

Upper bounds on the growth rates of hard squares and related models via corner transfer matrices

Yao-ban Chan^{1†}

¹*School of Mathematics and Physics, The University of Queensland, Brisbane, Australia*

Abstract. We study the growth rate of the hard squares lattice gas, equivalent to the number of independent sets on the square lattice, and two related models — non-attacking kings and read-write isolated memory. We use an assortment of techniques from combinatorics, statistical mechanics and linear algebra to prove upper bounds on these growth rates. We start from Calkin and Wilf’s transfer matrix eigenvalue bound, then bound that with the Collatz-Wielandt formula from linear algebra. To obtain an approximate eigenvector, we use an ansatz from Baxter’s corner transfer matrix formalism, optimised with Nishino and Okunishi’s corner transfer matrix renormalisation group method. This results in an upper bound algorithm which no longer requires exponential memory and so is much faster to calculate than a direct evaluation of the Calkin-Wilf bound. Furthermore, it is extremely parallelisable and so allows us to make dramatic improvements to the previous best known upper bounds. In all cases we reduce the gap between upper and lower bounds by 4-6 orders of magnitude.

Résumé. Nous étudions le taux de croissance du système de particules dur sur un réseau carré. Ce taux est équivalent au nombre d’ensembles indépendants sur le réseau carré. Nous étudions également deux modèles qui lui sont reliés : les rois non-attaquants et la mémoire isolée d’écriture-réécriture. Nous utilisons techniques diverses issues de la combinatoire, de la mécanique statistique et de l’algèbre linéaire pour prouver des bornes supérieures sur ces taux de croissances. Nous partons de la borne de Calkin et Wilf sur les valeurs propres des matrices de transfert, que nous bornons à l’aide de la formule de Collatz-Wielandt issue de l’algèbre linéaire. Pour obtenir une valeur approchée d’un vecteur propre, nous utilisons un ansatz du formalisme de Baxter sur les matrices de transfert de coin, que nous optimisons avec la méthode de Nishino et Okunishi qui exploite ces matrices. Il en résulte un algorithme pour calculer la borne supérieure qui n’est plus exponentiel en mémoire et est ainsi beaucoup plus rapide qu’une évaluation directe de la borne de Calkin-Wilf. De plus, cet algorithme est extrêmement parallélisable et permet ainsi une nette amélioration des meilleurs bornes supérieures existantes. Dans tous les cas l’écart entre les bornes supérieures et inférieures s’en trouve réduit de 4 à 6 ordres de grandeur.

Keywords: Transfer matrix, growth rate, hard squares, upper bound

1 Introduction

We study the growth rate of the statistical mechanical model of a hard square lattice gas. In this model, the vertices of a two-dimensional square lattice are either occupied by a gas particle or vacant, denoted by ● and ○ respectively. However, no two occupied vertices may be immediate neighbours. We can think of each occupied vertex as being covered by a square (rotated 45° and

[†]E-mail: yaoban.chan@uq.edu.au

with side-length $\sqrt{2}$). The constraint forbids these squares from overlapping, giving rise to the name hard squares.

In one dimension, the hard square lattice gas can be used as a simple model of data storage on a magnetic tape. As the tape is read from left to right, a vacant vertex indicates that the stored field is unchanged, while an occupied vertex indicates that the field flips. To avoid potential *intersymbol interference*, we forbid field flips from occurring in close succession — hence $\bullet\bullet$ is forbidden [21]. Because of this connection with magnetic fields, the vertices and their states are often referred to as *spins*.

The storage capacity of a length of tape is dependent on the number of legal configurations of the spins it contains. Accordingly, we define the *partition function*, which is simply the total number of possible configurations:

$$Z_N = \# \text{ of legal configurations of } N \text{ spins.} \quad (1)$$

The partition function typically grows exponentially with respect to the number of vertices, and so a more appropriate measure for the number of possible configurations is the *partition function per site* or *growth rate*

$$\kappa = \lim_{N \rightarrow \infty} Z_N^{1/N}. \quad (2)$$

The *capacity* of the system is the number of independent bits of storage per vertex in the system. This is directly related to κ by

$$\text{capacity} = \log_2 \kappa. \quad (3)$$

While growth rates for one-dimensional systems can be readily computed as the dominant eigenvalue of a finite transfer matrix⁽ⁱ⁾, the corresponding problem for two-dimensional systems is typically very hard. Despite considerable effort, exact growth rates are only known for a very small number of systems⁽ⁱⁱ⁾. Indeed, it has been shown by Berger [5] that there exist systems for which it is undecidable if there exist any valid configurations at all. Because of this, computing growth rates has become something of a numerical challenge and has been studied extensively for many different models under numerous guises.

