alternating plaquettes with high probability. Numerical studies suggest that the plaquette phase emerges immediately to the left of the RK point and transitions into the columnar phase via a first-order phase transition. All ordered phases, including the columnar, staggered, and plaquette phases, exhibit monomer confinement. The RK point serves as the critical point separating the plaquette and staggered phases. It is the only (isolated) point in the phase diagram where deconfinement occurs, which is why it is referred to as the deconfined critical point.

Cantor deconfinement : is a phenomenon proposed to emerge in certain field-theoretic descriptions claimed to model QDMs on 2D bipartite lattices, particularly when these models are perturbed near the RK point [82]. In the unperturbed QDM, an order parameter known as the ‘tilt’ (or winding number per unit length) jumps discontinuously from 0 to 1 as the system moves from the plaquette to the staggered phase.

Introducing perturbations to the system leads to the emergence of new crystalline phases between the plaquette and staggered states. Consequently, the tilt increases from 0 to 1 in a continuous manner and has zero derivative almost everywhere, except on a Cantor set where the derivative is undefined. This function resembles a staircase, commonly referred to as the ‘devil’s staircase’, and is predominantly observed near the RK point.

Within this devil’s staircase structure, regions corresponding to rational tilt values exhibit confined monomers, and regions associated with irrational tilt values display monomers that are deconfined. This intricate interplay between confined and nearly deconfined states in the system is termed Cantor deconfinement.

This challenges the direct application of Polyakov’s result to QDMs, which asserts that in compact $U(1)$ gauge theories in $2 + 1$ dimensions — such as the 2D bipartite QDM — a truly deconfined phase cannot exist at $T = 0$ [20]. But this also suggests that the confinement-deconfinement transition in QDMs is more nuanced. This implies that while Polyakov’s result provides a general framework, the specific characteristics of QDMs, especially near the RK point with perturbations, lead to a richer and more intricate phase structure than the theorem might initially suggest.




Figure 3.3. A configuration of fully packed loops (thick black lines) on a square lattice (dotted lines). The loops are non-intersecting, self-avoiding, and visit every vertex of the lattice exactly once.

3.3 The $O(n)$ Loop Model

The $O(n)$ model in statistical physics is a versatile framework that describes systems of n -component spins, $\mathbf{s} = (s_1, s_2, \dots, s_n)$, interacting with their nearest neighbors on a lattice [83, 84]. This model encompasses various well-known cases: $n = 1$ corresponds to the Ising model, $n = 2$ to the XY model, and $n = 3$ to the Heisenberg model. The $O(n)$ model is widely employed to study critical phenomena associated with second-order phase transitions.

Many spin models, such as the q -state Potts model, can be reformulated as random cluster models [85], where cluster boundaries in two dimensions form self-avoiding loops. Similarly, the $O(n)$ model can be mapped to a loop-gas model on lattices [86, 87]. In this formulation, the partition function is expressed as:

$$Z_{O(n)} = \sum_C t^{N_e} n^{N_l}, \quad (3.4)$$

where the summation is over all self-avoiding loop configurations C . In each configuration, N_e represents the number of empty vertices not visited by loops, each contributing a weight t , and N_l is the number of loops, each assigned a weight n . Here, t physically corresponds to temperature, and in the zero-temperature limit ($t \rightarrow 0$), the $O(n)$ model reduces to the fully packed loop (FPL) model in Fig. 3.3, where all vertices are visited by self-avoiding loops [87, 88].

When the loop weight $n \rightarrow 0$, the loop model further simplifies to a single Hamiltonian cycle, a self-avoiding loop that visits each vertex exactly once [88]. This close relationship links Hamiltonian cycles, the FPL model, and the broader $O(n)$ loop model [89]. Studying Hamiltonian cycles on specific lattices thus provides valuable insights into the $O(n)$ loop model on the same systems.

Moreover, the FPL model has been applied to investigate the configurational statistics of dense polymers [90], demonstrating its utility beyond traditional lattice spin models. These connections underscore the $O(n)$ loop model's significance in understanding critical phenomena and phase behavior in statistical mechanics.

Energetics: The fully packed loop (FPL) model studied in Chapter 8 includes an aligning interaction energy term that accounts for parallel loop segments on a single plaquette, similar to the classical dimer model. The

partition function and energy for the interacting FPL model on AB tiling are defined as follows:

$$Z = \sum_c e^{-\beta E_c}, \quad (3.5)$$

$$E_c = V \cdot N \left(\begin{array}{|c|} \hline \square \\ \hline \end{array} + \begin{array}{|c|} \hline \square \\ \hline \end{array} \right)$$

In the partition function Z , the summation runs over all possible fully packed loop configurations on the AB tiling. The parameter $\beta = 1/T$ represents the inverse temperature. The interaction energy E_c for each configuration counts the number (N) of flippable plaquettes with parallel loop segments. Note that, no energy is assigned to a plaquette containing more than two loop segments.

For this model, the interaction strength V is set to -1 , signifying attractive interactions between parallel loop segments. This choice highlights the tendency of the system to favor configurations with maximum flippable plaquettes, thereby influencing the statistical and energetic properties of the FPL model on any tiling.

3.4 Height Model Framework

A height representation establishes a one-to-one correspondence between the microstates of models like dimer and loop models and the configurations of a related height interface model. Visually, this is similar to constructing a surface of discrete heights over the lattice, where the height values encode the local structure of the underlying microstate. This concept has been pivotal in analyzing critical phenomena on two-dimensional lattices, enabling the mapping of complex interactions onto a more tractable scalar field framework. In these systems, correlation functions typically exhibit power-law decay with distance, a hallmark of criticality. Models such as dimer and loop models, which are prevalent in statistical physics, can be reformulated using height variables, revealing deep connections among diverse physical systems. This reformulation facilitates the application of conformal field theory techniques, aiding in the computation of correlation functions and critical exponents, and thereby deepening our understanding of their critical behavior. Furthermore, height representations illuminate roughening transitions and phase behaviors, making them a valuable analytical tool in statistical mechanics. Pioneering contributions by researchers such as [91, 92] have laid the foundation for applying height representations to a wide variety of models in this field.

To construct a height representation for a perfect dimer configuration on any bipartite lattice with all vertices having degree d , consider partitioning the lattice sites into two bipartite sets, A and B . The height field can be assigned as follows (see Fig. 3.4):

- Assign zero height to an initial plaquette, say $h(\vec{x}) = 0$.
- Select any corner of this plaquette. If the site belongs to the B (A) sublattice, move to the adjacent (next) plaquette in a counterclockwise (clockwise) direction around that site.
- Assign a height $h(\vec{x}') = h(\vec{x}) + 1$ to the next plaquette (following the chosen direction around the selected corner as described above) if there is no dimer between the two plaquettes. If a dimer is present, assign $h(\vec{x}') = h(\vec{x}) - C_r$, where $C_r = d - 1$ is the compactification radius, and d denotes the degree of the vertices.
- Select the plaquette that has just been assigned a height and repeat steps 2 and 3. Continue this procedure until every plaquette has an assigned height.




Figure 3.4. (a) In the square lattice, the sites can be partitioned into two bipartite sets: A (blue) and B (red). (b-c) Two examples of height representations in the fully packed dimer model on a square lattice, constructed using the rules described in Sec. 3.4. The construction begins from the top leftmost plaquette, which is assigned a height zero. These rules apply to any 2D bipartite lattice.

These steps creates a discrete height surface above the lattice. It is important to note that the height field assignment is non-unique, as the initial plaquette can be assigned any arbitrary integer height. Nevertheless, the differences in height values between neighboring plaquettes, which encode the physical properties of the dimer system, remain well-defined. This representation underscores the utility of height variables in providing a unified framework for studying the critical and geometric properties of dimer models.

The conformal field theory (CFT) of loop models can be explicitly constructed using Liouville field theory [88]. This formulation maps the loop model to a model of a fluctuating height surface, with the loops serving as height contour lines [93]. Heights are defined according to the following set of rules (see Fig. 3.5):

- Randomly assign orientations (clockwise or anti-clockwise) to all fully packed loops in a given configuration.
- Start from any arbitrary plaquette and assign it an initial height, $h(\vec{x}) = 0$, where \vec{x} is the position of the center of the plaquette.
- Hop to neighboring plaquettes: Let the common edge between two plaquettes be e , and let the center of the next plaquette be \vec{x}' . The height model is then defined as follows:
 - $h(\vec{x}') = h(\vec{x}) + 1$ if e is part of a loop and is directed to the right when viewed from \vec{x} .
 - $h(\vec{x}') = h(\vec{x}) - 1$ if e is part of a loop and is directed to the left when viewed from \vec{x} .
 - $h(\vec{x}') = h(\vec{x})$ if e is not part of any loop.

The heights $h(\vec{x})$ obtained through the steps outlined above, are discrete and reside on the plaquettes of the lattice, or equivalently, on the vertices of the dual lattice. These discrete heights can be coarse-grained to form a *continuous height surface* over the lattice. To do this, define a fine grid over the tiling, with as many points as necessary to achieve the desired resolution. For each grid point, assign the height value to be equal to the height of the plaquette in which the point lies. For points on the grid outside the tiling, heights can be computed by extending the height values from the nearest boundary plaquette.

We can study this continuous height surface by looking at its variance and using a Fourier transform. This helps us understand the type of fluctuations in the height field and whether the surface is ‘rough’ or ‘smooth’. A rough height surface means the system behaves critically, similar to many two-dimensional lattice models




Figure 3.5. Example of height representations in the fully packed loop model on a square lattice, constructed using the rules described in Sec. 3.4. These rules apply to any 2D bipartite lattice.

[92, 94, 93]. This roughness is characterized by the Fourier transform of the height fluctuations, which, after coarse-graining, follows the relation [93]:

$$\langle |h(\vec{p})|^2 \rangle \sim \frac{K}{|\vec{p}|^2}. \quad (3.6)$$

Here, $h(\vec{p})$ is the Fourier transform of the height function $h(\vec{x})$, and K is a constant related to the system's stiffness. This indicates that long-wavelength fluctuations are significant, leading to a rough surface profile.

Conversely, in the ‘smooth’ phase, the height field exhibits only short-range fluctuations, and the Fourier transform behaves as: $\langle |h(\vec{p})|^2 \rangle \sim \text{const}$, when $\vec{p} \rightarrow \vec{0}$. This suggests that the system is not critical, as the model tends to select a mean height, resulting in minimal height variations that do not scale with the system size.

The behavior of the height field can also be characterized by examining its net height variance, W^2 , defined as:

$$W^2 \equiv \langle h(\vec{x})^2 \rangle - \langle h(\vec{x}) \rangle^2. \quad (3.7)$$

Here, the averages are taken over an ensemble of various loop or dimer configurations. In the thermodynamic limit, if W^2 approaches a finite value, the system is in a smooth phase, indicating that the height field selects a particular mean height with minimal fluctuations. Conversely, in a rough phase, W^2 exhibits a logarithmic divergence with system linear size L , typically expressed as below, where K is a constant related to the stiffness of the height field:

$$W^2 = \text{const} + \frac{1}{2\pi K} \ln L. \quad (3.8)$$

NP-Hard Problems on the AB tiling

Optimization problems are central to decision-making across numerous disciplines, such as engineering, computer science, economics, physics, and logistics. These problems require identifying the best solution from a set of feasible options, considering certain constraints and objectives. However, as the size and complexity of these problems increase, discovering optimal solutions becomes more challenging and often demands substantial computational resources. This convergence of optimization and computational complexity has led to the classification of problems into various complexity classes, particularly NP -complete and NP -hard [49]. These classifications shed light on the intrinsic difficulty of optimization tasks and the computational resources necessary to solve them.

NP -complete (NPC) problems are a class of decision problems that yield a simple ‘yes’ or ‘no’ answer. The term ‘ NP ’ stands for ‘nondeterministic polynomial time’, implying that while it may not be possible to find a solution in polynomial time, a given solution can be verified within polynomial time. Every problem in the NP complexity class can be reduced to an NP -complete problem in polynomial time. This implies that if any NP -complete problem can be solved in polynomial time, then all problems in NP can also be solved with the same efficiency. This characteristic is known as ‘completeness’. Examples of NP -complete problems include the traveling salesman problem, graph coloring problem, and subset sum problem, as well as various combinatorial optimization problems [46, 49].

Conversely, NP -hard problems may or may not be decision problems and often require finding a solution if one exists. They represent a class of problems that are at least as challenging as the hardest NP problems but may not necessarily have a binary solution. Unlike NP -complete problems, there is no requirement for solutions of NP -hard problems to be verified in polynomial time. These problems serve as a benchmark for assessing the computational complexity of optimization tasks.

A well-known NP -complete problem with significant applications in physics, computer science, and biology is the Hamiltonian cycle problem. A Hamiltonian cycle is a closed, self-avoiding loop that visits every vertex of a graph exactly once. The study of Hamiltonian cycles dates back to the 9th century CE when the Indian poet Rudrata wrote a poem based on a ‘Knight’s tour’ of a chessboard. The relevance of these cycles spans multiple scientific fields. For example, in protein folding [18, 95], the compact state of a folded protein is well approximated by such cycles. This is because the self-avoiding property of these cycles models the short-range repulsion between amino acid chains. Similarly, Hamiltonian cycles have been used in studying the dynamics of polymer melting [96, 97, 17], where they provide insights into critical scaling exponents of polymer chains,

especially in 2D.

Also, such loop structures frequently emerge as critical elements in models describing phenomena like topological phases and dynamic domain boundaries in Ising magnets. These loops manifest in systems such as the trajectories of quantum particles through space-time and the basis-state wavefunctions for the Toric quantum error correcting code [98]. The $O(n)$ model in statistical physics, which describes the spins of n components that interact with their nearest neighbors, can be mapped to a self-avoiding loop problem [83, 84]. In the zero-temperature limit and for $n \rightarrow 0$, this model simplifies to configurations of Hamiltonian cycles [88, 86]. As such, studying Hamiltonian cycles offers valuable information on various physical models, mainly studied to date in the context of periodic regular lattices [84, 97, 86, 88, 90, 99].

However, many complex structures modeled by Hamiltonian cycles, such as proteins and polymers, lack translational symmetry and exhibit a disordered growth. To better represent these intricate structures, it is essential to investigate Hamiltonian cycles in non-periodic settings. Such exploration could provide valuable insights into optimization tasks specific to these disordered environments [100]. In fact, the lack of periodicity in certain systems, such as quasicrystals, often leads to novel behaviors in form of exotic criticality [38, 39, 40] to charge order [61, 101] to topology [102, 103]. This makes the study of Hamiltonian cycles in aperiodic settings particularly intriguing.

In this chapter, I present a polynomial-time algorithm for constructing Hamiltonian cycles on arbitrarily large finite subgraphs of infinite Ammann-Beenker (AB) tilings defined in Sec. 2.1. These 2D aperiodic tilings lack translational symmetry but still feature long-range order. This long-range order allows for exact analytical results to be proven, and their infinite extent offers a pathway to explore the thermodynamic limit, which is crucial in condensed matter physics, where quasicrystals with AB symmetries are observed [41, 42]. Recent experiments in systems like twisted trilayer graphene [43] and optical lattices [104] further underscore the relevance of AB tilings in physical phenomena.

While efficient Hamiltonian cycle constructions exist for specific classes of graphs [105, 106, 107], AB tilings do not belong to any such category. What makes the existence of Hamiltonian cycles on AB tilings even more surprising is that their close relative, the Penrose tiling, cannot admit Hamiltonian cycles. Despite sharing many similar properties with AB tilings, Penrose tilings lack perfect matchings, a necessary condition for the existence of Hamiltonian cycles in any graph [4].

Through results in this chapter, I demonstrate that AB tilings are a previously unknown special case of the Hamiltonian cycle problem that falls within the P complexity class, rather than NPC . This discovery does not suggest that all NP problems can be solved in polynomial time in general settings, but it opens new avenues for tackling other non-trivial optimisation problems in such unusual aperiodic settings. I will also explain how these findings can have potential applications in areas such as atomic imaging, chemical catalysis, and the Potts model.

This chapter is based on the work in Ref. [3], authored by me, Jerome Lloyd, and Felix Flicker. The proof in Sec. 4.2 was derived by the three of us together, based on earlier work by Jerome and Felix [10]. The proofs in Sec. 4.3 were derived by me, except for the three-coloring problem proof, which was derived by Felix and me together.

4.1 Ammann-Beenker Tilings - Key Properties

From Sec. 2.1, we can see that a set of AB tilings generated using inflation rules is infinite. There are two key consequences of this inflation structure relevant to our analysis:




Figure 4.1. The seven vertex configurations and their transformations under inflation (Fig. 4). Each vertex maps to an 8-vertex after at most two inflation steps. (Following Ref.[10])

Definition II.1 (Linear repetitivity): There exists a finite constant $C > 0$ such that for any set of vertices V within a disk of radius r , every disk of radius Cr contains a copy of V . AB tilings are linearly repetitive.

Definition II.2 (Local isomorphism): Two tilings are locally isomorphic if any finite patch in one tiling appears in the other. All AB tilings are locally isomorphic, and unless otherwise specified, ‘AB’ refers to the entire set of these tilings. Due to linear repetitivity and local isomorphism, any reference to a vertex set V in a specific AB tiling is considered modulo translations.

In a locally isomorphic set, arbitrary rescalings of tilings are not distinguished. For instance, the edges and tiles surrounding each vertex in the tiling can be classified into one of seven unique configurations, called ‘m-vertices’, where ‘m’ refers to the vertex connectivity. The AB tiling configurations include 3-, 4-, 5_A -, 5_B -, 6-, 7-, and 8-vertices (shown in Fig. 4.1), each appearing with a frequency determined by the silver ratio δ_S . Configurations with five edges are distinguished by their inflation behavior.

The 8-vertices are especially significant because any vertex configuration can transform into an 8-vertex after at most two inflations. This allows us to draw edges between 8-vertices to form a new AB tiling, highlighting the discrete scale symmetry that contributes to the distinctive physical properties of quasicrystals.

An 8₀-vertex becomes an 8_n-vertex after $2n$ inflations. I denote the local empire of a 8₀-vertex as W_0 , which exhibits discrete eightfold rotational (plus reflection) symmetry (D_8 in Schönflies notation) [108]. The local empire of a vertex configuration is the simply connected set of tiles that invariably appears around that configuration wherever it occurs in the tiling [109]. The inflation rule maps 8_n-vertices to 8_{n+1}-vertices, enlarging their local empire. This D_8 symmetry persists under inflation, with 8_n-vertices having local empires W_n that grow approximately as $R_n \approx \delta_S^{2n}$.

The empire W_n contains all $W_{0 \leq m \leq n}$ empires. Due to linear repetitivity, any finite vertex set in AB can be found within W_n for sufficiently large n . Hence, I will focus on W_n to derive more general results. A modified version of the AB tiling, with all 8-vertices removed, is known as the AB* tiling [10]. Individual squares and rhombuses are classified as level L_0 , while twice-inflated ones are level L_1 . Since any vertex becomes an 8-vertex after two inflations but not after one, the once-inflated tiles are designated as level $L_{1/2}$.

This chapter considers graphs constructed from the vertices and edges of AB tilings. I will employ the shorthand notation $G = AB$ to denote the case where G represents the graph derived from any infinite AB tiling. Additionally, I may refer to graphs formed from finite sections of AB, which will be evident from the context.

4.2 Hamiltonian Cycle Problem

This *NPC* decision problem asks whether a given graph G admits a Hamiltonian cycle or not. If the answer is yes, the corresponding optimization problem seeks to find such a cycle on G . For the Ammann-Beenker

(AB) tiling, I will present a constructive proof of the following theorem:

Theorem : Given an AB tiling and a finite set of vertices $V \subset AB$, there exists a set U_n , where $V \subseteq U_n \subset AB$, such that U_n contains a Hamiltonian cycle H .

Proof : I will focus on the finite tile regions, denoted as W_n , generated by applying the twice-inflation process n times to the local empire of the 8_0 -vertex (see Fig. 4.3). These regions exhibit a D_8 symmetry. For sufficiently large n , any finite vertex set $V \subset AB$ will be contained within this infinite hierarchy of W_n regions.

In Section 4.2.1, I will explain how we found a set of fully packed loops (FPL) on the AB^* tiling, i.e., the AB tiling with all 8_0 -vertices removed. This FPL set will consist of edges such that every vertex (except for the 8_0 -vertices) is touched by two such edges.

In Section 4.2.2, I outline our method for reconnecting these fully packed loops on the AB^* tiling, ensuring that all 8_0 -vertices in W_n are included. First, we create twice-inflated versions of FPLs on AB^* for all $m \leq n$. Then, we identify a process to connect these hierarchies of FPL sets into a loop that visits every vertex in W_n , except the central 8_n -vertex. The final step is to include the central 8_n -vertex and demonstrate how certain loops can be merged into a single loop that also visits the central 8_n -vertex, forming the Hamiltonian cycle. The resulting vertex set, visited by this H-cycle, is denoted as U_n . As U_n grows exponentially with n , it eventually covers increasingly larger sections of the AB tiling. By the property of linear repetitivity, any finite vertex set V will be included in U_n for sufficiently large n .

4.2.1 Constructing Fully Packed Loops on AB^*

Consider the twice inflation of the rhombus and square prototiles shown in Fig. 4.2, along with a set of thick black edges, denoted as e_0 , highlighted on the L_1 tiles. These e_0 edges connect to every vertex of the L_1 tiles, except at the corners. Additionally, each vertex within the L_1 tiles is touched by exactly two e_0 edges, except for those on the boundary. When these L_1 tiles are combined to form a valid patch of the Ammann-Beenker (AB) tiling, the e_0 edges cover all the vertices in the patch, except those at the corners of the L_1 tiles. These corner vertices correspond to the 8_0 -vertices of the L_0 tiling. Since every vertex other than the 8_0 -vertices has exactly two e_0 edges connected to it, the e_0 edges naturally form loops around all vertices, excluding the 8_0 -vertices. Thus, these edges create a set of fully packed loops on the AB^* tiling (see Fig. 4.3).

We picked this particular choice of the set of e_0 edges as canonical, as it forms closed loops that respect the D_8 local symmetry of the AB and AB^* tilings. However, it is possible to find other sets of edges with these same properties.

4.2.2 Constructing Hamiltonian cycles on AB

In this section, I present our method to construct the Hamiltonian cycle on AB (for any level n) by joining the fully packed loops (FPLs) on AB^* and incorporating the missing 8_0 vertices into a single loop. To achieve this, consider the red alternating path in Fig. 4.2, consisting of alternating e_0 edges. Augmenting this path will swap the e_0 edges, covering the other set of alternate edges on the path that are currently uncovered. This augmentation has three effects: (i) the two corner vertices of the L_1 tiles (i.e., the terminating vertices of the red path) will now be covered by e_0 edges, (ii) the total number of e_0 edges will increase by one, and (iii) any vertex previously visited by a pair of e_0 edges will still be visited by a pair of e_0 edges.

While other alternating paths with these three properties can exist, we identified this red path as our canonical choice for the following reasons: (i) the same path works for both L_1 tiles, producing the three




Figure 4.2. Twice-inflations of the two AB prototiles. These represents level one (L_1) of the tiling. Each small tile is at level zero (L_0). The thick black edges visit all but the corner vertices. These edges are called e_0 edges. Since the 8-vertices at L_0 corresponds to the tile corners at L_1 , the union of e_0 constitutes fully packed loops on AB^* (AB without the 8-vertices). Augmenting the red path (switching covered and uncovered edges) places e_0 edges on the two corner vertices while still visiting the original vertices. The red path can be thought of as the twice-inflation of an e_0 edge, and denoted as e_1 . It has the same effect along any tile edge. (Following Ref. [3])




Figure 4.3. A patch of AB tiling showing the eightfold symmetric region W_1 , and W_0 in gray. W_1 is the twice inflation of W_0 , formed from the double-inflated tiles shown in Fig. 4.2. Purple lines show the boundary of these L_1 tiles (and do not belong to the tiling itself). The thick black edges form a fully packed loop set on the AB^* patch. (Following Ref. [3])




Figure 4.4. The red alternating path in Fig. 4.2 can be thought of as a twice-inflation of an e_0 edge from which loops are constructed at level zero (L_0). It connects nearest neighbors in the L_1 tiling (8-vertices in the L_0 tiling) while following an alternating path in the canonical choice of e_0 placements. I therefore denote this path an e_1 edge. Alternating paths connecting higher-order 8_n -vertices can be constructed by concatenating e_{n+1} edges in the same way that e_1 was formed from a concatenation of e_0 ; levels e_2 and e_3 are shown here. The orientation of e_n is defined according to the arrow. (This figure was made by my supervisor.)

aforementioned effects, (ii) it does so while remaining entirely within the tiles, ensuring that the augmentation does not interfere with any neighboring L_1 tiles, and (iii) the augmentation can be undone by performing the same operation along the path a second time, which proves to be a crucial feature.

The red alternating path follows the boundary edges of L_1 tiles as closely as possible and terminates at the corners of L_1 tiles, so essentially it can be considered as an edge of either L_1 tiles, denoted by e_1 , similar to e_0 edge for L_0 tiles. So naturally an e_1 edge can be thought of as a twice-inflation of an e_0 edge. The edge inflation rules are defined by continuing this pattern: an e_2 edge (twice-inflation of e_1) is constructed by joining e_1 edges in the same manner as an e_1 edge is constructed from e_0 edges. The first three levels of this process are depicted in Fig. 4.4. Also note that, e_0 edges connects all non-8-vertices at the L_0 level, while e_1 edges connect 8_0 -vertices. Higher inflations will follow this pattern, so e_{n+1} edges will connect vertices of type 8_n .

The fractal nature of e_n structure at infinite inflations can be captured by calculating its fractal dimension. Going from e_0 to e_1 , the original side length of each prototile is subdivided into nine segments, each of length δ_S^{-2} . This same scaling pattern holds for all subsequent inflations, leading to the fractal dimension of e_∞ being expressed as:

$$\dim(e_\infty) = \lim_{n \rightarrow \infty} \frac{\log(9^n)}{\log(\delta_S^{2n})} = \frac{1}{\log_3(1 + \sqrt{2})} \approx 1.246. \quad (4.1)$$

Although e_1 edges must follow the boundaries of L_1 tiles, each edge is shared by two tiles, allowing two possible orientations for every e_1 . These orientations can be chosen freely without affecting any results. However, we adopt a natural choice where e_1 edges are oriented as shown in Fig. 4.4. In L_1 , we alternate the direction of e_1 edges along a loop, pointing alternately inward and outward. Since the AB tiling is bipartite, all loops have even lengths, ensuring no inconsistencies.

To finalize the proof, I will now demonstrate how all the FPLs on AB^* and all 8-vertices can be incorporated into a single closed loop (Hamiltonian cycle), for any large n . Fig. 4.5 - Fig. 4.8 show all the next steps for level

$n = 2$ of AB tiling, which is 4-times inflated version of W_0 .

The canonical e_0 edge arrangement ensures that all $8_{n<0}$ (vertices other than 8_0 -vertices) vertices in AB are connected by e_0 loops. To include all 8_0 vertices, we introduced $e_{\leq n}$ loops based on the canonical placement of edges, where e_n loops are $2n$ -inflation of e_0 loops (see Fig. 4.5). The canonical placement of such e_n loops remains consistent across all levels, as the twice inflation or deflation of any AB tiling results in another AB tiling, preserving the placement.

Replace all the e_n loops with their corresponding e_n edges from Fig. 4.4. We can see that augmenting the closed e_1 loops produces the three effects mentioned earlier, ensuring that all 8_0 -vertices along the loop are incorporated into the e_0 loops as required. Similarly by induction, 8_n -vertices can be added by connecting them through loops of e_{n+1} edges. After n steps, every vertex with $8_{\leq n}$ is included in a set of fully packed loops on AB defined on those vertices.

Note that the placement of e_n edges does not interfere with the established connections of e_{n-1} edges. When e_2 edge loops are placed, it is always possible to choose an orientation such that the overlapping e_1 edges in e_2 loop and e_1 loop will have the same orientation (see Fig. 4.6). Then augmenting those overlapping edges will not have any effect, hence can simply be deleted. Removing these overlapping edges will join e_1 and e_2 loops and rewire them into a single alternating path of e_0 edges (see Fig. 4.7). The rewiring works at any level by induction. Augmenting this single alternating loop will connect the underlying e_0 edges into a single loop visiting every vertex in W_n , but the central 8_0 vertex. It is not possible to include this vertex, because till now none of the $e_{\leq n}$ loops visits it as they admit a D_8 symmetry which forces these loops to circle the central 8_0 vertex. To include this vertex in the cycle we need to break this D_8 symmetry. Note that, even if we take the smallest e_0 loop that is a star around any 8_0 vertex in L_0 tiling, we can include the central vertex in this loop by just inverting one of the points of the star.

Based on this observation, a simple way to include central-8 vertex is to fold a single corner of the outermost e_n loop, $2n$ -fold inflation of the smallest star loop, encircling the 8_n -vertex. This connects every loop contained within it (and all those it passes through) into a single loop. Augmenting this one loop will result in a Hamiltonian cycle on the simply connected set of vertices visited by this deformed star. We denote this set U_n to distinguish it from the D_8 -symmetric set W_n . The cycle on U_2 is shown in Fig. 4.8.

Note that any vertex in V will be contained in some U_n , by linear repetitivity (see definition II.1). By this property, any vertex lying in a disk of radius r will also be in all disks of radius Cr . Define the radius of the largest disk within U_n to be Cr . Since the e_n star loop (like the blue loop in U_2 in Fig. 4.5) will be consistently present at any level n , we can consider the largest possible disk inside it with radius Cr as shown by the green disk in Fig. 4.5.

Let l_2 represent the green chord lengths, which are the edges of an L_2 rhombus. The radius of this disk can be expressed as $\frac{l_2 \cos(22.5^\circ)}{2} = l_2 \sqrt{\frac{\cos(45^\circ)+1}{2}} = 2^{-1/4} \delta_S^{-1/2} l_2$. Similarly, the largest equivalent disk in region U_n has a radius of $2^{-1/4} \delta_S^{-1/2} l_n$. Given that e_0 is a unit length, the inflation rules illustrated in Fig. 2.4 suggest that $l_n = \delta_S^{2n}$, and therefore, U_n encompasses a disk with a radius of $2^{-1/4} \delta_S^{2n-1/2}$. Consequently, linear repetitivity indicates that every U_n contains all vertex sets within a disk of radius $\frac{2^{-1/4} \delta_S^{2n-1/2}}{C}$. This radius increases exponentially with n , ensuring that any set V will be within U_n for sufficiently large n .

For any region U_n , a Hamiltonian cycle can be constructed using the method described in this section. Hence given any finite vertex set V , there exists a region U_n such that $V \subseteq U_n$. This allows U_n to extend




Figure 4.5. The canonical placements of e_0 edges (black) form e_0 -loops visiting every $8_{n<0}$ vertex (i.e. every vertex which is not an 8-vertex). The red and blue lines form e_1 - and e_2 -loops, respectively. The central 8-vertex is now added onto loops by folding one corner of the level ' n ' inflation (blue star) inwards to break the 8-fold symmetry. Here we have shown the process explained in Sec. 4.2.2 to order $n = 2$, but it can be iterated to any order n . The green outline represents the largest disk that fits within the blue loop, connecting three 8-vertices with green chords. (Following Ref. [3])




Figure 4.6. Next, e_1 edges from Fig. 4.4 are placed along the (red) e_1 loops to incorporate 8_0 -vertices into the loops. Subsequently, e_2 edges from Fig. 4.4 are added along the (blue) e_2 loop, with one corner folded, to include the central vertex and 8_1 -vertices into the loops. Note that some blue e_2 edges overlap with the red e_1 edges. This process can be iterated to any desired order. (Following Ref. [3])




Figure 4.7. The overlapping e_2 edges and e_1 edges as shown in Fig. 4.6 will augment to leave the L_0 tiling unchanged. Hence we can remove these overlapping e_2 edges and e_1 edges completely. This results in the rewiring of e_1 -loops and the e_2 -loop into a single loop as shown here. Finally augmenting this single loop, consisting of the remaining red and blue edges, will give the Hamiltonian cycle shown in Fig. 4.8. (Following Ref. [3])




Figure 4.8. The Hamiltonian cycle visiting all the vertices of a U_2 region (the AB tiles themselves are omitted for clarity). The image is obtained from Fig. 4.5 by placing e_1 and e_2 edges in alternate orientations along the red and blue loops respectively, as shown in Fig. 4.6, and augmenting. Note the present figure has been rotated through $1/16^{\text{th}}$ of a turn relative to other figures. (Following Ref. [3])

naturally to the thermodynamic limit, where $U_n \rightarrow \infty$ represents arbitrarily large finite patches. The algorithm operates with linear complexity in the number of vertices within U_n .

