

Auspicious tatami mat arrangements

Alejandro Erickson
ate@uvic.ca

Frank Ruskey
ruskey@uvic.ca

Jennifer Woodcock
jwoodcoc@uvic.ca

Department of Computer Science
University of Victoria
PO BOX 3055, STN CSC, Victoria BC, V8W 3P6, Canada

Mark Schurch
mschurch@uvic.ca

Department of Mathematics and Statistics
University of Victoria
PO BOX 3060, STN CSC, Victoria BC, V8W 3R4, Canada

Submitted: March 4, 2011; Accepted: XX; Published: XX
Mathematics Subject Classification: 05B45, 05A19

Abstract

An *auspicious tatami mat arrangement* is a tiling of a rectilinear region with two types of tiles, 1×2 tiles (dimers) and 1×1 tiles (monomers). The tiles must cover the region and satisfy the constraint that no four corners of the tiles meet; such tilings are called *tatami tilings*. The main focus of this paper is when the rectilinear region is a rectangle. We provide a structural characterization of rectangular tatami tilings and use it to prove that the tiling is completely determined by the tiles that are on its border. We prove that the number of tatami tilings of an $n \times n$ square with n monomers is $n2^{n-1}$. We also show that, for fixed-height, the generating function for the number of tatami tilings of a rectangle is a rational function, and outline an algorithm that produces the generating function.

Keywords: tatami, monomer-dimer tiling, rational generating function

1 What is a tatami tiling?

Traditionally, a tatami mat is made from a rice straw core, with a covering of woven soft rush straw. Originally intended for nobility in Japan, they are now available in mass-market stores. The typical tatami mat occurs in a 1×2 aspect ratio and various configurations of them are used to cover floors in houses and temples. By parity considerations it may be necessary to introduce mats with a 1×1 aspect ratio in order to cover the floor of a room. Such a covering is said to be “auspicious” if no four corners of mats




Figure 1: (a) Vertical bond pattern. (b) Horizontal bond pattern. (c) Herringbone pattern.




Figure 2: What is the least number of monomers among all tatami tilings of this region? The answer is provided at the end of the paper in Figure 21.

meet at a point. Hereafter, we only consider auspicious arrangements, since without this constraint the problem is the classical and well-studied dimer tiling problem ([6], [10]). Following Knuth ([7]), we will call the auspicious tatami arrangements, *tatami tilings*. The fixed-height enumeration of tatami tilings that use only dimers (no monomers) was considered in [9], and results for the single monomer case were given in [1].

Perhaps the most commonly occurring instance of tatami tilings is in paving stone layouts of driveways and sidewalks, where the most frequently used paver has a rectangular shape with a 1×2 aspect ratio. Two of the most common patterns, the “herringbone” and the “running bond,” shown in Figure 1, have the tatami property. Consider a driveway of the shape in Figure 2. How can it be tatami tiled with the least possible number of monomers? The answer to this question could be interesting both because of aesthetic appeal, and because it could save work, since to make a monomer a worker typically cuts a 1×2 paver in half.

Before attempting to study tatami tilings in general orthogonal regions it is crucial to understand them in rectangles, and our results are primarily about tatami tilings of rectangles.

1.1 Outline

In Section 2 we determine the structure of tatami tilings in a rectangle. Our structural characterization has important algorithmic implications, for example, it reduces the size of the description of a tiling from $\Theta(rc)$ to $O(\max\{r, c\})$ and may be used to generate tilings quickly. The three theorems in Section 3 are the main results of the paper and are also stated here. The first of these concerns the maximum possible number of monomers. Let $T(r, c, m)$ be the number of tilings of the $r \times c$ grid, with m monomers (and the other tiles being horizontal or vertical dimers).

Theorem 1. *If $T(r, c, m) > 0$, then m has the same parity as rc and $m \leq \max(r+1, c+1)$.*

Following this we prove a counting result for maximum-monomer tilings of square grids.

Theorem 2. *The number of $n \times n$ tilings with n monomers, $n2^{n-1}$.*

Our final result concerns fixed-height tilings with an unrestricted number of monomers.

Theorem 3. *For a fixed number of rows r , the ordinary generating function of the number of tilings of an $r \times n$ rectangle is a rational function.*

We also provide an algorithm which outputs this generating function for a given r and explicitly give the generating function for $r = 1, 2$ and 3 , along with the coefficients of the denominator for $1 \leq r \leq 11$. In Section 4 we return to the question of tatami tiling general orthogonal regions and introduce the “magnetic water strider problem” along with additional conjectures and open problems.

2 The structure of tatami tilings: T-diagrams

We show that all tatami tilings have an underlying structure which partitions the grid into regions, where each region is filled with either the vertical or horizontal running bond pattern (or is a monomer not touching the boundary). For example, in Figure 3 there are 11 regions, including the interior monomer. We will describe this structure precisely and prove some results for tilings of rectangular grids.

Whenever a horizontal and vertical dimer share an edge $\square\square$, either the placement of another dimer is forced to preserve the tatami condition, or the tiles make a T with the boundary of the grid $\perp\text{---}\perp$. In the former case, the placement of the new dimer again causes the sharing of an edge $\square\square$, and so on $\square\square$, until the boundary is reached.

The successive placement of dimers, described above gives rise to skinny herringbone formations, which we call *rays*. These are always directed in such a way that they propagate from their source to the boundary of the grid and cannot intersect one another.




Figure 3: A tiling showing all four types of sources. Coloured in magenta, from left to right they are, a clockwise vortex, a vertical bidimer, a loner, a vee, and another loner. Jagged edges are indicated by brackets.




Figure 4: These two types of sources must have their coloured tiles on a boundary, as shown, up to rotational symmetry.

Between the rays, there are only vertical or horizontal running bond patterns. The intersection of a running bond with the boundary is called a *segment*. This segment is said to be *jagged* if it consists of alternating monomers and dimers orthogonal to the boundary; otherwise it is said to be *smooth* because it consists of dimers that are aligned with the boundary. Every jagged segment is marked with square brackets in Figure 3.

We know that a ray, once it starts, propagates to the boundary. But how do they start? In a rectangular grid, we will show that a ray starts at one of four possible types of *sources*. In our discussion we use inline diagrams to depict the tiles that can cover the grid squares at the start of a ray. We need not consider the case where the innermost square (denoted by the circle) $\begin{array}{|c|} \hline \bullet \\ \hline \end{array}$ is covered by a vertical dimer $\begin{array}{|c|} \hline \blacksquare \\ \hline \end{array}$ because this would move the start of the ray.