One of the most-studied models in this area is the aforementioned hard squares model [6, 22, 14]. This model is well-studied in statistical mechanics⁽ⁱⁱⁱ⁾, not just for the combinatorics of the number of configurations, but also for its macroscopic behaviour as the relative weighting of \bullet and \circ spins changes. Obtaining very precise upper bounds for the growth rate of this model is the main focus of this paper.

We also consider two related models which forbid certain local configurations of \bullet spins (see Figure 1). The *non-attacking kings* model forbids horizontal, vertical or diagonal adjacency of \bullet spins — they can be considered as kings on a chessboard which cannot be placed in such a way that one attacks another. The *read-write isolated memory* model [10, 16] forbids two horizontally or diagonally adjacent \bullet spins (vertical adjacency is allowed). If we consider the two-dimensional array of spins as a horizontal line of spins evolving in time (each row separated by 1 time unit), then the RWIM constraint can be viewed as forbidding us from altering two adjacent spins in one time step.

⁽ⁱ⁾ This implies that κ for a 1d system is a Perron number. Lind [20] proved that any Perron number is also the growth rate of some similar 1d system.

⁽ⁱⁱ⁾ For example, the fully-packed dimer model [18, 27], 3-colourings of the square grid [19], and the “odd” model [22].

⁽ⁱⁱⁱ⁾ It is a close relative of the famous hard hexagons model which was solved exactly by Baxter [3].




Fig. 1: The models that we study in this paper. The above diagrams show the restrictions on adjacent \bullet spins — dashed red lines are disallowed and solid black lines are allowed. The lower diagrams show example configurations. From left to right: hard squares, non-attacking kings (NAK) and read-write isolated memory (RWIM).




Fig. 2: The transfer matrix V_w extends a system of N columns to a system of $N + 1$ columns by using the state of the N th column $(\sigma_1, \dots, \sigma_w)$. The transfer matrix V_w^o constructs the same system on a cylinder.

It is not at all clear that tractable closed-form expressions for κ exist for these models, and none are known. There is a significant body of work on finding rigorous bounds on κ , most of which are based on the analysis of transfer matrices (see Figure 2). Rather than writing down a transfer matrix for an infinite two-dimensional system — an obviously impossible task — we consider the system on an infinite strip of finite width, say w vertices. Let V_w be the column transfer matrix associated with this strip, and $\Lambda(w)$ its dominant eigenvalue. Then

$$\kappa = \lim_{w \rightarrow \infty} \Lambda(w)^{1/w}. \tag{4}$$

Lower bounds for κ can then be computed using a clever formula based on Rayleigh quotients due to Engel [12], and Calkin and Wilf [6] (we refer the reader to the latter paper for its proof):

$$\kappa^p \geq \frac{\Lambda(p + 2q)}{\Lambda(2q)} \quad \text{with } p, q > 0. \tag{5}$$

Calkin and Wilf also proved the following upper bound on κ , and we again refer the reader to their paper for its proof:

$$\kappa^{2p} \leq \Lambda^o(2p) \quad \text{with } p > 0, \tag{6}$$

where $\Lambda^o(w)$ is the dominant eigenvalue of a related matrix, V_w^o , the transfer matrix for the same system but with cylindrical boundary conditions (see Figure 2).

Almost all works which compute bounds on the growth rates of the hard squares and related models use the above two inequalities (two recent exceptions being [15] and previous work by the authors [9], both of which use methods from statistical physics). However, the practical application of these inequalities is quickly hampered by the exponential growth of the size of transfer matrix with respect to strip width.^(iv) There are creative methods to avoid storing the full matrix, perhaps the most successful being a matrix compression method due to Lundow and Markström [23]. Together with Friedland [14], they used this method to compute $\Lambda(28)$ and $\Lambda^o(36)$, which in turn allowed them to exactly determine the first 15 digits of κ for the hard squares model:

$$\kappa = 1.503\,048\,082\,475\,33\dots \quad (7)$$

While this compression method greatly decreases the time and memory needed, the requirements still grow exponentially with strip width.

