

1 Snipperclips: Cutting Tools into Desired
2 Polygons using Themselves*

3 Zachary Abel[†] Hugo Akitaya[‡] Man-Kwun Chiu[§]
4 Erik D. Demaine[†] Martin L. Demaine[†] Adam Hesterberg[†]
5 Matias Korman[¶] Jayson Lynch^{||} André van Renssen^{**}
6 Marcel Roeloffzen^{††}

7 **Abstract**

8 We study *Snipperclips*, a computer puzzle game whose objective is
9 to create a target shape with two tools. The tools start as constant-
10 complexity shapes, and each tool can snip (i.e., subtract its current shape
11 from) the other tool. We study the computational problem of, given a
12 target shape represented by a polygonal domain of n vertices, is it possible
13 to create it as one of the tools' shape via a sequence of snip operations? If
14 so, how many snip operations are required? We consider several variants
15 of the problem (such as allowing the tools to be disconnected and/or using
16 an undo operation) and bound the number of operations needed for each
17 of the variants.

18 **1 Introduction**

19 *Snipperclips: Cut It Out, Together!* [10] is a puzzle game developed by SFB
20 Games and published by Nintendo worldwide on March 3, 2017 for their new
21 console, Nintendo Switch. In the game, up to four players cooperate to solve

*An extended abstract of this paper appeared in the proceedings of the 29th Canadian Conference on Computational Geometry (CCCG 2017) [4]. M. C. was supported by ERC StG 757609. M. K. was partially supported by MEXT KAKENHI Nos. 12H00855, and 17K12635. M.-K. C., M. R. and A. v. R. were supported by JST ERATO Grant Number JPMJER1201, Japan.

[†]Massachusetts Institute of Technology, Cambridge, Massachusetts, USA, {zabel, edemaine, mdemaine, achester, }@mit.edu

[‡]University of Massachusetts Lowell, USA, hugo_akitaya@uml.edu

[§]Institut für Informatik, Freie Universität Berlin, chiumk@zedat.fu-berlin.de

[¶]Siemens EDA (formerly Mentor Graphics), OR, USA. matias_korman@mentor.com.

^{||}University of Waterloo, Ontario, Canada. jayson.lynch@uwaterloo.ca

^{**}University of Sydney, Sydney, Australia, andre.vanrenssen@sydney.edu.au

^{††}TU Eindhoven, Eindhoven, the Netherlands, m.j.m.roeloffzen@tue.nl

22 puzzles. Each player controls a character¹ whose shape starts as a rectangle in
23 which two corners have been rounded so that one short side becomes a semi-
24 circle. The main mechanic of the game is *snipping*: when two such characters
25 partially overlap, one character can *snip* the other character, i.e., subtract the
26 current shape of the first character from the current shape of the latter charac-
27 ter; see Figure 1 (top middle) where the yellow character snips the red character
28 subtracting from it their intersection (which is shown in green). In addition, a
29 *reset* operation allows a character to restore its original shape. Finally, an *undo*
30 operation allows a character to restore its shape to what it was before the prior
31 snip or reset operation. A more formal definition of these operations follows
32 in the next section. An unreleased 2015 version of this game, *Friendshapes* by
33 SFB Games, had the same mechanics, but supported only up to two players [6].

34 Puzzles in Snipperclips have varying goals, but an omnipresent subgoal is
35 to form one or more players into desired shape(s), so that they can carry out
36 required actions. In particular, a core puzzle type (“Shape Match”) has one
37 target shape which must be (approximately) formed by the union of the char-
38 acters’ shapes. In this paper, we study when this goal is attainable, and when
39 it is, analyze the minimum number of operations required.

40 2 Problem definition and results

41 For the remainder of the paper we consider the case of exactly two characters
42 or *tools* \mathcal{T}_1 and \mathcal{T}_2 . For geometric simplicity, we assume that the initial shape of
43 both tools is a unit square. Most of the results in this paper work for nice (in
44 particular, fat) constant-complexity initial shapes, such as the rounded rectangle
45 in Snipperclips, but would result in a more involved description.

