Decomposing Polygons Into Diameter Bounded Components

Chris Worman*

Abstract

A decomposition of a polygon P is a set of polygons whose geometric union is exactly P. We consider the problem of decomposing a polygon, which may contain holes, using subpolygons that have a bounded diameter. We show that this problem is NP-complete via a reduction from $Planar\ 3,4SAT$.

1 Introduction

Polygon decomposition problems arise in applications where objects represented by polygons need to be subdivided for the sake of tractability. Many variations of decomposition problems have recieved attention in the literature. The reader is directed towards [5] for a synopsis of recent polygon decomposition results. Of particular interest are those results concerning the decomposition of non-simple polygons. The problem of minimally decomposing a polygon that may contain holes has proven to be difficult, and is typically NP-hard.

Bounding box heuristics are commonly used in object intersection algorithms. It has been shown that these algorithms have better performance guarantees when the bounding boxes have similar sizes [6]. This result motivates Damian-Iordache [3] to explore the idea of restricting the diameter of the components in the decomposition of a polygon. Damian-Iordache is able to develop a polynomial time algorithm for partitioning a simple polygon into the minimum number of components that have a maximum diameter of α. Here α is a fixed real number that is part of the input to the partioning algorithm. The problem of decomposing a polygon, which may have holes, with the minimum number of diameter bounded components is conjectured to be NP-hard [3]. We confirm this conjecture by reducing $Planar\ 3,4SAT$ to the corresponding covering and partitioning decision problems.

2 Definitions

A polygon may or may not contain *holes*, which are non-overlapping simple polygons that are completely inside the polygon. The interiors of the holes are considered to be removed from the polygon that contains the holes. A polygon is *simple* if it does not contain holes and only adjacent edges intersect. A *decomposition* of a polygon P is a set of polygons whose union is exactly P. A *covering* of a polygon is a decomposition where the subpolygons are allowed to overlap. In a *partition* of a polygon, the subpolygons in the decomposition do not overlap. We insist that all decompositions are free of Steiner points, which means that subpolygons may only use vertices from P, where P is the polygon being covered.

We define the diameter of a polygon P to be the side length of the orthogonal bounding square for P[1]. A polygon is said to be α-*boundable* if its diameter is less than or equal to α. A decomposition is called an α-*decomposition* if every member of the decomposition is α-boundable. The $Minimum$ α-*cover* problem seeks to find a covering of minimum cardinality such that each member of the covering is α-boundable. The decision version of this problem, which we call *Decide* α-*cover*, asks whether or not there exists a set of k or less α-boundable polygons that cover P. The $Minimum$ α-*partition* and *Decide* α-*partition* problems are defined analogously for partitions.

3 Decomposing Non-Simple Polygons

As with many decomposition problems on polygons with holes, *Decide* α-*cover* and *Decide* α-*partition* are NP-complete. We begin with the covering version of the problem.

Theorem 1 *Decide* α-*cover is NP-complete for non-simple polygons.*

Proof. *Decide* α-*cover* is in NP since we can "guess" a set of k α-boundable subpolygons and verify that they cover P.

We now proceed to reduce $Planar\ 3,4SAT$[4] to *Decide* α-$cover$. $Planar\ 3,4SAT$ is very similar to $Planar\ 3SAT$. In the $Planar\ 3SAT$ problem we want to decide if a given boolean function ϕ is satisfiable. The problem stipulates that ϕ is in conjunctive normal form with exactly 3 literals per clause. The set of literals found in ϕ is referred to as U and the set of clauses is C. Another restriction is placed on ϕ in this problem: the graph $G = (V, E)$ is planar, where $V = U \cup C$ and

*Department of Computer Science, University of Saskatchewan, cmw133@mail.usask.ca

[1]The definition of diameter used herein is different then the one used in [3]. As we shall see in Section 4, we can use either definition of diameter in the NP-completeness proof without loss of correctness.

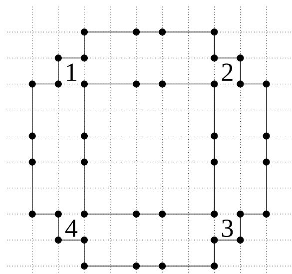

Figure 1: The variable polygon

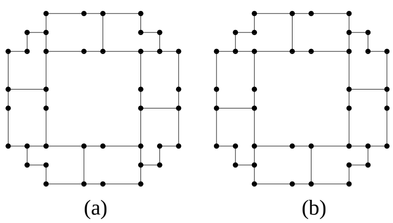

Figure 2: Representing truth assignments (a) True (b) False

$E = \{(u,c) | u \in U, c \in C, u \text{ } or \text{ } \bar{u} \text{ } is \text{ } a \text{ } literal \text{ } in \text{ } c\}$. In the *Planar* $3,4SAT$ problem there is one more restrcition that is imposed: all variables appear at most 4 times negated or unnegated within ϕ. The *Planar* $3,4SAT$ problem was shown to be NP-complete in [4].