In this work, we compute upper bounds on the growth rates of the models. We start with equation (6), but rather than computing $\Lambda^o(w)$ exactly, we compute an upper bound on it using an approximation of the dominant eigenvector of V_w^o and the Collatz-Wielandt formula [11, 28] (which we include as Lemma 2). To form the approximate eigenvector, we use corner transfer matrix formalism. This is a very powerful approach developed in statistical mechanics by Baxter [1, 2, 4] as a way to estimate the partition function of various lattice models, both numerically and via series expansions [4, 7, 8].

The CTM formalism allows us to express each component of the approximate eigenvector as the trace of a product of auxiliary matrices. This means that we are no longer required to store either the transfer matrix or vector; we can compute the vector component by component as the Collatz-Wielandt formula requires. To optimise the choice of vector (and the associated auxiliary matrices), we use an extension of CTM known as the corner transfer matrix renormalisation group (CTMRG) method, developed by Nishino and Okunishi [24, 25].

2 Upper bounds on upper bounds

We start by restating Calkin and Wilf's upper bound on the growth rate κ .

Lemma 1 ([6]) *Let V_w^o be the column transfer matrix for a system of width w with cylindrical boundary conditions, and let $\Lambda^o(w)$ be its dominant eigenvalue. Then for any even integer $m > 0$,*

$$\kappa^m \leq \Lambda^o(m). \quad (8)$$

As noted in the introduction, almost all upper bounds in the literature are derived using Lemma 1 by exact calculation of the dominant eigenvalue of V_w^o . As such they are restricted, both in time and memory, by the exponential growth of the transfer matrix size. We also start with this bound, but do not compute the eigenvalue exactly. Instead we find upper bounds for $\Lambda^o(m)$ using the Collatz-Wielandt formula [11, 28].

Lemma 2 (Collatz-Wielandt formula) *Let A be an irreducible square matrix with non-negative entries. Then for any vector $\mathbf{x} > 0$, the largest eigenvalue of A (denoted λ) is real and positive and is bounded by*

$$\min_i \frac{(A\mathbf{x})_i}{x_i} \leq \lambda \leq \max_i \frac{(A\mathbf{x})_i}{x_i}. \quad (9)$$

^(iv) For the three models we consider, the matrices grow as the Fibonacci numbers.




Fig. 3: Pictorial interpretation of (12), (13) and (14).

Proof: Let \mathbf{u} be the left eigenvector of A corresponding to λ . By the Perron-Frobenius theorem, λ is real and positive and \mathbf{u} has strictly positive entries. Then for any $\mathbf{x} > 0$,

$$\langle \mathbf{u} | A | \mathbf{x} \rangle = \lambda \langle \mathbf{u} | \mathbf{x} \rangle \tag{10}$$

and so

$$\langle \mathbf{u} | A\mathbf{x} - \lambda\mathbf{x} \rangle = \sum_{i=1}^n u_i x_i \left(\frac{(A\mathbf{x})_i}{x_i} - \lambda \right) = 0. \tag{11}$$

Some terms in the sum must be non-negative and some non-positive (although they may all be 0 if $\mathbf{x} = \mathbf{u}$), and therefore the maximum of the summands is non-negative and the minimum is non-positive. Since $\mathbf{u}, \mathbf{x} > 0$, the result follows. □

We use this lemma to find an upper bound $\Lambda^o(m)$. In order to calculate a tight bound, we need to get as close as possible to the dominant eigenvector of V_m^o . We do this using Baxter’s corner transfer matrix ansatz, which specifies to choose a vector ψ of the form

$$\psi_\sigma = \text{Tr} [F(\sigma_1, \sigma_2)F(\sigma_2, \sigma_3) \dots F(\sigma_m \sigma_1)], \tag{12}$$

where the index $\sigma = (\sigma_1, \dots, \sigma_m) \in \{\bullet, \circ\}^m$ runs over all legal configurations of one row of the cylinder, and $\{F(a, b)\}$ is a set of four matrices of size $n \times n$, indexed by two spin values a and b . Here n is an arbitrary positive integer which need not be related to m . The F matrices are calculated from corner transfer matrix methods; the process is quite involved, and we leave its description until the next section. We can interpret this ansatz pictorially by thinking of the F matrices as half-row transfer matrices which build up the state ψ_σ one row at a time — see Figure 3.