46 We view each tool as an open set of points that can be rotated and trans-
47 lated freely.² After any rigid transformation, if the two tools have nonempty
48 intersection, we can *snip* (or *cut*) one of them, i.e., remove from one of the tools
49 the closure of the intersection of the two tools (or equivalently, the closure of
50 the other tool, see Figure 2). Note that by removing the closure we preserve the
51 invariant that both tools remain open sets. In addition to the snip operation,
52 we can *reset* a tool, which returns it back to its original unit-square shape.

53 After a snip operation, the changed tool could become disconnected. There
54 are two natural variants on the problem of how to deal with disconnection. In
55 the *connected model*, we force each tool to be a single connected component.
56 Thus, if the snip operation disconnects a tool, the user can choose which com-
57 ponent to use as the new tool. In the *disconnected model*, we allow the tool
58 to become disconnected, viewing a tool as a set of points to which we apply
59 rigid transformations and the snip/reset operation. The Snipperclips game by

¹The game in fact allows one human to control up to two characters, with a button to switch between which character is being controlled.

²In the actual game, the tools’ translations are limited by gravity, jumping, crouching, stretching, standing on each other, etc., though in practice this is not a huge limitation. Rotation is indeed arbitrary.

60 Nintendo follows the disconnected model, but we find the connected model an
61 interesting alternative to consider.

62 The actual game has an additional *undo/redo* operation, allowing each tool
63 to return into its previous shape. For example, a heavily cut tool can reset
64 to the square, cut something in the other tool, and use the undo operation to
65 return to its previous cut shape. The game has an undo stack of size 1; we
66 consider a more general case in which the stack could have size 0, 1 or 2.

67 2.1 Results

68 Given two target shapes P_1 and P_2 , we would like to find a sequence of snip/reset
69 operations that transform tool \mathcal{T}_1 into P_1 and at the same time transform \mathcal{T}_2




Figure 1: Cropped screenshots of Snipperclips: snipping, resetting, and solving a Shape Match puzzle. Sprites copyright SFB/Nintendo and included here under Fair Use.




Figure 2: By translating and rotating the two tools we can make them partially overlap (left figure). On the right we see the resulting shape of both tools after the snip operation.

| Undo stack size | Connected Model | | Disconnected Model | |
|-----------------|-----------------|------------|--------------------|------------|
| | 1 shape | 2 shapes | 1 shape | 2 shapes |
| 0 | $O(n)$ | No | $O(n^2)$ | No |
| 1 | $O(n)$ | $O(n + m)$ | $O(n)$ | Yes |
| 2 | $O(n)$ | $O(n + m)$ | $O(n)$ | $O(n + m)$ |

Table 1: Number of operations required to carve out the target shapes of n and m vertices, respectively. A cell entry with “No” means that it is not always possible to do whereas “Yes” means it is possible (but the number of operations needed is not bounded by any function of n or m).

70 into P_2 . Because our initial shape is polygonal, and we allow only finitely many
 71 snips, the target shapes P_1 and P_2 must be polygonal domains of n and m ver-
 72 tices, respectively. Whenever possible, our aim is to transform the tools into the
 73 desired shapes using as few snip and reset operations as possible. Specifically,
 74 our aim is for the number of snip and reset operations to depend only on n
 75 and m (and not depend on other parameters such as the feature size of the
 76 target shape).

77 In Section 3, we prove some lower bound results. First we show in Section 3.1
 78 that, without an undo operation, it is not always possible to cut both tools into
 79 the desired shape, even when $P_1 = P_2$. Then we show lower bounds on the
 80 number of snips/undo/redo/reset operations required to make a single target
 81 shape P_1 . For the connected model, Section 3.2 proves an easy $\Omega(n)$ lower
 82 bound. For the disconnected model, Section 3.3 gives a family of shapes that
 83 need $\Omega(n)$ operations to carve in a natural 1D model, and gives a lower bound
 84 of $\Omega(\log n)$ for all shapes in the 2D model.

