Packing 2×2 unit squares into grid polygons is NP-complete

Dania El-Khechen* Muriel Dulieu†‡ John Iacono†§ Nikolaj van Omme¶

1 Introduction

In a packing problem, the goal is to put some small *objects* disjointly into a large *container*, while optimizing some objective function. The packing problem is very general, and a rich variety of objects and containers are possible. One major split in the taxonomy of packing concerns whether the objects to be packed are allowed to be different or are all identical. In the case of identical packing, probably the simplest non-trivial variant is the following problem, which we call the 2×2 *packing problem*: How many axis-aligned 2×2 squares can be packed inside a polygon P with n edges drawn on a unit grid, where the squares must be packed such that each occupies exactly four grid locations (i.e. rotation or fractional placement is forbidden)? Whether the problem when P is restricted to be simple (the 2×2 *simple packing problem*) is polynomial or NP-complete appears on the open problem project [1] as problem 56. There they cite [3] as proving the 2×2 packing problem NP-complete (where P must be allowed to have holes). This is not true, as in [3] they do prove NP-completeness, but only when all possible locations where the squares could be packed are explicitly provided in the input. Since the size of a normal representation of a grid polygon using binary integers can differ exponentially from the number of possible packing locations inside the polygon (e.g. a $k \times k$ square requires $\Theta(\log k)$ bits to represent but has $\Theta(k^2)$ packing locations), it does not follow from their result that the 2×2 packing problem is in NP if the input is simply a polygon. The result of [3] does prove NP-hardness for 2×2 packing problem. This paper is devoted to proving the 2×2 packing problem is in NP, and thus, combined with the work of [3] proves NP-completeness for the first time. Whether the 2×2 simple packing problem is polynomial or NP-complete remains an important open question, and remains open quite surprisingly even if P is restricted to be orthogonally convex. A simple greedy algorithm works for Manhattan Skyline polygons [2], the most complex class of polygons for which a polynomial algorithm is known.

*Concordia University, Montréal, Québec, Canada.
†Polytechnic Institute of NYU, Brooklyn NY USA.
‡Research partially supported by a grant from the U.S.-Israel Binational Science Foundation and by NSF Grant CCF-08-30691.
§Research partially supported by NSF grants OISE-0334653 and CCF-0430849 and by an Alfred P. Sloan research fellowship
¶CIRRELT, C.P. 6128, succursale Centre-ville, Montréal (Québec) H3C 3J7 CANADA.

A closely related problem is the *pallet loading problem*. It appears as problem 55 on the Open Problems Project [1]: "Given two pairs of numbers, (A, B) and (a, b), and a number n, decide whether n small rectangles of size $a \times b$, in either axis-parallel orientation, can be packed into a large rectangle of size $A \times B$." This problem is not known to be in NP. It shares many of the same issues as the 2×2 packing problem considered here, where the size of an explicit representation of the packing is not necessarily polynomial in the size of the input. The main impediment to showing pallet loading is in NP is the issue of orthogonal rotations, which in our 2×2 packing problem are thankfully not relevant.

2 Definitions

Definition 1 *Given a simple orthogonal grid polygon P with n edges, let* $\max(P)$ *be the maximum number of axis-aligned 2×2 squares that can be packed into P. We call any packing of P with* $\max(P)$ *axis-aligned 2×2 squares* maximum. *The decision version of the 2×2 packing problem is given P and a integer k, report if* $\max(P) \le k$.

We call a 2×2 square a *block*, while a 1×1 unit of the underlying grid we call a *square*. Blocks have one of four *types* bases on the parity of coordinates of their vertical and horizontal edges. By *adjacent* we mean edge-adjacant. We call a maximal adjacent set of blocks a *superblock*. In a packing, maximal adjacent squares not covered by blocks are called *holes*. Holes and blocks are classified as *interior* or *exterior* based on whether they are adjacent to the boundary of P. We call a 1×1 interior hole a *unit hole*. An interior unit hole is adjacent to four blocks, one of each type. There are two possible configurations of the four blocks around a unit hole, which we call the *spin* (See Figure 3); every hole has p-spin or n-spin. We call the edge where two blocks of different types meet a *skeleton* edge. (See Figure 4).