In order to use (9), we must also be able to compute $(V_m^o \psi)_\sigma$, and to do this we define the *face weight* ω . The weight of a face (a single cell of the lattice) is 1 when the spins around it form a legal configuration and 0 otherwise. The element of V_m^o which maps a column in state τ to a column in state σ is then given by the product

$$(V_m^o)_{\sigma, \tau} = \prod_{i=1}^m \omega \begin{pmatrix} \sigma_{i+1} & \tau_{i+1} \\ \sigma_i & \tau_i \end{pmatrix}, \quad \text{where } \begin{matrix} \sigma_{m+1} & \equiv & \sigma_1, & \text{and} \\ \tau_{m+1} & \equiv & \tau_1. \end{matrix} \tag{13}$$

This is shown pictorially in Figure 3. We can then write the action of V_m^o on ψ as

$$\begin{aligned} (V_m^o \psi)_\sigma &= \sum_{\tau} (V_m^o)_{\sigma, \tau} \psi_\tau = \sum_{\tau} \prod_{i=1}^m \omega \begin{pmatrix} \sigma_{i+1} & \tau_{i+1} \\ \sigma_i & \tau_i \end{pmatrix} \text{Tr} [F(\tau_1, \tau_2) F(\tau_2, \tau_3) \dots F(\tau_m, \tau_1)] \\ &= \text{Tr} \sum_{\tau} \prod_{i=1}^m \omega \begin{pmatrix} \sigma_{i+1} & \tau_{i+1} \\ \sigma_i & \tau_i \end{pmatrix} F(\tau_i, \tau_{i+1}) = \text{Tr} [F_l(\sigma_1, \sigma_2) F_l(\sigma_2, \sigma_3) \dots F_l(\sigma_m, \sigma_1)], \end{aligned} \quad (14)$$

where F_l is a $2n \times 2n$ matrix defined by the block matrix equation (shown pictorially in Figure 5)

$$F_l(c, a)|_{d, b} = \omega \begin{pmatrix} a & b \\ c & d \end{pmatrix} F(d, b). \quad (15)$$

We are now able to combine the above expression with (8) and (9) to derive

$$\kappa^m \leq \Lambda^o(m) \leq \max_{\sigma} \frac{(V_m^o \psi)_\sigma}{\psi_\sigma} = \max_{\sigma} \left\{ \frac{\text{Tr} [F_l(\sigma_1, \sigma_2) F_l(\sigma_2, \sigma_3) \dots F_l(\sigma_m, \sigma_1)]}{\text{Tr} [F(\sigma_1, \sigma_2) F(\sigma_2, \sigma_3) \dots F(\sigma_m, \sigma_1)]} \right\}. \quad (16)$$

This is the upper bound that we use for κ , and it is valid for any m , n , and F , as long as we satisfy the following conditions:

- V_m^o is non-negative. In fact V_m^o is a 0-1 matrix for the considered models, so this follows immediately.
- V_m^o is irreducible. This is equivalent to showing that every set of states σ on a cut can be reached from any other. This is easy to show for our models, as there is no restriction on adjacency to \circ spins — so any σ can be adjacent to $\{\circ, \circ, \dots, \circ\}$, and thus can reach any other set of states.
- ψ_σ is positive. This does not immediately follow from (12). We must show this by explicitly computing (12) for every σ and verifying that it is positive. However, this does not result in any extra work since we already need these values to compute (16).

Putting all of this together, we find an upper bound as follows:

1. Select $m, n > 0$ with m even.
2. Calculate a set of $n \times n$ matrices $\{F(a, b)\}$ as specified in the next section.
3. For all possible cut states σ of m spins:
 - (a) Verify $\psi_\sigma > 0$, where ψ_σ is given by (12).
 - (b) Calculate $\left(\frac{(V_m^o \psi)_\sigma}{\psi_\sigma} \right)^{1/m}$.
4. The upper bound for κ is the maximum of all values calculated in step 3b.