85 On the positive side, we first consider the problem without the undo operation
 86 in Section 4. We give linear and quadratic constructive algorithms to carve
 87 a single shape P_1 in both the connected and disconnected models, respectively.

88 In Section 5 we introduce the undo operation. We first show that even a

89 stack of one undo allows us to cut both tools into the target shapes, although
90 the number of snip operations is unbounded if we use the disconnected model.
91 We then show that by increasing the undo stack size, we can reduce the number
92 of operations needed to linear. A summarizing table of the number of snips
93 needed depending on the model is shown in Table 1.

94 2.2 Related Work

95 Computational geometry has considered a variety of problems related to cutting
96 out a desired shape using a tool such as circular saw [3], hot wire [7], and
97 glass cutting [8, 9]. The Snipperclips model is unusual in that the tools are
98 themselves the material manipulated by the tools. This type of model arises
99 in real-world manufacturing, for example, when using physical objects to guide
100 the cutting/stamping of other objects—a feature supported by the popular new
101 *Glowforge* laser cutter [1] via a camera system.

102 Our problem can also be seen as finding the optimal Constructive Solid
103 Geometry (CSG) [5] expression tree, where leaves represent base shapes (in our
104 model, rectangles), internal nodes represent shape subtraction, and the root
105 should evaluate to the target shape, such that the tree can be evaluated using
106 only two registers. Applegate et al. [2] studied a rectilinear version of this
107 problem (with union and subtraction, and a different register limitation).

108 3 Lower Bounds

109 In this section, we first prove that some pairs of target shapes cannot be realized
110 in both tools simultaneously, using only snip and reset operations. Then we
111 focus on achieving only one target shape. In the connected model, we give a
112 linear lower bound (with respect to the number n of vertices of the target shape)
113 on the number of operations to construct the target shape. In the disconnected
114 model, we give a logarithmic lower bound, and give a linear lower bound in a
115 natural 1D version of Snipperclips.

116 3.1 Impossibility

117 We begin with the intuitive observation that not all combinations of target
118 shapes can be constructed when restricted to the snip and reset operations.

119 **Observation 1.** *In both the connected and disconnected models, there is a target*
120 *shape that cannot be realized by both tools at the same time using only snip and*
121 *reset operations.*

122 *Proof.* Consider the target shape shown in Figure 3: a unit square in which we
123 have removed a very thin rectangle, creating a sort of thick “U”. First observe
124 that, if we perform no resets, neither tool has space to spare to construct a thin
125 auxiliary shape to carve out the rectangular gap of the other tool. Thus, after




Figure 3: A target shape that cannot be realized by both tools at the same time.

126 we have completed carving one tool, the other one would need to reset. This
 127 implies that we cannot have the target shape in both tools at the same time.

128 Now assume that we can transform both tools into the target shape by
 129 performing a sequence of snips and resets. Consider the state of the tools just
 130 after the last reset operation. One of the two shapes is the unit square and thus
 131 we still need to remove the thin hole using the other shape. However, because no
 132 more resets are executed, the other tool is currently and must remain a superset
 133 of the target shape. In particular, it can differ from the square only in the thin
 134 hole, so it does not have any thin portions that can carve out the hole of the
 135 other tool.

136 Because the above argument is based solely on the shape of the figure, it
 137 holds in both the connected and disconnected model. \square

138 3.2 Connected Model

139 Next it is easy to see that in the connected model a target shape with $\Theta(n)$
 140 holes requires $\Omega(n)$ operations.