The region U_∞ , without any surrounding tiles, can be viewed as an infinite AB tiling. It forms an infinite, simply connected set of tiles generated by inflating a valid AB patch (with surrounding tiles in W_∞ removed).

While an AB tiling can have at most one global center of D_8 symmetry (since multiple centers would violate the crystallographic restriction theorem [110]), the set of such tilings with a global D_8 center is of measure zero among all AB tilings [52]. However, our definition of the thermodynamic limit does not require the presence of a global D_8 center; it is sufficient that at least one region ($U_{n \rightarrow \infty}$) is included, and all AB tilings contain infinitely many such regions.

4.3 Solution to other NP -Hard problems on AB

There are several other problems classified as NP -Complete, similar to the Hamiltonian cycle problem [49, 46]. These are decision problems that inquire whether a solution exists for a general class of graphs. The corresponding optimization problems for these are NP -hard, requiring the discovery of a solution if one exists. In this section, I will present solutions to five such challenging optimization problems on the AB lattice, utilizing discrete scale symmetry and/or our Hamiltonian cycle constructions. This demonstrates once again that the AB lattice offers a unique and simplified case. I will also discuss the physical applications of the solutions to the first three problems, while the last two problems are inspired purely by mathematical considerations.

4.3.1 Equal-weight travelling salesperson problem

Problem Statement [49, 46] - Given a number of cities N , unit distances between each pair of connected cities, and an integer k , does there exist a route shorter than k which visits every city exactly once and returns back to the original city? The corresponding optimization problem is to find such a route.

Solution on AB - If cities are the vertices of $U_{n \geq 0}$, yes if and only if $k > |U_n|$.

Proof - In graph theory, each city can be represented as a vertex, and an edge connects cities with an edge weight indicating the distance between them. The problem of finding the shortest route that visits each city exactly once and returns to the starting city is equivalent to finding a Hamiltonian cycle with the minimum sum of edge weights.

For an unweighted AB graph, where all edge weights are set to one, this problem simplifies to the equal-weight traveling salesperson problem (EWTSP), which in turn is the Hamiltonian cycle problem itself. Consequently, the Hamiltonian cycle constructed in the previous section provides a solution to the EWTSP on U_n . For any integer k , we can have an $U_{n \geq 0}$ region such that $|U_n| < k$, and the Hamiltonian cycle created on such U_n using the steps in previous sections will be the solution of this problem.

4.3.2 Application: ultra - high resolution atomic imaging

Atomic imaging of material surfaces is typically done using techniques like scanning tunneling microscopy (STM) and magnetic force microscopy (MFM). In STM, a voltage between a sharp tip and the surface enables electron tunneling, producing a current that reflects local electronic structure [111, 112]. MFM, by contrast, measures magnetic forces with a magnetized tip to image magnetic domains [113]. One application of the traveling salesperson problem is to find the most efficient route for such tips to image every atom on a surface. For simple periodic cases, such as atoms arranged in a square lattice, a snake-like Hamiltonian cycle that visits

each site once in a regular sequence provides a straightforward solution. However, even with the knowledge of such an obvious most efficient route, the high resolution imaging of a $100 \times 100 \text{ nm}^2$ area of a surface with square lattice symmetry can take up to a month.

While STM and MFM are commonly used for crystalline surfaces, they can also image physical quasicrystals, such as those with the symmetries of AB tilings [41]. Recent experiments have demonstrated tunable quasicrystal geometries in twisted trilayer graphene [43] and eightfold symmetric structures have also been created in optical lattices [104, 114].

In such cases, the optimal route for the STM tip is not immediately apparent. Our solution to the traveling salesperson optimization problem on AB tilings provides the most efficient route for the STM tip to take while imaging any such aperiodic quasicrystal surfaces. Although the surface needs to be scanned initially to define U_n regions, this route ensures maximum efficiency for repeated scans under varying conditions, such as changes in temperature or magnetic field.

4.3.3 Longest Path Problem

Problem statement [49, 46] - Given an unweighted graph G and an integer k , does G contain a path of length at least k ? The path may not revisit vertices. The corresponding optimization problem is to find a maximum length path.

Solution - Yes, if $G = U_{n \geq 0}$ and $k \leq |U_n| - 1$.

Proof: A Hamiltonian path can be derived by removing any single edge from a Hamiltonian cycle. Thus the results from Sec. 4.2.2 confirms the existence of such a path of length $|U_n| - 1$ within any region U_n .

Comment: Shorter paths can be obtained by removing additional contiguous edges. Since every AB tiling contains regions U_n for any n , this implies that paths of any desired length can be found within the tiling.

4.3.4 Application: Adsorption and catalysis

Adsorption has significant practical applications across various industries. In catalysis, for instance, the adsorption of molecules onto a surface can lower the activation energy required for reactions. While the efficient packing of molecules has been extensively studied in periodic crystal structures, such as square lattices with limited bond angle options, quasicrystalline surfaces—such as those featuring AB tiling—provide a broader range of bond angles. This flexibility may better accommodate the adsorption of more adaptable or flexible molecules. Potential applications of this phenomenon include carbon sequestration, storage, and protein adsorption.

As mentioned earlier, the ‘dimer model’ was originally inspired by studies of the packing statistics of short, linear molecules adsorbed onto crystal surfaces [13, 14, 115]. However, it can also extend to the atomic-scale symmetries observed in AB tiling within physical quasicrystals. The existence of a longest path visiting every vertex in these systems suggests that long, flexible molecules, such as polymers, could be efficiently packed onto selected surfaces of such materials. These paths can be divided into smaller segments, indicating that molecules of any length can pack efficiently. Furthermore, the diverse range of rotational symmetries in quasicrystals enables them to more accurately approximate real-space configurations and accommodate a wider variety of bond angles, potentially enhancing the adsorption of folded molecules. This suggests that surfaces with quasicrystalline symmetries could serve as superior adsorbing materials for applications in catalysis, carbon capture, or fertilizer production [116, 13, 115].




Figure 4.9. (a),(b) Tiles used for 3-coloring of the AB tiling. (c) The structure with three (shaded) rhombuses present around every ladder. (Following Ref. [3])

4.3.5 Three-coloring Problem

Historically, the five-color theorem was proven in the 19th century and it states that any political map can be colored with at most five colors to avoid neighboring regions sharing the same color [117]. The four-color theorem, which posed a greater challenge, was eventually proven in 1976 using theorem-proving software [118]. The three-coloring problem is known to be NP -complete on general graphs, though certain special cases can be solved. For instance, a polynomial-time algorithm for three-coloring the rhombic Penrose tiling was found in 2000 [119], covering all rhombus tilings, including AB tilings. Our solution to the three-coloring problem on AB tilings follows a distinct approach and generates a different three-coloring.

Problem statement [49, 46] - Can all the faces of a planar graph G be colored using three colors such that no faces sharing an edge share a color? The corresponding optimization problem is to find such a three-coloring.

Solution - Yes, if $G \subseteq AB$.

Proof - The AB tiling is a bipartite graph, meaning its vertices can be divided into two disjoint sets where no edge connects vertices from the same set. Assign two opposite bipartite charges to the vertices based on the set they belong to.

I define two distinct three-colored tile sets, labeled 1 and 2 as shown in Fig. 4.9 (a,b). These two tile sets can tile the entire AB tiling as follows:

- (i) Place tile 1 on all 8-vertices with the same bipartite charge.
- (ii) Place the mirror image of tile 1 on all 8-vertices with the opposite charge.
- (iii) Place tile 2 or its mirror image such that the three (shaded) rhombuses (shown in Fig. 4.9(c)) around each ladder match in color.

This consistent placement is guaranteed by the structure of the tiling itself. After filling an AB tiling in this way, the only parts that will not be colored are ladders (see Fig. 4.10), which can be colored consistently based on their surrounding tiles so that no two adjacent faces share the same color. The resulting three-coloring of the AB tiling is displayed in Fig. 4.11. Note that this 3-coloring solution on AB is not unique.

This proof does not rely directly on the existence of Hamiltonian cycles. Instead, it utilizes an intermediate step in their construction — specifically, the fully packed loops on AB^* discussed in Sec. 4.2.1 — along with the




Figure 4.10. A small patch of 3-colored AB tiling using the steps as mentioned in Sec 4.3.5. Tile 1 and tile 2 are outlined in gold. The only remaining portions of the tiling are segments of ladders that can be 3-colored consistently based on their surrounding tiles. (Following Ref. [3])

discrete scale symmetry of AB. It demonstrates that intermediate steps in the Hamiltonian cycle construction can also solve other *NP*-complete problems on AB without needing to rely on the final result.

4.3.6 Application: The Potts model and protein folding

The q -state Potts model, with $q \in \mathbb{N}$, generalizes the Ising model to spins σ_i that can take one of q values [120]. It is defined by the Hamiltonian:

$$H = -J \sum_{\langle ij \rangle} \delta_{\sigma_i, \sigma_j} \quad (4.2)$$

where the sum runs over nearest neighbors, and δ is the Kronecker delta. The model can be defined on the AB tiling by placing a spin at the center of each face, with neighbors defined by spins on adjacent faces across shared edges.

The Potts model finds applications across statistical physics, showing first- and second-order phase transitions, as well as BKT transitions, depending on J and q . It is used to study the random cluster model [121], percolation [122], and chromatic polynomials [123], with physical applications in quark confinement [124], foams [122], and biological morphogenesis [125]. While primarily studied on periodic lattices, some work on Penrose tilings suggests universal behavior similar to the periodic case [126, 127].

For $J > 0$, the model exhibits a ferromagnetic ground state with all spins aligned along one of their q directions. The behavior for $J < 0$, however, is more intricate. For $q = 2$, the model reduces to the Ising model, but since a two-coloring of the faces of AB is impossible (due to the presence of three-vertices), the antiferromagnetic ground state on AB becomes geometrically frustrated. This frustration arises because the connectivity of sites prevents spins from simultaneously minimizing their energies. In contrast, the existence of a face-three-coloring demonstrates that for $q \geq 3$, the $J < 0$ state is no longer frustrated. For $q = 3$, three-colorings define the possible ground states, while for $q \geq 4$, they represent a subset of the ground states.

The connection between the Potts model and protein folding has been extensively explored in polymer physics and computational biology [128, 129, 130]. The q -state Potts model, in particular, has been employed to




Figure 4.11. A fully 3-colored AB tiling, featuring tile 1 and tile 2 (both outlined in gold), with all ladders consistently colored as described in Sec. 4.3.5. (Following Ref. [3])

study the thermodynamics of protein folding [131]. In this framework, the lattice represents a solvent surrounding the protein, which is a chain of amino acids [132]. Upon dissolution, the protein may achieve a stable, typically compact configuration that facilitates its biological functions [133]. This folding process is governed by pairwise interactions between amino acids and the solvent [17, 131, 128]. If interactions between similar amino acids are attractive, they form a compact folded structure. If the interactions are repulsive, similar amino acids disperse across the solvent sites, producing a folded but less compact structure [129].

In this context, the Hamiltonian in Eq. 4.2 represents the interactions between amino acids $\{\sigma_i\}$. For $J < 0$, similar amino acids exhibit repulsive interactions in the solvent. When considering protein chains composed of only three types of amino acids (each represented by a color), the three-coloring of AB predicts a stable, minimum-energy folded configuration. In this configuration, each amino acid type occupies the center of a tile with its corresponding color.

It is reasonable to propose that modeling the folding of long protein chains could be improved by representing the solvent as a quasicrystalline surface instead of a simple periodic lattice. Quasicrystals offer an aperiodic surface with a wider range of bond angles, as opposed to the limited flexibility of regular lattices. For example, square and hexagonal lattices have only one bond angle each (90° and 120°) and a single vertex type. In contrast, the AB tiling provides three bond angles — 45° , 90° , and 135° — and a variety of local environments generated by seven vertex types. Additionally, the eightfold rotational symmetry of AB patches surpasses the maximum sixfold symmetry allowed by crystallographic restrictions, making AB tiling a closer approximation to the continuous rotational symmetry of real space.

4.3.7 Minimum dominating set problem




Figure 4.12. A portion of the AB tiling showing the solution to the minimum dominating set problem in red, which, as proved in Sec. 4.3.7, are also vertices of a single inflated AB tiling. (Following Ref. [3])

A dominating set D of a graph G is a subset of vertices such that every vertex not in D is adjacent to at least one vertex in D . The minimum dominating set (MDS) is the smallest such set, and its cardinality is called the domination number, $\gamma(G)$.

Problem Statement [49, 46] - Given a graph G and an integer k , is $\gamma(G) \leq k$? The corresponding optimization problem is to find the MDS.

Solution: If G is constructed from N copies of $L_{1/2}$ tiles, the answer is yes if and only if $N \leq k$.

Proof: Fig. 2.4 illustrates the single-inflation rules for the rhombus and square at level $L_{1/2}$. The set of red vertices forms an MDS for all vertices within each $L_{1/2}$ tile. Since these inflated tiles cover the entire tiling (due to discrete scale symmetry), the union of the minimum dominating sets for all inflated tiles constitutes a dominating set for any patch of the AB tiling or for the full tiling in the thermodynamic limit. The absence of redundancy in the placement of red vertices in Fig. 2.4, where no red vertex neighbors another, demonstrates that this set is indeed an MDS for the entire AB tiling. A larger region of the MDS is shown in Fig. 4.12. Note that this MDS is not an unique solution.

4.3.8 Induced Path Problem




Figure 4.13. FPLs on AB^* created using the method in Sec. 4.2.1 also forms induced paths on AB . The total region shown here is W_1 , with W_0 highlighted in grey. The corresponding induced paths are shown in gold. (Following Ref. [3])

An induced path in an undirected graph G is a sequence of vertices such that two vertices in the sequence are adjacent if and only if they share an edge in G . An induced cycle is an induced path that forms a closed loop.

Problem Statement [49, 46] - For a given graph G and a positive integer k , does G contain an induced cycle of length at least k ? The associated optimization problem is to find an induced cycle of length k .

Solution: Yes, if $G = AB$.

Proof: In this case, it is convenient to examine the D_8 -symmetric regions W_n . Fig. 4.13 illustrates the loops of the Fully Packed Loop (FPL) model on AB^* , as constructed in Sec. 4.2.1. We observe that any pair of nonadjacent vertices on these loops lacks an edge in G , while adjacent vertices are connected by edges. Therefore, these loops constitute induced cycles. Since any AB tiling contains W_∞ , the answer to the problem is ‘yes’ for any k , whether finite or infinite.

4.4 Conclusion

This chapter provides an exact solution to various NP-hard optimization problems for instances defined on the aperiodic AB tiling. It is possible that other similar problems can also be solved on AB tilings. Our solution to the Hamiltonian cycle problem leverages the long-range order inherent in the AB tiling, which arises from its discrete scale symmetry. This raises the question of whether we can apply the same approach to solve the Hamiltonian cycle problem on other quasicrystal tilings. For Hamiltonian cycles to exist on any graph, a necessary but not sufficient condition is the presence of a perfect matching. This is because if we delete every second edge of a Hamiltonian cycle, we are left with a perfect matching. The work in [10] demonstrates that a perfect matching exists on the AB tiling. However, my preliminary checks on other tilings—by examining whether the number of defects tends to zero in the thermodynamic limit—indicate that other tilings with 7-, 10-, 12-, or 16-fold symmetry cannot be perfectly matched. In these cases, the number of defects does not tend to zero as it does in the case of the AB tiling. This suggests that the AB tiling is indeed a unique and special tiling, even among its close cousins in the family of aperiodic tilings.

On the other hand, preliminary checks on the recently discovered ‘Hat’ and ‘Spectre’ monotilings suggest they also admit perfect dimer matchings, even when vertices are added to make the graphs bipartite. However, the structure of these matchings implies that a continuation to fully packed loops is not possible. All the examples so far consider the simplest case where the vertices and edges of the graph align with the corresponding tiling. Fig. 4.14 shows the modified trivalent penrose tiling (introduced in Sec. 2.2.1) forming fully packed loops (purple) by construction. In Fig. 4.14(b), if we augment the alternating cycle around each red cross, resulting in a red cycle that visits every vertex on adjacent Penrose tiles, forming a Hamiltonian cycle on the relevant subgraph. This construction mirrors our method used for AB tilings in Sec. 4.2.1 and 4.2.2: first create fully packed loops and then connect them into a single cycle.

We anticipate that these methods may generalize to a broader range of structures. Our Hamiltonian cycle construction demonstrated the existence of fully packed loops on arbitrarily large finite subgraphs of AB, where each vertex is visited by a loop, though not necessarily the same one [134]. The FPL model is crucial for understanding the ground states of frustrated magnetic materials like spin ices [135]. Loops played a significant role in Onsager’s exact solution to the Ising model on the square lattice, a major result in statistical physics [15, 16].

Hamiltonian cycles and loop models are often used to approximate protein folding by representing molecules on periodic lattices [17, 97]. This approach simplifies the problem, allowing statistical physics tools to simulate protein folding [17, 136, 137]. Quasicrystals may provide more accurate models for such studies, as they offer higher rotational symmetry than crystals, better approximating the continuous rotational symmetry of real space [51]. Quasicrystals also feature a wider range of nearest-neighbor bond angles and more complex structures with additional neighbors. Although they lack discrete translational symmetry—a characteristic of lattices that could limit their ability to approximate continuous space—they possess discrete scale invariance, which could help




Figure 4.14. (a) Penrose tiling shown by gray lines. The modified trivalent penrose tiling (black lines) decorated to admit fully packed loops (purple) by construction. (b) Augmenting the alternating cycle surrounding each red cross leads to a cycle (red) which visits every graph vertex on any Penrose tile touching a cross. (This figure was made by my supervisor.)

address these challenges. Random graphs share some advantages but lack the long-range order necessary for exact solutions.

In the context of NP -complete problems, exact solutions to Hamiltonian cycles are used as benchmarks for testing new algorithms [138, 139]. Even though NP -complete problems are generally difficult, many algorithms work well on a variety of graphs [138, 140, 141]. Identifying Hamiltonian cycle instances that are particularly challenging for these algorithms is important for real-world applications and can provide valuable insights into the broader NPC problem. A future direction is to explore how existing Hamiltonian cycle algorithms perform on AB regions, compared to the approach in our algorithm. Since the AB tiling combines elements of both regular and random graphs, it has unique structures not found in either. Additionally, our Hamiltonian cycle construction is nonlocal and takes advantage of the discrete scale symmetry of the tiling, which could make it a challenging test case. Understanding which aspects of AB make it difficult for solvers would be a valuable problem to tackle.

In conclusion, while AB tilings may initially appear unsuitable for exact results due to their lack of periodicity, their inherent long-range order makes it possible. While periodic lattices with discrete translational symmetry have been central to fields like condensed matter physics [110], Floquet theory [142], and lattice QCD [143, 144], the discrete scale invariance offered by quasicrystal lattices like AB also provides a unique and powerful simplification. This characteristic opens new avenues for solving complex problems and offers an exciting potential for future research.

Classical Dimer Model on Random Planar Graphs

The classical dimer model, as introduced in Section 3.1, has been extensively studied on various periodic lattices due to its rich physical and mathematical properties. Here, I summarize the major known results for this model on regular periodic lattices and discuss how I aim to generalize these findings to random graphs.

The combinatorial problem of determining the exact number of dimer coverings was solved in the early 1960s for planar two-dimensional (2D) lattices using Pfaffian techniques [67, 68, 27]. These methods have been extended to calculate dimer-dimer and monomer (i.e., unpaired sites)-monomer correlation functions [80]. Studying classical correlations also provides insights into the corresponding quantum problem, as at the Rokhsar-Kivelson (RK) point, the quantum ground state is an equal-amplitude superposition of all classical perfect matchings.

Some earlier works have considered an Hamiltonian that assigns different energies (a_i) to geometrically distinct bonds of a lattice rather than to flippable plaquettes. On the honeycomb lattice (Fig. 5.1(a)), if we adopt such a Hamiltonian and assign different weights (z_i) to the three geometrically distinct edges (e_i) — such that $z_i = \exp(-a_i/(\beta T))$ with $a_1 < a_2 < a_3$ — a Kasteleyn-type phase transition emerges. This transition is characterized by a highly asymmetric specific heat around the critical temperature T_c [145, 146, 147]. Specifically, the specific heat is zero everywhere below T_c , while above T_c , it varies as $C_v \propto |T - T_c|^{-1/2}$. Above T_c , correlations decay as a power law, indicating a deconfined phase at high temperatures.

A similar Hamiltonian on the square-octagon lattice (Fig. 5.1(b)), also known as the 4-8 lattice, does not exhibit a phase transition at any $T > 0$ when equal weights are assigned to both square and octagon edges. However, assigning different weights — z_{sq} for square edges and z_o for octagon edges, with $z_i = \exp(-a_i/(\beta T))$ and $a_o < a_{sq}$ — induces a phase transition similar to that observed in the 2D Ising model at $T > 0$. This transition is marked by a highly symmetric, logarithmically diverging specific heat, varying as $C_v \propto -\ln(T - T_c)$ [145]. In this case, correlations decay exponentially to zero at all temperatures except at the critical temperature T_c , where they follow a power law. Thus, the dimer model on the square-octagon lattice exhibits confined monomers at all temperatures (with $v = -1$) except at T_c , where monomers become deconfined.

In the square lattice, with aligning interactions given by Eq. 3.1 where $v = -1$, the system exhibits a high-temperature critical disordered phase that transitions to a low-temperature ordered columnar phase via




Figure 5.1. (a) Honeycomb lattice. (b) Square-octagon lattice.

a Kosterlitz-Thouless (KT) transition [26]. The critical phase at high temperatures have monomer-monomer correlations decaying as a power law. The decay exponent (α_d) varies with temperature, with an exact solution of $\alpha_d = 1/2$ at $T = \infty$. In contrast, the low-temperature ordered phase features confined monomers with exponentially decaying monomer correlations. This Hamiltonian with aligning interactions on flippable plaquettes, as defined in Eq. 3.1, is more widely studied compared to one based on weighted edges. Therefore, in my study, I employ the Hamiltonian given below:

$$Z = \sum_c \exp\left(-\frac{E_c}{T}\right),$$

and

$$E_c = v \cdot N_{fp}, \quad (5.1)$$

where N_{fp} denotes the total number of flippable plaquettes in a given perfect dimer configuration c . This is similar to Eq. 3.1.

Field theories associated with dimer models often rely on an underlying discrete structure, such as a crystal lattice in condensed matter physics or a high-energy cutoff in particle physics. This structure is typically assumed to be periodic for ease of calculation. However, realistic materials often exhibit amorphous structures that lack both translational and long-range order, resulting in a more random atomic arrangement. These materials are finite-sized and usually contain defects or impurities. Additionally, the only space-time discretizations compatible with Poincaré symmetry are random. Thus, dimer models on random lattices may more accurately describe the physics of such systems.

Exploring these generalized dimer models, rather than strictly periodic ones, and their associated field theories could be particularly compelling in high-energy physics. They offer a natural framework for discretizing quantum field theories. Unlike regular lattices, which inherently break Lorentz invariance and preclude fully Poincaré invariant models, aperiodic lattices may provide a way to overcome this limitation.

Quasicrystals, which also lack periodic lattice discretizations, occur naturally and represent cases that lie between periodic and random graphs — aperiodic graphs. While quasicrystals lack translational symmetry, they exhibit long-range order and possess discrete scale symmetry. Studying dimer models within the broader context of aperiodic or random lattices, such as those found in quasicrystals and amorphous materials, can yield insights into systems that deviate from traditional periodicity.

Building upon established results on periodic lattices, my research aims to generalize the understanding of dimer models to random and aperiodic graphs using the Hamiltonian defined in Eq. 3.1, as it is more widely studied. This involves exploring how the absence of regular periodic lattice structures influences phase transitions, correlation functions, and other critical phenomena inherent to dimer models.

Our analysis of these lattices is based on three fundamental assumptions: planarity, bipartiteness, and regularity. A graph is called k -regular if every vertex has exactly k edges incident to it. k -regularity is essential for ensuring the locality of the field theory across various lattices [148], while regularity, together with bipartiteness, is crucial for the existence of perfect dimer matchings on any graph as given by Hall’s Marriage Theorem [149].

As discussed in Section 3.4, bipartite k -regular planar graphs facilitate integer-valued height models that correspond directly to perfect dimer matchings.

5.1 Random Graphs

I will now explain the method to construct k -regular planar bipartite random graphs. The first step is to determine the maximum value of k that can be used to create such graphs. One of my collaborators on this work, Dr Peru d’Ornellas, proved that no k -regular bipartite graph can be planar for $k \geq 4$. To understand why, consider a graph with V vertices, E edges, and F faces, which must satisfy the following two relations:

$$2E = kV, \quad (5.2)$$

where the first relation comes from the k -regularity of the graph, and

$$E \geq 2F, \quad (5.3)$$

where the second relation arises from the fact that the minimum number of edges for each face in a bipartite graph is four.

Substituting these two conditions into Euler’s formula, $\chi = V - E + F$, leads to the following inequality:

$$(4 - k)V \geq 4\chi. \quad (5.4)$$

Since $V \geq 0$, a planar solution (which must have $\chi = 2$) is only possible for $k \leq 3$. Therefore, we focus on 3-regular planar bipartite (3RPB) random graphs for our numerical analysis of dimer models on them. Bipartite graphs with all vertices of degree 3 are called bicubic graphs. Thus, 3RPB graphs can also be referred to as planar bicubic graphs.

Currently, no algorithm exists for efficiently sampling random cubic bipartite graphs*, however, there are algorithms available for generating all such graphs. The challenge lies in the rapid growth of the number of these graphs with increasing node count, as documented in the OEIS sequence A008325 [150]. While the exact count of planar cubic graphs remains unknown, the number of trivalent 3-connected† bipartite planar graphs is recorded in OEIS sequence A007083 [151], which also exhibits rapid growth with system size. Hence, due to these computational limitations, we must rely on a few existing code packages that can generate all possible 3RPB random graphs for small vertex counts. For our study, we used the `plantri` package, which can efficiently

*Private communication with Prof. Brendan McKay.

†A graph is said to be 3-connected if it has more than 3 vertices and remains connected upon the removal of fewer than 3 vertices.




Figure 5.2. Modulus heights represented on a 3RPB random graph (generated by `plantri`). The purple edges form a random perfect dimer matching. (a) Shows the heights obtained using the rules in Sec. 3.4, while (b) demonstrates that taking the modulus-3 of these heights corresponds to a 3-coloring of the graph, with only three possible values $\{0, 1, 2\}$, each represented by a different color. Note that the exterior tile of the graph has a height of 2.

generate all such 3RPB random graphs with vertex counts up to 30 [152, 153, 154]. A sample of such a graph is shown in Fig. 5.2. While `plantri` is capable of generating graphs with more than 30 vertices, it requires significant computational power and time.

If we take a random perfect dimer matching on such graphs and assign heights according to the rules mentioned in Sec. 3.4, then take the modulus 3 (since the graphs are 3-regular) of these heights, the resulting modulus heights will be restricted to $\{0, 1, 2\}$, ensuring that no two adjacent plaquettes share the same mod height. By assigning a unique color to each modulus height value, this height representation effectively produces a 3-coloring of the graph. One of my collaborators on this work, Lizzy Rieth, made this observation. Fig. 5.2 illustrates this for a sample 3RPB random graph obtained using `plantri`. We can categorize the plaquettes into three distinct types based on the 3-coloring. Note that the exterior of the graph is also considered a tile. Although this imposes a strong constraint, it ensures that the graphs have no boundaries and are effectively defined on a sphere rather than a plane.

Since purely random 3RPB graphs are not feasible to generate for large sizes, we designed an algorithm to create large-scale 3RPB graphs. However, there is a caveat: one of the three plaquette types (considering the 3-coloring) will always be a square. The algorithm is outlined in Fig. 5.3 and proceeds as follows:

- Begin by generating a random walk for a given number of steps, then connect the start and end points with a straight line to form a closed loop. This random loop is shown in red in Fig. 5.3(a), with black dots indicating its individual step points.
- Treat the loop as an infinitesimally thin line. Each self-intersection point of this loop represents a vertex of degree 4. Any two such vertices are connected by a segment of the loop, forming an edge between them.




Figure 5.3. Schematic of the algorithm for constructing large 3RPB graphs. (a) A closed random walk loop is shown in red, with thick black dots representing its step points. The self-intersection points of the loop are marked by black and green circles. These points form the nodes of a 4-regular graph, with edges given by the loop segments connecting them. The green circled node has a self-loop; removing this node and connecting its adjacent nodes results in a new 4-regular graph, shown in black, with black circle nodes only. (b) A square is drawn around each node of the 4-regular graph, shown in blue, such that each corner of a square lies on the original walk. (c) The resulting 3-regular planar bipartite (3RPB) graph is obtained by removing all self-intersection points inside the squares and their associated edges. Connections between squares are simplified as straight lines by removing intermediate random walk step points.

Together, these vertices and edges define a 4-regular graph, shown in red in Fig. 5.3(a), with its nodes marked by black/green circles. Note that this graph may not be bipartite.

- Remove all nodes in this 4-regular graph that contain a self-loop, such as the node indicated by the green circle in Fig. 5.3(a). After removing each such node, connect its adjacent nodes directly. This again yields a 4-regular graph, shown in black in Fig. 5.3(a). This step is performed to avoid multiple edges between a pair of nodes in the final Step 5.
- Construct a square around each self-intersection such that each of its four vertices lies on the loop. These squares are shown in blue in Fig. 5.3(b).
- Remove the intersection points inside these squares along with their associated edges. This transforms the 4-regular graph into a bipartite 3-regular graph. The resulting graph is shown in Fig. 5.3(c), with simplified straight edge connections between the squares. At this step, the vertices of the 3-regular graph correspond to the corners of the introduced squares. Since a square is inherently bipartite, this guarantees that the resulting graph is also bipartite.

In the following sections, graphs generated using the `plantri` package will be referred to as *random graphs*, while large-scale 3RPB graphs generated by our algorithm will be referred to as *height gems*. The term reflects the fact that our construction corresponds to forming a Graph-Encoded Map (GEM) from twigless, self-loop-free random graphs with the same embedding [155], and that these resulting graphs admit a height model. For the height gems, when modulus 3 heights are assigned, all squares will always have the same value (either 0, 1, or 2), meaning that all squares will always be assigned the same color in a 3-coloring. This is because the pattern of modulus heights around a vertex with a given bipartite charge can be consistently repeated for every other vertex in the same bipartite class.




Figure 5.4. (a) A height gem generated using our random walk algorithm in Fig. 5.3. (b) Trivalent Penrose: Thick black lines represent the modified Penrose tiling, where all vertices have degree 3, derived from the underlying standard Penrose tiling (thin gray lines). (c) A purely random 3RPB graph generated using the `plantri` package. Note that all these graphs are represented using the Tutte planar embedding.

5.2 Modified (Trivalent) Penrose Tiling

We also examine a class of graphs that lies between periodic and random structures — namely, aperiodic graphs. These graphs lack the translational symmetry of periodic lattices but are not entirely random; they exhibit discrete scale symmetry and thus retain long-range order. However, a key challenge with aperiodic tilings is that their vertices typically have varying degrees, which poses a problem for constructing a height field model due to its locality constraints. To address this, we developed a modified version of the Penrose tiling where all vertices have degree 3 while preserving essential properties such as bipartiteness, fivefold symmetry, and aperiodicity. This ensures that the long-range order stays while maintaining a uniform vertex degree. The construction of the trivalent Penrose tiling follows these steps: 1. Insert a square onto each face of the Penrose tiling. 2. For two adjacent faces sharing an edge, introduce an edge between the nearest vertices of their corresponding squares.