3 Regularizing

Lemma 1 *For any P, there is a maximum packing of P where the interior holes are unit. We call such a packing a partially-regularized maximal packing of P.*

Proof. We say two blocks in a packing are *aligned* if their enclosing rectangle has width or height 2 and there are no other blocks in the enclosing rectangle. The *alignment score* of a packing is the number of aligned pairs of blocks. The *gravity score* of a packing is the sum of the x and y coordinates of the centers of all blocks in the packing. Given two packings, the *better* one is the one with more blocks. If two packings have the same number of blocks, the better one is the one with the higher alignment score. If two packings have the same number of blocks and the same alignment score, the better one is the one with smaller gravity score. Let K be (one of the) best packings of P. We argue that K only has unit holes in the interior; if it does not we show there is a better packing, a contradiction. If K does not have only unit holes in the interior it must have one 1×2 (or by symmetry a 2×1) hole in the interior. In Figure 2, starting with the assumption of a 1×2 hole we enumerate a number of cases. For each of these cases we show in Figures 5-9 that there is a better packing. □

Lemma 2 *Let K be an arbitrary optimal packing of P with a minimum area of interior holes. The packing K must only have unit holes in the interior.*

Proof. We now use the cases developed for the previous lemma in an algorithmic manner. Suppose K has a 1×2 hole. Determine which case you are in according to Figure 2 and for each case a new, better, equal-sized packing is obtained or it is concluded that the packing was not optimal. In all cases where a new packing is obtained, there is a new 1×2 non-covered region. This new non-covered region is part of a hole that is either interior or exterior. If it is exterior, that contradicts that K had the minimum area of interior holes (since an area of two was transferred from an interior hole to and exterior one). If it is interior, apply this procedure to the new packing. This process must terminate since at each step we get a better packing. And the only ways to terminate is to contradict that K is a optimal packing or that K has the minimum area of interior holes. Thus the assumption that K has a 1×2 hole is false and K only has 1×1 holes. □

In a partially regularized maximum packing of a polygon P, we call a sequence of vertical skeleton edges disjoint only at interior unit holes a vertical *microbone*[1]. All vertical microbones in a partially regularized maximum packing begin either on an edge of P or at a vertical distance 1 away from P; this is because given two adjacent boxes of different type with a vertical skeleton edge between them, boxes of different type must continue to be on either side of the vertical extension of the skeleton edge until a hole or the exterior of the polygon is encountered. The union of a microbone, its extensions adjacent to unit holes, and the possible two unit-length grid edges to connect to P forms the *bone*. Note that the union of all bones partition P into regions such that each region is either a hole or contains blocks of only one type.

Given a vertical bone b in a maximum packing of a polygon P with the minimum area of interior holes, call $l(b)$ the set of squares in the packing to the left of b and adjacent to b. Look at a bone which has a top and bottom which are not incident to or within distance 2 of a vertex of P. Call this a floating bone. Given a floating bone in a maximum packing of a polygon P with the minimum area of interior holes, define the shift right operation as follows: remove the squares $l(b)$ from the packing and insert a copy of $r(b)$ shifted two units to the left (See Figure 11). We need to argue that this gives a valid packing; this is done by observing that in all cases an arbitrary square in $r(b)$ will have a vacant space immediately to its left after the removal of $l(b)$. If $|l(b)| = |r(b)|$ then a left or right shift on b will preserve the size of the packing; if they are not equal than either a left or right shift will increase the size of the packing by $|l(b) - r(b)|$ and show that the original was not optimal. By the same logic the number of interior holes must remain the same. Also note that after a shift, the bone has moved two units in the direction of the shift. The shifting of a bone can not cause any 1×2 holes to appear (for example by moving a hole next to the hole in another bone), as this would by Lemma 2 indicate that the packing does not have a minimum number of interior holes.