141 **Theorem 2.** *There are target shapes that require $\Omega(n)$ operations (snip, reset,*
 142 *undo and redo) to construct in the connected model.*

143 *Proof.* Consider the target shape to be a square with $n/3$ triangular holes. Since
 144 we consider the connected model, the cutting tool created by any operations is
 145 connected and it can only carve out one hole at a time. \square

146 3.3 Disconnected Model

147 In the disconnected model, we conjecture that most shapes require $\Omega(n)$ snip
 148 operations to produce (see Conjecture 4), but such a proof or explicit shape
 149 remains elusive. The challenge is that a cutting tool may be reused many times,
 150 which for some shapes leads to an exponential speedup. Indeed, we prove in
 151 Theorem 5 that *every* shape requires $\Omega(\log n)$ snips. As a step toward a linear
 152 lower bound, we prove that a natural 1D version of the disconnected Snipperclips
 153 model has a linear lower bound.

154 Define the *disconnected 1D Snipperclips* model (with arbitrarily many tools)
 155 as follows. A *1D tool* is a disjoint set of intervals in \mathbb{R} . A *1D snip* operation takes

156 a translation of one tool, optionally reflects it around the origin, and subtracts
 157 it from another tool, producing a new tool.

158 The main difference with the disconnected model that we consider is that
 159 we allow for arbitrarily many tools. Alternatively, this can be done with two
 160 tools if you can recall any shape that has been created in the past (i.e., having
 161 infinitely many undo, redo, and reset operations).

162 **Theorem 3.** *For M a positive integer, consider the set of all 1D tools consisting*
 163 *of n disjoint intervals having integer endpoints between 0 and M . For all positive*
 164 *integers n and all $\epsilon \in (0, 1)$, for all sufficiently large M , almost all such tools*
 165 *(at least a $1 - \epsilon$ fraction of them) require at least $2n$ 1D snip operations to build*
 166 *from a single 1D tool consisting of a single interval.*

167 *Proof.* Starting from $k = 1$, the k th snip operation is determined by:

- 168 1. A choice of the $k + 1$ existing tools for the cutting tool T ;
- 169 2. A choice of the $k + 1$ existing tools for the cut tool U ;
- 170 3. An offset x_k of U relative to T .

171 If T has interval endpoints t_0, t_1, \dots and U has interval endpoints $u_0, u_1,$
 172 \dots , then each interval endpoint of the tool created by the k th operation is either
 173 t_j or $x_k + u_j$. The first tools have interval endpoints 0 and $x_0 = M$ (the board
 174 width), so by induction on k , each interval endpoint of the tool created by the
 175 k th operation is of the form $\sum_{i \in I} x_i$ for some $I \subset \{0, \dots, k\}$. Therefore, if we
 176 make, with $k < 2n - 1$ operations, a tool with endpoints y_0, \dots, y_{2n-1} , then

177 there is a $(2n - 1) \times (k + 1)$ 0-1 matrix A such that $A \begin{bmatrix} x_0 \\ \vdots \\ x_k \end{bmatrix} = \begin{bmatrix} y_0 \\ \vdots \\ y_{2n-1} \end{bmatrix}$.

178 If this matrix has rank $rk(A) < k + 1$, set $k + 1 - rk(A)$ of the x_i to be
 179 0 such that it still has a solution, and choose $rk(A)$ of the y_i such that the
 180 $rk(A) \times rk(A)$ square matrix B formed by restricting to the rows corresponding
 181 to nonzero x_i and columns corresponding to those y_i is full-rank. $\det(B)$ is a
 182 sum, over $rk(A)!$ permutations σ , of a product of entries of A (or its negative).
 183 The entries of A are 0 or 1, so $0 < |\det(B)| \leq rk(A)! \leq (k + 1)!$. All the y_i
 184 are integers, so for all i , $\det(B)x_i$ is an integer, since we have that $Bx = y$, so
 185 $x = B^{-1}y$, and B^{-1} is $1/\det(B)$ times the cofactor matrix of B (which has only
 186 integer entries).