This transformation results in a 3RPB-modified Penrose graph, structurally similar to our height gem construction (see Fig. 5.4). Like height gems, this tiling also admits a three-coloring. When modulus-3 heights are assigned, all squares will again take the same value (either 0, 1, or 2), ensuring that all squares are assigned the same color in a three-coloring. In the rest of the text, I will refer to this modified tiling as the Penrose graph.

For all three graph types — height gem, Penrose, and random — we numerically calculated the Tutte embedding using an algorithm developed by my collaborator, Dr Peru d’Ornellas [156], to obtain a planar representation of the graphs as shown in Fig. 5.4.

5.3 Numerical methods: Worm algorithm

I used Monte Carlo simulations to numerically study the classical dimer model and correlation functions at finite temperatures across the different categories of graphs defined above. I also performed exact calculations of correlation functions at $T = \infty$ using the Fisher–Kasteleyn–Temperley (FKT) algorithm [67, 27]. The FKT algorithm efficiently counts the number of perfect matchings in planar graphs by transforming the problem into the computation of the Pfaffian — a specialized determinant — of a skew-symmetric matrix derived from the planar graph embedding. This Pfaffian can then be evaluated efficiently using standard determinant algorithms.

I used a Monte Carlo (MC) directed loop or ‘worm’ algorithm which allows for non-local updates by flipping dimers along an alternating closed loop (or worm) that can span a significant portion of the lattice at once [26]. This method enables us to efficiently sample the configurations of large systems at any given v and T . Furthermore, the algorithm naturally incorporates test monomers, which allows for the computation of monomer-monomer correlation functions, which will be discussed in later sections. The algorithm’s performance is strongly influenced by the physical properties of the test monomers in various phases.

For clarity, I briefly outline the algorithm, with modifications specific to our implementation. I run the algorithm for multiple sweeps until the energy converges at any given v and T . Each MC sweep consists of multiple worms such that all the edges of the lattice are, on average, visited at least once by a worm. The algorithm proceeds in the following three steps:

1. The ‘worm’, composed of two monomers (head and tail), is initially placed at a random site $i = i_0$ in the current dimer configuration.
2. The site i is connected to a neighboring site j by a dimer in the background configuration, denoted as (i, j) . The dimer (i, j) is removed, and the head of the worm moves to site j , leaving j without a dimer. One of the three neighboring sites of j (since the lattice is 3-regular), say site k , is selected based on a local detailed balance rule (explained below). A new dimer is then placed between sites j and k .
3. If $k = i_0$, the worm is completed, resulting in a new valid dimer configuration. Otherwise, we rename $i = k$ and return to step 2.

To select site k in step 2, we assign weights (based on Eq. 3.1) to all empty edges (e_0, e_1, e_2) at site j , such that $W_{e_i} = \exp(-vN_{e_i}/T)$, where N_{e_i} represents the number of plaquettes that would become flippable ($N_{e_i} \in \{0, 1, 2\}$) if a dimer were placed on edge e_i . Let edge e_2 be the reverse direction (j, i) . We then calculate the probability for each empty edge at j as $P_{e_i} = \frac{W_{e_i}}{\sum_{e_i} W_{e_i}}$. A random real number $n \in (0, 1]$ (uniformly distributed) is generated, and to select the edge (j, k) , we place the dimer on edge e_0 if $n < P_{e_0}$, on e_1 if $n < P_{e_0} + P_{e_1}$, and on e_2 (the reverse direction) if $n < P_{e_0} + P_{e_1} + P_{e_2} = 1$.

5.4 Results

For the Hamiltonian defined in Eq. 5.1, I first analyzed how the number of flippable plaquettes (NFP) varies with temperature at a fixed negative value of $v = -1$. I analysed results only for negative v , since in the positive v region the worm algorithm struggles to reach the true staggered state when lowering T , as these states lie in a different topological sector. For any given (v, T) , I initialized the system in a random perfect dimer state and executed the Monte Carlo (MC) algorithm described earlier for a number of sweeps to ensure energy convergence. Each MC sweep comprised multiple worm updates, ensuring that every edge of the lattice was visited at least once. The MC simulations were run for 1000 sweeps, which was sufficient for the energy to converge. After achieving convergence, I computed the expectation values of the ratio of the number of flippable plaquettes to the total number of plaquettes, $\langle NFP \rangle$, the energy cumulant, EC , and the specific heat per site, C_v/N , with the detailed definitions of these quantities provided later, by averaging over their values obtained from different ground-state samples at a given (v, T) . Here, different ground-state samples refer to the distinct perfect dimer matchings on a lattice that minimize the energy (as defined in Eq. 5.1) for any given (v, T) . The data and plots presented below are for graph sizes as follows: Height gem with $N = 884$ vertices, Penrose with $N = 1720$ vertices, and random with $N = 36$ vertices. For comparison, I have also included data for the periodic square-octagon lattice with $N = 1152$ sites.




Figure 5.5. Expectation value of the ratio of flippable plaquettes to total plaquettes, $\langle NFP \rangle$, as a function of temperature T for different lattices at $v = -1$. For this case, the system evolves from a disordered state with intermediate flippability at high temperature to a maximally flippable state as $T \rightarrow 0$. The dashed gray line indicates the analytical prediction for the maximum number of flippable plaquettes in the height gem, random, and Penrose lattices as $T \rightarrow 0$ at $v = -1$.

Fig. 5.5 presents the plot of $\langle NFP \rangle$ as a function of T for all three graph types at $v = -1$. The plot indicates that, as $T \rightarrow 0$, the system reaches a highly flippable ordered state. In fact, this corresponds to the maximum flippable state possible for these graphs. We can determine the maximum number of flippable plaquettes for each graph type as follows: The 3-coloring of these graphs naturally divides them into three distinct plaquette types. In the thermodynamic limit, these three tile types appear with equal frequency. However, for a finite lattice, there is typically some imbalance, meaning that two of the three tile types will have the highest count in total. A maximally flippable state can then be constructed by placing dimers only along edges that connect the two majority tile types.

This generic rule holds true for the height gem and random graphs, but the Penrose graph, which has discrete scale symmetry, follows a different rule. For the Penrose tiling, our Monte Carlo simulations instead converged to a ground state where all squares and hexagons are flippable at $(v = -1, T \rightarrow 0)$. This behavior is due to two key reasons: 1) In the original Penrose tiling, the number of 3-degree vertices is always greater than the total number of vertices of any other degree [4]. So, in the modified trivalent Penrose tiling, the number of hexagons will always exceed the total number of other faces of any size. 2) Each hexagon has two possible flippable states. Choosing one of these states will make all squares flippable, while choosing the other will prevent any of the squares from being flippable. Therefore, by combining these two arguments, it becomes clear why a configuration with all squares and hexagons flippable has the lowest energy at $(v = -1, T \rightarrow 0)$.

Thus, by numerically computing the modulus-3 heights and counting the number of plaquettes belonging

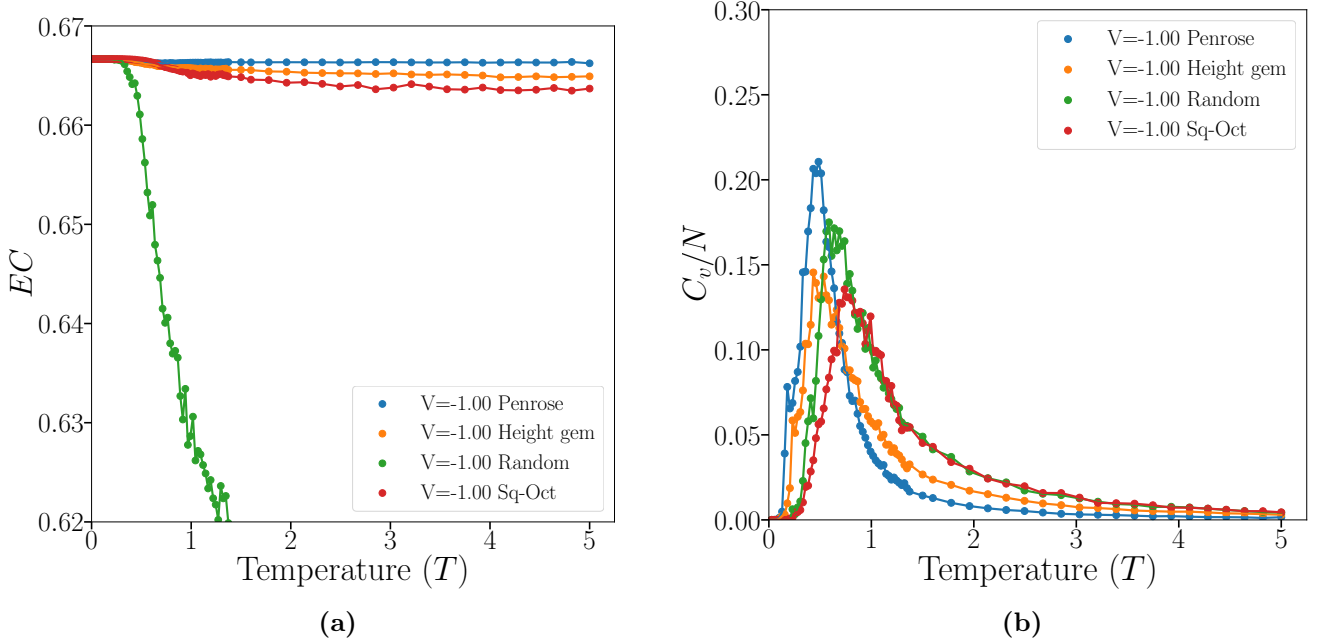


Figure 5.6. (a) Energy cumulant (EC) as a function of temperature at $v = -1$. (b) Specific heat per site (C_v/N) as a function of temperature at $v = -1$.

to the two majority tile types, we obtain the maximum possible number of flippable plaquettes for the height gem and random cases. For the Penrose case, I calculated the total number of squares and hexagons. This theoretical maximum is indicated by the gray dashed lines in Fig. 5.5 for the three graphs, and hence confirms that the system is converging to the maximum flippability ordered state at $(v = -1, T \rightarrow 0)$. This behavior is analogous to the square lattice reaching a maximally flippable columnar state as $T \rightarrow 0$.

As the temperature increases, thermal fluctuations gradually disrupt the ordered configurations, causing the system to transition into a disordered regime with no long-range order. In this regime, the flippability stabilizes at approximately 50% for $v = -1$.

However, the $\langle NFP \rangle$ plot does not clearly indicate whether there is a true phase transition. To investigate further, I compute additional quantities, focusing on the $v < 0$ regime and considering the case $v = -1$.

Energy Cumulant, $EC = 1 - \frac{\langle E^4 \rangle}{3\langle E^2 \rangle^2}$. In both the ordered and disordered phases, EC converges to a value of $2/3$ in the thermodynamic limit. This is also true for a critical point of a second-order phase transition, though for finite systems, it shows a slight dip at the critical temperature T_c . In the case of a first-order phase transition, EC exhibits a non-trivial minimum different from $2/3$ in the thermodynamic limit [157].

Fig. 5.6(a) shows that across all lattice types, the value of EC converges to $2/3$ over the entire temperature range. This suggests that there is no first-order phase transition at any $T \geq 0$. Note that, in the case of random graphs, EC shows a drop due to the effects of the small system size.

Specific heat per site, $\frac{C_v}{N} = \frac{\langle E^2 \rangle - \langle E \rangle^2}{NT^2}$. For a second-order phase transition, C_v/N is expected to either diverge or exhibit a cusp at the critical temperature T_c . However, as shown in Fig. 5.6(b), specific heat does not exhibit any divergence or cusp for any of the graphs at any $T > 0$. Instead, a broad (pronounced) peak

is observed in the specific heat plot which might correspond to a crossover from disordered states to ordered columnar states (maximal flippability) as $T \rightarrow 0$. Hence, the specific heat C_v is featureless at T_c , which might indicate a Kosterlitz–Thouless (KT) transition. However, in a KT transition the system typically evolves from a high-temperature phase with power-law correlations to a low-temperature phase with exponentially decaying correlations. Therefore, to confirm that our case indeed corresponds to a KT transition, further tests are required.

The specific-heat behavior we observe is consistent with that of the periodic square-octagon lattice under the same conditions, as seen in Figs. 5.5 and 5.6. And it is also similar to the square lattice classical dimer model with the same Hamiltonian, which exhibits a high-temperature critical disordered phase that transitions into a low-temperature ordered columnar phase. In Ref. [25], the specific heat was also found to be featureless, hinting at a Kosterlitz–Thouless (KT) transition. This was confirmed by observing power-law behavior in monomer correlations at high temperature, consistent with the presence of a critical phase.

Monomer-Monomer Correlations: I also examined the two-point correlation function for monomers, which are sites unpaired by a dimer. Although monomers do not appear naturally in the model, their behavior can be studied by introducing two test monomers into a dimer-covered background and computing the monomer-monomer correlation function [158]:

$$M(r) = \langle M(x)M(0) \rangle - \langle M(x) \rangle \langle M(0) \rangle \quad (5.5)$$

where $M(x)M(0)$ represents the number of valid dimer configurations with monomers fixed at x and $(0, 0)$, while $M(x)$ corresponds to the number of configurations with a single monomer at x . The parameter $r = |\vec{x} - \vec{0}|$ denotes the real-space distance between the two monomers. On a bipartite lattice, $M(x) = 0$ when only one monomer is present, or if both monomers belong to the same sublattice, since the remaining lattice cannot be perfectly matched. As a result, $M(x, 0) = \langle M(x)M(0) \rangle$, which focuses exclusively on configurations where the two monomers are positioned on opposite sublattices.

In Monte Carlo (MC) simulations, snapshots taken during intermediate steps of worm construction reveal two test monomers, allowing the worm algorithm to compute $M(r)$ at any (v, T) without explicitly introducing monomers into the model. After convergence, $M(r)$ is obtained as a histogram of monomer separation distances over multiple MC sweeps [159]. The behavior of the worms provides insight into monomer confinement: Confined monomers - Short worms, updating a small number of dimers. Deconfined monomers - Long worms, updating a large number of dimers.

If $M(r)$ decays exponentially or faster to zero as $r \rightarrow \infty$, the system is in a confined phase, where monomers cannot be separated to infinite distance without an increasing energy penalty. If $M(r)$ decays to zero as $r \rightarrow \infty$ not exponentially but as a power law, the system is said to be in a critical phase. Conversely, if $M(r)$ approaches a nonzero constant as $r \rightarrow \infty$, the system is in a deconfined phase, where monomers can be separated to infinite distance at a finite energy cost.

At infinite temperature ($T = \infty$), the partition function simplifies to an unweighted sum over all possible dimer coverings. At this limit, $M(r)$ can be exactly computed using the FKT algorithm, which has been used to derive the exact result $M(r) \propto r^{-1/2}$ for the square lattice [80]. This power-law dependence confirms a critical phase for the square-lattice dimer model at $T = \infty$. The FKT algorithm efficiently counts all perfect matchings in planar graphs via Pfaffian techniques. I developed an implementation of the FKT algorithm for arbitrary planar lattices, which is also used in the next chapter for numerical calculations.

Fig. 5.7(a) presents the exact calculation of $M(r)$ as a function of real-space distance $r = |\vec{x} - \vec{x}_0|$ at

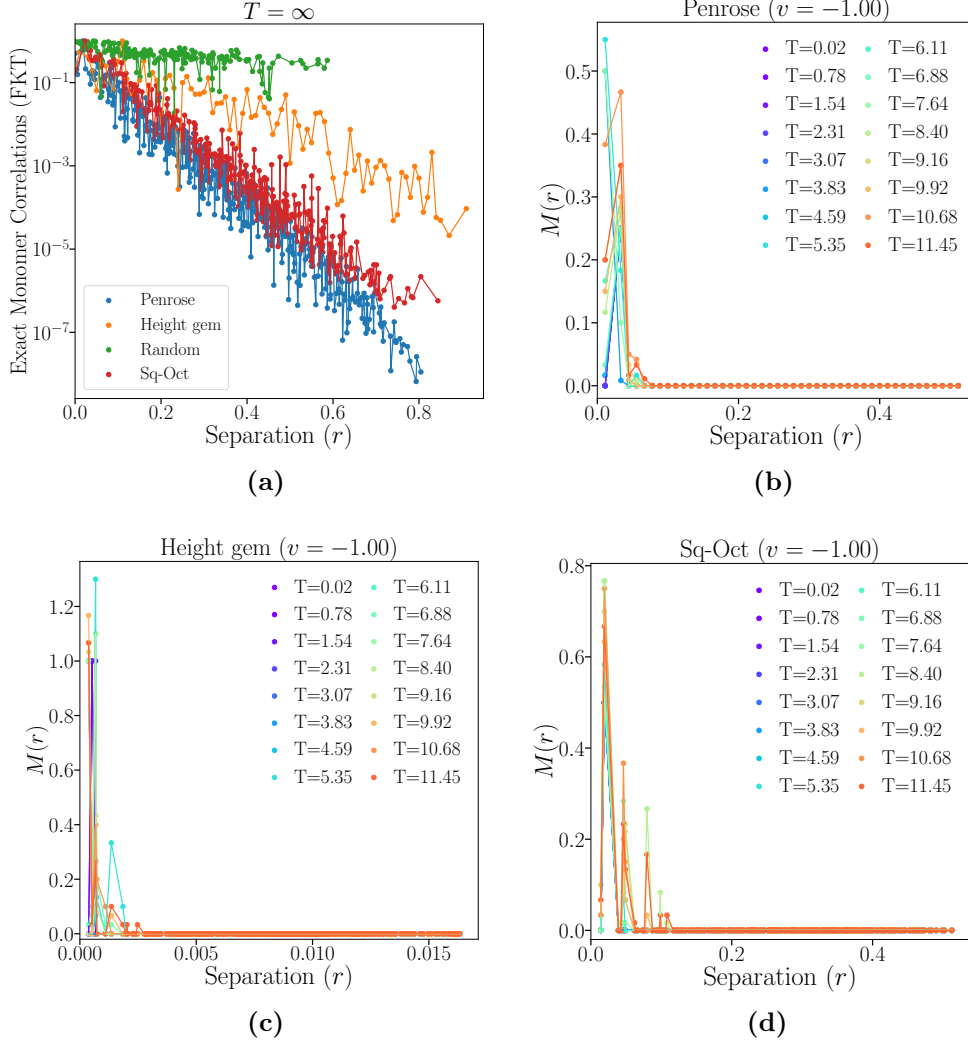


Figure 5.7. Monomer-monomer correlations, $M(r)$, as a function of monomer separation r at $v = -1$. (a) Exact calculations of $M(r)$ at $T = \infty$ using the FKT algorithm, with a logarithmic y -axis, showing a straight-line behavior indicative of exponential decay, $\log(y) = -mx$ where $m > 0$. (b, c, d) Rapid decay of $M(r)$ at finite temperatures (T) for the Penrose, height gem, and square-octagon lattices, respectively, indicating confinement at all $T > 0$ for each lattice type. The separation distance (r) on the x-axis is normalized according to the Tutte planar embedding, which maps the entire lattice within a disk of radius 0.5 centered at $(0.5, 0.5)$, ensuring consistent comparison across different graph structures.

$T \rightarrow \infty$. Here, one monomer is fixed at position \vec{x}_0 , while the other monomer is placed at different positions x on the complementary bipartite set. The total number of perfect dimer matchings for each monomer pair (x, x_0) is computed using the FKT algorithm. These values from the FKT results for a given distance r are then averaged over different monomer pairs to obtain the corresponding $M(r)$. The resulting plot exhibits a linear trend on a log-linear scale, indicating an exponential decay of monomer-monomer correlations at $T = \infty$, with

$$\log M(r) = -mr, \quad m > 0. \quad (5.6)$$

This result implies a confined phase at $T = \infty, v = -1$, in contrast to the square lattice, where correlations exhibit power-law decay [26], indicating a critical phase at high temperatures [80].

At finite temperatures, as mentioned before, the locally detailed balance of the worm algorithm ensures that the histogram of the distance r between the worm's head and tail is proportional (with a small correction) to the monomer-monomer correlation function $M(r)$ [159]. Fig. 5.7(b,c,d) shows that $M(r)$ decays rapidly to zero at all $T > 0$ for the Penrose, height gem, and square-octagon lattices, respectively. This indicates a confined phase throughout all $T > 0, v = -1$, again in contrast with the square lattice. Due to the limited size of random graphs, I was unable to plot $M(r)$ as a function of separation r for this case.

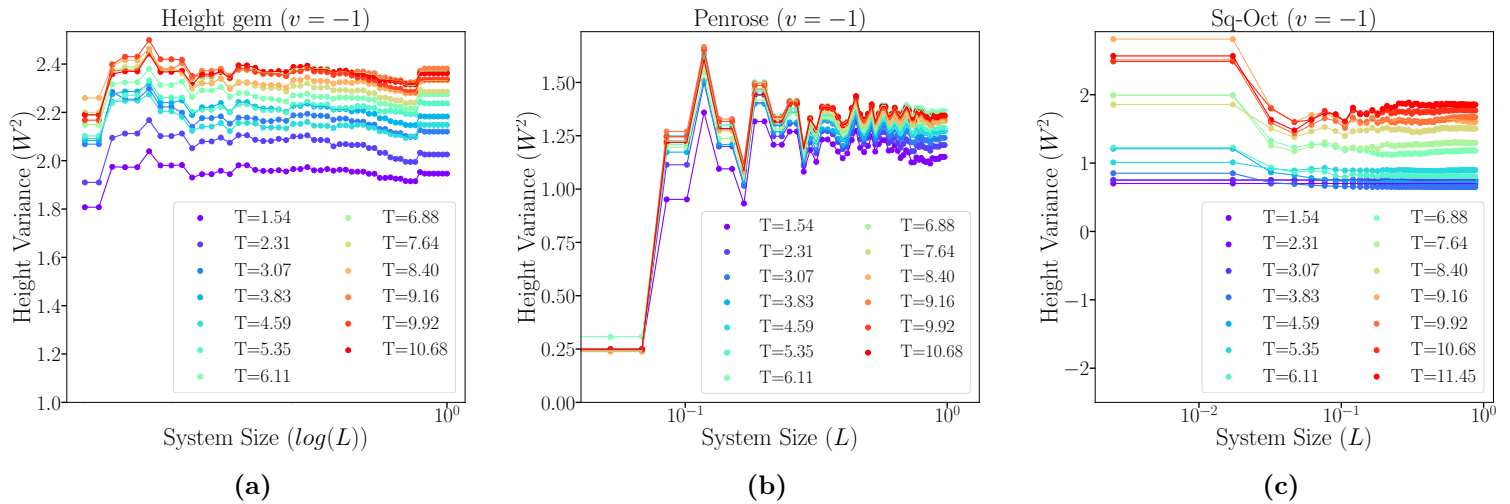


Figure 5.8. Height variance (W^2) as a function of linear system size at finite temperatures with $v = -1$ for the (a) height gem, (b) Penrose, and (c) square-octagon cases. The x-axis is log-scaled and normalized according to the Tutte planar embedding. The results show that W^2 converges to a finite value in the thermodynamic limit, confirming a smooth confining phase.

Height Variance: To further confirm the confinement at all $T > 0$, I also calculated the height variance (W^2) in real space for each plaquette, as given in Eq. 3.7. To determine whether it log-diverges (in the thermodynamic limit) with the linear system size as in Eq. 3.8 — indicating a rough (deconfined) phase — or remains finite, indicating a smooth (confined) phase, I plotted W^2 vs. $\ln(L)$, where L is the linear system size.

For a large lattice system, I constructed concentric circles of increasing radius from the center to the outer boundary. Each enclosed region effectively represents a lattice of increasing size, with the linear system size L corresponding to the radius of each region. I then averaged the variance of all plaquettes within each region to obtain W^2 for a given system size L .

Fig. 5.8(a,b) shows the plot of W^2 versus $\ln(L)$, revealing that the height variance remains finite as the

system size increases for all $T > 0, v = -1$. This result is consistent with the monomer-monomer correlation findings, confirming a smooth (confined) phase at all $T > 0$. Due to the limited size of random graphs, we could not plot the height variance against system size for this case.

5.5 Conclusion

I have numerically studied the classical dimer model on the generalized lattices considered — specifically, height gems, random graphs, and aperiodic planar cubic graphs — to investigate how these systems behave as a function of temperature when $v < 0$, with the aligning interaction Hamiltonian given in Eq. 5.1. There is a crossover from a high-temperature disordered regime to an ordered regime near $T = 0$, as shown in Fig. 5.5. The specific heat and energy cumulant results indicate that there might be a KT phase transition due to C_v/N being featureless and EC approaching the value $2/3$. But the system remains confined at all $T > 0$ from a field-theoretic perspective. This is because we observe that monomer-monomer correlations decay rapidly to zero and the height variance remains finite in the thermodynamic limit. So, there seems to be no critical phase at high T , which is not typical of a KT transition. Hence, more work is required to confirm a KT phase transition. We also see similar behavior in the square-octagon lattice as well (as shown in the figures). So the square-octagon lattice might represent a general case for the classical dimer model with aligning interactions. The analogous classical dimer model on the square lattice is said to have a KT phase transition because not only is the specific heat featureless and the energy cumulant value tends to $2/3$, but the finite-temperature phases exhibit power-law decays of correlation functions with exponents that vary continuously with temperature [26]. Hence, the square lattice exhibits a critical or algebraically confined phase at high temperatures. By contrast, in our cases of height gems, random graphs, and aperiodic lattices, we observe exponentially confined phases at both low and high temperatures. Also note that, in our current construction of height gems, one type of plaquette is necessarily a square. If we can reduce the number of such squares to zero and construct fully random 3-regular bipartite planar graphs, it will provide the most generalized perspective on this class of classical dimer models. In such a scenario, we may also observe a high-temperature critical deconfined phase, characterized by power-law correlations decaying to a finite non-zero value. This will form the next step of this project.

Dimer models on the Spectre tiling

The classical dimer model, as introduced in Chapter 3, examines all possible perfect dimer matchings on a planar lattice. The exact analytical solution for an N -vertex planar dimer model involves finding all perfect matchings through the partition function:

$$\mathcal{Z}_N[w] = \sum_{\mathcal{M}_i \in \mathcal{M}} \prod_{e \in \mathcal{M}_i} w(e), \quad (6.1)$$

where \mathcal{M}_i is a perfect matching in the set of all perfect matchings \mathcal{M} , and e represents the edges of the graph. When the weight $w = 1$ is assigned to all edges, $\mathcal{Z}[1]$ counts the total number of perfect matchings. From \mathcal{Z} , thermodynamic functions such as the free energy per dimer can be derived [79] for an N -vertex graph:

$$f_N[w] = \frac{1}{N/2} \ln(\mathcal{Z}_N[w]). \quad (6.2)$$

Determining the total number of perfect matchings for graphs of arbitrary size is computation-ally challenging. The Fisher-Kasteleyn-Temperley (FKT) algorithm provides an efficient numerical method for counting the total number of perfect matchings in planar graphs [67, 27, 68, 78]. While FKT is polynomial in theory, it becomes impractical for graphs with millions of nodes due to significant time and memory demands. Even for regular periodic graphs, the analytical solution to the dimer model has been achieved only in limited cases [28, 78, 79, 71].

In contrast, aperiodic tilings, which encapsulate the symmetries of physical quasicrystals, introduce an even greater degree of frustration in dimer arrangements due to their irregular vertex degrees. This suggests that solving the dimer model on aperiodically ordered planar tilings is significantly more complex than on periodic ones. However, the unique long-range order and discrete scale invariance of these tilings, as discussed in Chapter 2, often lead to surprising simplifications, enabling analytical results even in cases where periodic tilings remain unsolved.

In this chapter, I will present a mathematical proof for the exact analytical solution of the classical dimer model on the recently discovered aperiodic monotile known as the spectre tiling (see Section 2.3.2). I will also compute the free energy per dimer numerically in the thermodynamic limit using the FKT algorithm and compare these results with analytical calculations. Building on these classical results, I will also present exact analytical solution for the quantum dimer model on the Spectre tiling.

6.1 Exact Solution to the Classical Dimer Model

A single spectre at the zeroth level of inflation (S_0) contains a total of 14 vertices, including the gold vertex shown in Fig. 6.1(a). This gold vertex is included even when it is not implied by the edge connections, such as on the outermost boundary. This ensures that the spectre tiling is bipartite and that every spectre has an even number of vertices — a necessary condition for perfect matching since each dimer connects a pair of vertices.

Recall from Sec. 2.3.2 that two green tiles together form the “Mystic,” denoted as M_0 . The dark green tile, referred to as the Upper Mystic, is denoted as M_0^+ , where 0 represents the zeroth level of inflation (no inflation). Similarly, M_1^+ represents the first inflation of M_0^+ , and so on.

In a Mystic M_0 , out of the four internal edges, exactly one of the central two edges must be covered with a dimer in any perfect matching (shown by the dark red and dark blue edges in Fig. 6.1(a)). The dark red edge enforces the presence of pink and purple dimers, while the dark blue edge enforces light blue and purple dimers. The purple dimers are fixed regardless of whether a dimer is placed on the dark red or dark blue edge.

So a M_0^+ tile can only have two distinct types of perfect matchings, as illustrated in Fig. 6.1(a) with red and blue colors. According to the inflation or substitution rules for the Spectre tiling discussed in Sec. 2.3.2, Fig. 6.1(a) shows that at the first inflation level, S_1 and M_1 tiles inherit a fixed set of purple dimers on their internal edges enforced by either chosen matching of M_0^+ . At this stage, the only degrees of freedom exist in the two-fold matching choice on M_0^+ and along the boundaries of S_1 and M_1 . Notably, only one special purple dimer in either S_1 or M_1 touches the boundary, appearing in M_1 (highlighted with a green circle in Fig. 6.1(a)).

By applying the inflation rules again, the S_1 and M_1 tiles combine to form S_2 , as shown in Fig. 6.1(b). Here, pink tiles represent M_1 , while the remaining tiles belong to different S_1 regions. The boundaries between these regions are marked by thick gray paths, with the special dimer highlighted in gold. At this stage, S_2 is perfectly matched except along the thick gray boundary. The gold dimer enforces the positions of nearby green dimers, which in turn constrain all other green dimers in Fig. 6.1(b). Consequently, all purple and green dimers — covering the internal edges of S_2 outside M_0^+ tiles — are fixed and perfectly matched.

This behavior generalizes to all S_N regions. Reference [6] enumerates all possible configurations where S_1 and M_1 regions meet. The boundaries of S_1 and M_1 contain only bivalent or trivalent vertices. Specifically, the bottom vertex of the Mystic (the rightmost of the two lowest vertices of M_0 in Fig. 6.1(a)) connects exactly once to an internal dimer of S_1 or M_1 . It does so at a trivalent vertex and touches the (gold) boundary dimer of M_1 . In all other cases, it connects to the boundary of S_1 or M_1 . A single (gold) dimer forces the absence of dimers on the two other edges. Therefore, a single forced dimer along the network of S_1 and M_1 boundaries constrains all remaining dimers, except on the outermost boundary of the entire connected tile set, S_N .

The path of green dimers always connects to the outermost boundary via a degree-3 vertex such that the internal edge associated with this vertex does not contain a dimer. This also introduces a two-fold degree of freedom on the outermost boundary. Thus, for any finite patch of the spectre tiling, the twofold degree of dimer freedom exists only on M_0^+ tiles and the outermost boundary. Accordingly, the total number of dimer matchings is given by:

$$\mathcal{Z}_N[1] = 2^{N_{\text{Mystic}} + 1}, \quad (6.3)$$

where N_{Mystic} is the number of Mystic tiles M_0 (equal to the number of Upper Mystic tiles M_0^+).