If it is covered by a horizontal dimer $\begin{array}{|c|c|} \hline \blacksquare \blacksquare \\ \hline \end{array}$, the source, which consists of the two dimers that share a long edge, is called a *bimer*. Otherwise it is covered by a monomer $\begin{array}{|c|} \hline \blacksquare \\ \hline \end{array}$ in which case we consider the grid square beside it $\begin{array}{|c|} \hline \bullet \\ \hline \end{array}$. If it is covered by a monomer the source is called a *vee* $\begin{array}{|c|c|} \hline \blacksquare \blacksquare \\ \hline \end{array}$; if it is covered by a vertical dimer the source is called a *vortex* $\begin{array}{|c|} \hline \blacksquare \\ \hline \end{array}$; if it is covered by a horizontal dimer it is called a *loner* $\begin{array}{|c|c|} \hline \blacksquare \blacksquare \\ \hline \end{array}$. Each of these four types of sources forces at least one ray in the tiling and all rays begin at either a bidimer, vee, vortex or loner. The different types of features are depicted in Figures 4-6.

The coloured tiles in Figures 4-6 characterize the four types of sources. A bidimer or vortex may appear anywhere in a tiling, as long as the coloured tiles are within its boundaries. The vees and loners, on the other hand, must appear along a boundary, as shown in Figure 4.




Figure 5: A vertical and a horizontal *bidimer* source. A bidimer may appear anywhere in a tiling provided that the coloured tiles are within the boundaries of the grid.




Figure 6: A counter clockwise and a clockwise *vortex* source. A vortex may appear anywhere in a tiling provided that the coloured tiles are within the boundaries of the grid.

The collection of bold staircase-shaped curves in each of the four types of source-ray drawings in Figures 4-6, is called a *feature*. These features do not intersect when drawn on a tatami tiling because rays cannot intersect. A *feature-diagram* refers to a set of non-intersecting features drawn in a grid. Not every feature-diagram admits a tatami tiling; those that do are called *T-diagrams*. See Figure 7.




Figure 7: (a) The T-diagram of Figure 3. (b) A feature diagram that is not a T-diagram.

Recall that a tatami tiling consists of regions of horizontal and vertical running bond patterns. A feature-diagram is a T-diagram if and only if each pair of rays bounding the same region admit bond patterns of the same orientation and the distance between them has the correct parity. The precise conditions are stated in Lemma 1.

Features decompose into four types of rays, to which we assign the symbols *NW*, *NE*, *SW*, and *SE*, indicating the direction of propagation. Two rays are said to be *adjacent* if they can be connected by a horizontal or vertical line segment which intersects no other ray. If (α, β) is an adjacent pair, then α is on the left when considering horizontally adjacent pairs and on the bottom when considering vertically adjacent pairs.

Lemma 1. *A feature diagram is a T-diagram if and only if the following four conditions hold.*

Horizontal Conditions:

- (H1) *There are no horizontal $(\alpha E, \beta E)$ -adjacencies, nor are there horizontal $(\alpha W, \beta W)$ -adjacencies, where α and β are either N or S (Figure 8);*
- (H2) *all distances are even, except for horizontal (NE, NW) -distances and horizontal (SE, SW) -distances, which are odd (Figure 9).*

Vertical Conditions:

- (V1) *There are no vertical $(S\alpha, S\beta)$ -adjacencies, nor are there any vertical $(N\alpha, N\beta)$ -adjacencies, where α and β are either E or W;*
- (V2) *all distances are even, except for vertical (NW, SW) -distances and vertical (NE, SE) -distances, which are odd.*




Figure 8: Incompatible pairs of adjacent rays. The region between the adjacent rays would have to contain both horizontal and vertical dimers.




Figure 9: If the size of the gap between adjacent rays has the correct parity then it can be properly tiled, as shown on the left. On the right, the red regions cannot be tiled because the gaps have the wrong parities.

This characterization has some implications for the space and time complexity of a tiling.

Lemma 2. *Let G be an $r \times c$ grid, with $r < c$.*

- (i) *A tatami tiling of G is uniquely determined by the tiles on its boundary.*
- (ii) *The storage requirement for a tatami tiling of G is $O(c)$; that is, a tatami tiling can be recovered from $O(c)$ bits.*

(iii) Whether a feature diagram in G is a T -diagram can be determined in time $O(c)$.

Proof. To prove (i), we need to show that we can recover the T -diagram from the tiles that touch the boundary. Those portions of the T -diagram corresponding to vees and loners, as well as bidimers whose source tiles are both on the boundary $\sqcup\sqcup$, are easy to recover. The black rays in Figure 10 show their recovery. Imagine filling in the remaining red rays, whose ends look like $\sqcup\sqcup$, by following them naïvely, backwards from their endings to the boundary. The ends of the four rays emanating from a bidimer or vortex will always form exactly one of the four patterns illustrated in Figure 11; in each case, it is straightforward to recover the position and type of source. This proves (ii).

Part (ii) follows from (i), because we can use a ternary encoding for the perimeter squares.




Figure 10: The same tiling as in Figure 3 with only the boundary tiles showing. Rays emanating from sources on the boundary are in black and otherwise, they are drawn naïvely in red, to be matched with a candidate source from Figure 11.

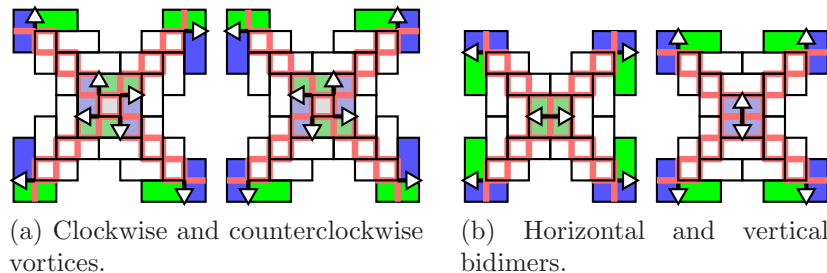


Figure 11: The four types of vortices and bidimers are recoverable from the ends of their rays, at the boundary of the grid. Extending the rays naïvely, backwards from the boundary, we form one of the two patterns in the red overlay. One occurs only for bidimers and the other for vortices. Successively placing tiles, working from the ends of the rays towards the central configuration, we also find the orientation of the source, as shown in the figure.