187 Therefore, the k th snip operation has at most $(k + 1)^2$ choices for the cutting
 188 tools and $(k + 1)!M < (k + 1)^{k+1}M$ choices for the offset x_k , so the number of
 189 choices for operations up to the $(k - 1)$ st is at most $M^{k-1}k^{k^2}$. On the other
 190 hand, the number of 1D tools consisting of n intervals with integer endpoints $y_0,$
 191 \dots, y_{2n-1} between 0 and M is $\binom{M}{2n} > (M - 2n)^{2n} > (\frac{M}{2})^{2n}$. If $k - 1 < 2n$, then
 192 the total number of integer-endpoint tools with n intervals is asymptotically (for
 193 large M) at most $O(M^{-1})$ times the number of integer-endpoint tools we can
 194 build in $k - 1$ steps, so almost all integer-endpoint tools with n intervals require at

195 least $2n$ steps, as claimed. In particular, if $\epsilon \in (0, 1)$ and $M > 2^{2n}(2n)^{(2n)^2}\epsilon^{-1}$,
 196 then at most an ϵ fraction of such tools can be built in fewer than $2n$ snip
 197 operations, as claimed. \square

198 We conjecture that the same linear lower bound applies to the 2D (discon-
 199 nected) model of Snipperclips as well:

200 **Conjecture 4.** *For M a positive integer, consider the collection of all possible*
 201 *2D “comb” tools consisting of a $1 \times M$ rectangle with n disjoint $1 \times t_i$ “teeth”*
 202 *attached above it (by its side of length t_i), where each tooth has integer coordi-*
 203 *ates and $1 \leq t_i \leq M$ so that the construction fits a $2 \times M$ rectangle. For*
 204 *all positive integers n and all $\epsilon \in (0, 1)$, for all sufficiently large M , almost*
 205 *all such tools (a $1 - \epsilon$ fraction of them) require $\Omega(n)$ snip operations to build,*
 206 *even with arbitrarily many tools (and thus with arbitrary undo, redo, and reset*
 207 *operations).*




Figure 4: Illustration of Conjecture 4. Note that because the teeth are disjoint and have integer coordinates, they are at least one unit apart.

208 Unfortunately, a reduction from 2D Snipperclips to 1D Snipperclips remains
 209 elusive. A natural approach is to view a 2D tool T as a set of 1D tools, one for
 210 each direction that has perpendicular edges in T . But in this view, it is possible
 211 in a linear number of snips to construct a 2D tool containing exponentially many
 212 1D tools, by repeated generic rotation and snipping of the tool by itself. The
 213 information-theoretic argument of Theorem 3 might still apply, but given the
 214 exponential number of tool choices in each step, it would give only a logarithmic
 215 lower bound on the number of snips. We can instead prove such a bound holds
 216 for *all* shapes:

217 **Theorem 5.** *Every tool shape with n edges requires $\Omega(\log n)$ snip operations to*
 218 *build from initial shapes of $O(1)$ edges in the disconnected model. This result*
 219 *holds even with arbitrarily many tools (and thus with arbitrary undo, redo, and*
 220 *reset operations).*

221 *Proof.* Each snip operation involving two tools with n_1 and n_2 edges, respec-
 222 tively, produces a shape with at most $n_1 + n_2$ edges. Thus, if we start with tools
 223 having $c = O(1)$ edges, then in k snips we can produce a shape having at most
 224 c^k edges, proving a lower bound of $k \geq \log_c n$. \square

225 **4 Making one shape with snips and resets**

226 **4.1 Connected Model**

227 In the connected model, the shapes must remain connected. Whenever the snip
 228 operation would break a tool into multiple pieces, we can choose one piece to
 229 keep. In this model, we show that $O(n)$ snips suffice to create any polygonal
 230 shape of n vertices.

231 **Theorem 6.** *We can cut one of the tools into any target polygonal domain P_1
 232 of n vertices using $O(n)$ snip operations (and no reset or undo operations) in
 233 the connected model.*

234 *Proof.* The idea is that we can shape \mathcal{T}_2 into a very narrow triangle, a *needle*,
 235 and use that to cut along the edges of the target shape P_1 . Whenever a snip
 236 disconnects the shape, we simply keep the one containing the target shape.
 237 Initially, we start with a long needle to cut the long edges of \mathcal{T}_2 and we gradually
 238 shrink the needle to cut the smaller edges.