Figure 6.1. (a) The dark and light green tiles, collectively called the Mystic (M_0), with the dark green tile referred to as the Upper Mystic (M_0^+). This shows the first inflation of a spectre S_0 and Mystic M_0 . The Upper Mystic can have only two types of perfect matchings, shown in red and blue. Choosing either red or blue matching enforces the same set of purple dimers, thus fixing the purple dimers. The internal edges of S_1 and M_1 tiles allow only two dimer matchings per Mystic. (b) The twice-inflated spectre region, S_2 . Mystic regions are highlighted in pink, with the special gold dimer (arrow) enforcing all green dimers. The two-fold dimer freedom is preserved on the Upper Mystic tiles and the outermost boundary, validating Eq. (6.3). Source:[11]

Using Eq. 6.2 and Eq. 6.3, the free energy per dimer for an S_N region can be calculated as:

$$f_N [1] = \frac{\ln (\mathcal{Z}_N [1])}{N/2} = \frac{N_{\text{Mystic}} + 1}{N/2} \cdot \ln (2) . \quad (6.4)$$

In the thermodynamic limit $S_{N \rightarrow \infty}$, the free energy per dimer can be calculated analytically using spectre inflation rules and series expansion methods. These analytical calculations were performed by my supervisor and are detailed in our work here, Ref.[11]*. The final analytical result is given below:

$$f_{\lim_{N \rightarrow \infty}} [1] = \frac{1}{5 + \sqrt{15}} \cdot \frac{\ln (2)}{3} \approx 0.02604. \quad (6.5)$$

6.2 Numerical Results

I calculated the free energy per dimer numerically for finite patches S_2 to S_6 using the Fisher-Kasteleyn-Temperley (FKT) algorithm, which I implemented from scratch. The FKT algorithm counts all possible perfect matchings in a planar graph in polynomial time, a task that is NP -complete for general graphs [67, 27, 68, 78]. The algorithm achieves this by transforming the problem into a Pfaffian computation of a skew-symmetric matrix derived from a planar embedding of the graph. The Pfaffian is efficiently computed using standard determinant

*See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevB.109.L220303> for details of the analytical calculations.

algorithms. Below are the key steps of the FKT algorithm:

- Take a planar graph G .
- Compute a spanning tree T_1 of the input graph G . A spanning tree T is a connected, acyclic subgraph that includes all vertices of G .
- Assign an arbitrary orientation to each edge in G that is also in T_1 .
- Create an undirected graph T_2 with the same vertex set as the dual graph of G .
- Create an edge in T_2 between two vertices if their corresponding faces in G share an edge that is not in T_1 . Note that T_2 is also a tree.
- For each leaf V in T_2 (excluding the root):
 - Let e be the lone edge of G in the face corresponding to V that does not yet have an orientation.
 - Orient e such that the number of edges oriented clockwise is odd.
 - Remove V from T_2 .
- Return the absolute value of the Pfaffian of the adjacency matrix of G , which is the square root of its determinant.

The returned Pfaffian provides the value for $\mathcal{Z}_N[1]$. Using these $\mathcal{Z}_N[1]$ values from the FKT algorithm, I calculated f_N for regions S_2 to S_6 . The results are shown in Fig. 6.2.

Calculating f_N for larger regions involves computing the determinant of increasingly large matrices, which becomes computationally prohibitive once the number of nodes reaches the order of a million.

In Fig. 6.2, the raw data demonstrate the convergence of the numerically calculated f_N values to the analytical result for $f_{N \rightarrow \infty}$ derived in the previous section. The dashed line represents the analytical value, while square points correspond to numerical results for S_2 through S_6 . It is evident that f_N converges exceptionally slowly as N increases. This slow convergence arises from the fractal nature of the boundary in the S_N regions. As N grows, the boundary becomes increasingly intricate, resulting in a disproportionately large number of vertices residing on the boundary. Since boundary vertices contribute differently to the free energy compared to bulk vertices, this imbalance slows the approach of f_N to its thermodynamic limit $f_{N \rightarrow \infty}$.

Addressing this challenge requires advanced numerical techniques, such as series acceleration methods or extrapolation strategies, to account for the boundary effects and improve the accuracy of f_N for finite N . We employed these methods, and the calculations are detailed in our work, Ref. [11]. However, I am not elaborating on them here, as they were primarily carried out by my supervisor.

The free energy per dimer in the spectre tiling (≈ 0.02604) is significantly smaller than the corresponding values observed in periodic lattices [18], such as the square lattice (0.583), honeycomb lattice (0.323), triangular lattice (0.857), and Kagome lattice (0.462). This aligns with the observation that most bulk dimers, apart from those on M_0^+ tiles, are entirely constrained.

In periodic lattices, the higher free energy per dimer reflects a greater number of possible configurations due to the relatively uniform distribution of constraints across the lattice. In contrast, the spectre tiling's hierarchical and quasiperiodic structure introduces strong constraints on dimer arrangements, significantly reducing the




Figure 6.2. The free energy per dimer f_N for spectre patches containing N vertices, shown as purple points, corresponds to S_2 through S_6 spectre regions. These data points are calculated numerically using the FKT algorithm. The dashed line represents the analytical result from Eq. 6.5, valid as $N \rightarrow \infty$. The slow convergence is due to the fractal nature of the boundary.

number of allowable configurations. These constraints are particularly pronounced for dimers outside M_0^+ tiles, where the freedom of placement is almost entirely suppressed, leading to the low free energy per dimer.

6.3 Exact Solution to the Quantum Dimer Model

In Section 3.2 and existing literature [20], it is noted that the ground states for Quantum Dimer Models (QDM) have not been exactly solved for any two-dimensional bipartite lattice across the entire parameter space. Even in simple cases like the square lattice, the phase diagram region between the columnar and plaquette phases remains unknown. However, these studies primarily focus on periodic lattices. In this section, we aim to extend these studies to aperiodic cases, we were surprised to find that we can exactly solve the QDM on the newly discovered spectre monotiling and determine its ground and excited states in a surprisingly simple way.

Recall the Hamiltonian for the Quantum Dimer Model (QDM) on the basic square lattice, as mentioned in Eq. 3.2. The summation is over all the plaquettes, and the kinetic term describes flipping between two possible dimer states on a single square plaquette. There is also a potential term corresponding to each of these dimer states.

To write a similar Hamiltonian for the QDM on the spectre tiling (S_N) with N nodes, we sum over all the spectre plaquettes ($S_{0,i}$), where ‘0’ denotes level 0 of the tiling with no inflation. For each spectre plaquette ($S_{0,i}$), there are only two possible dimer states, as shown in red and blue in Fig. 6.1. These will be the two quantum states, $|r_i\rangle$ and $|b_i\rangle$, for quantum dimers on any $S_{0,i}$. So, the QDM Hamiltonian on S_N will be:

$$\hat{H} = \sum_{S_{0,i} \in S_N} -J (|r_i\rangle\langle b_i| + |b_i\rangle\langle r_i|) + v (|r_i\rangle\langle r_i| + |b_i\rangle\langle b_i|) \quad (6.6)$$

where J and v are real, and J is positive. Results from the previous chapter count the total number of perfect matchings possible on an S_N region and also tell us that any possible such perfect dimer matchings on any S_N region will have all dimers fixed except the two degrees of freedom on special dark green upper mystic tiles (M_0^+). This means that the only flippable plaquettes, which enact ‘flips’ $|r\rangle \leftrightarrow |b\rangle$, possible in the spectre tiling are the upper mystic M_0^+ tiles. Hence, the QDM Hamiltonian for the spectre tiling will be modified as below:

$$\hat{H} = \sum_{M_{0,i}^+ \in S_N} -J (|r_i\rangle\langle b_i| + |b_i\rangle\langle r_i|) + v (|r_i\rangle\langle r_i| + |b_i\rangle\langle b_i|) \quad (6.7)$$




Figure 6.3. Monomer Deconfinement on Spectre Tiling: If a dimer is removed, creating two monomers (represented by red circles), one monomer will always follow the path of green dimers along the boundary of the S_1/M_1 tiles and move toward the outermost boundary. The other monomer, however, will become trapped on an upper mystic dark green tile. These possible paths for the two red monomers are shown in yellow—one reaching the outer boundary and the other confined to an upper mystic tile. Similarly any pair of monomers can be separated to infinite distances at zero energy cost, as their movement does not require crossing any M_0^+ tiles during separation. (Following Ref. [11])

with the summation over only flippable M_0^+ plaquettes in a given S_N region. Since all these M_0^+ are non-intersecting, the problem of finding ground states of QDM on the entire spectre tiling actually decouples into the problem of finding the ground state of QDM on each of these M_0^+ independently. For each M_0^+ tile, there will be two energy eigenstates possible as below:

$$|\pm_i\rangle = (|r_i\rangle \pm |b_i\rangle) / \sqrt{2} \quad (6.8)$$

with corresponding energies $v \mp J$. The ground state of Eq. 6.7 for the whole S_N spectre tiling is therefore $|+_i\rangle$ states on all M_0^+ , with the total ground state energy being $(v - J)N_{\text{Mystic}}$:

$$\hat{H} \prod_{M_{0,i}^+} |+_i\rangle = (v - J)N_{\text{Mystic}} \prod_{M_{0,i}^+} |+_i\rangle \quad (6.9)$$

All excited states can be formed by swapping individual $|+_i\rangle$ for $|-_i\rangle$ at a cost of $2J$ energy per swap. Note that in the square lattice, since all plaquettes are flippable and coupled with one another, the sign of v/J either attempts to maximize or minimize the number of flippable plaquettes. But in the spectre tiling, all the flippable plaquettes are decoupled and are fixed in number.

Additionally, I did some preliminary numerical checks on the Hat monotiling which involved overlaying two random perfect matchings of the hat tiling and removing only the common dimers. The remaining dimers were located only on reflected hat tiles (forming a loop). This suggests that the hat tiling may exhibit behavior similar to the Spectre tiling in terms of classical and quantum dimer model solutions, with a two-fold degree of freedom on certain special decoupled tiles, specifically the mirror (reflected) hat tiles shown in dark blue in Fig. 2.10.




Figure 6.4. Possible Escape Routes for a Monomer: All possible paths (gold) by which a monomer (black), created on any upper mystic tile, can escape to the boundary of the S_1/M_1 tiles — and ultimately to the outermost boundary — by following the green dimer paths shown in Fig. 6.3 [11]. (This figure was made by my supervisor.)

6.3.1 Monomer Deconfinement

An intriguing result we have found regarding the ground and excited states of the Quantum Dimer Model (QDM) on the spectre tiling is that any pair of monomers created will be deconfined for all values of J and v . I will now explain how this works and why it is surprising.

The spectre tiling has an interesting feature: in any perfect dimer matching, if we remove a dimer to create two monomers, using alternating paths, one monomer will always get stuck on an upper mystic tile, while the other monomer can always hop to the outermost boundary using alternating paths of green dimers, as shown in Fig. 6.3.

For example, in Fig. 6.4, removing a dimer on an upper mystic tile creates a black and a white monomer. If you try moving the white monomer around using alternating dimers, you will find that it remains confined to the upper mystic tile. In contrast, the black monomer can travel to the boundary of the S_1 (or M_1) region using different types of alternating paths shown in gold. Once this black monomer reaches the S_1/M_1 boundary, which consists of green dimers, it can always follow the green dimer path further to reach the outermost boundary for any S_N region. Notably, this green dimer path does not cross (or flip) any upper mystic tile along the way. Hence moving the black monomer to the outermost boundary preserves the number of flippable plaquettes and does not require any extra energy because only upper mystic tiles contribute to the energy for any S_N region. This is possible for any dimer in any S_N region in the thermodynamic limit.

Hence, this shows that any pair of monomers created on any state of the QDM on the spectre tiling can be separated to an infinite distance with zero energy cost at any v and J . Therefore, there is a deconfined phase over the entire parameter space.

This is surprising because Polyakov’s result states that deconfined phases cannot exist in compact (2+1)-dimensional U(1) gauge theories [160]. In general, QDMs on bipartite graphs map to compact (matter-free) quantum electrodynamics [74, 82]. Hence, QDMs on 2D bipartite lattices cannot have an extended deconfined

phase. This aligns with QDM results on previously studied planar bipartite graphs, such as the square lattice, where the Rokhsar-Kivelson (RK) point $J = v$ constitutes a deconfined quantum critical point (not a phase) between ordered phases, as discussed in Sec. 3.2. Surprisingly, the spectre aperiodic tiling seems to contradict this, as it appears to have a deconfined phase over the entire parameter space.

We think that this happens because there seems to be no obvious mapping of the spectre dimer model to a compact $U(1)$ gauge theory, since the vertices in the spectre tiling connect to variable numbers of edges. As explained in earlier chapters, a necessary condition to obtain a compact $U(1)$ gauge theory like a height model is that the lattice be regular.

Additionally, noting that the eigenstates are constant across v/J , it would be reasonable to suggest that the line v/J is a simple extension of the deconfined RK point by interactions that do not affect the physics, rather than a true deconfined phase. For example, setting the RK model on a graph consisting of disconnected squares would lead to similar phenomenology. However, there is a testable difference between this artificially simple scenario and ours. Since the simplification in our results came about naturally, moving off the line v/J by tuning some other physical parameter should lead to different behaviours at different v/J . If so, v/J would represent a limiting behavior within a physically meaningful larger parameter space. In this sense we would argue that it constitutes a true deconfined phase. One such larger space to consider is that of a softened dimer constraint, allowing monomer creation (equivalently, weakening the spin interactions from which the RK model emerges in the infinite-coupling limit). A full quantum treatment is beyond the present study, but it is sufficient to demonstrate that the behaviour in the classical limits $v/J \rightarrow \pm\infty$ must differ. To see this, allow monomer pairs to be created with energy $2\Delta \ll v$. For $v/J \rightarrow \infty$ the flippable antispectres are energetically unfavourable, so the ground state will have monomers on these plaquettes to stop them flipping. For $v/J \rightarrow -\infty$ flippability is favoured, and so the ground state will be monomer free. Hence, while our eigenstates are constant along v/J , they are nevertheless measurably distinct when factoring in the surrounding parameter space.

Another difference from previous QDM studies is that all previously known bipartite RK points were characterized by algebraic dimer correlations, which means power-law decay, making the RK point a critical point [20]. However, in the spectre tiling, dimer correlations are expected to be completely uncorrelated and decay to zero immediately due to decoupled M_0^+ tiles spread in a field of fixed dimers.

While no material has yet been identified with the symmetries of the spectre tiling, it seems plausible that such a structure could be manufactured or discovered in the future. At its core, the spectre tiling represents a novel symmetry of nature, inherently deserving of study. This is particularly compelling given that this work shows how it can offer unexpected simplifications in physical models, enabling exact results that are unattainable in periodic or other aperiodic tilings.

The tiling's ability to simplify physical models raises intriguing questions about whether a deeper, underlying property of its structure is responsible for this behavior.

Future research could investigate whether the spectre tiling's unique characteristics extend to other physical phenomena, such as electronic, optical, or vibrational properties, or if its structural features inspire new paradigms in designing materials with tailored functionalities.

The Quantum dimer model on the AB tiling

In classical physics, common examples of long-range magnetic order include ferromagnetic and antiferromagnetic arrangements. However, both theoretical and experimental studies, like Refs. [161, 74, 162, 163], have shown a wide variety of long-range orders influenced by quantum effects. Hence, the study of frustrated quantum magnetism has become a rapidly evolving area within condensed matter physics. Traditionally, researchers have explored quantum magnetism using simplified spin models, such as the Ising or Heisenberg models.

In 1988, Rokhsar and Kivelson introduced a model of interacting hard-core dimers to describe the resonating-valence bond (RVB) phase of spin-1/2 systems, particularly in the context of high-temperature cuprate superconductors. In this framework, hard-core dimers represent $SU(2)$ invariant spin singlets [20]. This model is known as the quantum dimer model (QDM), as introduced in Sec.3.2. A distinction between spin models and the QDM lies in the definition of their basic components. While spin degrees of freedom are located at the lattice sites, dimer degrees of freedom are associated with the bonds between these sites. Also, there is a strict local constraint that only one dimer can emerge from a given site. As a result, the associated Hilbert spaces consist of the complete set of perfectly packed dimer configurations. One benefit of these constraints is that the Hilbert space grows more slowly with the number of sites N compared to 2^N for a spin-1/2 system of size N . For example, in the square lattice dimer model, the Hilbert space grows asymptotically as $(1.3385\dots)^N$ [67, 27].

Beyond the QDM, other constrained models have gained attention over the past decades. For instance, the hard-boson model introduced by Fendley et al. [164] and models of interacting anyons [165] are notable examples. These constrained models are valuable as they serve as simplified representations for various physical problems. By focusing on a particular constrained basis, one can concentrate on the essential properties of the low-energy sector, effectively disregarding less significant higher-energy effects.

The QDM has been studied on various two-dimensional bipartite periodic lattices, leading to a generally agreed-upon phase diagram, as detailed in the Sec.3.2.1. Additionally, such strongly correlated models have been used to create topologically protected qubit systems, where an isolated two-fold degenerate quantum dimer liquid ground state can be used in the construction of topologically stable qubits protected from decoherence by exploiting topological stability [166]. The QDM has also been studied on quasi-one-dimensional systems like the square ladder, revealing a rich phase diagram with a rung-singlet, a leg-singlet, and a period-2 phase, as well

as an Ising transition between the first two phases. Various other works have explored the QDM on different lattices [167].

In the previous chapter, I found the exact analytical solution for the ground state of the quantum dimer model on the Spectre tiling. This was possible due to the specific graph connectivity of the Spectre tiling, which simplified the entire problem by allowing it to decouple into disconnected upper-mystic tiles. This kind of special property is usually absent in other tilings, making it hard to find an exact ground state in most cases. This challenge is central to modern quantum many-body physics. Strongly correlated quantum systems on low-dimensional lattices continue to attract significant interest among physicists. However, both analytically and numerically, these systems are difficult to study. Exact analytical solutions exist only in rare cases, such as those solvable via the Bethe ansatz in one dimension [168].

Numerical methods commonly used in these cases include Exact Diagonalization (ED), which fails for large system sizes due to the exponential growth of the Hilbert space, and Quantum Monte Carlo (QMC). In addition, variational numerical algorithms are also employed to study the ground states of strongly correlated many-body systems in low dimensions. One such algorithm is the Density Matrix Renormalization Group (DMRG), which has become the most powerful and widely used technique for studying one-dimensional quantum lattices [12, 167].

In the DMRG framework, the ground states of any 1D quantum system can be simulated using a chain of tensors called Matrix Product States (MPS). For the simplest case — a 1D Ising chain with a quantum spin at each site — the Hilbert space dimension grows exponentially as $d^L = 2^L$, where $d = 2$ is the local Hilbert space dimension of a single spin. Finding the ground state in such a vast space is like searching for a needle in a haystack. However, it has been proven that, at least for local Hamiltonians with a gap between the ground state and the first excited state, the effective search space is infinitesimally small compared to the full Hilbert space [169, 170]. This relevant corner of the Hilbert space can be efficiently parameterized using numerical techniques to solve fundamental questions in quantum physics. This parameterization is precisely what Matrix Product States (MPS) provide [12].

Our work aims to extend existing numerical techniques, particularly Density Matrix Renormalization Group (DMRG), to study the ground state phases of the Quantum Dimer Model (QDM) on large-scale two-dimensional (2D) systems. We focus on a specific example of a quasicrystal, which lacks periodicity but exhibits discrete scale invariance. In this chapter, I present results on the ground state phase diagram of the QDM defined on the two-dimensional AB* tiling — an AB tiling with all 8_0 -vertices removed, as introduced in Chapter 4. These results are obtained for a region containing approximately 10^6 sites, using Matrix Product States (MPS), the DMRG algorithm, and leveraging a unique structural property of the tiling that arises from the interplay between dimer constraints and aperiodicity.

7.1 Studying QDM on the AB* tiling using MPS-DMRG

DMRG has demonstrated remarkable accuracy in finding the ground states of various 1D systems. However, it is not as successful for 2D or higher-dimensional systems because the computational resources required grow exponentially with system size, making large 2D lattices inaccessible. The different behavior of DMRG in one and two dimensions is closely related to the scaling of quantum entanglement in these systems, as dictated by the so-called area laws [12, 169]. If a system is divided into two subsystems, A and B, the entanglement entropy of A is given by:

$$S_A = -\text{tr}(\rho_A \log \rho_A), \quad (7.1)$$

where $\rho_A = \text{Tr}_B(\rho)$ is the reduced density matrix of subsystem A , obtained by tracing out the degrees of freedom of subsystem B from the full system density matrix ρ . For a random state in the Hilbert space of any Hamiltonian, entanglement entropy follows the volume law:

$$S_A \propto L^d \quad (7.2)$$

where d is the dimension of the system. However, ground states of local gapped Hamiltonians are less entangled than random states in the Hilbert space [169], as they follow the area law of entanglement entropy:

$$S_A \propto L^{d-1} \quad (7.3)$$

rather than the volume law. For locally gapped systems, which include the RK Hamiltonian [169], we have: in 1D ($d = 1$): $S_A = \text{const}$ and in 2D ($d = 2$): $S_A \propto L$. This explains why DMRG works well for large system sizes in 1D, as entanglement does not grow. However, in 2D, entanglement increases with linear system size, and the computational resources required to capture this growth become overwhelming, leading to the failure of DMRG for large 2D systems.

The area law also explains why MPS can be used to model 1D quantum spin chains or ladders. Since MPS is essentially a chain of tensors, if we break it at any point to create two subsystems A and B , the entanglement between them depends only on the bond at the break-point, not on the entire length of the chain. Hence, $S_A = \text{const}$ is an in-built property of a MPS tensor network. Similarly, in the DMRG framework, the Hamiltonian can also be written as a chain of tensors, known as Matrix Product Operators (MPO).

As discussed, DMRG has been highly successful in studying 1D quantum many-body systems efficiently and accurately [12, 167, 171]. However, it has not been as effective for 2D or higher-dimensional systems on a large scale due to increasing entanglement. The AB* tiling is constructed by removing the degree-8 vertices and their connecting edges from the AB tiling, as shown in Fig. 7.1. It is important to note that the AB* tiling still remains a connected 2D graph. So studying systems like the quantum dimer model (QDM) on the AB* tiling in the thermodynamic limit should not be feasible using DMRG. This is the primary challenge addressed in this chapter. Again, the symmetries and graph connectivity of the AB* tiling provide a crucial simplification. From results on the classical dimer model on the AB* tiling, as discussed in Sec. V of Ref. [10], we see that any possible ground state will not have dimers on the gray edges in Fig. 7.1. The 8-star loops, shown as purple edges, have only two trivial choices of dimer placements. The only non-trivial dimer phases appear on the 1D ladders, shown with blue or yellow tiles.

In the quantum dimer model, ground states are essentially superpositions of ground states from the classical dimer model with Hamiltonian given in Eq. 3.1 (or as aligning interactions in Sec.V of Ref. [10]). Therefore, the only non-trivial phases for QDM on AB* will also appear on the 1D ladders that constitute the AB* tiling. The rest of the tiling behaves trivially, as in the classical case. This allows us to simplify the entire QDM on the 2D AB* tiling into a QDM on decoupled 1D ladders. Therefore, to determine the phase diagram of the ground state for the entire 2D tiling, we only need to study the QDM on all the individual ladders contained within the tiling, which is computationally feasible using the MPS-DMRG algorithm.

The quantum dimer model has a hard-core constraint, meaning dimers cannot overlap. The quantum many-body systems studied with DMRG so far either lack such constraints, or these constraints have not been explicitly encoded in the DMRG algorithm. Instead, they are typically included through an additional term in the Hamiltonian that energetically penalizes states that do not satisfy the local constraints [165, 172, 173].




Figure 7.1. A patch of AB* Tiling, which is an AB tiling with all 8-vertices and their connected edges removed. The ground states for the classical dimer model on AB*, with Hamiltonian given in Eq. 3.1 (or defined as aligning interactions in Sec.V of Ref. [10]), will never hold a dimer on gray edges. Consequently, for both the classical and quantum dimer models, the AB* tiling can be decomposed into independent loops and ladders, represented by thick purple and black edges, respectively. The purple edges form 8-star loops, which allow only two trivial dimer configurations. The only remaining portion of the tiling with nontrivial phases consists of ladders made up of yellow and blue tiles, outlined here with thick black edges.

This approach does not take advantage of the fact that the Hilbert space of such constrained models grows more slowly with system size. To leverage this property, these constraints need to be explicitly encoded in the MPS-DMRG algorithm. This was first achieved in Ref. [167], where the ground-state phase diagram of the QDM on a 1D square ladder was studied. In that case, the Hilbert space grew as 1.6^N instead of 2^N . They demonstrated how to implement hard dimer constraints directly into this variational technique. In this work, I will employ similar techniques to study the QDM on AB ladders.

7.1.1 Methods: MPS-DMRG on AB ladders

In this section, I will explain how I extended the method of applying DMRG with quantum dimer constraints, used in Ref. [167], from square ladders to AB ladders. Any ladder on the AB tiling consists of three components shown in Fig. 7.2: rungs (vertical edges) of the square plaquettes, a rhombus with four outward edges connected to it (set of blue tiles), and a double-rhombus structure with four outward edges connected to it (set of yellow tiles). I have labeled these as s , A , and B respectively.

For example, the smallest ladder, shown entirely in blue in Fig. 7.1, can be written using these labels as $sAsAsAsAsAsAsAsA$ with periodic boundary conditions. These ladders exhibit long-range order, meaning the




Figure 7.2. These three components are the fundamental building blocks that make up ladders of any size in an AB tiling. The black circles show the nodes included in each of the components. For instance, the largest ladder in Fig. 7.1 can be expressed in terms of these components as $[..sAssBssAssBssAssBssAssBs..]$, with periodic boundary.

structure of any higher-order ladder of any size can be predicted from the smallest ladder. Let's denote the smallest ladder as order Λ_0 , with higher-order ladders subsequently labeled as $\{\Lambda_1, \Lambda_2, \Lambda_3, \dots\}$. This long-range order arises from the discrete scale invariance of the AB tiling. Consequently, these ladders follow a set of inflation rules:

$$s \rightarrow s, \quad A \rightarrow AssBs, \quad ssB \rightarrow AsAsAsAsAssB \quad (7.4)$$

as described in Ref. [10]. A ladder of any order (Λ_n) can be generated by applying these rules iteratively to the smallest ladder n times.

$$\begin{aligned} \Lambda_0: & \quad sAsAsAsAsAsAsA \\ & \quad \downarrow \\ \Lambda_1: & \quad sAssBssAssBssAssBssAssBssAssBssAssBssAssBssAssBs \end{aligned}$$

This matches the second ladder in Fig. 7.1.

Any quantum dimer state on an AB ladder will be a superposition of dimer basis states possible on the three ladder components (s, A, B). All possible basis states on each of these components are shown in Fig. 7.3. The basis states for the 's' component simply represent the presence and absence of a rung dimer, respectively. For components 'A' and 'B', the basis states include all possible dimer configurations that cover every node within them.

MPS Construction: Any arbitrary pure quantum state of a 1D lattice with N sites and d -dimensional local state spaces $\{\sigma_n; \quad 1 \leq n \leq N, \quad n \in \mathbb{Z}\}$ can be exactly represented in the matrix product state (MPS) form through an iterative process involving a sequence of Singular value decompositions (SVDs):

$$|\Psi\rangle = \sum_{\sigma_1 \dots \sigma_N} c_{\sigma_1 \dots \sigma_N} |\sigma_1 \dots \sigma_N\rangle = \sum_{a_1, \dots, a_N} L_{a_1}^{\sigma_1} L_{a_1, a_2}^{\sigma_2} \dots L_{a_{N-2}, a_{N-1}}^{\sigma_{N-1}} L_{a_{N-1}, a_N}^{\sigma_N} |\sigma_1 \dots \sigma_N\rangle. \quad (7.5)$$

The detailed steps for deriving Eq. 7.5 are provided in Ref. [174]. The iterative construction of an exact MPS representation of an arbitrary quantum state through a sequence of SVDs is illustrated in Fig. 7.4. On every site n in the bulk, L^{σ_n} is a three-legged left-normalized tensor, as shown graphically in Fig. 7.5. It has one physical leg, denoted by σ_n , representing the local state space of dimension d . For Ising spins, $d = 2$ (up or down spins), while for the quantum dimer model, $d = 2$ (presence or absence of a dimer). The other two legs in L^{σ_n} are auxiliary legs, whose dimensions match their neighboring tensors.




Figure 7.3. All possible basis states for dimer configurations on each of the constituent ladder components. Green and purple edges represent rung and leg dimers, respectively. The reason behind assigning specific edges on rhombus tiles as rungs or legs will be explained later in the text. The matrix representation of the basis states for component ‘s’ is given by $h_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $h_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. Similarly, each basis state in $\{r_i^A; i = 1, 2, 3, 4\}$ and $\{r_i^B; i = 1, 2, 3, 4, 5\}$ is represented as a column vector of size 4×1 and 5×1 , respectively, with the $(i - 1)$ -th entry equal to 1 and all other entries set to 0.

The MPS state can be left- (right-) canonical depending on whether the tensors L_n (R_n) at each site n satisfy the following conditions:

$$\sum_{\sigma_n} L^{\sigma_n} L^{\sigma_n \dagger} = I, \quad (\text{Left canonical form}) \quad (7.6)$$

$$\sum_{\sigma_n} R^{\sigma_n} R^{\sigma_n \dagger} = I, \quad (\text{Right canonical form}), \quad (7.7)$$

where on every site n in the bulk, R^{σ_n} is a three-legged right-normalized tensor.

For any AB ladder, the tensors L^{σ_n} are placed on components (s, A, B) , with the dimension of the physical leg, $d = 2, 4$, and 5 , respectively corresponding to the number of local basis states for each component. In DMRG, the entire system is divided into left and right environments with a central part, usually containing two sites. These three parts are contracted to form a network known as the effective Hamiltonian and the states violating the quantum dimer model (QDM) constraints must be removed. The effective Hamiltonian is then diagonalized in the reduced basis states.

Since the quantum dimer constraints are purely local, Ref. [167] proposed selecting the valid basis states at each level of the algorithm. For the MPS, I only retain states allowed by the QDM constraints, which gives the onsite tensors L^{σ_n} a block-diagonal form. This approach eliminates the need to construct the full tensor for the unconstrained Hilbert space, instead directly selecting blocks of allowed states and performing SVD decompositions separately for each block. This significantly reduces computational costs, allowing access to larger system sizes for ladders.

To implement the quantum dimer constraints and obtain the block-diagonal form of the MPS tensors,

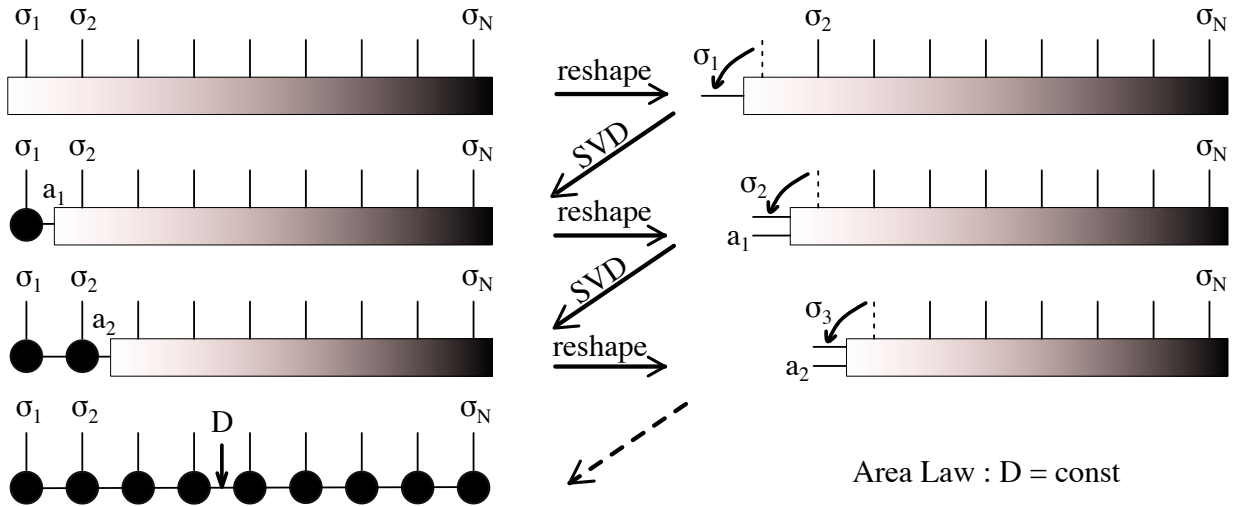


Figure 7.4. Illustration of the step-by-step process used to construct an exact Matrix Product State (MPS) representation of any quantum state, utilizing a sequence of singular value decompositions.