Claim (iii) is true provided that Lemma 1 only needs to be applied to $O(c)$ ray-adjacencies. Notice that a pair of rays can be adjacent and yet not be adjacent on the boundary. For example, it happens in Figure 7.

Each ray bounds exactly two regions, each of which is bounded by at most three other rays, and two rays must bound the same region to be adjacent. Thus, a ray is adjacent

to at most six other rays. Let the ray-adjacencies be the edges of a graph $G = (V, E)$ whose vertex set is the set of rays, so that G has maximum degree at most 6. Therefore, the number of ray-adjacencies, $|E|$, and hence applications of Lemma 1, is linear in the number of rays, $|V|$, which is at most four times the number of features, which is in $O(c)$. This proves (iii). \square

The T-diagram structure is a useful tool for enumerating and generating tatami tilings as will be illustrated in the following sections.

3 Counting results

Let $T(r, c, m)$ be the number of tatami tilings of a rectangular grid with r rows, c columns, and m monomers. Also, $T(r, c)$ will denote the sum

$$T(r, c) = \sum_{m \geq 0} T(r, c, m).$$

We begin by giving necessary conditions for $T(r, c, m)$ to be non-zero.

Theorem 1. *If $T(r, c, m) > 0$, then m has the same parity as rc and $m \leq \max(r+1, c+1)$.*

Proof. Let r, c and m be such that $T(r, c, m) > 0$ and let d be the number of grid squares covered by dimers in an $r \times c$ tatami tiling so that $m = rc - d$. Since d is even, m must have the same parity as rc .

It suffices to assume that $r \leq c$, and prove that $m \leq c + 1$. The proof proceeds in two steps. First, we will show that a monomer on a vertical boundary of any tiling can be *mapped* to the top or bottom, without altering the position of any other monomer. Then we can restrict our attention to tilings where all monomers appear on the top or bottom boundaries, or in the interior. Secondly, we will show that there can be at most $c + 1$ monomers on the combined horizontal boundaries.

Let T be a tatami tiling of the $r \times c$ grid with a monomer μ on the left boundary, touching neither the bottom nor the top boundary. The monomer μ is (a) part of a vee or a loner, or (b) is on a jagged segment of a region of horizontal bond. Define a *diagonal* to be μ together with a set of dimers in this region which form a stairway shape from μ to either the top or bottom of the grid as shown in purple in Figure 12a. If such a diagonal exists, a *diagonal flip* can be applied, which changes the orientation of its dimers and maps μ to the other end of the diagonal. In case (a) a diagonal clearly exists since it is a source and its ray will hit a horizontal boundary because $r \leq c$.

If μ is on a jagged segment, then we argue by contradiction. Suppose neither diagonal exists, then they must each be impeded by a distinct ray. Such rays have this horizontal region to the left so the upper one is directed *SE* and the lower *NE* and they meet the right boundary (before intersecting). Referring to Figure 12c,

$$\begin{aligned} \alpha + \beta + j &= \gamma + \delta + 1 \leq r \\ &\leq c \leq c' = \alpha + \gamma = \beta + \delta, \end{aligned}$$




Figure 12: (c) If both diagonals are blocked, then $c < r$. The tiling is at least this tall and at most this wide.

where j is some odd number. Thus $\alpha + \beta + j \leq \alpha + \gamma$ implying that $\beta < \gamma$. On the other hand,

$$\gamma + \delta + 1 = r \leq c \leq c' = \beta + \delta$$

implies that $\gamma < \beta$, which is a contradiction. Therefore at least one of the diagonals exists and the monomer can be mapped to a horizontal boundary.

We may now assume that there are no monomers strictly on the vertical boundaries of the tiling, and therefore all monomers are either in the top or bottom rows or in vortices. Let v be the number of vortices. Encode the bottom and top rows of the tiling by length c binary sequences P and Q , respectively. In the sequences, 1s represent monomers and 0s represent squares in dimers.

If Q contains a consecutive pair of 1s, this represents a vee in the top row; the vee has a region of horizontal dimers directly below it. This region of horizontal bond must reach the bottom row somewhere, otherwise, by an argument similar to one given previously, we would have $c < r$ (see Figure 13a). Therefore, there must be a consecutive pair of 0s in P unique to these 1s in Q . Modify Q so that the 11 becomes a 1, and modify P so that the 00 becomes 010. Do this for each consecutive pair of 1s in P and Q , as illustrated in Figure 13b. The updated P and Q sequences contain no consecutive pairs of 1s, but the total number of 1s remains unchanged.

Now we show that for each of the v vortices, a 00 can be removed from each of P and Q and there will still be no consecutive pairs of 1s in the new sequences. Each vortex generates 4 rays; at least one of these rays will hit the top boundary, and at least one will hit the bottom boundary. In fact, a horizontal and a vertical dimer on either side of the ray will lie on the upper boundary, and similarly for the bottom boundary. Figure 13a is helpful in seeing why this is true. These dimers on either side of the ray induce a 000 in P and another 000 in Q . (Although not used in this proof, note that the comments above also apply to bidimers.) In total, there are at least v distinct triples of 0s in each sequence, one for each vortex. Now remove 00 from each triple as in Figure 13b. The updated P and Q sequences have a combined length of $2c - 4v$ and neither of them contains a 11. Thus the total number of 1s is at most $\lceil |P|/2 \rceil + \lceil |Q|/2 \rceil$, which is at most $c - 2v + 1$. Adding back the v vortex monomers, we conclude that there are at most $c - v + 1$ monomers in total, which finishes the proof.

Note that, to achieve the bound of $c + 1$, we must have $v = 0$, and that the maximum is achieved by a vertical bond pattern. \square




Figure 13: Each vortex and vee is associated with segments of monomer-free grid squares shown in purple. (a) Segments associated with vortices have length at least three. Those associated with vees have at least two 0s. (b) The two types of updates to sequences P and Q . The upper sequences are before the updates and the lower are after updates. The symbol \times represents a deletion from the sequence.

The converse of Theorem 1 is false, for example, Alhazov et al. ([1]) show that $T(9, 13, 1) = 0$. We now state a couple of consequences of Theorem 1.