Lemma 3 *For any P, there is a partially regularized maximum packing of P where all bones have at least one endpoint on P within 2 units from a vertex of P. We call such a packing* fully regularized.

Proof. Suppose all maximum packings of P with a minimum number of interior holes have at least l bones that are at least $k \geq 3$ units away from a vertex of P. By performing a left or right shift on a bone of distance k from a vertex, a packing of the same size is obtained that has a bone that is $k - 2$ away from the closest vertex, a contradiction. Thus the lemma follows. □

Lemma 4 *In a fully regularized packing of P, all unit holes have x and y coordinates within $O(n)$ of that of two vertexes of P. Thus for any P there are only $O(n^4)$ possible unit hole locations.*

Proof. By Lemma 3, there are only $O(n)$ bones in a fully regularized packing. Since any unit hole is at the intersection of two bones, there are only $O(n)$ unit holes on a bone, and $O(n^2)$ unit holes total in a regular packing. The horizontal extent of a vertical bone (or vice-versa) is given by the number of unit holes on it and is thus $O(n)$. Since unit holes can only appear at the

[1]From here on consider all definitions and claims to include the symmetric and rotated variants

intersection of bones, their x and y-coordinates must be within $O(n)$ of that of a vertex of P. □

Call a bone in a regularized packing of P that is not incident to a unit hole a *chordal bone*. By Lemma 3 there are only $O(n)$ possible locations for a chordal bone in a regularized packing.

4 Algorithm

Lemma 5 *Given a polygon P with n vertexes presented as a array of binary numbers with a total of N bits, a skeleton which partitions P into m regions, and a type assignment for each region, the size of the packing can be determined in time polynomial in N and m.*

Proof. For any polygonal region with k edges, determining the number of blocks of one particular type that can be packed in is simple to do (e.g. via line sweep) and can be done with $k \log k$ coordinate arithmetic operations. Since each coordinate is a number with at most N bits, coordinate arithmetic can be done in time polynomial in N. For any one region we know $k = O(n)$ and $n \leq N$. The count operation can be repeated for each of the m regions to obtain a total count. □

Theorem 6 *Given a simple orthogonal grid polygon P with n vertexes presented as a array of binary numbers with a total of N bits and a positive integer k, the decision problem* $\max(P) \leq k$ *is in NP.*

Proof. There are only $O(n^4)$ possible locations for unit holes and $O(n)$ locations for chordal bones in a regularized packing of a polygon P. A subset of these can be chosen nondeterministically, and the spin of the holes can be chosen nondeterministically as well. From this the complete skeleton of size at most $O(n^4)$ can be computed in polynomial time by extending four skeleton lines form every unit hole in a manner consistent with its parity. If two lines of the skeleton cross not at a unit hole, it is not a valid skeleton. This leaves a partition of the plane into regions. Each unit hole, together with its spin, assigns a type to each of the four incident regions; if there is any conflict it is not valid. Finally, regions incident only to chordal bones (and not unit holes) do not have a type assigned. Since there are only four types and $O(n)$ such regions, this can be done in *NP*. Finally by Lemma 5, the size of the packing can be computed. If it is at most k then "yes" is returned by that branch of the nondeterministic algorithm. □

Corollary 7 *Given a simple orthogonal grid polygon P with n vertexes presented as a array of binary numbers with a total of N bits and a positive integer k, the decision problem* $\max(P) \leq k$ *is NP-complete.*

Proof. Follows from Theorem 6 and the proof of *NP*-hardness in [3]. □

5 Acknowledgments

We would like to thank Boris Aronov and Thomas Fevens for their helpful conversations on this problem.