Figure 7.5. Schematic representation of L - tensors at the boundaries and within the bulk of MPS chains: the left diagram illustrates $L_{1,a_1}^{\sigma_1}$, the row vector at the left boundary; the right diagram shows $L_{a_{N-1},1}^{\sigma_N}$, the column vector at the right boundary. The central diagram depicts $L_{a_{l-1},a_l}^{\sigma_l}$ for the bulk. The σ_n are the physical legs, while the a_n are the auxiliary legs of the MPS tensors.

binary labels ‘0’ and ‘1’ are introduced to distinguish two sectors of the Hilbert space. Consider the construction of the left-canonical MPS on an AB ladder. At any bond in the ladder, if the ladder is split into two subsystems, the left subsystem is labeled ‘0’ if all its nodes are covered by dimers and ‘1’ if its end nodes are uncovered. These labels define the fusion rules for constructing the block-diagonal form of the MPS, starting from the left end of the ladder. The fusion rules and their shorthand notations are detailed in Fig. 7.6. These can be straightforwardly generalized to the right-canonical form.

These fusion rules essentially select blocks of allowed states in MPS tensors. Starting from the left, for a given L^{σ_n} , the labels 0 and 1 of the left subsystems are fed in and combine with the basis states, producing the same labels as output, which are then fed into the next site tensor, as illustrated in Fig. 7.6. A similar process can be repeated from the right end of any ladder to obtain the right-canonical MPS state.

I will now explain how the RK-Hamiltonian is modified for AB ladders and present the corresponding Matrix Product Operators (MPO). In the case of a simple square ladder, there are two types of edges on a square plaquette: vertical edges, called ‘rungs’, and horizontal edges, called ‘legs’. However, since AB ladders contain rhombuses in components ‘A’ and ‘B’, we cannot define rungs and legs as the usual vertical and horizontal dimers. Instead, I defined the edges for each of the three components, as shown in Fig. 7.7. Thick black edges represent rungs, while thin edges represent legs. This choice follows the intuition that, in rhombuses, lower two edges will be rungs since they comes in the path of the usual two-leg ladder, while remaining edges should be legs.

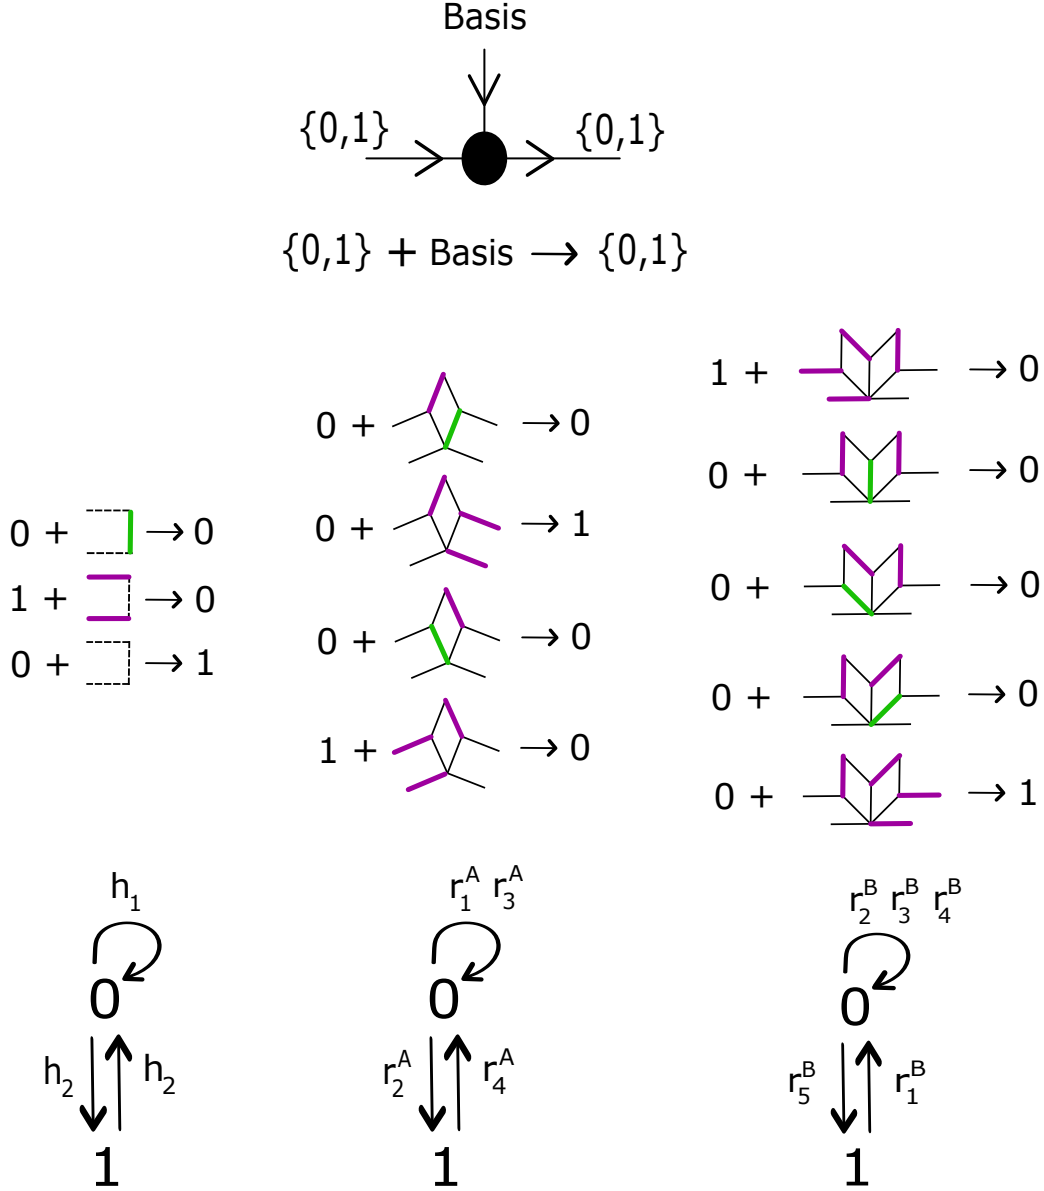


Figure 7.6. Fusion rules and dimer constraints: The *top row* illustrates how the left environment is constructed during a left-to-right MPS sweep. Incoming labels $\{0,1\}$ at each site (black circles) encode the left subsystem state and fuse with the onsite basis—depending on whether the site is of type s , A , or B —to produce outgoing labels. These define the updated subsystem state and are passed recursively to the next site. A similar process builds the right environment. The *middle row* shows all valid fusion rules for any ladder, detailing allowed combinations of incoming labels and onsite basis states. The *bottom row* presents their shorthand notation.




Figure 7.7. The thick and thin black edges represent the assigned rungs and legs, respectively, in the three fundamental ladder components. The rationale behind this assignment is that edges forming a standard square-ladder equivalent are designated as rungs, while all remaining edges are classified as legs.

Notably, in square ladders, a valid plaquette basis state can only have either all rung dimers or all leg dimers. However, in AB ladders, as seen in Fig. 7.3, some ‘A’ and ‘B’ component basis states contain both rung and leg dimers. This is evident in the basis states $|r_1^A\rangle, |r_3^A\rangle, |r_2^B\rangle, |r_3^B\rangle$, and $|r_4^B\rangle$.

As explained earlier, in the MPS representation of AB ladders, tensors are placed for components s , A , and ‘B’. The modified RK-Hamiltonian for AB ladders is given by:

For sites with component ‘A’:

$$H = H_A + H_{QDM} \quad (7.8)$$

$$H_A = \sum_i -J(|r_1\rangle\langle r_3| + |r_3\rangle\langle r_1|) + v_l(|r_2\rangle\langle r_2| + |r_4\rangle\langle r_4|) + v_{lr}(|r_1\rangle\langle r_1| + |r_3\rangle\langle r_3|)$$

For sites with component ‘B’:

$$H = H_B + H_{QDM} \quad (7.9)$$

$$H_B = \sum_i -J(|r_2\rangle\langle r_3| + |r_3\rangle\langle r_2| + |r_2\rangle\langle r_4| + |r_4\rangle\langle r_2|) + v_l(|r_1\rangle\langle r_1| + |r_5\rangle\langle r_5|)$$

$$+ v_{lr}(|r_3\rangle\langle r_3| + |r_2\rangle\langle r_2| + |r_4\rangle\langle r_4|)$$

For sites with component ‘s’:

$$H = H_{QDM} \quad (7.10)$$

Here, $J > 0$, v_l , and v_{lr} are real numbers, and H_{QDM} is given in Eq. 3.2. The terms H_A and H_B represent the onsite Hamiltonian terms for ‘A’ and ‘B’ components of a ladder. H_A contains a kinetic term with J , allowing flipping between local basis states $r_1^A \leftrightarrow r_3^A$. It also includes a new potential term with coefficient v_{lr} for $\{r_1^A, r_3^A\}$, since these basis states contain a dimer on both rungs and legs. Additionally, a full-leg potential term appears with v_l for r_2^A and r_4^A , as these basis states contain dimers only on their legs. Similarly, H_B contains a kinetic term with J , allowing flipping between $\{r_2^B \leftrightarrow r_3^B, r_2^B \leftrightarrow r_4^B\}$. A new potential term with coefficient v_{lr} is introduced for $\{r_2^B, r_3^B, r_4^B\}$, since these basis states contain dimers on both rungs and legs. Additionally, a full-leg potential term with v_l appears for r_1^B and r_5^B , as these basis states contain dimers only on their legs. The interaction term H_{QDM} , defined in Eq. 3.2, governs the interactions between adjacent components $\{s, A, B\}$ located on neighboring sites $\{i, i+1\}$. All interacting pairs, $\{ss, sA, As, sB, Bs\}$, form four-sided square plaquettes between nearest-neighbor components. Thus, H_{QDM} effectively serves as the interaction term for these plaquettes.

Matrix Product Operators (MPO) Construction: To compute the ground state using the DMRG algorithm and MPS parametrization, we must express our Hamiltonian as a Matrix Product Operator (MPO). An MPO consists of a chain of four-legged tensors, W_n , for each site of the MPS, as shown in Fig. 7.8. The left




Figure 7.8. (a) Onsite matrix product operators W_n with physical indices σ, σ^* and horizontal matrix indices. (b) A chain of MPOs acting on an MPS, with contracted horizontal indices and vertical contractions applying the MPO to the MPS.

and right legs serve as auxiliary indices, maintaining the same dimension as their adjacent MPO tensors. The upper and lower legs correspond to the physical indices and contract with the physical legs of MPS tensors L_n^* and L_n , respectively. (Here $*$ represents conjugate of a tensor.)

For AB ladders, the Hamiltonian at each site depends on its component, making the construction of MPO tensors W_n more challenging compared to the square ladder case, where all MPS sites share the same MPO structure. The MPO tensors W_n for sites $\{s, A, B\}$ can be constructed using the standard procedure outlined in Ref. [167, 174]. The explicit form of the Hamiltonian for AB ladders is given in Eq. 7.1.1.

$$W_s = \begin{pmatrix} I_2 & 0 & 0 & 0 & 0 & 0 \\ S_m & 0 & 0 & 0 & 0 & 0 \\ S_p & 0 & 0 & 0 & 0 & 0 \\ I_2 - n & 0 & 0 & 0 & 0 & 0 \\ n & 0 & 0 & 0 & 0 & 0 \\ 0 & -JS_m & -JS_p & v_l(I_2 - n) & v_r n & I_2 \end{pmatrix} \quad (7.11)$$

$$W_A = \begin{pmatrix} I_4 & 0 & 0 & 0 & 0 & 0 \\ S_{34}^A & 0 & 0 & 0 & 0 & 0 \\ S_{43}^A & 0 & 0 & 0 & 0 & 0 \\ S_{44}^A & 0 & 0 & 0 & 0 & 0 \\ S_{33}^A & 0 & 0 & 0 & 0 & 0 \\ E & -JS_{12}^A & -JS_{21}^A & v_l S_{22}^A & v_r S_{11}^A & I_4 \end{pmatrix}$$

$$W_B = \begin{pmatrix} I_5 & 0 & 0 & 0 & 0 & 0 \\ S_{31}^B & 0 & 0 & 0 & 0 & 0 \\ S_{13}^B & 0 & 0 & 0 & 0 & 0 \\ S_{11}^B & 0 & 0 & 0 & 0 & 0 \\ S_{33}^B & 0 & 0 & 0 & 0 & 0 \\ F & -JS_{45}^B & -JS_{54}^B & v_l S_{55}^B & v_r S_{44}^B & I_5 \end{pmatrix}$$

and,

$$E = -J \cdot (S_{31}^A + S_{13}^A) + v_l \cdot (S_{22}^A + S_{44}^A) + v_{lr} \cdot (S_{11}^A + S_{33}^A)$$

$$F = -J \cdot ((S_{32}^B + S_{23}^B + S_{42}^B + S_{24}^B) + v_l \cdot ((S_{11}^B + S_{55}^B) + v_{lr} \cdot ((S_{22}^B + S_{33}^B + S_{44}^B)$$

$$W_{(n=0)} = \begin{pmatrix} 0 & -JS_m & -JS_p & v_l(I_2 - n) & v_r n & I_2 \end{pmatrix}; \quad W_{(n=N-1)} = \begin{pmatrix} 0 \\ S_m \\ S_p \\ (I_2 - n) \\ n \\ 0 \end{pmatrix}$$

Here, the matrices are defined as follows: $S_p = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$ is the dimer creation operator. $S_m = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}$ is the dimer annihilation operator. S_{ij}^A is a 4×4 matrix for component ‘A’, with all elements set to 0 except $S_{ij}^A[j-1][i-1] = 1$, which swaps basis states $r_i^A \rightarrow r_j^A$. S_{ij}^B is a 5×5 matrix for component ‘B’, with all elements set to 0 except $S_{ij}^B[j-1][i-1] = 1$, which swaps basis states $r_i^B \rightarrow r_j^B$. I_a is the identity matrix of dimension a . The MPO tensor dimensions $W^i = (b_L, b_R, \sigma_L, \sigma_L^*)$ are: $W_s = (6, 6, 2, 2)$, $W_A = (6, 6, 4, 4)$, $W_B = (6, 6, 5, 5)$. At the first and last sites of the ladder, the dimensions are: $W_{(n=0)} = (1, 6, 2, 2)$ and $W_{(n=N-1)} = (6, 1, 2, 2)$.

AB ladders form closed loops, but for the DMRG algorithm, we open the ladder at a single point, ensuring that the first and last sites contain the s component. This allows $W_{(n=0)}$ and $W_{(n=N-1)}$ to remain fixed for all higher-order ladders. In the smallest (Λ_0) ladder, this is not possible, as it ends with a component ‘A’. Therefore, in this case, $W_{(n=N-1)}$ is equal to the first column of the W_A tensor in Eq. 7.1.1. After applying DMRG, periodic boundary conditions are naturally satisfied in the ground states on every ladder due to their inherent structure.

MPS-DMRG: an iterative ground state search: In this section, I’ll provide a concise overview of the Density Matrix Renormalization Group (DMRG) algorithm using Matrix Product States (MPS) for quantum states and Matrix Product Operators (MPO) for Hamiltonians. For a more detailed discussion, see Ref. [174]. To begin, we start with a randomly initialized MPS, represented by matrices M^σ of appropriate dimensions, which are either left- or right-normalized. The goal is to find the ground state of a given Hamiltonian H , meaning we need to determine an MPS $|\Psi\rangle$ that minimizes the energy functional:

$$E = \frac{\langle \Psi | H | \Psi \rangle}{\langle \Psi | \Psi \rangle}$$

This translates to extremizing the quantity:

$$\langle \Psi | H | \Psi \rangle - \lambda \langle \Psi | \Psi \rangle \quad (7.12)$$

where λ is a Lagrange multiplier enforcing normalization. The corresponding tensor network representation for this equation is shown in Fig. 7.9. DMRG relies on two essential principles:

- **Site-by-Site Optimization:** The minimization is carried out iteratively, optimizing one MPS tensor at a




Figure 7.9. Tensor network representation of Eq. 7.12, which is then optimized to determine the ground state and its corresponding energy. The left side illustrates the expectation value $\langle \Psi | H | \Psi \rangle$, while the right side depicts the squared norm $\langle \Psi | \Psi \rangle$. (Following Ref. [12])




Figure 7.10. Tensor network representation corresponding to the equation $Hv - \lambda v = 0$, derived by differentiating the equation in Fig. (7.9) to optimize a single site $v = L_{a_{l-1}, a_l}^{\sigma_l}$, which is circled in both the left and right networks. This formulation leads to a generalized eigenvalue problem, where the Hamiltonian H is represented by the tensor network enclosed in the blue box. The networks in the two orange boxes depict the left and right environments of the current site l , labeled as LE and RE, respectively. Meanwhile, the tensor network within the purple boxes contracts to yield the identity.

time while keeping all others fixed. At each step, the equation in Fig. 7.9 simplifies to a quadratic form on a single tensor, making it a straightforward linear optimization problem. This progressively reduces the energy and improves the approximation of the ground state. However, since each update is only locally optimal, repeated sweeps through all sites (left-to-right and right-to-left) are necessary until convergence is reached.

- **Truncation via Singular Value Decomposition (SVD):** After optimizing an MPS tensor at a given site, we perform an SVD to restore normalization and retain only the most significant D singular values. Here D is the bond dimension of the MPS. This step is crucial for computational efficiency, as MPS tensors can grow exponentially in size. Fortunately, in 1D systems, the eigenvalue spectrum often decays exponentially, allowing accurate truncation with minimal loss of precision.

During a sweep, optimizing a single site — i.e., differentiating the tensor network equation shown in Fig. 7.9 — leads to a standard eigenvalue problem, illustrated in Fig. 7.10. The key tensor to be optimized, denoted by v , satisfies the equation:

$$Hv - \lambda v = 0 \quad (7.13)$$

where H is the effective Hamiltonian at that site, represented by the tensor network in the blue box of Fig. 7.10. The surrounding tensors (purple boxes) contract to the identity due to MPS normalization, simplifying the problem to the above eigenvalue equation.

DMRG using MPS is a powerful iterative method for finding the ground state of quantum systems. By optimizing site-by-site and enforcing truncation via SVD, it balances accuracy and computational feasibility. The

sweeping procedure (left-to-right and right-to-left) systematically improves the state until the energy converges. This approach is particularly effective in 1D systems, where entanglement remains limited and allows for efficient compression.

7.1.2 Results: Ground-State Phase Diagram

The ground state phase diagram for the quantum dimer model on AB ladders for the RK-Hamiltonian given in Eq. 7.10, is shown in Fig. 7.12. This two-dimensional phase diagram is constructed by varying the parameters v_r and v_l , while first keeping $v_{lr} = 0.01$ fixed to minimize leg-rung interaction. The phases displayed here exist for all AB ladders, which indicates that the non-trivial ground state phases found on these ladders, when combined with the trivial phases present on the purple and gray edges in Fig. 7.1, provide the phase diagram for the entire 2D AB* tiling.

I present results here for AB ladders of order $\{\Lambda_4, \Lambda_5\}$, with corresponding lengths $L = \{432, 1296\}$. In all simulations, I retain all states with Schmidt values greater than 10^{-8} . Typically, only 1-2 full sweeps are required to achieve excellent energy convergence for these large-size ladders. However, when running DMRG for smaller ladders of order Λ_0, Λ_1 , it is necessary to increase the number of sweeps. Additionally, the bond dimension is increased by 100 after every half sweep.

Fig. 7.12 shows a total of five phases, with the corresponding probability density for the dimer basis on each MPS site. This provides insight into how these QDM phases look. To calculate the probability density observable for basis states on each site, we first need to define operators for these observables.

For an MPS site with the ‘s’ component of the ladder, the operator for calculating the probability of having a dimer is:

$$\hat{n}^s = S_p S_m = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad (7.14)$$

and the operator for not having a dimer is:

$$\hat{O}^s = I_2 - \hat{n}^s \quad (7.15)$$

For an MPS site with the ‘A’ component of the ladder, the set of operators used to calculate the probability of each basis state $\{r_j^A; j = 1, 2, 3, 4\}$ is represented by a set of 4×4 diagonal matrices. Each operator \hat{O}_j^A is defined such that:

$$\hat{O}_j^A[j-1][j-1] = 1, \quad \text{and all other entries are zero,} \quad (7.16)$$

where $j = 1, 2, 3, 4$.

For an MPS site with the ‘B’ component of the ladder, the set of operators used to calculate the probability of each basis state $\{r_j^B; j = 1, 2, 3, 4, 5\}$ is represented by a set of 5×5 diagonal matrices. Each operator \hat{O}_j^B is defined such that:

$$\hat{O}_j^B[j-1][j-1] = 1, \quad \text{and all other entries are zero,} \quad (7.17)$$

where $j = 1, 2, 3, 4, 5$.

Using these on-site probability operators, the probability density for each basis state can be calculated via tensor network contraction, as shown in Fig. 7.11, depending on whether the left or right canonical MPS tensor is used at site k .

In Fig. 7.13 and Fig. 7.14, the blue plot points collectively illustrate probability of forming either just rung or leg-rung dimers on any site. For ‘s’ sites it shows $\langle \Psi | \hat{n}^s | \Psi \rangle$, which gives the probability of having a rung




Figure 7.11. Tensor network for calculating the onsite expectation value of any operator O^i . This is also used to calculate the onsite probability density for different dimer basis states. The tensors in green and blue represent the left and right environments, respectively, while the red diamond-shaped tensors are the Schmidt matrices on any bond.




Figure 7.12. Ground state phase diagram of the QDM on the AB ladder, as defined by Eq. 7.10, plotted as a function of the parameters $\{v_r, v_l\}$ with $v_{lr} = 0.01$. The diagram displays five distinct phases described in the text, along with their phase boundaries. The thick black line denotes the isotropic line where $v_r = v_l$. Solid (pink, green) lines indicate first-order phase transitions, while dashed (pink, green, blue) lines correspond to second-order transitions. The orange and black circles mark the values of v_r at which various observables are evaluated to study the nature of the transitions in Sec. 7.1.3. The purple circle highlights the region where all five phases converge, suggesting a transitional regime.

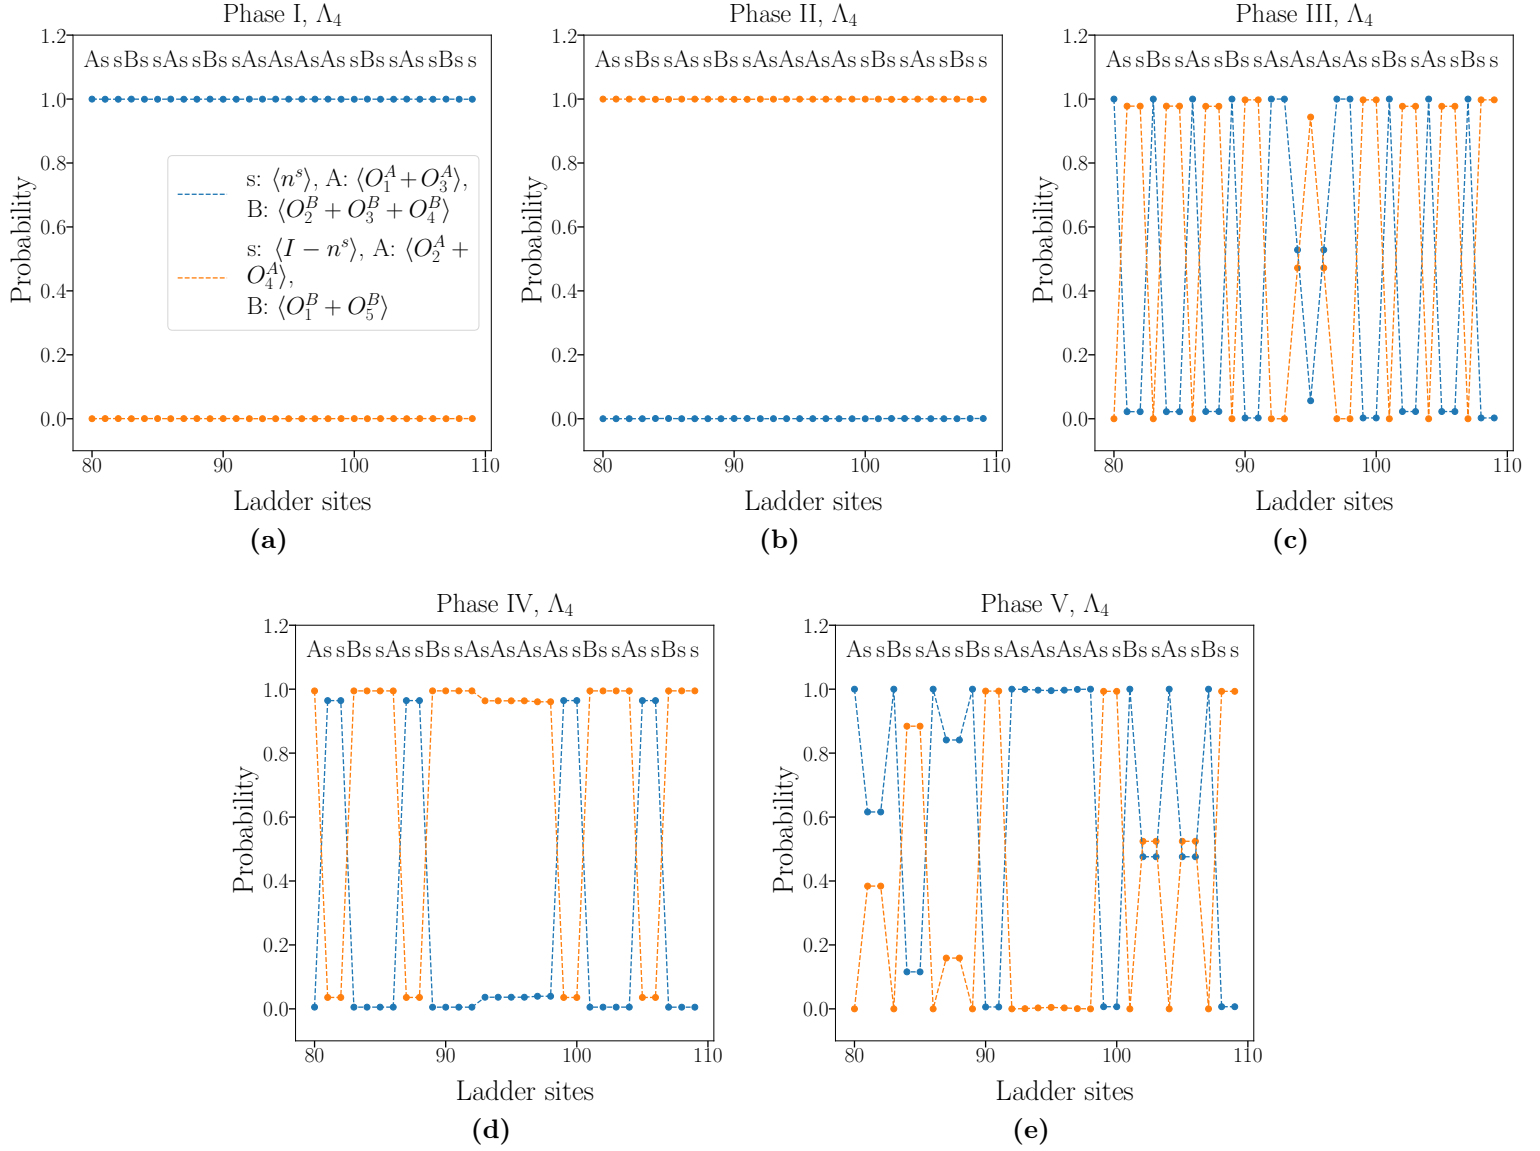


Figure 7.13. Probability density profiles for the five distinct phases on the AB ladder Λ_4 with 432 sites. A small segment of the ladder is shown for illustration, along with its site components $\{s, A, B\}$. Blue points indicate the probability of either a rung dimer on s sites or a leg-rung dimer basis on $\{A, B\}$ sites. Orange points represent the probability of no rung dimer on s sites or an all-leg dimer basis on $\{A, B\}$ sites. The legend in (a) applies to all subfigures.

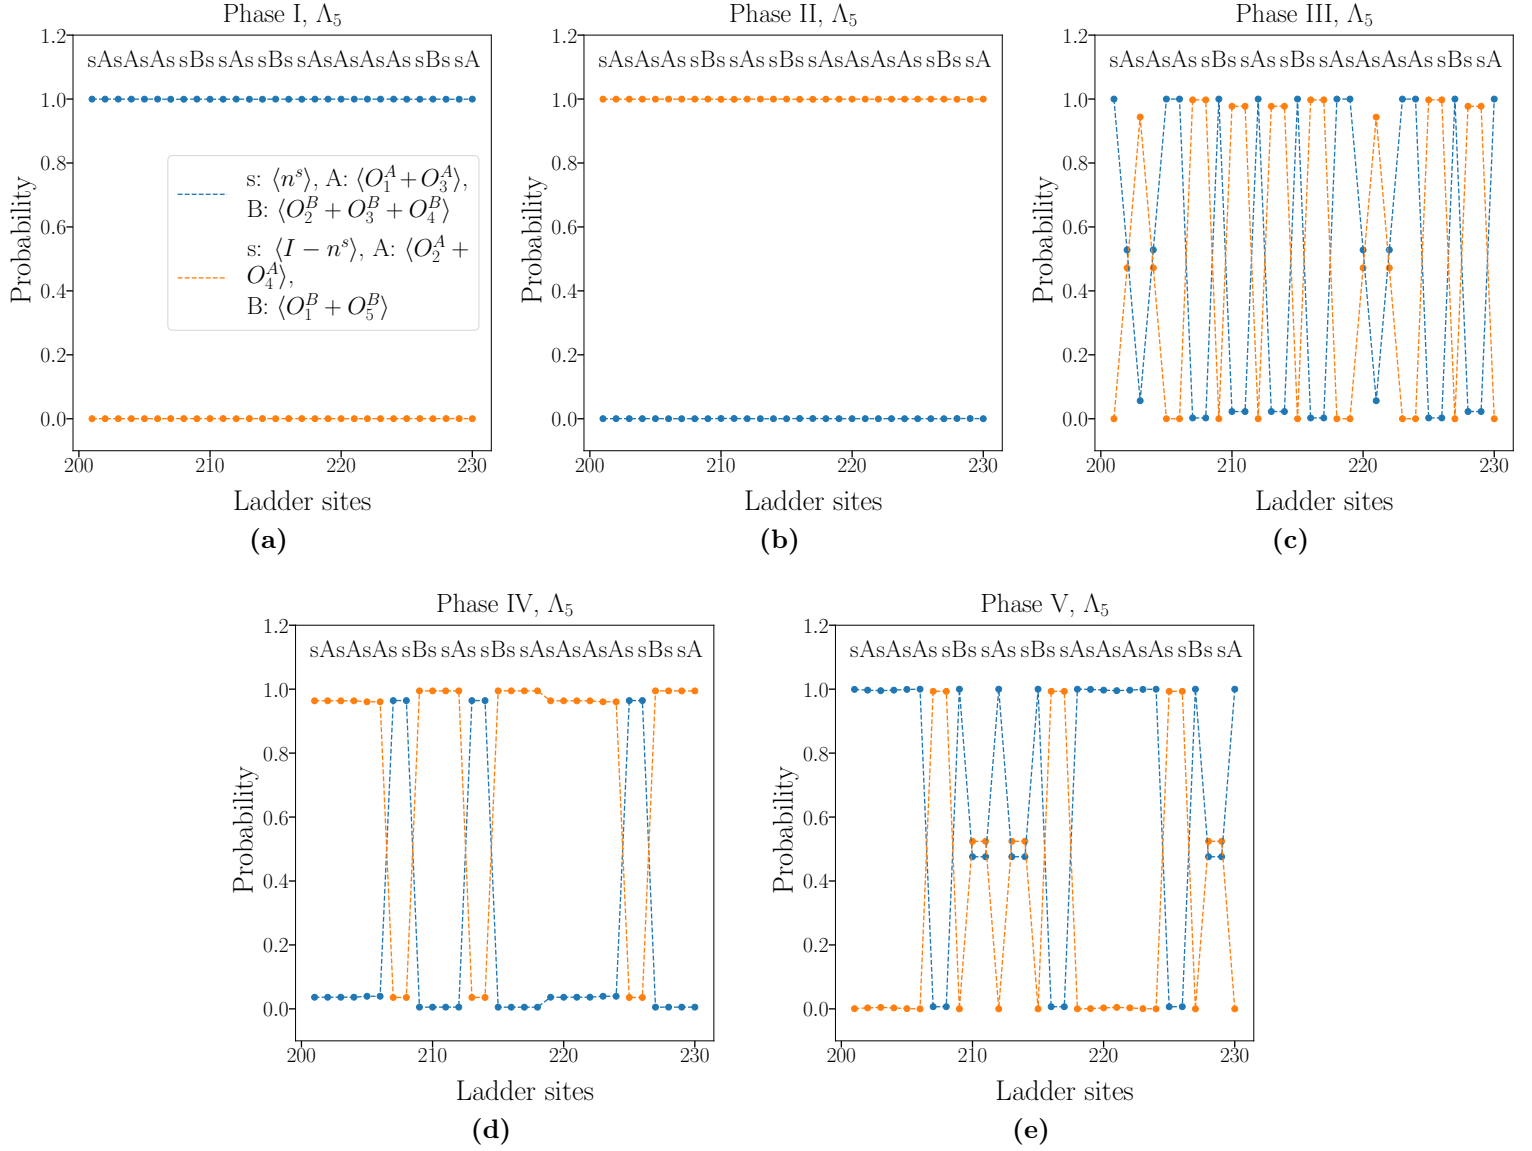


Figure 7.14. Probability density profiles for the five distinct phases on the AB ladder Λ_5 with 1296 sites. A small segment of the ladder is shown for illustration, along with its site components $\{s, A, B\}$. Blue points indicate the probability of either a rung dimer on s sites or a leg-rung dimer basis on $\{A, B\}$ sites. Orange points represent the probability of no rung dimer on s sites or an all-leg dimer basis on $\{A, B\}$ sites. The legend in (a) applies to all subfigures.