Corollary 1. *The following three statements are true for tatami tilings of an $r \times c$ grid with $r \leq c$.*

- (i) *The maximum number of monomers in a $r \times c$ grid is $c+1$ if r is even and c is odd; otherwise it is c . There is a tatami tiling achieving this maximum.*
- (ii) *A tatami tiling with the maximum number of monomers has no vortices.*
- (iii) *A tatami tiling with the maximum number of monomers has no bidimers.*

Proof. (i) That this is the correct maximum value can be inferred from Theorem 1. A tiling consisting only of vertical running bond achieves it, for example.

(ii) This was noted at the end of the proof of Theorem 1.

(iii) We can again use the same sort of reasoning that was used for vortices in Theorem 1, but there is no need to “add back” the monomers, since bidimers do not contain one. \square

3.1 Square tatami tilings

In this section, we show that $T(n, n, n) = n2^{n-1}$. Theorem 2 relies on the following lemma and corollary.

Lemma 3. *For each $n \times n$ tiling with n monomers, a trivial tiling can be obtained via a finite sequence of diagonal flips in which each monomer moves at most once. Reversing this sequence returns us to the original tiling.*




Figure 14: In the tiling T from Lemma 3, the ray ρ belongs to the corner γ and it is associated with the diagonal δ . The area $A(T)$ counts the grid squares that are not between any ray and its corner. The monomer that is moved in the diagonal flip becomes part of $A(T')$ and is therefore moved only once in the sequence. The corner monomers are never moved.

Proof. Let T be the T-diagram of an $n \times n$ tiling with n monomers. Each ray ρ in T touches two adjacent boundaries which form a corner γ , so ρ and γ are said to belong to each other. For each corner γ , choose the ray which belongs to it and is farthest away from it; if a corner does not have a ray, then choose the corner itself. Between the four chosen rays/corners, our tiling can only contain either horizontal or vertical running bond (by Lemma 1). Let $A(T)$ be the area of this *central running bond*.

We begin a sequence of diagonal flips by choosing one ray ρ that is farthest from its corner and flipping the diagonal δ touching ρ that is between ρ and its corner. Let T' be the resulting T-diagram. In T , δ is not part of the central running bond and in T' , it is; thus $A(T') > A(T)$. Continuing this process yields a trivial tiling via a finite sequence of diagonal flips.

□

Corollary 2. *Every $n \times n$ tiling with n monomers has two corner monomers and they are in adjacent corners.*

Proof. The sequence of diagonals chosen for diagonal flips described in Lemma 3 never includes a diagonal containing a corner monomer because such a diagonal is never between a ray and its associated corner. As such, the corner monomers are fixed throughout the sequence of diagonal flips yielding a trivial tiling. Since a trivial $n \times n$ tiling with n monomers has two monomers in adjacent corners, then, so must every other $n \times n$ tiling with n monomers.

□

Corollary 2 show that the four rotations of any $n \times n$ tiling with n monomers are distinct. We call the rotation with monomers in the top two corners the *canonical case*.

Theorem 2. *The number of $n \times n$ tilings with n monomers, $T(n, n, n)$, is $n2^{n-1}$.*

Proof. We count the $n \times n$ tilings with n monomers up to rotational symmetry by counting the canonical cases only. Let $S(n) = T(n, n, n)/4$. We will give a combinatorial proof that $S(n)$ satisfies the following recurrence:

$$S(n) = 2^{n-2} + 4S(n-2), \text{ where } S(1) = \frac{1}{4} \text{ and } S(2) = 1. \quad (1)$$

The solution to the recurrence (1) is $S(n) = n2^{n-3}$.

In Theorem 1 we defined a diagonal flip which results in a monomer μ moving up or down depending on the orientation of its diagonal. As such, we simplify our terminology by referring to flipping a monomer in a particular direction (up, down, left, or right).

We treat the even and odd cases separately, though the proofs are naturally similar. In both cases, we begin with the canonical trivial case and consider all possible sequences of flips in which each monomer is moved at most once and the corner monomers are fixed. By Lemma 3 and its corollary, this counts the canonical tilings.

The canonical trivial case for even n , shown in Figure 15a for $n = 8$, is a horizontal running bond tiling with fixed (black) monomers in the top corners and $n/2$ (red and yellow) monomers on both the left and right boundaries. We classify the tilings according to what happens to the bottom (yellow) monomer on each of these boundaries, which we will call w and e .




Figure 15: (a) Canonical trivial case for an 8×8 square with 8 monomers. (b) Flipping w up. (c) 180 degree rotation of the canonical trivial case for a 6×6 square with 6 monomers. (d) An 8×8 tiling with its associated 6×6 tiling.

First, suppose $\mu \in \{w, e\}$ is flipped up as shown in Figure 15b. Because our tiling is square, this flip inhibits any orthogonal diagonal flips and thus the monomers that shared a boundary with μ before it was flipped up can only be flipped up and monomers on the opposite boundary can only be flipped down. There are $n - 3$ such monomers that are not fixed and can be flipped independently of each other. This gives 2^{n-3} possibilities when either w or e is flipped up, resulting in a total of 2^{n-2} tilings.

If neither w nor e is flipped up, these monomers can be flipped (or not) independently of each other and of other non-fixed monomers, as shown in Figure 16. As such, we can now ignore what happens to w and e and consider them fixed, keeping in mind that for each such tiling, there are three others with w and e in different positions.

Consider canonical trivial case for $(n - 2) \times (n - 2)$ tilings with $n - 2$ monomers, and rotate it 180 degrees so that its fixed (black) monomers are in the bottom corners, as




Figure 16: The four possibilities for flipping w and/or e down.

shown in Figure 15c for $n - 2 = 6$. We map the $n - 4 = (n - 2) - 2$ non-fixed (red) monomers of this tiling to the $n - 4$ non-fixed (red) monomers of the $n \times n$ canonical trivial case in a natural way: pairing those in the same position relative to the bottom fixed (black/yellow) monomers. In Figure 15, for example, the bottom left red monomer in (a) maps to the bottom left red monomer in (c), and the top left red monomer in (a) maps to the top left red monomer in (c) (the right side is analogous). Further, the diagonals angling upwards from a mapped pair of monomers are associated, as are those angling downwards.

Next, we prove that any tiling obtained by a sequence of flips of the non-fixed (red) monomers in the rotated $(n - 2) \times (n - 2)$ canonical trivial case corresponds to a tiling obtained by the same sequence of flips of the associated non-fixed (red) monomers in the $n \times n$ canonical trivial case, and vice versa. An example of such a correspondence is shown in Figure 15d.