In our reduction, various polygon components will be constructed in a manner that is similar to that presented in [1, 2]. These polygons will have a direct correspondance to the components of the graph G that is described above. The polygons will be joined together to form one large polygon P that will have a minimum α-covering of size k iff ϕ is satisfiable. The value of k will be determined during the construction of P. For the remainder of this discussion we will fix α at 3 for reasons that will become apparent.

Variable Polygons: The polygon used to represent a variable, called the *variable polygon*, is given in Figure 1. Recall that we have fixed α at 3. This polygon can be minimally covered by 8 polygons in exactly 2 ways (Figure 2). One of the coverings will represent the variable being set to true and the other will represent false. Wires will attach to variable polygons at 1 of the 4 labeled terminals.

Wire Polygons: The truth value of a variable will be "transmitted" from a variable polygon to a clause polygon using a sequence of one or more *wire polygons*, or simply *wires* (Figure 3). Wire polygons can be attached to variable polygons, other wires, and clause polygons. A sequence of wires connecting a variable polygon to a clause polygon will represent an edge from G. Wires need to be slimmed down near the terminals of variable polygons so they can attach properly (Figure 4). The orientation of the attachment will determine whether the variable is to be negated in the connecting clause. Figure 4(a) shows a variable polygon that has been set to true. The covering subpolygons within the variable polygon that overlap with terminals have been shaded in. Notice that these shaded subpolygons can extend over the tip of the outgoing wires. This is because the outgoing wires are all in the unegated orientation. This is not the case in Figure 4(b), where the outgoing wires are in the negated position. Here the shaded polygons cannot extend over the tip of the wires since we have set α to 3. This illustrates how truth values travel along wires. Wires that carry true will have the top portion covered by a subpolygon from a variable polygon. When a wire is connected to another wire the truth value will propagate to the next wire. Thus wires that carry true will have the "advantage" of being coverable by one less polygon. Clause polygons will be constructed in a way that exploits this. For a wire to connect a variable polygon to a clause polygon it may have to be bent, shifted, or offset (Figure 5). A wire may need to be bent or shifted in order to avoid other components, and it may need to be offset if the clause it is connecting to is 1 unit too close to the wire. If we use one of components from Figure 5 then k must be updated accordingly. If we use an offset component, for example, then we must increase k by 4. The remaining segment of a wire carrying false will be covered by a clause polygon.

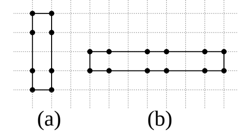

Figure 3: (a) A wire polygon (b) 2 wire polygons that are connected

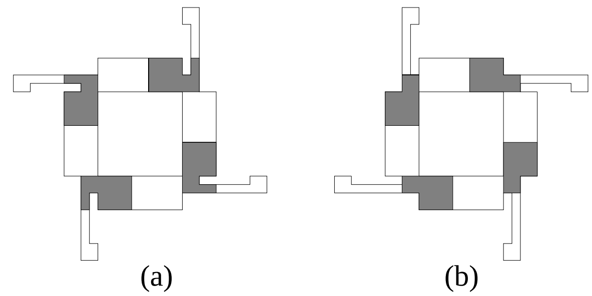

Figure 4: (a) A variable polygon set to true with outgoing wires orientated in the unegated position (b) A vairable polygon set to true with outgoing wires orientated in the negated position

Clause Polygons: The *clause polygon* is shown in Figure 6(a). A clause polygon has 3 terminals where incoming wires will be attached. Since each clause has 3

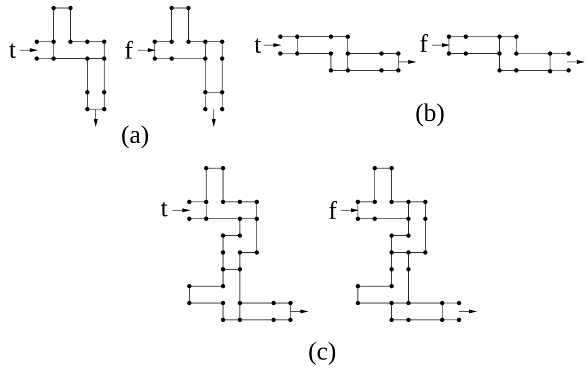

Figure 5: (a) Bending a wire (b) Shifting a wire (c) Offsetting a wire

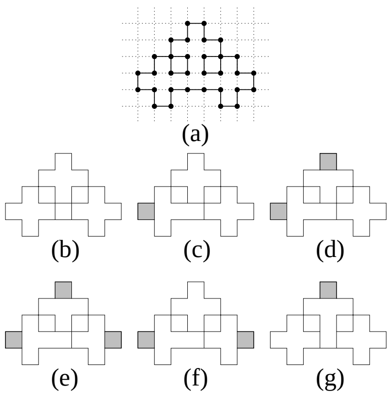

Figure 6: (a) The clause polygon (The symetric cases are not shown). (b) false, false, false (c) true, false, false (d) true, true, false (e) true, true, true (f) true, false, true (g) false, true, false