Figure 7.15. Pictorial representation of phase I on the Λ_1 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. Sites ‘s’ always host a rung dimer (green). Sites ‘A’ are in an equal superposition of $\{|r_1^A\rangle, |r_3^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$. Sites ‘B’ are in a superposition of $\{|r_2^B\rangle, |r_3^B\rangle, |r_4^B\rangle\}$ with coefficients depending on v_r : $(d, e, f) = (0, 1/\sqrt{2}, 1/\sqrt{2})$ for $v_r \ll 0$, $(d, e, f) = (1/\sqrt{3}, 1/\sqrt{3}, 1/\sqrt{3})$ for $v_r < 0$, $(d, e, f) = (1/\sqrt{2}, 1/2, 1/2)$ for $v_r > 0$, and $(d, e, f) = (1, 0, 0)$ for $v_r \gg 0$.

dimer. For ‘A’ sites it shows the sum $\langle \Psi | \hat{O}_1^A | \Psi \rangle + \langle \Psi | \hat{O}_3^A | \Psi \rangle$, indicating the combined probability of having a leg-rung dimer configuration. For ‘B’ sites it shows the sum $\langle \Psi | \hat{O}_2^B | \Psi \rangle + \langle \Psi | \hat{O}_3^B | \Psi \rangle + \langle \Psi | \hat{O}_4^B | \Psi \rangle$, representing the combined probability of having a leg-rung dimer configuration.

The orange plot points provide the complementary probabilities: For ‘s’ sites it shows $\langle \Psi | I_2 - \hat{n}^s | \Psi \rangle$, which gives the probability of not having a rung dimer (i.e., having a leg dimers instead). For ‘A’ sites it shows the sum $\langle \Psi | \hat{O}_2^A | \Psi \rangle + \langle \Psi | \hat{O}_4^A | \Psi \rangle$, which represents the probability of only having leg dimers. For ‘B’ sites it shows the sum $\langle \Psi | \hat{O}_1^B | \Psi \rangle + \langle \Psi | \hat{O}_5^B | \Psi \rangle$, which again gives the probability of only leg dimers on ‘B’ sites.

This visualization clearly distinguishes between (rung + leg-rung) or just legs dimer configurations across different sites, helping to characterize the quantum dimer phases in the AB ladder system. From the probability density for each basis state across all sites, we can now understand the mean-field quantum dimer configuration of the phases in Fig. 7.12 in detail:

Phase I: This phase is equivalent to the rung-dimer phase observed in Ref. [167] for the square ladder, where each plaquette favors rung dimers. An illustration is provided in Fig. 7.15. So, the ‘s’ sites always host dimers. The ‘A’ sites are in the state $(|r_1^A\rangle + |r_3^A\rangle)/\sqrt{2}$. The state on ‘B’ sites varies along the v_r -axis:

- At large negative v_r , the ‘B’ sites are in the state $(|r_3^B\rangle + |r_4^B\rangle)/\sqrt{2}$.
- As v_r approaches 0 from the left, the probability of r_2^B increases while those of r_3^B and r_4^B decrease. Near $v_r = 0$, the ‘B’ sites are in the state $(|r_2^B\rangle + |r_3^B\rangle + |r_4^B\rangle)/\sqrt{3}$.
- For slightly positive v_r , the ‘B’ sites are in the state $|r_2^B\rangle/\sqrt{2} + (|r_3^B\rangle + |r_4^B\rangle)/2$.
- As v_r increases further, the ‘B’ sites approach the pure state $|r_2^B\rangle$.

Phase II: This phase is equivalent to the columnar phase observed in Ref. [167] for the square ladder, where each plaquette contains only leg dimers. An illustration is provided in Fig. 7.16. So, the ‘s’ sites never host dimers. In any AB ladder, the system will exhibit one of two configurations: (i) ‘A’ and ‘B’ sites in the pure states $|r_4^A\rangle$ and $|r_5^B\rangle$, or (ii) ‘A’ and ‘B’ sites in the pure states $|r_2^A\rangle$ and $|r_1^B\rangle$. To satisfy the hard dimer constraint, dimers on any square plaquette can only reside on horizontal legs or be entirely absent, depending on the configuration of neighboring sites.




Figure 7.16. Pictorial representation of phase II on the Λ_1 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges indicate the absence of a dimer. Sites ‘s’ never host a dimer. In any AB ladder, the system exhibits one of two configurations: (i) ‘A’ and ‘B’ sites in the pure states $|r_2^A\rangle$ and $|r_1^B\rangle$, or (ii) ‘A’ and ‘B’ sites in the pure states $|r_4^A\rangle$ and $|r_5^B\rangle$.




Figure 7.17. Pictorial representation of phase III on the Λ_2 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. For ladder segments of the form $(sAsAsA\dots)$, the ‘A’ sites alternate between equal superpositions of $\{|r_3^A\rangle, |r_4^A\rangle\}$ and $\{|r_1^A\rangle, |r_2^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$. The ‘s’ sites alternate between hosting a rung dimer and being empty, depending on the neighboring ‘A’ state, in order to satisfy the hard-core dimer constraint. For ladder segments of the form $(AssB\dots)$, the ‘ss’ sites do not host rung dimers but only leg dimers. The ‘A’ sites are in an equal superposition of $\{|r_1^A\rangle, |r_3^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$. The ‘B’ sites are in a superposition of $\{|r_2^B\rangle, |r_3^B\rangle, |r_4^B\rangle\}$ with coefficients $(d, e, f) = (1/\sqrt{2}, 1/2, 1/2)$.

Phase III: For Λ_0 , with a ladder structure $[sAsAsAsAsAsAsA]$, the alternating ‘A’ sites are in the state $\frac{|r_1^A\rangle + |r_2^A\rangle}{\sqrt{2}}$ or $\frac{|r_3^A\rangle + |r_4^A\rangle}{\sqrt{2}}$. The ‘s’ sites alternate between hosting dimers and being empty.

For Λ_1 , the ‘A’ sites are in the state $\frac{|r_1^A\rangle + |r_3^A\rangle}{\sqrt{2}}$. The ‘B’ sites are in the state $\frac{|r_2^B\rangle}{\sqrt{2}} + \frac{|r_3^B\rangle + |r_4^B\rangle}{2}$. The ‘s’ sites do not host dimers. To satisfy the hard dimer constraint, dimers on any square plaquette can only reside on horizontal legs or be entirely absent, depending on the configuration of neighboring sites.

From Fig. 7.1 and Eq. 7.4, note that any higher-order ladder ($\Lambda_n, n > 1$) structure is a combination of Λ_0 ladder segments of the form $(sAsA\dots)$ and Λ_1 ladder segments of the form $(AssB\dots)$.

Our results indicate that the ground states of all higher ladders also arise as combinations of dimer configurations from these two base ladder structures, $\{\Lambda_0, \Lambda_1\}$. The key distinction arises at the junction where a segment of type $(sAsA\dots)$ connects with a segment of type $(AssB\dots)$. At this junction, the A sites are in either of the pure states $|r_1^A\rangle$ or $|r_3^A\rangle$. An illustration is provided in Fig. 7.17.

Phase IV: For Λ_0 , the ground states remain the same as in Phase II. Similarly, for Λ_1 , the states on ‘A’ and ‘B’ sites are identical to those in Phase II. In this phase, alternating ‘ss’ components on square plaquettes




Figure 7.18. Pictorial representation of phase IV on the Λ_2 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. The states on ‘A’ and ‘B’ sites are identical to those in phase II. In this phase, alternating ‘ss’ components on square plaquettes host two rung dimers, depending on their neighbouring configurations.




Figure 7.19. Pictorial representation of phase V on the Λ_2 ladder segment, for instance. Thick green (purple) edges denote rung (leg) dimers, while thin black edges denote the absence of a dimer. The ‘A’ sites are in an equal superposition of $\{|r_1^A\rangle, |r_3^A\rangle\}$ with coefficients $(a, b) = (1/\sqrt{2}, 1/\sqrt{2})$, while the ‘B’ sites are in a superposition of $\{|r_2^B\rangle, |r_3^B\rangle, |r_4^B\rangle\}$ with coefficients $(d, e, f) = (1/\sqrt{2}, 1/2, 1/2)$. In this phase, ‘ss’ components on square plaquettes alternately host either two rung dimers or two leg dimers.

host either two rung dimers or two leg dimers, depending on their neighbouring configurations. Any higher-order ladder ($\Lambda_n, n > 1$) structure is a combination of Λ_0 ladder segments of the form $(sAsA\dots)$ and Λ_1 ladder segments of the form $(AssB\dots)$. Our results indicate that the ground states of all higher ladders also arise as combinations of dimer configurations from these two base ladder structures, $\{\Lambda_0, \Lambda_1\}$. An illustration is provided in Fig. 7.18.

Phase V: For Λ_0 , the ground states remain identical to those in Phase I. Similarly, for Λ_1 , the basis states of ‘A’ and ‘B’ sites are the same as in Phase I. Alternating ‘ss’ components host either two rung dimers or two leg dimers, based on their neighboring configurations. Any higher-order ladder ($\Lambda_n, n > 1$) structure is a combination of Λ_0 ladder segments of the form $(sAsA\dots)$ and Λ_1 ladder segments of the form $(AssB\dots)$. Our results indicate that the ground states of all higher ladders also arise as combinations of dimer configurations from these two base ladder structures, $\{\Lambda_0, \Lambda_1\}$. The key distinction arises at the junction where a segment of type $(sAsA\dots)$ connects with a segment of type $(AssB\dots)$. At this junction, the A sites are in either of the pure states $|r_1^A\rangle$ or $|r_3^A\rangle$. An illustration is provided in Fig. 7.19.

The region enclosed by the purple circle, with a radius of approximately 1.5 centered at $(-1.5, 0)$, marks an area where all five phases converge, resulting in ground states characterized by distinct quantum superpositions.

In this region, all ‘A’ sites are found in the state $(|r_1^A\rangle + |r_3^A\rangle)/\sqrt{2}$, while all ‘B’ sites are in the state $|r_2^B\rangle/\sqrt{2} + (|r_3^B\rangle + |r_4^B\rangle)/2$. Any ‘s’ component flanked by ‘A’ sites on both sides (as in the case of Λ_0) hosts a rung dimer with a probability of approximately 0.75. Any square plaquette (such as those found in the middle of configurations like $..AssB..$) exhibits a plaquette phase, where dimers appear with equal probability on all four edges of the corresponding square plaquettes. Thus, this region appears to host transitional ground states, smoothly connecting multiple phases as the system moves across phase boundaries.

7.1.3 Phase Transitions

In this section, I analyze the nature of quantum phase transitions at the boundaries between different phases in Fig. 7.12. This analysis is based on two key quantities:

1. **Order Parameter (OP):** This quantity captures symmetry breaking between two phases. A first-order phase transition is characterized by a discontinuous jump in the OP (from 0 to 1) when crossing the boundary, while a second-order phase transition exhibits a continuous and gradual increase in the OP from 0 to 1 as the system transitions between phases.

2. **Ground-State Energy (E_N/N):** A first-order phase transition is typically marked by a kink in the ground-state energy per site, which becomes more pronounced in the thermodynamic limit. In contrast, a second-order phase transition is indicated by a smooth, continuous variation of E_N/N between the two phases.

First-Order Phase Transitions: I observe that two pairs of phases, $\{\text{II}, \text{III}\}$ and $\{\text{IV}, \text{I}\}$, are separated by first-order phase transitions. Fig. 7.20 illustrates the variation of the OP and E_N/N as a function of v_l for different fixed values of v_r . These plots are shown for two ladder sizes, $\{\Lambda_4, \Lambda_5\}$.

Fig. 7.20(a,c) depicts the transition between phases II and III. From the earlier description of these phases, we see that moving from phase II to phase III leads to a change in the basis states on ‘B’ components. Specifically, the basis states shift from a quantum superposition of all leg dimer states $\{r_1^B, r_5^B\}$ in phase II to a quantum superposition of rung-leg dimer states $\{r_2^B, r_3^B, r_4^B\}$ in phase III.

To capture this symmetry breaking, we define the order parameter as the probability of finding leg dimers on any ‘B’ component: $OP = \langle \Psi | \hat{O}_1^B | \Psi \rangle + \langle \Psi | \hat{O}_5^B | \Psi \rangle$. Fig. 7.20(a) shows a sharp discontinuous jump in the OP from 1 to 0 as we increase v_l for a given v_r , indicating a first-order transition. Similarly, Fig. 7.20(c) shows a distinct kink in the ground-state energy E_N/N as a function of v_l , further confirming the first-order nature of the phase transition. The kinks in E_N/N and the abrupt jump in OP occur at $v_l = -0.56, -0.58$ for $v_r = 8, 20$, respectively. These transition points are marked as black circles in Fig. 7.12.

Fig. 7.20(b,d) illustrates the transition between phases IV and I. As described earlier, moving from phase IV to phase I leads to a change in the basis states on ‘B’ components, similar to the transition from phase II to III.

To capture this symmetry breaking, we can again define the order parameter as the probability of finding leg dimers on any ‘B’ component: $OP = \langle \Psi | \hat{O}_1^B | \Psi \rangle + \langle \Psi | \hat{O}_5^B | \Psi \rangle$. Fig. 7.20(b) shows a sharp discontinuous jump in the OP from 1 to 0 as v_l increases for a fixed v_r , indicating a first-order phase transition. Similarly, Fig. 7.20(d) displays a distinct kink in the ground-state energy E_N/N as a function of v_l , further confirming the first-order nature of the transition. The kinks in E_N/N and the abrupt jump in OP occur at $v_l = -24.2, -3.75$ for $v_r = -40, -6$, respectively. These transition points are marked as black circles in Fig. 7.12.

Second-Order Phase Transitions: The remaining three phase boundaries — between $\{\text{III}, \text{V}\}$, $\{\text{II}, \text{IV}\}$,

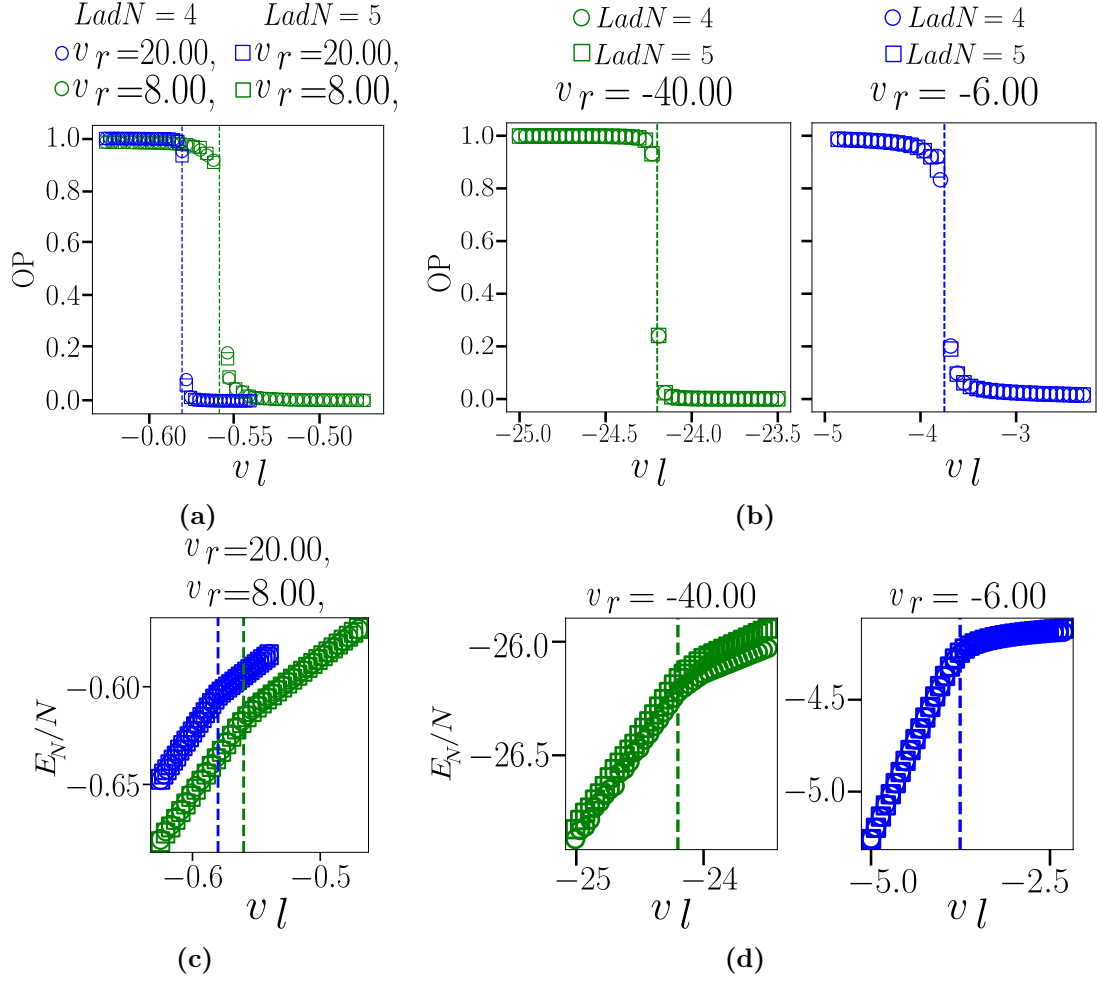


Figure 7.20. First-Order Phase Transitions: These transitions are characterized by the variation of the order parameter (OP) and the ground state energy per site (E_N/N) as functions of v_l , near the boundaries between different phases, for different fixed values of v_r . Subfigures (a,c) show the transition from Phase II to Phase III, while (b,d) correspond to the transition from Phase IV to Phase I. In all cases, the OP exhibits a discontinuous jump from 0 to 1, and the ground state energy per site (E_N/N) shows a distinct kink. This discontinuous behavior of both the order parameter and the energy confirms the first-order nature of these transitions. Vertical dashed lines indicate the location of the transitions.

and $\{V, I\}$ — exhibit continuous phase transitions. Fig. 7.21 illustrates the variation of the order parameter (OP) and ground-state energy per site E_N/N as functions of v_l for different fixed values of v_r . These plots are shown for two ladder sizes, $\{\Lambda_4, \Lambda_5\}$.

Fig. 7.21(a,b) focuses on the transition between phases III and V. As described earlier, moving from phase III to phase V results in a structural change in the basis states on alternate ‘ss’ components. Specifically, phase III has no dimers on ‘ss’ components, whereas phase V exhibits two rung dimers on alternating ‘ss’ components.

To characterize this transition, we define the order parameter as:

$$OP = \frac{2 \times \sum(\text{probability of finding rung dimers on any 'ss' pair})}{\text{total number of 'ss' pairs}}. \quad (7.18)$$

By construction, this parameter is 0 in phase III and 1 in phase V. Fig. 7.21(a) shows a continuous increase in the OP from 0 to 1 as v_l increases for a fixed $v_r = \{9, 20\}$, confirming a second-order phase transition. Additionally, Fig. 7.21(b) reveals the absence of a kink in E_N/N around the transition, further supporting the second-order nature of this phase change. The approximate quantum transition points are marked as red circles in Fig. 7.12.

Fig. 7.21(c,d) examines the transition between phases II and IV. As shown earlier, this transition also involves a structural change in the basis states on alternating ‘ss’ components, similar to the transition between phases III and V. Specifically, phase II has no dimers on ‘ss’ components, whereas phase IV features two rung dimers on alternating ‘ss’ components.

To characterize this transition, we use the same order parameter (OP) as in Eq. 7.18. By definition, this parameter is 0 in phase II and 1 in phase IV. Fig. 7.21(c) shows a continuous increase in the OP from 0 to 1 as v_l increases for a fixed $v_r = \{-20, -8\}$, confirming a second-order phase transition. Additionally, Fig. 7.21(d) shows no kink in the ground-state energy per site E_N/N , further supporting the second-order nature of the transition. The approximate quantum transition points are marked as red circles in Fig. 7.12.

Fig. 7.21(e,f) examines the transition between phases V and I. As described earlier, this transition involves a structural change in the basis states on alternating ‘ss’ components. Specifically, phase V features two rung dimers on alternating ‘ss’ components, whereas phase I has two rung dimers on all ‘ss’ components.

To characterize this transition, we define the order parameter (OP) as:

$$OP = \frac{2 \times \sum(\text{probability of finding rung dimers on any 'ss' pair})}{\text{total number of 'ss' pairs}} - 1.$$

By construction, this parameter is 0 in phase V and 1 in phase I. Fig. 7.21(e) shows a continuous increase in the OP from 0 to 1 as v_l increases for a fixed $v_r = \{9, 18\}$, confirming a second-order phase transition. Additionally, Fig. 7.21(f) shows no kink in the ground-state energy per site E_N/N , further supporting the second-order nature of the transition. The approximate quantum transition points are marked as red circles in Fig. 7.12.

7.2 Conclusion

Hence, the DMRG results obtained on AB* ladders, combined with the trivial phases known on 8-start loops, reveal how the ground state of the full AB* tiling quantum dimer model defined in Eq. 7.10 evolves across

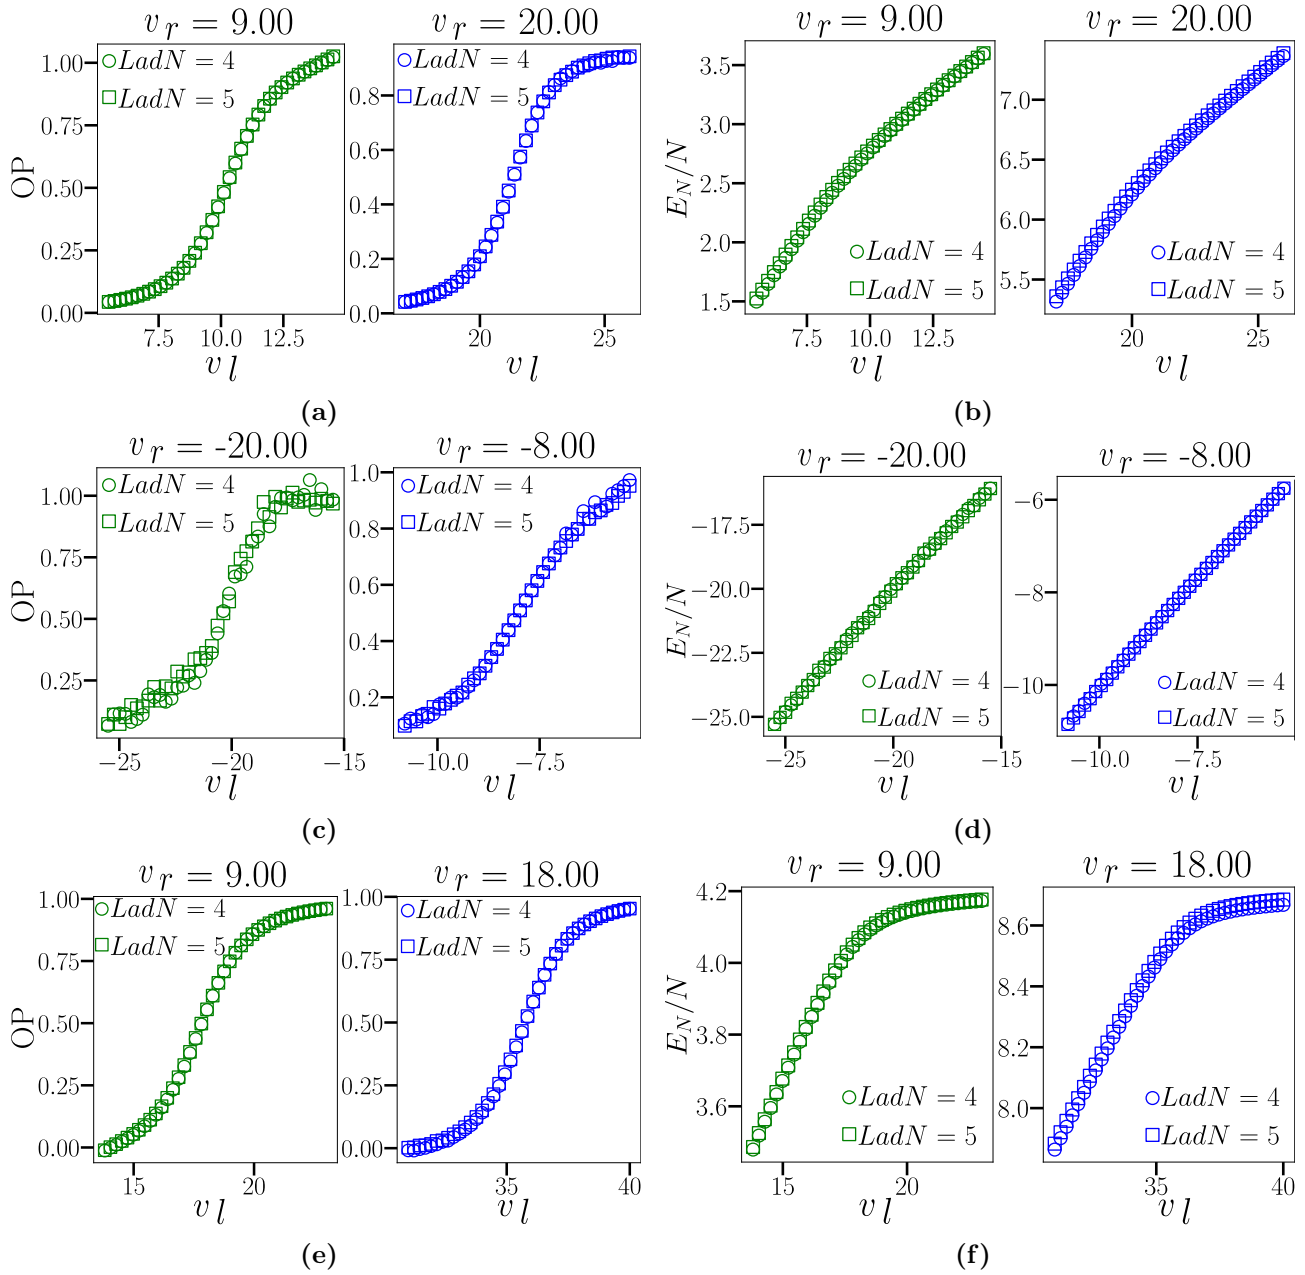


Figure 7.21. Second-Order Phase Transitions: These transitions are characterized by the variation of the order parameter (OP) and the ground state energy per site (E_N/N) as a function of v_l , near the boundaries between different phases, for different fixed values of v_r . Subfigures (a,b) show the transition from Phase III to Phase V, (c,d) correspond to the transition from Phase II to Phase IV, and (e,f) illustrate the transition from Phase V to Phase I. In all cases, the OP increases continuously from 0 to 1, and the ground state energy per site (E_N/N) varies smoothly without any kink. The continuous behavior of both the order parameter and the energy confirms the second-order nature of these transitions.

different values of $\{v_r, v_l\}$ at fixed $v_{lr} = 0.01$, and uncover the various types of phase transitions that emerge between the identified phases. We have thus successfully applied the DMRG method to study a 2D system based on the AB* tiling, leveraging its unique connectivity and discrete scale invariance properties.

Studying large-scale 2D quantum systems using DMRG is typically very challenging due to rapidly increasing entanglement entropy and computational constraints. For instance, one of the largest full 2D kagome clusters, with 48 sites, was studied using exact diagonalization (ED) rather than DMRG [175]. By contrast, state-of-the-art DMRG studies on the kagome Heisenberg antiferromagnet are typically performed on cylindrical geometries with circumferences up to about 12 sites [176, 177]. In this chapter, we present DMRG results for AB* ladders that contain up to $N = 1296$ sites. For a patch of tiling where this AB* ladder is the largest and wraps around a roughly circular geometry with perimeter $N = 2\pi R$. The enclosed area is then $\pi R^2 = \frac{N^2}{4\pi}$. For $N \sim 10^3$, this corresponds to an area containing approximately 10^6 sites, effectively accessing a two-dimensional system of the order of a million sites.

These results offer a promising perspective: the limitations of DMRG in handling large 2D systems may be circumvented by exploiting the combination of aperiodicity and local constraints in the QDM. While regular, periodic lattices impose strong entanglement constraints, aperiodic lattices such as the AB* tiling — with their non-trivial connectivity and discrete scale invariance — can, when coupled with the hard-dimer constraint, provide access to much larger 2D quantum systems using DMRG. This opens new possibilities for exploring quantum many-body physics in complex geometries beyond conventional periodic systems.

The Fully packed loop Model on AB tiling

The Fully Packed Loop (FPL) model, as introduced in Sec. 3.3, is a fundamental model in statistical mechanics, providing insights into critical phenomena and phase transitions. Extensive studies on regular lattices, particularly the square lattice, have enhanced our understanding of universality classes, critical exponents, and conformal invariance [99, 178, 87, 159].

In this chapter, I extend the exploration of the FPL model to Ammann-Beenker (AB) tilings — non-standard, aperiodic lattice structures — to investigate how lattice geometry influences these fundamental concepts. Universality suggests that systems with differing microscopic details can exhibit identical critical behavior, characterized by the same critical exponents and scaling functions. Critical exponents describe how physical quantities diverge near critical points, providing a quantitative measure of phase transitions.

Our objective is to compute the critical exponents on the AB tiling using various loop statistics parameters and correlations, and compare them with those from the square lattice. By implementing the model on the AB tiling, I aim to determine whether the universality class remains consistent or if the aperiodic nature of the lattice induces novel critical behavior.