To do this, we show that compatibility between diagonal flips is preserved between the smaller and larger cases: that is, if two diagonals cannot both be flipped in the larger square, the same is true for the corresponding diagonals in the smaller square, and vice versa. In both cases, two monomers on the same boundary can both be flipped if and only if they are either flipped in the same direction or the top one is flipped up and the bottom one is flipped down; compatibility is preserved.

For a pair of monomers on opposite boundaries, observe that a conflict between flips can only occur if we try to flip them both in the same direction. Further, conflict depends entirely on the distance of the monomers from the horizontal centerline of the grid. Let d_w and d_e respectively be the distances from the horizontal centerline, with negative values below the line and positive values above. If $d_w + d_e > 0$, then the two monomers cannot both be flipped down, and similarly, if $d_w + d_e < 0$, they cannot both be flipped up. This distance is preserved between the associated monomers in the larger and smaller squares and thus compatibility is also preserved.

There are $S(n - 2)$ ways of flipping the monomers of the (rotated) $(n - 2) \times (n - 2)$ canonical trivial case, and thus $S(n - 2)$ ways of flipping the corresponding monomers of the $n \times n$ canonical trivial case. This yields $4S(n - 2)$ tilings, one for each way of positioning w and e and establishes (1) for even n .

The canonical trivial case for odd n , shown in Figure 17a for $n = 7$, is a vertical running bond tiling with (black) monomers in the top corners. It has $\lceil n/2 \rceil$ monomers on the top boundary and $\lfloor n/2 \rfloor$ monomers on the bottom boundary. Label the bottom left and bottom right monomers w and e respectively.

Similar to the even case, if either w is flipped right (as in Figure 17b) or e is flipped




Figure 17: (a) Canonical trivial case for a 7×7 square with 7 monomers. (b) Flipping w to the right. (c) Canonical trivial case for a 5×5 square with 5 monomers. (d) An 7×7 tiling with its associated 5×5 tiling.




Figure 18: The four possibilities for flipping (or not flipping) w and e to the left and right respectively.

left, there are $n - 3$ monomers which can be flipped independently to obtain other tilings and this yields 2^{n-2} tilings.

Otherwise, w and e can be flipped left and right (respectively) independently of each other and of other monomers, as shown in Figure 18. Again we fix w and e , keeping in mind that for each such tiling, there are three others with w and e in different positions. We will find a similar one-to-one correspondence to the one in the even case.

There are $n - 4$ (red) monomers on the top and bottom boundaries of the canonical trivial case that we have not fixed. Once again, we associate these monomers in a natural way with those of the 180 degree rotation of the $(n - 2) \times (n - 2)$ canonical trivial case with $n - 2$ monomers which has fixed (black) monomers in the bottom corners. Figure 17a and Figure 17c show the two trivial cases we are considering, and Figure 17d shows an example of a correspondence between an $(n - 2) \times (n - 2)$ and an $n \times n$ tiling.

Arguing as in the even case, with a vertical centerline rather than a horizontal one, we conclude that a pair of monomers in the $n \times n$ tiling can be flipped if and only if the corresponding flips can be made in the $(n - 2) \times (n - 2)$ tiling, and vice versa. Again this yields $4(S(n - 2))$ tilings and establishes (1) for odd n . \square

3.2 Fixed height tatami tilings

In this section we show that for a fixed number of rows r , the ordinary generating function of the number of tilings of an $r \times c$ rectangle is a rational function. We will show that, for each value of r , the number of fixed-height tilings satisfies a system of linear recurrences with constant coefficients. We will derive the recurrences for small values of r and then discuss an algorithm which can be used for larger values of r .

Let $T_r(z)$ denote the generating function

$$T_r(z) = \sum_{c \geq 0} T(r, c) z^c.$$

For $c \geq 2$, a tatami tiling of a $1 \times c$ rectangle begins with either a monomer or a dimer. Thus, $T(1, c) = T(1, c - 1) + T(1, c - 2)$ for $c \geq 2$, where $T(1, 0) = 1$ and $T(1, 1) = 1$. This is the well known Fibonacci recurrence. Since it is a linear recurrence with constant coefficients, it is not a difficult task to verify that

$$T_1(z) = \frac{1 + z}{1 - z - z^2}.$$

For each $r \geq 2$ we derive a recurrence based on partial tilings which can be solved using mathematical software such as Maple. A partial tiling of an $r \times c$ grid is a *minimal $r \times k$ tiling* if and only if the first k columns are covered and no tile lies entirely outside of these columns. The r and k may sometimes be omitted. Let \mathcal{S}_r be the set of configurations which correspond to a minimal $r \times 1$ tiling. For $s_v \in \mathcal{S}_r$, let v be a ternary r -tuple whose elements correspond to the grid squares of the first column, ordered from top to bottom. The elements 0, 1, and 2, each represent a grid square covered by a vertical dimer, monomer, or horizontal dimer, respectively. Note that 0s always appear in consecutive pairs. For example, $s_{0012002} \in \mathcal{S}_7$ corresponds to the minimal 7×1 tiling shown in Figure 19a.




Figure 19: (a) The minimal 7×1 tiling corresponding to $s_{0012002}$. (b) A possible extension of the minimal tiling in (a). (c) Removing the first column yields a new minimal tiling, represented by $s_{2001211}$.

For $c \geq 1$, let $t_{r,v}(c)$ be the number of tilings of an $r \times c$ rectangle that begin with the minimal $r \times 1$ tiling s_v . Note that $t_{r,v}(1) = 1$ if v does not contain a 2, and is zero otherwise. To derive a recurrence we determine all ways of extending each configuration in s_v to a minimal $r \times 2$ tiling. By taking each of these minimal $r \times 2$ tilings and chopping off the first column we can match these tilings to an element in \mathcal{S}_r which will define a recurrence. Figure 19b shows an extension of the tiling $s_{0012002}$ and Figure 19c shows that this extension corresponds to the configuration $s_{2001211}$. Notice that Figure 19c can only be extended once more.

Lemma 4. $T_2(z) = \frac{1+2z^2-z^3}{1-2z-2z^3+z^4}$.