This method uses very little memory, because we do not need to keep the entire ψ vector in memory. Each component can be computed independently as it is needed from the F and F_l matrices. These matrices are tiny compared to ψ and V_m^o , so this is a very modest requirement. Furthermore, the calculation in step 3 for any particular component does not depend on any other component, so can be done in parallel (and ratios compared afterwards). This step is by far the most time-consuming part of the process — the number of traces required grows exponentially with m . Hence parallelisation creates a huge real-time saving. Details of our implementation and some optimisations are given in Section 4 below.




Fig. 4: Pictorial interpretation of the CTM equations.

3 Approximate eigenvectors from corner transfer matrices

So far, we have not described how to calculate the F matrices, and this is clearly critical in order to obtain a good upper bound. We need to choose the matrices in such a way that the approximate eigenvector ψ given by (12) is close to the true eigenvector of V_m^o . We accomplish this through the use of corner transfer matrix formalism, which we briefly outline here; we direct the reader to a more detailed discussion in [9].

The expression (12) is the starting point of Baxter’s corner transfer matrix formalism. It can be represented pictorially by considering $F(a, b)$ in the limit $n \rightarrow \infty$ as a half-row transfer matrix; it takes a half-infinite row of spins ending with spin a , and transfers it along one row to a half-infinite row of spins ending with spin b . Using this representation, we can see that the expression for ψ represents a half-plane partition function (normalised in some way), which is unchanged under the action of V_m^o save for a constant factor — *i.e.*, it is an eigenvector of V_m^o . This representation is shown in Figure 3.

These infinite-dimensional half-row transfer matrices can be shown [2, 9] to satisfy the *CTM equations*

$$\xi A^2(a) = \sum_b F(a, b) A^2(b) F(b, a), \tag{17}$$

$$\eta A(a) F(a, b) A(b) = \sum_{c, d} \omega \begin{pmatrix} a & b \\ c & d \end{pmatrix} F(a, c) A(c) F(c, d) A(d) F(d, b), \tag{18}$$

where $A(a)$ are matrices of the same size as F , indexed by a single spin value. The A matrices are the eponymous *corner transfer matrices*; they transfer a half-infinite row of spins ending with spin a around an angle of $\pi/2$, sweeping out a quarter of the lattice. The CTM equations can then be interpreted as equating half-plane transfer matrices (see Figure 4). The matrices on either side of the equality differ only by an extra row, which results in the scalar multipliers ξ and η .

In the infinite-dimensional limit, the CTM equations directly yield the solution of the model via $\kappa = \eta/\xi$. We can approximate the true solution by taking F matrices of finite size $n \times n$. The CTM equations can be solved for finite-dimensional matrices, and we take these finite sized solutions for our F matrices in our bound finding procedure.

Solving the CTM equations at finite size is quite non-trivial. For this purpose, we use the corner transfer matrix renormalisation group method developed by Nishino and Okunishi [24, 25]. In this method, we start with some trial values for $A(a)$ and $F(a, b)$ which are then “polished” into solutions. We do this by expanding A and F into the $2n \times 2n$ matrices A_l and F_l , using the block matrix equations (15) and

$$A_l(c)|_{d, a} = \sum_b \omega \begin{pmatrix} a & b \\ c & d \end{pmatrix} F(d, b) A(b) F(b, a). \tag{19}$$




Fig. 5: Pictorial interpretation of A_l and F_l .

These can be interpreted as “larger” versions of the corner and half-row transfer matrices (see Figure 5).

We then reduce the large matrices to generate new iterations of A and F . To do this, we diagonalise $A_l(a)$, producing orthogonal diagonalising matrices $P(a)$. We take the transformations

$$A(a) \rightarrow P^T(a)A(a)P(a), \quad F(a, b) \rightarrow P^T(a)F(a, b)P(b), \quad (20)$$

which leave the CTM equations invariant. We then truncate the matrices to the original $n \times n$ size by keeping the n largest eigenvalues of A_l and performing a consistent transformation on F_l . This has the effect of making the matrices intuitively “cover as large an area as possible”, so that they are close to the infinite-size solution of the CTM equations.