References

[1] Erik D. Demaine, Joseph S. B. Mitchell, and Joseph O'Rourke. The open problems project. `http://maven.smith.edu/~orourke/TOPP/`.

[2] Dania El-Khechen. *Decomposing and Packing Polygons.* PhD thesis, Concorida University, 2009.

[3] Robert J. Fowler, Mike Paterson, and Steven L. Tanimoto. Optimal packing and covering in the plane are NP-complete. *Inf. Process. Lett.*, 12(3):133–137, 1981.

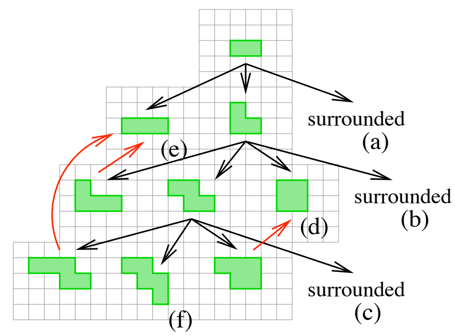

Figure 2: When there is a 1×2 uncovered region in an interior hole, this figure shows a decision process to determine which case should be used, excluding symmetries. For example, look at the root and its three children. At the root is a 1×2 uncovered region, and there are three choices. This region is either surrounded (it is a hole with block on all sides) or there is at least one additional uncovered square adjacent to its boundary. This could result in one of two (excluding symmetries) possible three-square uncovered regions. The surrounded case is (a) and is illustrated in Figure 5, while the case of three empty square in a row is (e) which is considered in Figure 9. However, the 2×2 L-shaped case requires additional subcases determined in a similar manner to that described for the subcases of the root. The red arrows indicate when we reach a case that properly contains another case.

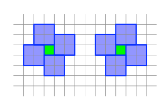

Figure 3: Spin

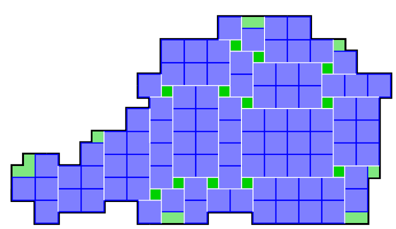

Figure 4: Illustration to help the reader understand the terminology. The green regions are holes. The dark green regions are internal 1×1 holes, or unit holes. The light green regions are external holes. The white lines are the bones, and separate blocks (blue squares) of differing parity. The bones do not intersect and meet internally only at unit holes. The two leftmost bones, which span the polygon without encountering a unit hole, are chordal bones.

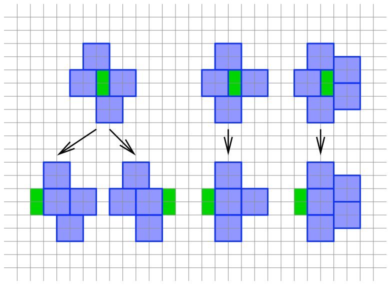

Figure 5: Case (a). Given a 1×2 hole completely surrounded by blocks, there are three possible local configurations, excluding symmetries (top). By moving one block on each of the right two subcases to create the figures on the bottom, the packing size stays the same but the alignment score increases by two. In the left subcase, the alignment score does not decrease, but one of the two moves illustrated will result in a decrease of gravity, depending on the direction of gravity. All subcases result in two new uncovered adjacent holes.

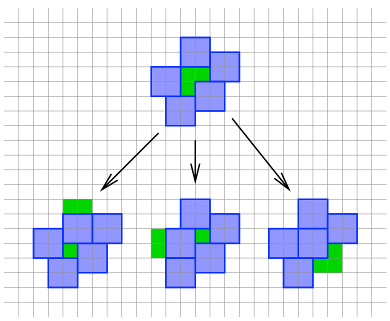

Figure 6: Case (b). Given a 2×2 L shape hole completely surrounded by blocks, there is only one possible local configuration, excusing symmetries (top). There are three possible reconfigurations illustrated (bottom). All of them yield packings of the same size; none of them will ever have a decrease in alignment score. Depending on the direction of gravity, one of the three will have a reduction in the gravity score. All result in two new uncovered adjacent holes.