Proof. For $r = 2$, we have $\mathcal{S}_2 = \{s_{00}, s_{11}, s_{12}, s_{21}, s_{22}\}$. Since \mathcal{S}_2 contains all possible ways to start a tiling of an $r \times c$ rectangle, with $c \geq 2$, it follows that

$$T(2, c) = t_{00}(c) + t_{11}(c) + t_{12}(c) + t_{21}(c) + t_{22}(c). \quad (2)$$

The initial conditions are $t_{00}(1) = 1, t_{11}(1) = 1, t_{12}(1) = 0, t_{21}(1) = 0$, and $t_{22}(1) = 0$. To derive the recurrence we consider the ways of extending each of the minimal 2×1 tilings in \mathcal{S}_2 to a minimal 2×2 tiling.

	Start	Extensions	Recurrences
S_{00}			
		$t_{00}(c) = t_{00}(c-1) + t_{11}(c-1) + t_{12}(c-1) + t_{21}(c-1) + t_{22}(c-1)$	
S_{11}			
		$t_{11}(c) = t_{00}(c-1)$	
S_{12}			
		$t_{12}(c) = t_{11}(c-1) + t_{21}(c-1)$	
S_{21}			
		$t_{21}(c) = t_{11}(c-1) + t_{12}(c-1)$	
S_{22}			
		$t_{22}(c) = t_{11}(c-1)$	

By solving the system of recurrences defined by these five cases and Equation (2) we arrive at the stated result for $r = 2$. \square

The process outlined in the proof of Lemma 4 can be implemented in an algorithm. We determine the set \mathcal{S}_r with an exhaustive search. Then, for each element $s_v \in \mathcal{S}_r$, we list all extensions to a minimal $r \times 2$ tiling. Each extension of s_v produces a unique sum-term in the recurrence for $t_v(c)$. Once again, the initial conditions are

$$t_v(1) = \begin{cases} 1, & \text{if } v \text{ contains a } 2, \\ 0, & \text{otherwise.} \end{cases}$$

We may reduce the number of equations in the system of recurrences by ignoring elements of \mathcal{S}_r which cannot be extended to a minimal $r \times 2$ tiling. This may be taken further by determining necessary conditions for an element of \mathcal{S}_r to be extendible to an $r \times c$ tiling.

The algorithm produces a system of linear recurrences with constant coefficients. This proves the following result.

Theorem 3. *For a fixed number of rows r , the ordinary generating function for the number of tilings of an $r \times n$ rectangle is a rational generating function.*

The output of our algorithm for $r = 3$ gives the following generating function:

$$T_3(z) = \frac{1 + 2z + 8z^2 + 3z^3 - 6z^4 - 3z^5 - 4z^6 + 2z^7 + z^8}{1 - z - 2z^2 - 2z^4 + z^5 + z^6}.$$

It is impractical to include the complete generating function for any larger values of r . The degrees for the numerators and denominators, however, as well as the coefficients in the denominators are given in Table 1 for $r = 1, 2, \dots, 11$. The salient patterns in these coefficients are summarized in Conjectures 1 and 2. Note that Conjecture 1(a) implies $g(z)$ is a self-reciprocal polynomial for $r \equiv 2 \pmod{4}$.

Conjecture 1. *Let $T_r(z) = \frac{f(z)}{g(z)}$, where $f(z)$ and $g(z)$ are relatively prime polynomials, and $\deg(g(z)) = n$, and $r \geq 1$. Then,*

$$g(z) = \begin{cases} -z^n g\left(\frac{1}{z}\right), & \text{if } r \equiv 0 \pmod{4}, \\ -z^n g\left(-\frac{1}{z}\right), & \text{if } r \equiv 1 \pmod{4}, \\ z^n g\left(\frac{1}{z}\right), & \text{if } r \equiv 2 \pmod{4}, \\ z^n g\left(-\frac{1}{z}\right), & \text{if } r \equiv 3 \pmod{4}. \end{cases}$$

A mod 4 pattern also seems to occur in the degrees of the denominators of $T_r(z)$. The rigid structure we encounter in tatami tilings prompts us to infer this pattern upon all values as well.

Conjecture 2. *Let $g(z)$ be the denominator of $T_r(z)$. Then,*

$$\deg(g(z)) = \begin{cases} 8m^2 + 2m + 1, & \text{if } r \equiv 0 \pmod{4}, \\ 8m^2 + 4m + 2, & \text{if } r \equiv 1 \pmod{4}, \\ 8m^2 + 10m + 4, & \text{if } r \equiv 2 \pmod{4}, \\ 8m^2 + 8m + 6, & \text{if } r \equiv 3 \pmod{4}. \end{cases}$$

We have not been able to devise a uniform presentation of the generating function of $T_r(z)$ similar to what was done in [9].

It would also be interesting to consider the generating function $T_r(x, y, z)$, in which the coefficient of $x^h y^v z^m$ is the number of tilings with h vertical dimers, v horizontal dimers, and m monomers.

r	p	q	Coefficients of $g(z)$ ordered from left to right by ascending degree and then folded like these arrows; \mapsto for $r \leq 3$, \curvearrowright for $r = 4, 5, 6, 7$, and \curvearrowleft for $r \geq 8$.
1	1	2	1, -1, 1
2	3	4	1, -2, 0, -2, 1
3	8	6	1, -1, -2, 0, -2, 1, 1
4	14	11	-1, 1, 1, 1, -1, 7 1, -1, -1, -1, 1, -7
5	18	14	-1, 1, 1, -1, 3, -1, 5, -2 1, 1, -1, -1, -3, -1, -5
6	27	22	1, -1, -1, 1, -1, -2, 2, -10, 9, -1, 4, 6 1, -1, -1, 1, -1, -2, 2, -10, 9, -1, 4
7	28	22	1, -1, -3, 3, 4, -4, -9, 7, 6, -5, 2, 0 1, 1, -3, -3, 4, 4, -9, -7, 6, 5, 2
8	44	37	-1, 1, 1, -1, 1, -1, 1, 3, -3, 13, -12 1, -1, -1, 1, -1, 1, -1, -3, 3, -13, 12 34, -2, 6, 20, -6, 12, 0, 0 -34, 2, -6, -20, 6, -12, 0, 0
9	50	42	-1, 1, 1, -1, 1, -1, 1, -1, 5, -3, 11, -8 1, 1, -1, -1, -1, -1, -1, -1, -5, -3, -11, -8 10, 24, 2, 28, 2, 20, 8, 14, 4, 6 -24, 2, -28, 2, -20, 8, -14, 4, -6
10	65	56	1, -1, -1, 1, -1, 1, -1, 1, -1, -4, 4, -16, 15, 1, -1 1, -1, -1, 1, -1, 1, -1, 1, -1, -4, 4, -16, 15, 1, -1 -120, 68, -78, -18, 18, -66, 66, -2, 7, 41, -23, 33, -17, 17 68, -78, -18, 18, -66, 66, -2, 7, 41, -23, 33, -17, 17
11	64	54	1, -1, -5, 5, 13, -13, -27, 27, 48, -48, -83, 81, 125, -120, -160 1, 1, -5, -5, 13, 13, -27, -27, 48, 48, -83, -81, 125, 120, -160 -34, 83, 89, -156, -165, 199, 210, -202, -206, 185, 193, -154 -34, -83, 89, 156, -165, -199, 210, 202, -206, -185, 193, 154