The expanding and reducing procedures are iterated, and it is observed that the matrices eventually converge to the finite-size solution of the CTM equations. The initial values for the A and F matrices can be taken to be of some small size, typically 1×1 or 2×2 . They can then be “built up” to the desired $n \times n$ size by sometimes keeping extra eigenvalues at the reduction step, resulting in larger matrices until the desired size is reached. We again direct the reader to [9] for a more detailed description of this process.

4 Results

The method was implemented in 2 distinct steps: computing the F matrices, and then computing the trace ratios for all legal states σ . Both parts were implemented in C++ using the Eigen numerical linear algebra library [17]. We used Eigen since it readily supports multiple precision computation through the MPFR library [13]. High precision is necessary in the first step because the eigenvalues of corner transfer matrices range over many orders of magnitude, and in the second step so that the trace ratios are also of high precision — though less precision is needed. The first step requires only modest computing resources, and all F matrices were generated within a few hours on a modest Linux laptop using only a few megabytes of memory.

The second step — computing the trace ratios — can be implemented quite naïvely and still give good results. However we use a number of simple and very effective optimisations. Firstly, traces are invariant under cyclic permutations, so we only have to consider legal cut states modulo cyclic permutations for σ . Traces are also invariant under transposes, and the F matrices satisfy $F(a, b) = F^T(b, a)$, so we can also disregard reflections. Thus we consider bracelets which are legal for the system (generated using the methods in [26]). This reduces the number of computations by a factor of approximately $2m$. Additionally, we can gain a little more speed by using similarity transforms to diagonalise one of the F matrices. Finally, we split the set of bracelets into batches of equal size across about 300 cores on the WestGrid computing cluster. Each job (with $n = 30, m = 50$) took approximately 1 day of CPU time. Note that final upper bounds were computed at high precision and then rounded up.

Model	m	n	Upper bound
Hard squares	50	30	<u>1.503 048 082 475 332 264 519</u>
NAK	50	30	<u>1.342 643 951 124 602 238</u>
RWIM	50	30	<u>1.448 957 372 102</u>

Tab. 1: Upper bounds on the growth rates of hard squares, NAK and RWIM. The previous best known bounds [14, 22] are underlined.

Model	Lower bounds [9]	Exact digits
Hard squares	<u>1.503 048 082 475 332 264 322 066 329 475 553 689 385 781 038 610</u> 305 062 028 101 735 933 850 396 923 440 380 463 299 47	19
NAK	<u>1.342 643 951 124 601 297 851 730 161 875 740 395 719 438 196 938</u> 393 943 434 885 455 0	15
RWIM	<u>1.448 957 371 775 608 489 872 231 406 108 136 686 434 371</u>	9

Tab. 2: Lower bounds on the growth rates of hard squares, NAK and RWIM from [9]; combining those lower bounds with the new upper bounds, the underlined digits are known rigorously. Note that the lower bounds agree the best estimates [9] in all but their last couple of digits.


Our results are given in Table 1. For each of the three models, we computed the bound using F matrices of size 30×30 on cylindrical systems of circumference 50 vertices. In this table, we also indicate the previous best upper bounds. In Table 2, we combine these new upper bounds with the best known lower bounds to show the digits of the growth rates which we now know rigorously. We extend the number of rigorously known digits from previous work by between 4 and 6 digits for our models.

We note here that the use of floating point precision for the trace-ratio computations made our code about 50-100 times slower than if we had been able to use hardware floating point. Unfortunately most commercially available processors, including those we had access to, support 80-bit double-extended precision rather than true 128-bit or higher precision. If such high precision hardware floating point were available, then we could have increased m by about 8, which would perhaps give an extra 2 or 3 digits exactly. The previous best upper bounds [14, 22] used machine floating point.


Due to a lack of rigorous results on the theoretical convergence of the CTMRG, it is difficult to predict how good our bounds will be for given m and n . Thus we must observe their empirical behaviour for various m and n . We first note that m is the computational bottleneck parameter, as the time required is exponential in m (while the time and memory requirements depend polynomially on n). Thus we wish to increase m to the limit of our computational power, and then select an optimal n for that value of m . We show the results in the context of the hard squares model; other models behave in much the same way.