239 Let α and h be two small numbers to be determined. Our needle will be an
 240 isosceles triangle, with the two equal-length edges making an angle of α and the
 241 base edge with length at most h . We refer to the *length* of the needle as the
 242 length of the equal-length edges. We will choose α small enough so that (i) the
 243 needle can fit into all reflex vertices, and we choose h small enough so that, (ii)
 244 when placed on an edge of the target polygon, the needle does not intersect a
 245 non-adjacent edge.




Figure 5: (a) The needle is an isosceles triangle with apex at most α and a base edge of length at most h . The equal-length edges have length at most 1 so that the whole triangle can fit inside a tool. (b) Dashed blue lines denote $\text{Ch}(P_1)$. The choice of h guarantees that there is a segment of length $\geq h$ contained on the boundary of \mathcal{T}_1 that can be used to shrink the needle \mathcal{T}_2 .

246 Refer to Figure 5. Let v be an arbitrary vertex on the convex hull of P_1 ,
 247 denoted $\text{Ch}(P_1)$, and let e_1 and e_2 be its incident edges. By the definition of
 248 convex hull, at least one edge in $\{e_1, e_2\}$ has the property that its normal vector
 249 at v is outside of $\text{Ch}(P_1)$. Without loss of generality, let that be e_1 and let u be
 250 the vertex of P_1 whose orthogonal projection u' on e_1 is closest to v and u lies on
 251 the closed half-plane defined by the supporting line of e_1 containing the normal

252 vector. Note that u might be also in the convex hull and then $u = u'$. We first
 253 make \mathcal{T}_2 into a needle of length $1/2$ using 2 snips. Fix a rigid transformation of
 254 P_1 so that it is entirely contained in \mathcal{T}_1 . We no longer move \mathcal{T}_1 . Use the needle
 255 to cut off a 90° wedge at u' containing the segment $u'v$ on its boundary and so
 256 that we do not cut off any point in the interior of P_1 . This is done with at most
 257 4 snips due to the length of the needle.

258 Now we group all edges of P_1 into sets based on their length. Let \mathcal{E} denote
 259 the full set of edges defining P_1 and let \mathcal{E}_i , for $0 \leq i$, be the set of edges whose
 260 length is between 2^{-i-1} and 2^{-i} . To cut along the edges of \mathcal{E}_i , we use a needle
 261 where the equal-length edges have length 2^{-i-2} . Such a needle can cut each
 262 edge in \mathcal{E}_i using at most four snips; see Figure 5 (a). For an edge e , its nearest
 263 other features of P_1 are its two adjacent edges, the vertices closest to the edge,
 264 and the edges closest to its endpoints. We avoid cutting into the adjacent edges
 265 by placing the tip of the needle at the vertex when cutting near a vertex. By
 266 Properties (i)–(ii), we can make e an edge of \mathcal{T}_1 without removing any point in
 267 the interior of P_1 .

268 By making the cuts along the edges in the sets \mathcal{E}_i in increasing order of i the
 269 needle has to only shrink, which is easily done by using the segment $u'v$ in the
 270 perimeter of \mathcal{T}_1 to shorten the needle by placing the short edge of the needle
 271 parallel to $u'v$. This is possible as long as (iii) $h < \|u'v\|$ where $\|\cdot\|$ denotes
 272 Euclidean norm. We are now ready to set α and h . Property (i) is achieved if
 273 α is smaller than every external angle in P_1 . Property (ii) is achieved if h is
 274 smaller than the shortest distance between an edge and a nonincident vertex.
 275 We also have that the length of the initial needle is $1/2$ and thus $\sin(\alpha/2) \leq h$
 276 using the law of cosines.