We can also create height field representations for these loop models, as explained in Sec. 3.4. In Sec. 3.4, I explained how to construct a height model for fully packed loops. This formulation maps the loop model onto a fluctuating height surface, where the loops act as contour lines of the height field. This approach is a powerful tool for utilizing conformal field theory techniques and understanding roughening transitions present in two-dimensional critical systems.

In summary, this chapter will first explain how to construct various loop configurations on the AB tiling. Subsequently, I will calculate, at infinite temperature, various loop and height statistics to determine critical exponents.

8.1 Construction of Fully Packed Loops on AB

To construct fully packed loops (FPLs) on the Ammann-Beenker (AB) tiling, one can start with an initial configuration, such as a Hamiltonian cycle that visits every vertex exactly once, or by placing loops on every non-adjacent plaquette to ensure all vertices are covered. Subsequent configurations can be generated by performing ‘flip’ moves on flippable plaquettes. A ‘flip’ move exchanges dimers on a flippable plaquette from covered to uncovered edges, analogous to the Rokhsar-Kivelson (RK) flip in quantum dimer models [179].




Figure 8.1. An example of a fully packed loop configuration (shown by purple edges) on the U_1 AB tiling (shown by thin gray edges), resulting from multiple flip move sweeps. The corresponding discrete height field is also illustrated, with each plaquette labeled by its respective height, derived from the steps in Sec. 3.4. Loops with blue nodes have clockwise orientation, while those with green nodes have anti-clockwise orientation.

Implementing these flip moves to sample FPL configurations is equivalent to applying the Metropolis Monte Carlo algorithm at infinite temperature, where every flip is accepted with probability one. Typically, a single sweep consists of a number of flip iterations equal to the total number of plaquettes in the tiling. This approach explores the configuration space of FPLs on the AB tiling, allowing for the analysis of their statistical and physical properties.

An example of a fully packed loop configuration on the AB tiling is shown in Fig. 8.1, and the corresponding flip move is illustrated in Fig. 8.2.

8.2 Infinite Temperature - Criticality

To determine the critical exponents at infinite temperature for the AB tiling, we must first establish whether the system exhibits critical behavior at infinite temperature. This can be assessed using the height field model. In this framework, the nature of the height field determines the criticality of the system: if the heights are ‘rough’, the corresponding 2D model is critical. Conversely, smooth heights indicate a non-critical phase [93]. I will use the loop height models explained in Sec. 3.4.




Figure 8.2. The Rokhsar-Kivelson (RK) flip, which shifts dimers (highlighted in purple) from one set of alternating edges to the complementary set within a plaquette.

As detailed in Sec. 3.4, the roughness of the height field is characterized by a specific scaling relation in Eq. 3.6. For a critical system, after coarse-graining, fourier transform of the height fluctuations follows: $\langle |h(\mathbf{p})|^2 \rangle^{-1} \sim K|\mathbf{p}|^2$. Taking the logarithm on both sides, the relationship becomes: $-\log(\langle |h(\mathbf{p})|^2 \rangle) \sim \log(K) + 2\log(|\mathbf{p}|)$. This implies that a log-log plot of $-\log(\langle |h(\mathbf{p})|^2 \rangle)$ versus $\log(|\mathbf{p}|)$ should exhibit a linear slope of 2 in the rough phase. If the slope deviates from 2, the system resides in the smooth phase [93].

Numerical Implementation : To verify the criticality of the height model, I calculated $\langle |h(\mathbf{p})|^2 \rangle$ numerically by averaging over a large number of height configurations obtained from fully packed loop (FPL) samples. For clarity, the following log-log plots present one-dimensional cuts through reciprocal space in the (1,0) (x-axis) and (1,1) (diagonal) directions, similar to the approach in Ref. [93].

Fig. 8.3(a) shows the log-log plot for the height model on a 64×64 square lattice with periodic boundary conditions. The ensemble average $\langle \dots \rangle$ is computed over 10,000 flip sweeps (Each sweep consists of a number of flips equal to the total number of flippable plaquettes). The slopes of the linear sections in the plot are approximately 1.78 ± 0.028 and 1.86 ± 0.023 for the (1,1) and (1,0) directions, respectively. Since the slopes are not equal to 2, these results indicate that the fully packed loop's height model, where loops represent equal-height contours, resides in the smooth phase (at $T = \infty$) for the square lattice. This conclusion is consistent with the findings of Ref. [93], which describe two distinct height models for fully packed loops (FPLs) on square lattices: the z_1 model [159], characterized by rough behavior, and the z_2 model, exhibiting smooth behavior. Our height model, with loops acting as height contours, corresponds to the z_2 model. Additionally, the z_1 height model requires a regular lattice where all vertices have the same degree, a condition that is not necessary for the z_2 model. For this reason, I did not explore the z_1 model in this work, as our ultimate goal was to study loops on the AB tiling, where all vertices have varying degrees.

For the U_1 AB tiling, comprising 464 vertices, the corresponding log-log plot is shown in Fig. 8.3(b). Here, the ensemble average $\langle \dots \rangle$ is calculated over 4,000 flip sweeps. The slope of the linear fit is approximately 2.12 ± 0.035 , which is very close to 2, though not within the stated error margin. Since this error bar reflects only the standard error, it provides a rough estimate rather than a definitive bound. Nevertheless, the result indicates that the height model on the AB tiling resides in the rough phase and is, therefore, critical at $T = \infty$.

For the U_2 AB tiling, comprising 14,992 vertices, I also calculated the variance of the height field according to Eq. 3.7 and confirmed that it diverged logarithmically with system size L in the thermodynamic limit. The variance $W^2(\mathbf{x})$ in real space was calculated for each plaquette center using an ensemble of FPLs and height samples obtained from 2000 flip sweeps, as shown in Fig. 8.4. Here, darker shades indicate lower variance, while lighter shades represent higher variance. This visualization clearly shows that variance increases from the center toward the boundary.

To confirm logarithmically increasing variance, I plot W^2 against $\ln(L)$. A linear relationship in this plot would indicate logarithmic divergence, as described by Eq. 3.8. Starting from the center of the tiling, I considered concentric circles of increasing radius, effectively creating AB regions of increasing sizes, with the linear system size L corresponding to the radius of each region. Then I averaged the variance of all plaquettes within each

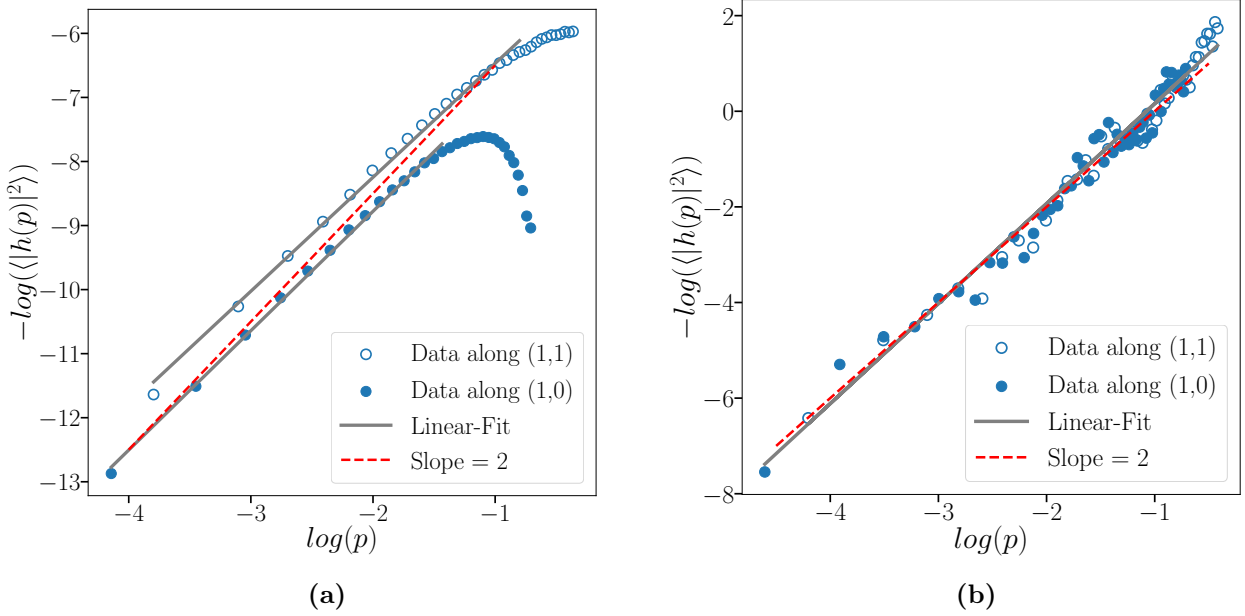


Figure 8.3. Plot of $[-\log(\langle |h(\mathbf{p})|^2 \rangle)]$ vs $[\log(|\mathbf{p}|)]$: (a) for the square lattice FPL model, where the slope of the linear fit (gray line) deviates from 2, confirming the smooth phase at $T = \infty$. (b) for the AB tiling FPL model, where the slope of the linear fit (gray line) is very close to 2, confirming rough phase or criticality at $T = \infty$.

region to obtain W^2 for a particular system size L . Fig. 8.5 shows the plot of W^2 versus $\ln(L)$, resulting in a straight line with slope $= 0.775 \pm 0.00226$, yielding $\frac{1}{2\pi K} = 0.77$ which implies the stiffness constant $K = 0.2067$.

Hence, this analysis confirms the presence of a rough phase, showing that the FPL model, with loops representing equal-height contours on the AB tiling, exhibits critical behavior at infinite temperature. This behavior contrasts with the smooth phase observed for the FPL model on the square lattice under the same conditions. This sets the stage for calculating critical exponents and further investigating the universality class of the AB tiling FPL model.

8.2.1 Geometric critical exponents of Loops

In many systems, especially in statistical mechanics, a system undergoes a phase transition at the critical point. The fixed point at criticality is a fixed point of the renormalization group (RG) flow and exhibits scale invariance. At this point, the system behaves similarly at all length scales, which is a key feature of critical phenomena. Systems that flow to the same fixed point under RG transformations are said to belong to the same universality class. This implies that, despite differences in microscopic details (such as lattice structure or interaction types), the large-scale behavior (e.g., critical exponents) of these systems will be the same.

For systems with continuous scale symmetry, at the critical fixed point, if we apply an RG transformation by scaling the system by a factor b , the observable should behave as,

$$O(x) = b^{-\alpha} O(bx). \quad (8.1)$$




Figure 8.4. Colormap of the net height variance (W^2) in real space (Eq. 3.7) for each plaquette, derived from an ensemble of loop configurations (at infinite T) for the U_2 AB tiling. Darker shades represent lower height variance, while lighter shades correspond to higher height variance. This spatial pattern indicates a rough phase, as the variance increases with system size when examining concentric regions from the center to the boundary.




Figure 8.5. Plot of W^2 versus $\ln(L)$, where L is the linear system size defined as the radius of the concentric circles in the U_2 region shown in (a). The linear fit demonstrates logarithmic divergence as predicted by Eq. 3.8, confirming the rough phase for the AB tiling at $T = \infty$.

In this case, b can take any value. This leads to a simple solution of Eq. 8.1: $O(x) = x^\alpha$, where the factor $b^{-\alpha}$ represents how much the observable reduces or increases depending on the scaling factor b . If $\alpha > 0$, this means that $O(x)$ decreases as we zoom out (increasing b) and increases as we zoom in (decreasing b).

Thus, at the critical point in 2D systems with continuous scale invariance, various physical quantities, such as correlation functions, follow power-law relationships characterized by critical exponents. These exponents are universal for systems within the same universality class, meaning they are independent of the specific details of the system.

For loop models, geometric critical exponents can be determined using scaling relations of various loop statistics [94]. These relations often take the form of power laws, such as:

- *Radius of Gyration (R) Scaling vs Loop Length (s):* The loop length (s) represents the number of edges in a loop, while R is defined as the radius of the smallest disk that completely encloses a loop. At the critical point, the scaling relation between these quantities is $s \sim R^{D_f}$, where D_f is the critical exponent corresponding to the fractal dimension of these geometric loop contours.
- *Area Density $P(s)$ of Loops vs Loop Length (s):* $P(s)$ denotes the total number of loops of size between s and $s + ds$, normalized by the lattice area and total number of loops. At the critical point, the scaling relation is $P(s) \sim s^{-\tau}$, where τ is the critical exponent. The exponents D_f and τ are related by the equation: $D_f(\tau - 1) = 2$.
- *Two-point Correlation Function $C(r)$:* $C(r)$ is the average probability that two sites separated by distance r belong to the same loop. At the critical point, the scaling relation is $C(r) \sim r^{-2x}$, where x is the critical exponent. The relationship between D_f , τ , and x is given by: $D_f(3 - \tau) = 2 - 2x$, which simplifies to $D_f = 2 - x$.




Figure 8.6. Log-log plot of area density $P(s)$ of loops versus loop length (s) for a square lattice of size 64×64 . The linear fit (gray line) indicates a power-law decay with a slope of $\tau \sim 2.14$.

Studies have shown that loop models on various 2D periodic lattices, including the square and honeycomb lattice, exhibit different geometric critical exponents [94]. Our motivation is to calculate these geometric critical exponents for fully packed loops on the Ammann-Beenker (AB) tiling.

For the square lattice, I successfully reproduced the $P(s)$ vs. s results, demonstrating that $P(s)$ follows a power-law decay, $P(s)_{SQ} \sim s^{-\tau}$, with $\tau = 2.136 \pm 0.0072$, as determined from its slope. Fig. 8.6 showcases this log-log behavior for a 64×64 square lattice with periodic boundary conditions. The data is based on an ensemble of fully packed loop configurations generated through 10,000 flip sweeps. This result is in close agreement with the established findings reported in Ref. [135] and aligns with the exact results provided in Refs. [94].

Now, if the system has discrete scale invariance like the AB tiling, the relevant fixed points are only invariant under a discrete, rather than continuous, set of scale transformations. Therefore, Eq. 8.1 will only be valid for a discrete set of scales, $b = \lambda^n$ for integer n . Substituting $O(x) = x^\omega$ gives [180]:

$$x^\omega = \lambda^{-n\alpha} \lambda^{n\omega} x^\omega \quad \rightarrow \quad \lambda^{n(\omega-\alpha)} = 1 = e^{2\pi i k} \quad \rightarrow \quad n\omega = n\alpha + \frac{i2\pi k}{\log(\lambda)}. \quad (8.2)$$

For simplicity, we set $n = 1$. Hence, ω is a complex critical exponent, with the imaginary part leading to log-periodic oscillations. So in systems with DSI, observables no longer follow pure power laws like $O(x) = x^\omega$. Instead, they behave like:

$$O(x) = x^{\alpha + i \frac{2\pi k}{\log \lambda}} = x^\alpha \cdot e^{i \frac{2\pi k \log x}{\log \lambda}}. \quad (8.3)$$

Now define $y = \frac{\log x}{\log \lambda}$, so the oscillatory factor becomes: $e^{i2\pi k y}$, which is explicitly 1-periodic in y . This is

why we can write the observable more generally as:

$$O(x) \sim x^\alpha \cdot P_O \left(\frac{\log x}{\log \lambda} \right), \quad (8.4)$$

where $P_O(y)$ is a periodic function of period 1. It captures the log-periodic oscillations that naturally arise in systems with discrete scale invariance.

This log-periodic structure is a distinctive signature of discrete scale invariance (DSI). The detailed explanation of these calculations and theory is also given in the review by Ref. [180]. Thus, for the AB tiling, observables such as $s(R)$, $P(s)$, and $C(r)$, defined above, will exhibit power-law forms with log-periodic modulations (as also derived here [181]). I will now numerically calculate these scaling relations for the AB tiling to determine its critical exponents.

As explained in Sec. 2.1, under a single inflation (σ), the edges of the AB tiling are scaled by the silver ratio $\delta_S = 1 + \sqrt{2}$. Both σ^{-1} and σ^{-2} generate discrete scale invariance (DSI), but σ^{-2} has a natural RG rule, which arises from its effect on vertex behavior.

The AB tiling features seven types of vertices, each defined by specific tile configurations: 3, 4, 5_A , 5_B , 6, 7, and 8, named after their coordination numbers. Under inflation, these vertices develop distinct tile neighborhoods. The action of σ alters vertex types and modifies edges, while σ^2 maps vertices of all coordination numbers onto 8-vertices. After two deflations of a AB tiling, only 8-vertices remain, forming a new tiling, $\sigma^{-2}(G)$, with edge lengths scaled by δ_S^2 . Thus, the process of removing all vertices except the 8-vertices yields a δ_S^2 -rescaled AB tiling and forms the basis for the RG transformation [181]. So, at the critical point with discrete scale symmetry, the AB system exhibits invariance under scaling by factors of $\lambda = \delta_S^2$. Importantly, this process is simpler under σ^{-2} , whereas such a simplification is not possible with σ^{-1} .

Using Eq. 8.4, the scaling of the observable $s(R)$ as a function of the radius R is described by:

$$s(R) = R^{D_f} \cdot P_{D_f} \left(\frac{\log R}{\log \delta_S^2} \right), \quad (8.5)$$

where P_{D_f} is a log-periodic function that encodes oscillatory behavior due to the discrete nature of the symmetry. This implies that $s(R)$ follows a power-law behavior modulated by a periodic function arising from the logarithmic scaling dependence. The finite-size scaling hypothesis extends this concept to finite systems with linear dimension L , as expressed in [181]:

$$s(R, L) = R^{D_f} \cdot F_{D_f}(R/L), \quad (8.6)$$

where $F_{D_f}(R/L)$ is the scaling function. For large systems ($R/L \ll 1$), $F_{D_f}(R/L)$ resembles the log-periodic function $P_{D_f} \left(\frac{\log R}{\log \delta_S^2} \right)$, reflecting critical scaling behavior. However, for small systems or near boundaries ($R/L \sim 1$), $F_{D_f}(R/L)$ accounts for finite-size and boundary effects, deviating from pure critical scaling. If we plot sR^{-D_f} for systems with different sizes L , these data sets should collapse onto the same scaling pattern only if the systems are related by discrete scaling transformations. Fig. 8.7 shows the histogram of the radius of gyration (R) versus loop lengths (s) for U_1 and U_2 regions of the AB tiling. The regions U_1 and U_2 are the twice- and four-times inflated AB tilings, containing 464 and 14,992 vertices, respectively. The results reveal a power-law behavior, with the slope of the linear plot yielding the fractal dimension $D_f = 1.568 \pm 0.00399$. The data for $s(R)$ are averaged over samples obtained from 10,000 and 100,000 flip sweeps for U_1 and U_2 , respectively. The linear sizes of the U_1 and U_2 regions of the AB tiling were determined by calculating the graph distance between the two farthest points along the diameter of the regions, resulting in $L = 13$ and $L = 76$, respectively.




Figure 8.7. The average loop size $s(R)$ plotted against its radius R on a log-log scale shows power-law behavior with weak periodic corrections. When $s(R)R^{-d_f}$ is plotted against R/L , the data for different system sizes collapse onto a finite-size scaling form (Eq. 8.6), with a fractal dimension $d_f = 1.568$, revealing weak log-periodic modulations to the power-law behavior.

In Fig. 8.7, plotting $s(R)R^{-D_f}$ against R/L , the data for different sizes collapse to a finite-size scaling form (Eq. 8.6) with a fractal dimension $D_f = 1.568$. This reveals weak log-periodic modulations superimposed on the power-law behavior.

We can apply a similar argument to derive the scaling form of the observable $P(s)$, which represents the area density of loops as a function of loop lengths s [181]. The number of loops N within a patch of area A must remain invariant under a rescaling $R \rightarrow Rb^{-1}$. From Eq. 8.5, this implies a scaling transformation $s \rightarrow s' = s/b^{D_f}$, where $b = \delta_s^2$ is the discrete scaling factor. Following this argument, we can derive:

$$P(s)A ds = P\left(\frac{s}{b^{D_f}}\right) Ab^{-2} \frac{ds}{b^{D_f}}, \quad (8.7)$$

$$P(s) = P\left(\frac{s}{b^{D_f}}\right) b^{-2-D_f}, \quad (8.8)$$

$$\text{Substituting } b = \delta_s^2 \implies P(s) = P\left(\frac{s}{\delta_s^{2D_f}}\right) (\delta_s^2)^{-2-D_f}. \quad (8.9)$$

Here, $P(s)$ is the area density of loops of size s , and the factor b^{-2-D_f} accounts for changes in both the area and scaling properties of the system. Comparing this with Eqs. 8.1 and 8.4, we observe that $P(s)$ follows a power-law behavior modulated by log-periodic oscillations:

$$P(s) \sim s^{-\tau} P_P\left(\frac{\log s}{\log \delta_s^{2D_f}}\right), \quad (8.10)$$

where the scaling relation $-2 - D_f = -\tau D_f$ leads to $D_f(\tau - 1) = 2$.

For finite systems of size L , the probability distribution $P(s)$ must depend on both s and L . Extending the




Figure 8.8. $P(s)$, the area density of loops plotted against their size s on a log-log scale, exhibits power-law behavior with periodic corrections. When $P(s)s^\tau$ is plotted against sL^{-D_f} , the data for different system sizes collapse onto a finite-size scaling form, as per Eq. 8.11, with $\tau = 2.282$ and $D_f = 1.56$, revealing log-periodic modulations to the power-law behavior.

scaling hypothesis [181], we write:

$$P(s, L) = s^{-\tau} F_P(sL^{-D_f}), \quad (8.11)$$

where $F_P(sL^{-D_f})$ captures finite-size effects. In the regime $sL^{-D_f} \ll 1$, F_P reduces to the log-periodic function $P_P\left(\frac{\log s}{\log \delta_s^{2D_f}}\right)$, consistent with the critical scaling behavior.

Fig. 8.8 presents the histogram of the area density of loops $P(s)$ versus loop lengths s for the U_1 and U_2 regions of the AB tiling. The results exhibit a power-law behavior, with the slope of the linear region of the plot yielding $\tau = 2.269 \pm 0.0842$. These data for $P(s)$ are averaged over samples obtained from 13,000 and 120,000 flip sweeps for U_1 and U_2 , respectively. In Fig. 8.8, by plotting $s^\tau P(s)$ against sL^{-D_f} , data from different system sizes collapse onto a universal finite-size scaling form (Eq. 8.6), with a fractal dimension $D_f = 1.568$ and $\tau = 2.269$. This scaling reveals weak log-periodic modulations superimposed on the power-law behavior. Moreover, using the analytical relation $\tau = \frac{2}{D_f} + 1$, I calculate $\tau = 2.282$, which is in agreement with the numerical estimate.

Fig. 8.9 presents the two-point Correlation Function $C(r)$ as a function of distance r for the U_1 and U_2 regions of the AB tiling. Using the relation $D_f = 2 - x$, I calculate the critical exponent as $x = 2 - 1.568 = 0.432$. The numerical results also demonstrate a power-law behavior, with the slope of the linear region of the plot yielding $2x = 0.7842 \pm 0.03$. The data for $C(r)$ are averaged over samples obtained from 4000 and 20,000 flip sweeps for U_1 and U_2 , respectively, giving $x = 0.3921 \pm 0.115$, which shows only a slight discrepancy compared to the analytical value.




Figure 8.9. Two-point correlation function $C(r)$ plotted on a log-log scale against the separation between two nodes r in real space for the U_1 and U_2 AB tiling regions. The linear fit gives the critical exponent x and reveals the power-law decay.

8.3 Conclusion

In this chapter, I demonstrate that the fully packed loop (FPL) model, with loops as height contours, on the aperiodic Ammann–Beenker (AB) tiling exhibits a critical rough phase at infinite temperature. A key result is that the geometric critical exponents for the FPL model on the AB tiling are $\tau = 2.282$, $D_f = 1.568$, and $x = 0.432$. In contrast, the analogous loop–height construction on the square lattice exhibits a smooth phase. This clear distinction provides strong evidence that the FPL model on the AB tiling belongs to a different universality class than its counterparts on periodic lattices.

These findings offer a deeper understanding of the interplay between loop configurations and height representations in aperiodic systems, and they highlight the critical behavior that emerges in non-periodic geometries such as the AB tiling.

Conclusion and Future Directions

This thesis, while exploring a range of concepts and seemingly diverse topics, is unified by a central objective: to extend the well-established research on strongly correlated physical models and mathematical problems from periodic systems to more general aperiodic and random settings. To pursue this aim, we have studied a variety of models in different contexts. In this chapter, I present a summary of the key findings and outline promising directions for future research inspired by this work.

9.1 Summary

We began by establishing the theoretical framework necessary for understanding various models studied and the aperiodic systems employed. In Chapter 2, I introduced the different aperiodic tilings explored in this thesis, detailing their construction methods, inflation rules, and key structural properties. In Chapter 3, I presented the various strongly correlated physical models studied in later chapters, including classical and quantum dimer models, loop models, and the corresponding height models.

Chapter 4 shifted the focus to a set of mathematical and computational problems in complexity theory. Although many of these problems have only been solved in small special classes of graphs, we extended the analysis to large-scale graphs, specifically the aperiodic Ammann–Beenker (AB) tiling. We provided a constructive proof that Hamiltonian cycles exist on the AB tiling for any finite size, thereby solving the problem in the thermodynamic limit. The discrete scale invariance and hierarchical structure of the AB tiling played a crucial role in enabling this proof. Building on the intermediate results of this construction and additional arguments, we further solved five more nontrivial NP-hard problems on the AB tiling. We also discussed potential physical applications of these results in ultra-high resolution atomic imaging, adsorption, catalysis, and protein folding.

In Chapter 5, I numerically studied the classical dimer model on general random and aperiodic graphs using the worm Monte Carlo algorithm. I described the method for constructing these random graphs and the specific implementation of the worm algorithm used in our simulations. The results for $\langle NFP \rangle$, energy cumulants, and specific heat indicate that no true phase transition occurs as a function of temperature. Instead, we observe only a crossover from a high-temperature disordered regime to a low-temperature ordered regime. The monomer–monomer correlation functions and the variance of the height field further suggest that, from a field-theoretic perspective, the system remains confined across the temperature range. Interestingly, these findings closely match those obtained for the square–octagon lattice and stand in sharp contrast to the square lattice case, where a Kosterlitz–Thouless (KT) phase transition is known to occur, as discussed in Refs. [26].

Hence, we conclude that the square–octagon lattice may actually represent the more generic case for the classical dimer model with aligning interactions, rather than the square lattice.

In Chapter 6, I explained how to obtain an exact solution of the classical and quantum dimer model on the recently discovered aperiodic Spectre monotiling. This is the first instance of exactly solving the classical dimer model on any aperiodic tiling, which is usually very challenging even on periodic lattices. In previous studies, only the QDM on the Kagome lattice—and more generally on corner-sharing lattices—has been studied by Misguich and collaborators [182]. They introduced an *arrow representation* that allows exact counting of dimer coverings; for Kagome, the number of configurations is $Z = 2^{N/3} + 1$. This shows that the special connectivity of these aperiodic tilings can lead to exact solutions of physical problems that remain unsolved even on most periodic tilings. We also observed an extended deconfined quantum phase on this Spectre monotiling.

In Chapter 7, I studied the ground state phase diagram of the quantum dimer model on the AB* tiling using the numerical variational method of DMRG. While DMRG is typically well-suited for one-dimensional systems, it generally fails for two-dimensional systems due to the area law of entanglement entropy. In Ref. [167], a modified DMRG approach was developed that explicitly incorporates quantum dimer constraints, allowing for the study of quasi-1D square ladders of significant size. However, applying DMRG to large-scale, truly two-dimensional quantum systems has remained a challenge. The AB* tiling, due to its unique connectivity and structural properties, can be decomposed into independent loops and ladder-like substructures. This decomposition enables the application of DMRG independently to each of these ladders, allowing us to construct the full phase diagram across a wide parameter range for the large AB* tiling system. In fact, we were able to set a benchmark by successfully applying DMRG on sufficiently large AB* ladders to obtain results for a two-dimensional AB* system comprising on the order of 10^6 sites. I also discussed the different phases identified and characterized the various types of phase transitions observed between them.

Finally, in Chapter 8, I studied the fully packed loop (FPL) model on the AB tiling. I first explained how to construct random FPL configurations on the AB tiling in the thermodynamic limit. Using height model analysis, we demonstrated that the FPL model is critical and in a rough phase at infinite temperature. So I computed the geometric critical exponents and compared them with those of the square lattice FPL model at infinite temperature. The differences indicate that the AB FPL model belongs to a distinct universality class.

9.2 Outlook

While significant progress has been made, the journey ahead remains exciting, with much more to explore and accomplish. As hinted at the end of Chapter 4, one intriguing direction is to investigate whether the constructive approach for Hamiltonian cycles on the AB tiling can be extended to other aperiodic infinite graphs that also exhibit discrete scale invariance. Moreover, the fully packed loop (FPL) model developed on the AB tiling suggests the possibility of enumerating these loops, which could, in turn, enable an exact solution of the Ising model on the AB tiling, with spins placed at the centers of AB faces—analogous to Onsager’s solution for the square lattice [15, 16].

On the experimental front, several promising avenues also emerge from this work, particularly in exploring applications such as adsorption and catalysis. Physical icosahedral quasicrystals (iQCs) exhibit properties like brittleness, thermal stability, and aperiodic surface structures, making them ideal candidates as adsorbing substrates [183]. Their non-periodic surface atoms allow for a wider range of bond angles, and iQC nanoparticles have been successfully coated with atomic monolayers — including catalytically active metals like silver — which inherit the quasicrystalline order [184, 34]. This opens a promising path for future experimental investigations.

Alternatively, the theoretical results presented here may find direct application in known two-dimensional ‘axial’ quasicrystals that do exhibit AB symmetries [41, 42, 185].

Building on the numerical results from Chapter 5, in our forthcoming analytical work, we aim to rigorously generalize Polyakov’s theorem to arbitrary spacetimes of the form $M^2 \times T$, independent of the genus of the two-dimensional manifold M^2 . This proof will show that the theorem relies solely on the regularity of the lattice — a crucial condition for preserving locality in the resulting field theory. Furthermore, we will derive the effective Sine-Gordon field theory (EFT) for the most general classical dimer model and develop a mathematical framework that captures the nature of columnar and staggered states in these generalized, disordered settings.

Continuing the work presented in Chapter 7, an important future direction is to determine the universality classes of the observed second-order phase transitions. Additionally, understanding how all five identified phases converge at the center of the phase diagram is crucial — particularly whether this convergence gives rise to an emergent higher-order symmetry. We can also investigate how varying the parameter v_{lr} affects the phase diagram, potentially leading to a three-dimensional phase diagram.

Bibliography

- [1] IR Fisher, Z Islam, J Zarestky, C Stassis, MJ Kramer, AI Goldman, and PC Canfield. Magnetic properties of icosahedral R–Mg–Zn quasicrystals (R= Y, Tb, Dy, Ho and Er). *Journal of Alloys and Compounds*, 303:223–227, 2000.
- [2] Dan Shechtman, Ilan Blech, Denis Gratias, and John W Cahn. Metallic phase with long-range orientational order and no translational symmetry. *Physical review letters*, 53(20):1951, 1984.
- [3] Shobhna Singh, Jerome Lloyd, and Felix Flicker. Hamiltonian cycles on Ammann-Beenker tilings. *Physical Review X*, 14(3):031005, 2024.
- [4] Felix Flicker, Steven H Simon, and SA Parameswaran. Classical dimers on Penrose tilings. *Physical Review X*, 10(1):011005, 2020.
- [5] David Smith, Joseph Samuel Myers, Craig S Kaplan, and Chaim Goodman-Strauss. An aperiodic monotile. *arXiv preprint arXiv:2303.10798*, 2023.
- [6] David Smith, Joseph Samuel Myers, Craig S Kaplan, and Chaim Goodman-Strauss. A chiral aperiodic monotile. *arXiv preprint arXiv:2305.17743*, 2023.
- [7] R Moessner, Shivaji Lal Sondhi, and P Chandra. Phase diagram of the hexagonal lattice quantum dimer model. *Physical Review B*, 64(14):144416, 2001.
- [8] Olav F Syljuåsen. Plaquette phase of the square-lattice quantum dimer model: Quantum Monte Carlo calculations. *Physical Review B—Condensed Matter and Materials Physics*, 73(24):245105, 2006.
- [9] Arnaud Ralko, Didier Poilblanc, and R Moessner. Generic mixed columnar-plaquette phases in Rokhsar-Kivelson models. *Physical review letters*, 100(3):037201, 2008.
- [10] Jerome Lloyd, Sounak Biswas, Steven H Simon, S A Parameswaran, and Felix Flicker. Statistical mechanics of dimers on quasiperiodic Ammann-Beenker tilings. *Physical Review B*, 106:094202, 2022.
- [11] Shobhna Singh and Felix Flicker. Exact solution to the quantum and classical dimer models on the spectre aperiodic monotiling. *Physical Review B*, 109(22):L220303, 2024.
- [12] Ulrich Schollwöck. The density-matrix renormalization group. *Reviews of modern physics*, 77(1):259–315, 2005.
- [13] J. K Roberts. *Proc. Roy. Soc. (London)*, A 152:469, 1935.
- [14] RH Fowler and GS Rushbrooke. An attempt to extend the statistical theory of perfect solutions. *Transactions of the Faraday Society*, 33:1272–1294, 1937.