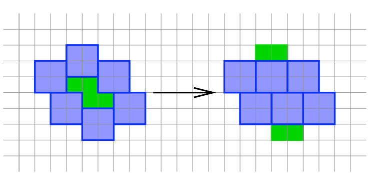

Figure 7: Case (c). Given the pictured S-shape hole (green) completely surrounded by blocks, there is only one possible local configuration, excusing symmetries (right). By moving the two blocks inwards, the size of the packing stays the same but the alignment score increases by 4. This creates two new pairs of uncovered adjacent holes

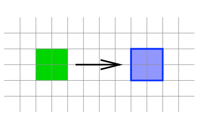

Figure 8: Case (d). Place a block in the hole. This gives a better packing, contradicting maximality.

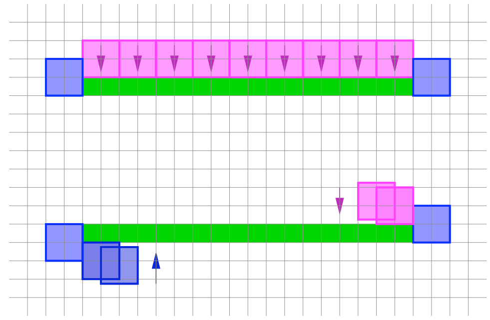

Figure 9: Case (e). Here we have a 1×3 hole. The green rectangle is the maximal $1 \times k$ hole, $k \geq 3$ that the 1×3 hole is contained in. There are two cases depending on whether the two blocks at the end of the green rectangle have the same (top) or different (bottom) vertical parities. In the top case, the pink blocks are all those blocks incident to the top edge of the green hole. There must be at least one or else there would be a 2×2 hole and case (d) could been applied. By moving all of these pink squares down one, the alignment score increases by 2. In the bottom case there must be a block in one of the two extreme right positions above the green rectangle (pink) and one block in one of the two extreme left positions below the green rectangle (blue); otherwise there would be a empty 2×2 hole, a contradiction by case (d). Depending on the direction of gravity, either moving the pink block down or the blue block up will give a packing of the same size and alignment score but with a smaller gravity score. All subcases result in two new uncovered adjacent squares.

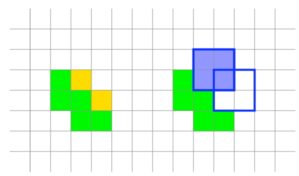

Figure 10: Case (f) Given the green three-step staircase as a hole or subset of a hole (left), there must be blocks in the two orange shaded locations, or else there would be a 2×2 hole and case (d) could be applied. But placing a block in either location (right) forces the other location to be empty, a contradiction.

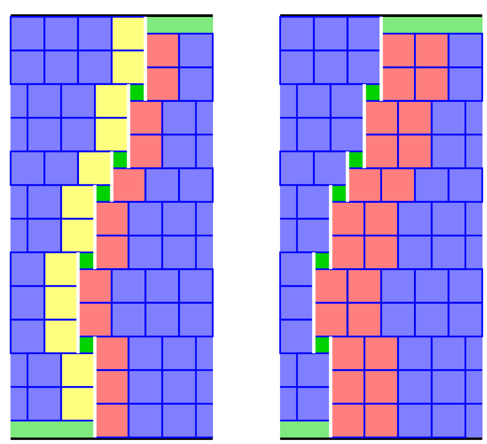

Figure 11: Illustration to help the reader understand the shift left operation. The packing before the shift is to the left. The bone b is illustrated in white, the blocks in $l(b)$ are shaded yellow while the blocks in $r(b)$ are shaded red. On the right is the result of the shift. The yellow blocks have been removed and a copy of the red blocks has been placed two units to the left. The bone has moved two units left.