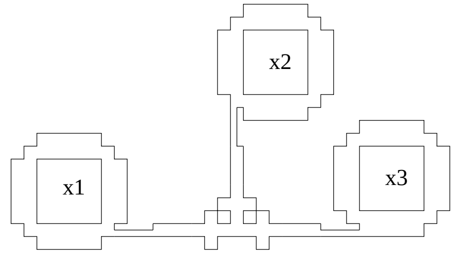

Figure 7: The polygon for the boolean expression $\phi = (x_1 \vee x_2 \vee \overline{x_3})$

literals, each clause polygon will have 3 incoming wires. The size of the minimum covering for some clause polygon will depend on the wires that are attached to it (Figure 6). If 1 or more of the incoming wires is carrying true then the clause polygon will require 3 polygons to cover it. If all incoming wires are carrying false then the polygon will require 4 polygons to be covered. When calculating k, each clause polygon will contribute 3 polygons.

Figure 7 shows the complete polygon P for the boolean expression $\phi = (x_1 \vee x_2 \vee \overline{x_3})$. To calculate k, each component must be accounted for. The 3 variable polygons contribute a total of $8 \times 3 = 24$ polygons, the 6 wires contribute 6, and the clause polygon adds 3. Thus we have $k = 24 + 6 + 3 = 33$. The minimum covering has cardinality 33 and this covering corresponds to a satisfying assignment for ϕ. We still must show that for an arbitrary instance of $Planar\ 3,4SAT$, P will have a minimum covering of size k if and only if ϕ is satisfiable.

Suppose that P has a minimum covering M of size k, where k is the value that was calculated during the construction of P. Consider the truth assignment associated with the minimum covering M. Variable polygons can only be minimally covered 2 ways, which ensures that each variable has a truth value. The value of k can be expressed as $8v + 1w + 3c$, where v is the number of variable polygons, w is the number of wire polygons, and c is the number of clause polygons. Since variable polygons can only be minimally covered 2 ways using 8 polygons we know that $8v$ of the polygons in M are used for covering variable polygons. A sequence of l wires will need l subpolygons to cover them regardless of the truth value that is being transmitted. If they are transmitting true, then they will need l subpolygons to cover the remaining portion of the wires that were not covered by the variable polygon attached to the wire. If the wires are transmitting false then l subpolygons will be necessary to cover all but the last unit of the last wire in the sequence. Thus we know that at least w of the polygons in M must be covering wire polygons. The remaining $3c$ polygons in M are used for covering clause polygons. Recall that a clause polygon needs at least 3 polygons to cover it. Since there are c clause polygons and only $3c$ polygons left in M, we know that each of the clause polygons was covered using 3 polygons. This corresponds to each clause being satisfied, and hence ϕ is satisfiable.

Assume that ϕ is satisfiable under some truth assignment T. We will show how to cover P with k polygons using T. First we cover the variable polygons with 8 polygons each according to T. Use w polygons to cover the wires, which may leave the tips of some wires uncovered. These tips will be within clause polygons. Since we know that T satisfies ϕ, every clause will have at least 1 incoming wire that does not have an uncovered tip. Since this is the case, each clause polygon can be covered with 3 polygons. Thus our covering has size $8c + w + 3c = k$. $\square$

Note that the above proof also works for the partitioning problem.

Theorem 2 *Decide α-partition is NP-complete for non-simple polygons.*

4 Conclusion

We have shown that *Decide α-cover* and *Decide α-partition* are NP-complete when considered on polygons with holes. These results can be extended in two ways. Consider these problems when the diameter of a polygon is computed as the diameter of the smallest circle that covers the polygon. Theorems 1 and 2 can be repeated for this problem if we fix α to be the diameter of the smallest bounding circle for a box whose sides are 3 units long[2]. Also, the polygon constructed in the reduction is orthogonal, and hence we can say that these problems are NP-complete for orthogonal polygons that may contain orthogonal holes. The time complexity of *Minimum α-covering* is still unknown for simple polygons.

5 Acknowledgements

The author would like to thank Mark Keil for useful discussions regarding this problem.

References

[1] L. J. Aupperle. Covering regions by squares. Master's thesis, University of Saskatchewan, July 1987.

[2] H. E. Conn and J. O'Rourke. Some restricted rectangle covering problems. Technical Report JHU87 /13, The John Hopkins University, Baltimore, MD, 1987.

[3] M. Damian-Iordache. *Shape Constrained Polygon Decomposition and Graph Domination Problems.* PhD thesis, University of Iowa, July 2000.

[4] K. Jansen and H. Müller. The minimum broadcast time problem for several processor networks. *Theoretical Computer Science*, 147(1–2):69–85, 1995.

[5] J. M. Keil. Polygon decomposition. In *Handbook of Computational Geometry*, chapter 11. Elsevier Science B. V, 1999.

[6] Y. Zhou and S. Suri. Analysis of a bounding box heuristic for object intersection. In *SODA: ACM-SIAM Symposium on Discrete Algorithms (A Conference on Theoretical and Experimental Analysis of Discrete Algorithms)*, 1999.

[2] Although the radius of such a circle will be $\sqrt{18}$, and hence irrational, we can use an approximate value, such as 4.242641, without loss of correctness.