Table 1: Summary of generating function attributes for fixed height tilings, $r = 1, \dots, 11$, and where p and q are the degrees of the numerator and denominator of $T_r(z)$, respectively. The degree ordering shows the patterns of Conjecture 1.

4 More conjectures and further research

The T-diagram structure removes much of the mystery from tatami tilings and motivates considerable future work. In this section we list some open problems and conjectures, beginning with another counting problem on rectangular grids.

4.1 Rectangular regions

Conjecture 3. For all $d \geq 0$ and $m \geq 1$ there is an n_0 such that, for all $n \geq n_0$.

$$T(n, n + d, m) = T(n_0, n_0 + d, m),$$

whenever $n(n + d)$ has the same parity as m (otherwise $T(n, n + d, m) = 0$, by Theorem 1).

Experimentally, it appears that the smallest n_0 is $m + d + 4$, if $d \geq 1$.

The easiest case occurs when $d = 0$ and $m = 1$. It is not hard to show that for all odd $n \geq 3$ we have $T(n, n, 1) = 10$ (the single monomer must go at a corner or in the center).

In a subsequent paper we will show that for $m < n$,

$$T(n, n, m) = m2^m + (m + 1)2^{m+1},$$

whenever m and n have the same parity.

Returning to the subject of generating functions, ignoring signs, it appears that the denominators of $T_r(z)$ in Section 3.2 are self-reciprocal. There must be a combinatorial explanation for this. Similar questions in the non-tatami case are considered in [2].

Generating functions also appear in Conjecture 4, inspired by conversations with Knuth. Let $T(n, z)$ be the generating polynomial for the number of $n \times n$ tilings with n monomers and i vertical dimers. Once again, to count such tilings we consider flipping diagonals, with the added precaution that the sum of the number of tiles in the flipped diagonals is a given constant. The relationship between this and subsets of $\{1, \dots, n\}$ which have a given sum is detailed in a subsequent publication.

Let $\phi_n(z)$ denote the n th cyclotomic polynomial. Recall that the roots of $\phi_n(z)$ are the primitive roots of unity. One of their more well-known properties is that

$$1 - z^n = \prod_{d|n} \phi_d(z). \quad (3)$$

Let $S_n(z)$ denote the ordinary generating function of subsets of $\{1, \dots, n\}$ which have a given sum. That is, $\langle z^k \rangle S_n(z)$ is the number of subsets A of $\{1, 2, \dots, n\}$ such that the sum of the numbers in A is k . It is not difficult to see that

$$S_n(z) = (1 + z)(1 + z^2) \cdots (1 + z^n) = \prod_{k=1}^n (1 + z^k). \quad (4)$$

Let $\nu(n)$ denote the number of 2s in the prime factorization of n and note that

$$1 + z^n = \frac{1 - z^{2n}}{1 - z^n} = \frac{\prod_{d|2n} \phi_d(z)}{\prod_{d|n} \phi_d(z)} = \prod_{\substack{d|2n \\ d \nmid n}} \phi_d(z) = \prod_{\substack{d|n \\ d \text{ odd}}} \phi_{2^{1+\nu(n)}d}(z).$$

When this latter expression is used in $S_n(z)$ some interesting simplification occurs.

Lemma 5. For all $n \geq 1$,

$$S_n(z) = \prod_{j=1}^n (\phi_{2j}(z))^{\lfloor \frac{n+j}{2j} \rfloor}$$

Proof. The index $2j$ will occur for those k s in Equation (4) for which $j = 2^{\nu(k)}d$ for some odd d where $d \mid k$. This equation is satisfied for $k = j, 3j, 5j, \dots$. There are $\lfloor (n+j)/(2j) \rfloor$ such k s that are less than or equal to n (See Figure 20). \square




Figure 20: A visual aid for the last line of the proof of Lemma 5. The pink dots represent the sequence, $j, 3j, 5j, \dots$, with $ij \leq n$. Adding j to n shows that the number of dots is $\lfloor (n+j)/(2j) \rfloor$.

Conjecture 4. The generating polynomial $T(n, z)$ has the factorization

$$T(n, z) = P(n, z) \prod_{j \geq 1} S_{\lfloor \frac{n-1}{2j} \rfloor}(z)$$

where $P(n, z)$ is an irreducible polynomial.

We return to the topic mentioned in the introduction: Tatami-tilings of orthogonal regions.

4.2 Orthogonal regions

We believe that the main structural components are the same as they were for rectangles, but there are a few subtleties to be clarified at inside corners, since a ray could begin at such a place.

What is the computational complexity of determining the least number of monomers that can be used to tile an orthogonal region given the segments that form the boundary of the region and the unit size of each dimer/monomer? In the rectangular grid this is




Figure 21: (a) The solution to the question posed in Figure 2; no monomers are required to tatami tile the region. (b) A legal configuration of six magnetic water striders in an orthogonal “pond”. Note that no further striders may be added.

answerable in polynomial time using T-diagrams, however, it appears to be NP-hard for an arbitrary number of segments.

The problem of minimizing the number of monomers in a tiling inspires what we call the “magnetic water strider problem”. This time the orthogonal region is a pond populated by water striders. A water strider is an insect that rides atop water in ponds by using surface tension. Its 4 longest legs jut out at 45 degrees from its body. In the fancifully named magnetic water strider problem, we require the body to be aligned north-south. Furthermore its legs support it, not by resting on the water, but by extending to the boundary of the pond. Naturally, the legs of the striders are not allowed to intersect. A legal configuration of magnetic water striders in an orthogonal pond is shown Figure 21b.