For any fixed m , the bound given by (16) decreases as n increases, but only up to a small limit, after which it stays level. Therefore, given a value of m , we can be assured of calculating the best bound for that m as long as n is sufficiently large. This is shown in Figure 6(a). The optimal n (*i.e.*, the value of n after which the bound does not decrease further) for any given m appears to follow a linear relationship with m . In Figure 6(b), we show this value for small m . For our final bound computation, we were able to run the method for $m = 50$. We estimate by extrapolating the fitted linear relationship that for this value of m , the optimal value of n is approximately 14. We have taken $n = 30$, and so are confident that this value is sufficient to achieve the best bound for our value of m .

Next, we look at the performance of the bound in relation to m . Because our bound is itself an upper bound on $\Lambda^o(m)^{1/m}$, we know that it will be inferior to an explicit evaluation of (8) for the



(a) Log error in the upper bounds generated by increasing n for $m = 12, 24$.



(b) Optimal n for various m . There appears to be a linear relationship with m .

Fig. 6: Choosing an n value for a given m .




Fig. 7: Log error in the upper bounds generated by (16) (red) and (8) (blue) for various m . Here we used $n = 20$ which always produces the best possible bound for these values of m .

same m . However, for sufficiently large n , the difference between (16) and (8) is very small. This is shown in Figure 7, where there is almost no discernible difference between the two bounds. Thus we expect that the accuracy of our bound will behave similarly to that of (8), which is conjectured to decrease to κ exponentially with respect to m . This behaviour can also be observed in Figure 7. This means that the algorithm, which is exponential-time in m , will also be exponential-time in the number of digits found. This is worse than the observed performance for the lower bound found in [9], although this is not unexpected.

As previously mentioned, (8) is the formula used to generate all previous best known upper bounds. The advantage of our method is that (16) is almost as accurate and much faster to compute, even though both methods are exponential in m . For example, in 8 the eigenvalue $\Lambda^o(36)$ was computed using over 200 gigabytes of memory and 40 days of cpu-time (and real-time). We were able to push to $m = 50$ using only a few megabytes of memory and about 300 days of cpu-time but only 1 day of real-time due to parallelisation. Thus we are able to achieve much better bounds.

5 Conclusion

In this paper, we have used transfer matrix analysis, the Collatz-Wielandt formula, and corner transfer matrix formalism to derive upper bounds on the growth rates of three lattice models motivated

by an information storage problem. Our bounds are a significant improvement on all previously known upper bounds, and together with the corresponding lower bounds from [9], fully determine the growth rates to a generous number of digits.

Our method is not limited to the models we have studied; nearly any lattice model which can be written in terms of local face weights and satisfies some elementary symmetry properties can be analysed in this way. It can be applied directly to colouring models, and it appears possible to exploit symmetries in this case to make it even more efficient. On the other hand, some other models lack irreducible transfer matrices, for example the “even” model in [9]. We are extending our method to account for these difficulties.

One weakness of our method is that it appears to be exponential-time in the number of digits required, due to the need to evaluate all the ratios in (16) (modulo symmetries). It seems intuitive that much of this is wasted work; since in the limit $n \rightarrow \infty$ it can be shown that the solution of the CTM equations gives the exact eigenvector in (12) and as such, all of the ratios should be very close to each other, and it should be unnecessary to calculate all 2^m of them. We observed, for example, that the traces were always positive but we have been unable to prove this. Similarly, we observed that the maximum or minimum trace ratio was frequently given by $\sigma = \circ^m$ or $\sigma = (\circ\bullet)^{m/2}$, but, again, we have been unable to prove this. Such results would allow us to compute bounds much more efficiently and we are currently working to prove them.

Acknowledgements

I would like to thank Andrew Rechnitzer, who contributed to this paper at all levels. I also would like to thank Ian Enting and Brian Marcus for many helpful and interesting discussions, and WestGrid for providing access their computer cluster.