- [15] Lars Onsager. Crystal statistics. I. A two-dimensional model with an order-disorder transition. *Physical Review*, 65(3-4):117, 1944.
- [16] Richard P Feynman. *Statistical mechanics: a set of lectures*. CRC press, 2018.
- [17] Hue Sun Chan and Ken A Dill. Compact polymers. *Macromolecules*, 22(12):4559–4573, 1989.
- [18] PJ Flory. Statistical thermodynamics of semi-flexible chain molecules. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 234(1196):60–73, 1956.
- [19] Daniel S. Rokhsar and Steven A. Kivelson. Superconductivity and the Quantum Hard-Core Dimer Gas. *Phys. Rev. Lett.*, 61:2376–2379, Nov 1988.
- [20] R. Moessner and K. S. Raman. *Quantum Dimer Models*. Lecture Notes, Trieste, 2007.
- [21] P. W. Anderson. The resonating valence bond state in La_2CuO_4 and superconductivity. *Science*, 235(4793):1196–1198, 1987.
- [22] Cyril Domb. *Phase transitions and critical phenomena*, volume 19. Elsevier, 2000.
- [23] Kenneth G Wilson. The renormalization group and critical phenomena. *Reviews of Modern Physics*, 55(3):583, 1983.
- [24] John Cardy. *Scaling and renormalization in statistical physics*, volume 5. Cambridge university press, 1996.
- [25] Fabien Alet, Jesper Lykke Jacobsen, Grégoire Misguich, Vincent Pasquier, Frédéric Mila, and Matthias Troyer. Interacting classical dimers on the square lattice. *Physical review letters*, 94(23):235702, 2005.
- [26] Fabien Alet, Yacine Ikhlef, Jesper Lykke Jacobsen, Grégoire Misguich, and Vincent Pasquier. Classical dimers with aligning interactions on the square lattice. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, 74(4):041124, 2006.
- [27] Michael E. Fisher. Statistical mechanics of dimers on a plane lattice. *Phys. Rev.*, 124:1664–1672, Dec 1961.
- [28] P. W. Kasteleyn. Dimer statistics and phase transitions. *Journal of Mathematical Physics*, 4(2):287–293, 1963.
- [29] Xiao-Ping Liu, Wen-Chao Chen, Yi-Fei Wang, and Chang-De Gong. Topological quantum phase transitions on the kagomé and square–octagon lattices. *Journal of Physics: Condensed Matter*, 25(30):305602, 2013.
- [30] Ruben Verresen, Frank Pollmann, and Roderich Moessner. Quantum dynamics of the square-lattice Heisenberg model. *Physical Review B*, 98(15):155102, 2018.
- [31] Lukas Janssen, Eric C Andrade, and Matthias Vojta. Honeycomb-lattice heisenberg-kitaev model in a magnetic field: Spin canting, metamagnetism, and vortex crystals. *Physical review letters*, 117(27):277202, 2016.
- [32] Cory R Dean, Lei Wang, P Maher, Carlos Forsythe, Fereshte Ghahari, Y Gao, Jyoti Katoch, Masa Ishigami, Pilkyung Moon, Mikito Koshino, et al. Hofstadter’s butterfly and the fractal quantum hall effect in moiré superlattices. *Nature*, 497(7451):598–602, 2013.
- [33] P.J. Steinhardt and S. Ostlund. *The Physics of Quasicrystals*. World Scientific (Singapore), 1987.

- [34] R McGrath, JA Smerdon, HR Sharma, Wolfgang Theis, and J Ledieu. The surface science of quasicrystals. *Journal of Physics: Condensed Matter*, 22(8):084022, 2010.
- [35] Jean-Marie Dubois. New prospects from potential applications of quasicrystalline materials. *Materials Science and Engineering: A*, 294:4–9, 2000.
- [36] Yu Zou, Pawel Kuczera, Alla Sologubenko, Takashi Sumigawa, Takayuki Kitamura, Walter Steurer, and Ralph Spolenak. Superior room-temperature ductility of typically brittle quasicrystals at small sizes. *Nature communications*, 7(1):12261, 2016.
- [37] Edy Widjaja. *Quasicrystalline thin films: Growth, structure and interface*. Northwestern University, 2004.
- [38] Grace M. Sommers, Michael J. Gullans, and David A. Huse. Self-dual quasiperiodic percolation. *arXiv:2206.11290 [cond-mat.stat-mech]*, 2023.
- [39] Doruk Efe Gökmen, Sounak Biswas, Sebastian D. Huber, Zohar Ringel, Felix Flicker, and Maciej Koch-Janusz. Machine learning assisted discovery of exotic criticality in a planar quasicrystal. *arXiv:2301.11934 [cond-mat.stat-mech]*, 2023.
- [40] Matteo Ciardi, Adriano Angelone, Fabio Mezzacapo, and Fabio Cinti. Quasicrystalline bose glass in the absence of disorder and quasidisorder. *Phys. Rev. Lett.*, 131:173402, Oct 2023.
- [41] Nanlin Wang, Hanyuan Chen, and KH Kuo. Two-dimensional quasicrystal with eightfold rotational symmetry. *Physical review letters*, 59(9):1010, 1987.
- [42] N. Wang, K. K. Fung, and K. H. Kuo. *Appl. Phys. Lett.*, 52:2120, 1988.
- [43] Aviram Uri, Sergio C de la Barrera, Mallika T Randeria, Daniel Rodan-Legrain, Trithep Devakul, Philip JD Crowley, Nisarga Paul, Kenji Watanabe, Takashi Taniguchi, Ron Lifshitz, et al. Superconductivity and strong interactions in a tunable moiré quasicrystal. *Nature*, 620(7975):762–767, 2023.
- [44] Konrad Viebahn, Matteo Sbroscia, Edward Carter, Jr-Chiun Yu, and Ulrich Schneider. Matter-wave diffraction from a quasicrystalline optical lattice. *Physical Review Letters*, 122(11):110404, 2019.
- [45] Stefan Wessel, Anuradha Jagannathan, and Stephan Haas. Quantum antiferromagnetism in quasicrystals. *Physical review letters*, 90(17):177205, 2003.
- [46] Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002.
- [47] Matthias Troyer and Uwe-Jens Wiese. Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations. *Phys. Rev. Lett.*, 94:170201, May 2005.
- [48] Cubitt, Toby S. and Pérez-García, David and Wolf, Michael M. Undecidability of the spectral gap. *Nature*, 528:207–211, 2015.
- [49] Michael R Garey, David S Johnson, and Larry Stockmeyer. Some simplified NP-complete problems. In *Proceedings of the sixth annual ACM symposium on Theory of computing*, pages 47–63, 1974.
- [50] Dov Levine and Paul Joseph Steinhardt. Quasicrystals: a new class of ordered structures. *Physical review letters*, 53(26):2477, 1984.
- [51] M. Senechal. *Quasicrystals and Geometry*. Cambridge University Press, 1995.

- [52] Michael Baake and Uwe Grimm. *Aperiodic order*, volume 1. Cambridge University Press, 2013.
- [53] B. Grünbaum and G. C. Shephard. *Tilings and Patterns*. W. H. Freeman and Company, New York, 1986.
- [54] Luca Bindi, Paul J Steinhardt, Nan Yao, and Peter J Lu. Natural quasicrystals. *science*, 324(5932):1306–1309, 2009.
- [55] Enrique MacIá. The role of aperiodic order in science and technology. *Reports on Progress in Physics*, 69(2):397, 2005.
- [56] Jan Fikar. Al-Cu-Fe quasicrystalline coatings and composites studied by mechanical spectroscopy. Technical report, EPFL, 2003.
- [57] Jin Jung Kweon, Hyo-Im Kim, Sang-hwa Lee, Jaeyong Kim, and Sung Keun Lee. Quantitative probing of hydrogen environments in quasicrystals by high-resolution nmr spectroscopy. *Acta Materialia*, 226:117657, 2022.
- [58] Matthew Kalman. The quasicrystal laureate. *Massachusetts Institute of Technology (MIT) Technology Review, Cambridge*, 2011.
- [59] H Bakhtiari, MR Rahimipour, M Farvizi, and MR Khanzadeh. An overview of quasicrystals, their types, preparation methods, properties. *Journal of Environmental Friendly Materials*, 5(1):69–76, 2021.
- [60] Doruk Efe Gökmen, Sounak Biswas, Sebastian D Huber, Zohar Ringel, Felix Flicker, and Maciej Koch-Janusz. Compression theory for inhomogeneous systems. *Nature Communications*, 15(1):10214, 2024.
- [61] F. Flicker and Jasper van Wezel. Quasiperiodicity and 2D topology in 1D charge-ordered materials. *Europhysics Letters*, 111(3):37008, 2015.
- [62] Robert Ammann, Branko Grünbaum, and Geoffrey C Shephard. Aperiodic tiles. *Discrete & Computational Geometry*, 8(1):1–25, 1992.
- [63] Rima Al Ajlouni. Octagon-based quasicrystalline formations in islamic architecture. In *Aperiodic crystals*, pages 49–57. Springer, 2013.
- [64] Luca Bindi, William Kolb, G Nelson Eby, Paul D Asimow, Terry C Wallace, and Paul J Steinhardt. Accidental synthesis of a previously unknown quasicrystal in the first atomic bomb test. *Proceedings of the National Academy of Sciences*, 118(22):e2101350118, 2021.
- [65] Martin Gardner. *Penrose tiles to trapdoor ciphers: And the return of Dr Matrix*. Cambridge University Press, 1997.
- [66] William Thomas Tutte and William Thomas Tutte. *Graph theory*, volume 21. Cambridge University Press, 2001.
- [67] P.W. Kasteleyn. The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice. *Physica*, 27(12):1209 – 1225, 1961.
- [68] H. N. V. Temperley and Michael E. Fisher. Dimer problem in statistical mechanics-an exact result. *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics*, 6(68):1061–1063, 1961.

- [69] H Christopher Longuet-Higgins. Some studies in molecular orbital theory I. Resonance structures and molecular orbitals in unsaturated hydrocarbons. *Journal of Chemical Physics*, 18(3):265–274, 1950.
- [70] Elliott H. Lieb. Two theorems on the Hubbard model. *Phys. Rev. Lett.*, 62:1201–1204, Mar 1989.
- [71] Rodney J Baxter. *Exactly solved models in statistical mechanics*. Elsevier, 2016.
- [72] R. Moessner and S. L. Sondhi. Resonating valence bond phase in the triangular lattice quantum dimer model. *Phys. Rev. Lett.*, 86:1881–1884, 2001.
- [73] R. Moessner and S. L. Sondhi. *Progress of Theoretical Physics Supplement*, 145:37–42, 2002.
- [74] Lucile Savary and Leon Balents. Quantum spin liquids: a review. *Reports on Progress in Physics*, 80(1):016502, 2017.
- [75] Christopher L Henley. The “coulomb phase” in frustrated systems. *Annu. Rev. Condens. Matter Phys.*, 1(1):179–210, 2010.
- [76] KJ Satzinger, Y-J Liu, A Smith, C Knapp, M Newman, C Jones, Z Chen, C Quintana, X Mi, A Dunsworth, et al. Realizing topologically ordered states on a quantum processor. *Science*, 374(6572):1237–1241, 2021.
- [77] Giulia Semeghini, Harry Levine, Alexander Keesling, Sepehr Ebadi, Tout T Wang, Dolev Bluvstein, Ruben Verresen, Hannes Pichler, Marcin Kalinowski, Rhine Samajdar, et al. Probing topological spin liquids on a programmable quantum simulator. *Science*, 374(6572):1242–1247, 2021.
- [78] P. Kasteleyn. Graph theory and crystal physics. *Graph Theory and Theoretical Physics*, pages 43–110, 1967.
- [79] Fa Y Wu. Dimers on two-dimensional lattices. *International Journal of Modern Physics B*, 20(32):5357–5371, 2006.
- [80] Michael E Fisher and John Stephenson. Statistical mechanics of dimers on a plane lattice. ii. dimer correlations and monomers. *Physical Review*, 132(4):1411, 1963.
- [81] Subir Sachdev. Spin-Peierls ground states of the quantum dimer model: A finite-size study. *Physical Review B*, 40(7):5204, 1989.
- [82] Eduardo Fradkin, David A. Huse, R. Moessner, V. Oganesyan, and S. L. Sondhi. Bipartite Rokhsar–Kivelson points and Cantor deconfinement. *Phys. Rev. B*, 69:224415, Jun 2004.
- [83] H Eugene Stanley. Dependence of critical properties on dimensionality of spins. *Physical Review Letters*, 20(12):589, 1968.
- [84] HWJ Blöte and B Nienhuis. Fully packed loop model on the honeycomb lattice. *Physical review letters*, 72(9):1372, 1994.
- [85] Geoffrey Grimmett. *The random-cluster model*, volume 333. Springer, 2006.
- [86] Jesper Lykke Jacobsen. Exact enumeration of Hamiltonian circuits, walks and chains in two and three dimensions. *Journal of Physics A: Mathematical and Theoretical*, 40(49):14667, 2007.
- [87] Jesper Lykke Jacobsen. On the universality of fully packed loop models. *Journal of Physics A: Mathematical and General*, 32(29):5445, 1999.

- [88] Jané Kondev. Liouville field theory of fluctuating loops. *Physical review letters*, 78(23):4320, 1997.
- [89] Murray T Batchelor, J Suzuki, and CM Yung. Exact results for Hamiltonian walks from the solution of the fully packed loop model on the honeycomb lattice. *Physical review letters*, 73(20):2646, 1994.
- [90] Junji Suzuki. Evaluation of the connectivity of Hamiltonian paths on regular lattices. *Journal of the Physical Society of Japan*, 57(3):687–690, 1988.
- [91] Rainer Liebmann. Statistical mechanics of periodic frustated ising systems. *Lecture notes in physics*, 251, 1986.
- [92] Jané Kondev and Christopher L Henley. Kac-moody symmetries of critical ground states. *Nuclear Physics B*, 464(3):540–575, 1996.
- [93] R Raghavan, Christopher L Henley, and Scott L Arouh. New two-color dimer models with critical ground states. *Journal of statistical physics*, 86:517–550, 1997.
- [94] Jané Kondev and Christopher L Henley. Geometrical exponents of contour loops on random gaussian surfaces. *Physical review letters*, 74(23):4580, 1995.
- [95] Hao Li, Robert Helling, Chao Tang, and Ned Wingreen. Emergence of Preferred Structures in a Simple Model of Protein Folding. *Science*, 273(5275):666–669, 1996.
- [96] Gérard Jannink. *Polymers in Solution: Their Modelling and Structure*. Clarendon Press, 1990.
- [97] Olga Bodroza-Pantic, Bojana Pantic, Ilija Pantic, and Marija Bodroza-Solarov. Enumeration of Hamiltonian cycles in some grid graphs. *MATCH Commun. Math. Comput. Chem*, 70(1):181–204, 2013.
- [98] A Yu Kitaev. Fault-tolerant quantum computation by anyons. *Annals of physics*, 303(1):2–30, 2003.
- [99] MT Batchelor, HWJ Blöte, B Nienhuis, and CM Yung. Critical behaviour of the fully packed loop model on the square lattice. *Journal of Physics A: Mathematical and General*, 29(16):L399, 1996.
- [100] Alexey V Melkikh and Dirk KF Meijer. On a generalized Levinthal’s paradox: The role of long-and short range interactions in complex bio-molecular reactions, including protein and DNA folding. *Progress in Biophysics and Molecular Biology*, 132:57–79, 2018.
- [101] F. Flicker and Jasper van Wezel. Natural 1D Quasicrystals from Incommensurate Charge Order. *Physical Review Letters*, 115:236401, 2015.
- [102] Yaacov E. Kraus, Yoav Lahini, Zohar Ringel, Mor Verbin, and Oded Zilberberg. Topological states and adiabatic pumping in quasicrystals. *Phys. Rev. Lett.*, 109:106402, Sep 2012.
- [103] Kevin A. Madsen, Emil J. Bergholtz, and Piet W. Brouwer. Topological equivalence of crystal and quasicrystal band structures. *Phys. Rev. B*, 88:125118, Sep 2013.
- [104] Konrad Viebahn, Matteo Sbroscia, Edward Carter, Jr-Chiun Yu, and Ulrich Schneider. Matter-Wave Diffraction from a Quasicrystalline Optical Lattice. *Phys. Rev. Lett.*, 122:110404, Mar 2019.
- [105] Martin Gardner. Mathematical games: About the remarkable similarity between the icosian game and the towers of Hanoi. *Sci. Amer*, 196:150–156, 1957.

- [106] John Adrian Bondy, Uppaluri Siva Ramachandra Murty, et al. *Graph theory with applications*, volume 290. Macmillan London, 1976.
- [107] Robert W. Robinson and Nicholas C. Wormald. Almost all cubic graphs are Hamiltonian. *Random Structures & Algorithms*, 3(2):117–125, 1992.
- [108] Michael Baake. A guide to mathematical quasicrystals. In *Quasicrystals: An introduction to structure, physical properties and applications*, pages 17–48. Springer, 2002.
- [109] Fang Fang, Sinziana Paduroiu, Dugan Hammock, and Klee Irwin. Non-local game of life in 2D quasicrystals. *Crystals*, 8(11):416, 2018.
- [110] N. W. Ashcroft and N. D. Mermin. *Solid State Physics*. Harcourt College Publishers, New York, 1976.
- [111] Gerd Binnig and Heinrich Rohrer. Scanning tunneling microscopy. *Surface science*, 126(1-3):236–244, 1983.
- [112] Paul K Hansma and Jerry Tersoff. Scanning tunneling microscopy. *Journal of Applied Physics*, 61(2):R1–R24, 1987.
- [113] Y Martin and H Kumar Wickramasinghe. Magnetic imaging by “force microscopy” with 1000 Å resolution. *Applied Physics Letters*, 50(20):1455–1457, 1987.
- [114] Matteo Sbroscia, Konrad Viebahn, Edward Carter, Jr-Chiun Yu, Alexander Gaunt, and Ulrich Schneider. Observing Localization in a 2D Quasicrystalline Optical Lattice. *Phys. Rev. Lett.*, 125:200604, Nov 2020.
- [115] O. J. Heilmann and E. H. Lieb. Theory of monomer-dimer systems. *Commun. Math. Phys.*, 25:190–232, 1972.
- [116] P. A. Webley and D. Danaci. *CO2 capture by adsorption processes*, pages 106–167. Number 26 in RSC Energy and Environment Series. The Royal Society of Chemistry, United Kingdom, 26 edition, 2020.
- [117] Percy John Heawood. Map color theorems. *Quant. J. Math.*, 24:332–338, 1890.
- [118] K. Appel and W. Haken. Every planar map is four colorable. *Contemporary Mathematics*, 98, 1989.
- [119] Tom Sibley and Stan Wagon. Rhombic Penrose tilings can be 3-colored. *The American Mathematical Monthly*, 107(3):251–253, 2000.
- [120] Fa-Yueh Wu. The Potts model. *Reviews of modern physics*, 54(1):235, 1982.
- [121] Vincent Beffara and Hugo Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probability Theory and Related Fields*, 153(3):511–542, 2012.
- [122] Walter Selke and David A Huse. Interfacial adsorption in planar potts models. *Zeitschrift für Physik B Condensed Matter*, 50(2):113–116, 1983.
- [123] Alan D Sokal and BS Webb. The multivariate Tutte polynomial (alias Potts model). *Surveys in Combinatorics*, 327(2005):173, 2005.
- [124] M Alford, S Chandrasekharan, J Cox, and U-J Wiese. Solution of the complex action problem in the Potts model for dense QCD. *Nuclear Physics B*, 602(1-2):61–86, 2001.
- [125] François Graner and James A Glazier. Simulation of biological cell sorting using a two-dimensional extended Potts model. *Physical review letters*, 69(13):2013, 1992.

- [126] William G Wilson and Chester A Vause. Evidence for universality of the Potts model on the two-dimensional Penrose lattice. *Physics Letters A*, 126(8-9):471–475, 1988.
- [127] Gang Xiong, Zhe-Hua Zhang, and De-Cheng Tian. Real-space renormalization group approach to the Potts model on the two-dimensional Penrose tiling. *Physica A: Statistical Mechanics and its Applications*, 265(3-4):547–556, 1999.
- [128] T Garel and H Orland. Mean-field model for protein folding. *Europhysics letters*, 6(4):307, 1988.
- [129] Theja N De Silva and Vattika Sivised. A statistical mechanics perspective for protein folding from q -state potts model. *arXiv preprint arXiv:1709.04813*, 2017.
- [130] Ronald M Levy, Allan Haldane, and William F Flynn. Potts Hamiltonian models of protein co-variation, free energy landscapes, and evolutionary fitness. *Current opinion in structural biology*, 43:55–62, 2017.
- [131] Renfrey Burnard Potts. Some generalized order-disorder transformations. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 48, pages 106–109. Cambridge University Press, 1952.
- [132] Thomas E Creighton. Protein folding. *Biochemical journal*, 270(1):1, 1990.
- [133] K. A. Dill, S. Bromberg, K Yue, H. S Chan, K. M. Ftebig, D. P. Yee, and P. D. Thomas. *Protein science*, 4:561–602.
- [134] Jan de Gier. Fully packed loop models on finite geometries. In *Polygons, Polyominoes and Polycubes*, pages 317–346. Springer, 2009.
- [135] LDC Jaubert, M Haque, and R Moessner. Analysis of a fully packed loop model arising in a magnetic coulomb phase. *Physical review letters*, 107(17):177202, 2011.
- [136] Hue Sun Chan and Ken A Dill. Intrachain loops in polymers: Effects of excluded volume. *The Journal of chemical physics*, 90(1):492–509, 1989.
- [137] Kit Fun Lau and Ken A Dill. A lattice statistical mechanics model of the conformational and sequence spaces of proteins. *Macromolecules*, 22(10):3986–3997, 1989.
- [138] Pouya Baniasadi, Vladimir Ejov, Jerzy A Filar, Michael Haythorpe, and Serguei Rossomakhine. Deterministic “Snakes and Ladders” Heuristic for the Hamiltonian cycle problem. *Mathematical Programming Computation*, 6(1):55–75, 2014.
- [139] Pouya Baniasadi, Vladimir Ejov, Michael Haythorpe, and Serguei Rossomakhine. A new benchmark set for Traveling salesman problem and Hamiltonian cycle problem. *arXiv preprint arXiv:1806.09285*, 2018.
- [140] Keld Helsgaun. An effective implementation of the lin–kernighan traveling salesman heuristic. *European journal of operational research*, 126(1):106–130, 2000.
- [141] David L Applegate, Robert E Bixby, Vašek Chvátal, and William J Cook. The traveling salesman problem. In *The Traveling Salesman Problem*. Princeton university press, 2011.
- [142] Marin Bukov, Luca D’Alessio, and Anatoli Polkovnikov. Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to Floquet engineering. *Advances in Physics*, 64(2):139–226, 2015.

- [143] Christof Gattringer and Christian Lang. *Quantum chromodynamics on the lattice: an introductory presentation*, volume 788. Springer Science & Business Media, 2009.
- [144] Heinz J Rothe. *Lattice gauge theories: an introduction*. World Scientific Publishing Company, 2012.
- [145] SR Salinas and JF Nagle. Theory of the phase transition in the layered hydrogen-bonded $\text{SnCl}_2 \cdot 2\text{H}_2\text{O}$ crystal. *Physical Review B*, 9(11):4920, 1974.
- [146] Helder L Casa Grande, SR Salinas, and FA da Costa. Fermionic representation of two-dimensional dimer models. *Brazilian Journal of Physics*, 41:86–93, 2011.
- [147] John F Nagle, Carlos SO Yokoi, and Somendra M Bhattacharjee. Dimer models on anisotropic lattices. *Phase Transitions*, 13:236–296, 1989.
- [148] Christopher L Henley. Classical height models with topological order. *Journal of Physics: Condensed Matter*, 23(16):164212, 2011.
- [149] Marshall Hall. *Combinatorial theory*, volume 71. John Wiley & Sons, 1998.
- [150] OEIS Foundation Inc. The on-line encyclopedia of integer sequences, sequence a123456. <https://oeis.org/A008325>, 2025.
- [151] OEIS Foundation Inc. The on-line encyclopedia of integer sequences, sequence a123456. <https://oeis.org/A007083>, 2025.
- [152] Gunnar Brinkmann, Sam Greenberg, Catherine Greenhill, Brendan D McKay, Robin Thomas, and Paul Wollan. Generation of simple quadrangulations of the sphere. *Discrete mathematics*, 305(1-3):33–54, 2005.
- [153] Gunnar Brinkmann, Brendan McKay, et al. *Fast generation of planar graphs*. University of Kragujevac, 2007.
- [154] Gunnar Brinkmann and Brendan McKay. plantri and fullgen. <https://users.cecs.anu.edu.au/~bdm/plantri/>, 2025.
- [155] C Paul Bonnington and Charles HC Little. *The foundations of topological graph theory*. Springer Science & Business Media, 2012.
- [156] W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, s3-13(1):743–767, 1963.
- [157] Murty SS Challa, David P Landau, and K Binder. Finite-size effects at temperature-driven first-order transitions. *Physical Review B*, 34(3):1841, 1986.
- [158] Anders W Sandvik and R Moessner. Correlations and confinement in nonplanar two-dimensional dimer models. *Physical Review B—Condensed Matter and Materials Physics*, 73(14):144504, 2006.
- [159] Bhupen Dabholkar, Xiaoxue Ran, Junchen Rong, Zheng Yan, GJ Sreejith, Zi Yang Meng, and Fabien Alet. Classical fully packed loop model with attractive interactions on the square lattice. *Physical Review B*, 108(12):125112, 2023.
- [160] A.M. Polyakov. Quark confinement and topology of gauge theories. *Nuclear Physics B*, 120(3):429 – 458, 1977.

- [161] Subir Sachdev. Quantum phase transitions. *Physics world*, 12(4):33, 1999.
- [162] Markus Greiner, Olaf Mandel, Tilman Esslinger, Theodor W Hänsch, and Immanuel Bloch. Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms. *nature*, 415(6867):39–44, 2002.
- [163] Chetan Nayak, Steven H. Simon, Ady Stern, Michael Freedman, and Sankar Das Sarma. Non-abelian anyons and topological quantum computation. *Rev. Mod. Phys.*, 80:1083–1159, 2008.
- [164] Paul Fendley, Krishnendu Sengupta, and Subir Sachdev. Competing density-wave orders in a one-dimensional hard-boson model. *Physical Review B*, 69(7):075106, 2004.
- [165] Adrian Feiguin, Simon Trebst, Andreas W. W. Ludwig, Matthias Troyer, Alexei Kitaev, Zhenghan Wang, and Michael H. Freedman. Interacting anyons in topological quantum liquids: The golden chain. *Phys. Rev. Lett.*, 98:160409, Apr 2007.
- [166] Lev B Ioffe, Mikhail V Feigel'man, Alexey Ioselevich, Dmitri Ivanov, Matthias Troyer, and Gianni Blatter. Topologically protected quantum bits using josephson junction arrays. *Nature*, 415(6871):503–506, 2002.
- [167] Natalia Chepiga and Frédéric Mila. DMRG investigation of constrained models: from quantum dimer and quantum loop ladders to hard-boson and Fibonacci anyon chains. *SciPost Physics*, 6(3):033, 2019.
- [168] Hans Bethe. Zur theorie der metalle: I. Eigenwerte und eigenfunktionen der linearen atomkette. *Zeitschrift für Physik*, 71(3):205–226, 1931.
- [169] Jens Eisert, Marcus Cramer, and Martin B Plenio. Colloquium: Area laws for the entanglement entropy. *Reviews of modern physics*, 82(1):277–306, 2010.
- [170] Alexander M Dalgell and Fernando GSL Brandão. Locally accurate mps approximations for ground states of one-dimensional gapped local hamiltonians. *Quantum*, 3:187, 2019.
- [171] Karen A Hallberg. New trends in density matrix renormalization. *Advances in Physics*, 55(5-6):477–526, 2006.
- [172] Junhyun Lee, Subir Sachdev, and Steven R White. Electronic quasiparticles in the quantum dimer model: Density matrix renormalization group results. *Physical Review B*, 94(11):115112, 2016.
- [173] Yin-Chen He, Fabian Grusdt, Adam Kaufman, Markus Greiner, and Ashvin Vishwanath. Realizing and adiabatically preparing bosonic integer and fractional quantum hall states in optical lattices. *Physical Review B*, 96(20):201103, 2017.
- [174] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of physics*, 326(1):96–192, 2011.
- [175] Andreas M Läuchli, Julien Sudan, and Roderich Moessner. $S = 1/2$ kagome Heisenberg antiferromagnet revisited. *Physical Review B*, 100(15):155142, 2019.
- [176] Simeng Yan, David A Huse, and Steven R White. Spin-liquid ground state of the $S=1/2$ kagome Heisenberg antiferromagnet. *Science*, 332(6034):1173–1176, 2011.
- [177] Sebastian Depenbrock, Ian P. McCulloch, and Ulrich Schollwöck. Nature of the Spin-Liquid Ground State of the $S=1/2$ Heisenberg Model on the Kagome Lattice. *Phys. Rev. Lett.*, 109:067201, 2012.

- [178] Bernard Nienhuis. Loop models. *Exact Methods in Low-dimensional Statistical Physics and Quantum Computing: Lecture Notes of the Les Houches Summer School: Volume 89, July 2008*, page 159, 2010.
- [179] Christopher L Henley. From classical to quantum dynamics at Rokhsar–Kivelson points. *Journal of Physics: Condensed Matter*, 16(11):S891, 2004.
- [180] Didier Sornette. Discrete-scale invariance and complex dimensions. *Physics reports*, 297(5):239–270, 1998.
- [181] Sounak Biswas and SA Parameswaran. Discrete scale invariant fixed point in a quasiperiodic classical dimer model. *arXiv preprint arXiv:2302.07879*, 2023.
- [182] Misguich, G and Serban, D and Pasquier, V. Quantum dimer model on the kagome lattice: Solvable dimer-liquid and Ising gauge theory. *Physical review letters*, 89(13):137202, 2002.
- [183] D.V. Louzguine-Luzgin and A. Inoue. Formation and properties of quasicrystals. *Annual Review of Materials Research*, 38(1):403–423, 2008.
- [184] R McGrath, J Ledieu, E.J Cox, S Haq, R.D Diehl, C.J Jenks, I Fisher, A.R Ross, and T.A Lograsso. Quasicrystal surfaces: potential as templates for molecular adsorption. *Journal of Alloys and Compounds*, 342(1):432–436, 2002. Proceedings from the ‘Quasicrystals 2001’ Conference.
- [185] D. A. Rabson, N. D. Mermin, D. S. Rokhsar, and D. C. Wright. *Reviews of Modern Physics*, 63(3), 1991.