There are two problems and a game here. The first is a packing problem: What is the largest number of magnetic waters striders that a pond can support? On the other hand, one can ask what is the minimum number that can be placed so that no more can be added. Placing and packing striders can be tricky, which gives rise to an adversarial game where players take turns placing striders in an orthogonal region. Brian Wyvill has kindly implemented a version of this game, available at <http://www.theory.cs.uvic.ca/~cos/tatami/>.

Interpreted as a matching problem on a subgraph of a grid graph G , a tatami tiling is a matching M with the property that $G - M$ contains no 4-cycles. Note that there is always such a matching (e.g., take the “running bond” layout on the infinite grid graph and then restrict it to G). However, if we insist on a perfect matching, then the problem is equivalent to our “perfect” driveway paving problem from the introduction.

More generally, a matching whose removal destroys k -cycles is called C_k -*transverse*. Ross Churchley proved that finding a C_k -transverse matching in an arbitrary graph is NP-hard when $k \geq 4$ (private communication [4]).

4.3 Combinatorial games

Consider the following game. Given an orthogonal region, players take turns placing dimers (or dimers and monomers); each placement must satisfy the tatami constraint and the last player who can move wins. This game, called Oku!, is reminiscent of the game called Nimm, in which players also win by making the last move, however a winning strategy for our game is unknown and there are grid sizes in which the second player can force a win. The name is a phonetic spelling of the Japanese word for “put”.

Another game applies tomography to rectangular tilings.

Tiling tomography is a rich and open area of complexity theory to which a good introduction can be found in [3]. The relevant question is as follows: Given $r + c$ triples of numbers (h, v, m) , one for each row and one for each column, is there a tatami tiling which has h horizontal dimers, v vertical dimers, and m monomers in the respective row or column?

Without the tatami condition this decision problem is NP-hard (Theorem 4, [5]). Hard or not, the tatami condition gives considerable information in practice, however, making the reconstruction of a tatami tiling an entertaining challenge. Erickson, A., has created an online computer game out of this called Tomoku. It is playable at <http://tomokupuzzle.com>, complete with music, countdown timers and high scores.

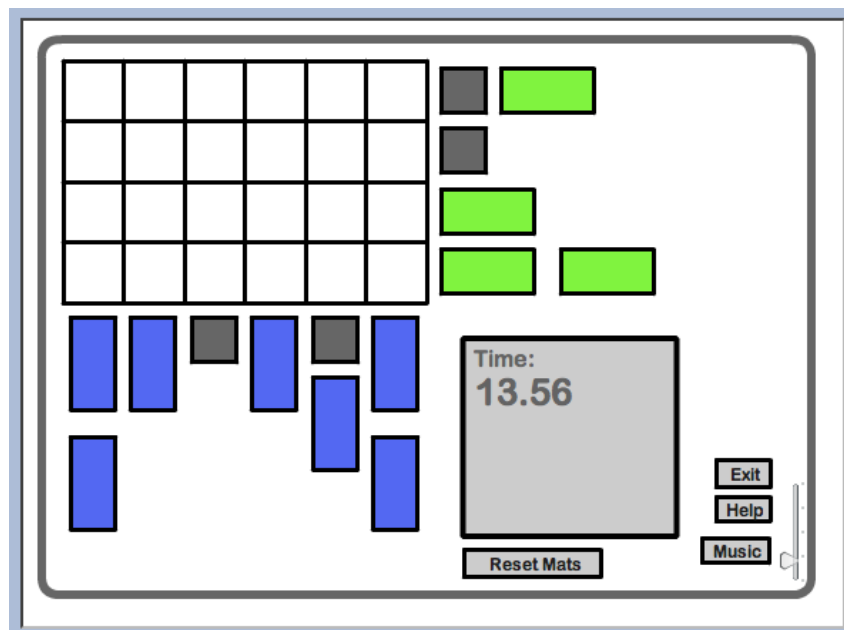


Figure 22: The Tomoku web game. The player is shown which tiles are completely contained in each row and column, and the object is to reconstruct the tiling. Note that each monomer appears twice in the projections.

Both of these games are at <http://www.theory.cs.uvic.ca/~cos/tatami/>, and it should be noted that they can also be played with a pencil and paper.

5 Acknowledgements

Thanks to Donald Knuth for his comments on an earlier draft of this paper and to Martin Matamala for pointing out the tomography problem.

References

- [1] Alhazov, A., Morita, K., Iwamoto, C.: A Note on Tatami Tilings. *Mathematical Foundation of Algorithms and Computer Science* (T. Tokuyama, Ed.), RIMS Kôkyûroku series, No. 1691, Research Institute for Mathematical Sciences, Kyoto, Japan, (2010), 1–7. <http://www.kurims.kyoto-u.ac.jp/~kyodo/kokyuroku/kokyuroku.html>
- [2] Anzalone, N., Baldwin, J., Bronshtein, I., Petersen, T.K.: A reciprocity theorem for monomer-dimer coverings. In Morvan, M., Rémila, É., eds.: *Discrete Models for Complex Systems, DMCS'03. Volume AB of DMTCS Proceedings.*, Discrete Mathematics and Theoretical Computer Science (2003) 179–194
- [3] Chrobak, M., Couperus, P., Dürr, C., Woeginger, G.: On tiling under tomographic constraints *Theoretical Computer Science* Vol: 290 Issue: 3 ISSN: 0304-3975 Date: 01/2003 Pages: 2125 - 2136
- [4] Churchley, R., Huang, J.: Private communication, November 2010
- [5] Dürr, C., Guíñez, F., Matamala, M.: Reconstructing 3-colored grids from horizontal and vertical projections is NP-hard. *Algorithms - ESA 2009 LNCS 5757* (2009) 776–787
- [6] Kenyon, R., Okounkov, A.: “What is ... a dimer?”. *Notices of the American Mathematical Society* **52** (2005) 342–343
- [7] Knuth, D. E.: *The Art of Computer Programming. Volume 4, fascicle 1B.* Addison-Wesley (2009)
- [8] Pachter, L.: Combinatorial approaches and conjectures for 2-divisibility problems concerning domino tilings of polyominoes. *Electronic Journal of Combinatorics* **4** (1997) 2–9
- [9] Ruskey, F., Woodcock, J.: Counting fixed-height tatami tilings. *Electronic Journal of Combinatorics* **16** (2009) 20
- [10] Stanley, R. P.: On dimer coverings of rectangles of fixed width. *Discrete Applied Mathematics* **12**(1) (1985) 81 – 87