References

- [1] R. J. Baxter. Dimers on a rectangular lattice. *J. Math. Phys.*, 9:650, 1968.
- [2] R. J. Baxter. Variational approximations for square lattice models in statistical mechanics. *J. Stat. Phys.*, 19(5):461–478, 1978.
- [3] R. J. Baxter. Hard hexagons: exact solution. *J. Phys. A: Math. Gen.*, 13:L61, 1980.
- [4] R. J. Baxter, I. G. Enting, and S. K. Tsang. Hard-square lattice gas. *J. Stat. Phys.*, 22(4):465–489, 1980.
- [5] R. Berger. *The undecidability of the domino problem*. Number 66 in Memoirs of the AMS. American Mathematical Society, 1966.
- [6] N. Calkin and H. Wilf. The number of independent sets in a grid graph. *SIAM J. Disc. Math.*, 11(1):54–60, 1998.
- [7] Y. Chan. Series expansions from the corner transfer matrix renormalization group method: the hard squares model. *J. Phys. A: Math. Theor.*, 45:085001, 2012.
- [8] Y. Chan. Series expansions from the corner transfer matrix renormalization group method: II. Asymmetry and high-density hard squares. *J. Phys. A: Math. Theor.*, 46(12):125009, 2013.
- [9] Y. Chan and A. Rechnitzer. Accurate lower bounds on 2-D constraint capacities from corner transfer matrices. *IEEE Trans. Info. Theory*, 60(7):3845–3858, 2014.
- [10] M. Cohn. On the channel capacity of read/write isolated memory. *Disc. Appl. Math.*, 56(1):1–8, 1995.
- [11] L. Collatz. Einschließungssatz für die charakteristischen Zahlen von Matrizen. *Math. Z.*, 48(1):221–226, 1942.
- [12] K. Engel. On the Fibonacci number of an $M \times N$ lattice. *Fibonacci Quart.*, 28(1):72–78, 1990.

- [13] L. Fousse, G. Hanrot, V. Lefèvre, P. Pélissier, and P. Zimmermann. MPFR: A multiple-precision binary floating-point library with correct rounding. *ACM Trans. Math. Software*, 33(2):13:1–13:15, June 2007.
- [14] S. Friedland, P. H. Lundow, and K. Markström. The 1-vertex transfer matrix and accurate estimation of channel capacity. *IEEE Trans. Inf. Theory*, 56(8):3692–3699, 2010.
- [15] D. Gamarnik and D. Katz. Sequential cavity method for computing free energy and surface pressure. *J. Stat. Phys.*, 137(2):205–232, 2009.
- [16] M. Golin, X. Yong, Y. Zhang, and L. Sheng. New upper and lower bounds on the channel capacity of read/write isolated memory. In *Information Theory, 2000. Proceedings. IEEE International Symposium on*, page 280. IEEE, 2000.
- [17] G. Guennebaud, B. Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
- [18] 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.
- [19] E. H. Lieb. Exact solution of the problem of the entropy of two-dimensional ice. *Phys. Rev. Lett.*, 18:692–694, Apr 1967.
- [20] D. Lind. The entropies of topological Markov shifts and a related class of algebraic integers. *Ergod. Theor. Dyn. Syst.*, 4:283–300, 6 1984.
- [21] D. Lind. *An introduction to symbolic dynamics and coding*. Cambridge University Press, 1995.
- [22] E. Louidor and B. Marcus. Improved lower bounds on capacities of symmetric 2D constraints using Rayleigh quotients. *IEEE Trans. Inf. Theory*, 56(4):1624–1639, 2010.
- [23] P. H. Lundow and K. Markström. Exact and approximate compression of transfer matrices for graph homomorphisms. *LMS J. Comput. Math*, 11:1–14, 2008.
- [24] T. Nishino and K. Okunishi. Corner transfer matrix renormalization group method. *J. Phys. Soc. Jpn.*, 65(4):891–894, 1996.
- [25] T. Nishino and K. Okunishi. Corner transfer matrix algorithm for classical renormalization group. *J. Phys. Soc. Jpn.*, 66(10):3040–3047, 1997.
- [26] F. Ruskey and J. Sawada. Generating necklaces and strings with forbidden substrings. In *Computing and Combinatorics*, pages 330–339. Springer, 2000.
- [27] H. Temperley and M. E. Fisher. Dimer problem in statistical mechanics — an exact result. *Philos. Mag.*, 6(68):1061–1063, 1961.
- [28] H. Wielandt. Unzerlegbare, nicht negative matrizen. *Math. Z.*, 52(1):642–648, 1950.