277 Recall that making the initial needle requires two snips, cutting each edge
 278 requires at most four snips and hence $O(n)$ snips in total, and reducing the
 279 needle length requires one snip per nonempty set \mathcal{E}_i of which there are at most
 280 $O(n)$. Thus, in total the required number of snips is $O(n)$. \square

281 4.2 Disconnected Model

282 We now consider the disconnected model. Recall that in this model we allow
 283 the tools to become disconnected. That is, when a snip would disconnect the
 284 tool, we keep all pieces. This is the actual version implemented in the game.
 285 Unfortunately, the method in the prior section will not work here. The first
 286 issue is that our tool must now remove the full area of the unwanted space
 287 rather than relying on separated components disappearing. The second issue is
 288 that we may cut up the boundary of our target in such a way that we can no
 289 longer ensure we have an exterior edge of sufficient size to efficiently trim our
 290 needle into the next needed shape. To solve these problems we end up using
 291 $O(n^2)$ snips and the reset operation which was not used in the previous section.
 292 The new algorithm works in phases where we only tackle an L-shaped portion
 293 of the shape at a time. This allows us to keep a solid square in the lower right
 294 which is sufficiently large to create the tools we need to carve out the desired
 295 shape. It also ensures that we can isolate the tool which we are using to carve




Figure 6: The squares S_1 and S_2 along with L-shaped region Q_1 and corner c .

296 the target region of the current phase. Thus each phase bounds how far into
 297 the target we must reach and ensures we have a block with which to alter our
 298 carving tool, allowing methods similar to those in Section 4.1 to complete each
 299 phase. We now give a formal description and proof of correctness.

300 In order to carve out a target shape P_1 , we virtually fix a location of P_1
 301 inside \mathcal{T}_1 , pick a corner c of \mathcal{T}_1 (say, the lower right one) and consider the set
 302 of distances $d_1, \dots, d_{n'}$ from each of the vertices in the fixed location of the
 303 target shape P_1 to c in decreasing order under the L_∞ -metric. For simplicity
 304 assume that all distances are distinct, and thus $n' = n$ (this can be achieved
 305 with symbolic perturbation). We refer to the part of \mathcal{T}_1 not in P_1 , i.e., $\mathcal{T}_1 \setminus P_1$,
 306 as the *free-space*. We will remove the free-space in n steps, where in each step
 307 i we remove the free-space from an L-shaped region Q_i that is the intersection
 308 of \mathcal{T}_1 and an annulus formed by removing the L_∞ -ball of radius d_i from the
 309 L_∞ -ball of radius d_{i-1} centered at c . We argue that in each step we will need
 310 $O(n)$ snips and resets, thus creating the target shape in $O(n^2)$ operations. Our
 311 inductive step is given in the following lemma.

312 **Lemma 7.** *The free-space in region Q_i can be removed in $O(n)$ snip and reset*
 313 *operations provided that $\bigcup_{j \geq i} Q_j$ is a square in \mathcal{T}_1 .*

314 *Proof.* Let S_i be the bounding square containing Q_i (see Figure 6) and let F_i be
 315 the set of faces created when removing the boundary edges of the target shape
 316 from Q_i . By definition all vertices of the target shape on Q_i must be on its inner
 317 or outer L-shaped boundary and all boundary segments must fully traverse Q_i ,
 318 i.e., they cannot have an endpoint inside Q_i . It then follows that the set F_i of
 319 faces consists of $O(n)$ constant complexity pieces. Now triangulate all faces of
 320 F_i and let T_i denote the resulting set of triangles (Figure 7). Note that our aim
 321 is to remove some of the triangles of T_i . We will show that we can remove any
 322 triangle that fits in $S_i \setminus S_{i+1}$ with a constant number of cuts.

323 For simplicity in the exposition we first consider the case in which S_{i+1} is
 324 **large**. That is, the side length of S_{i+1} is at least half the side length of S_i .




Figure 7: An L-shaped region Q_i , the edges of the target shape that cross it (thick edges) define F_i . We further triangulate each face (thin edges), and consider the corresponding dual graph (dotted edges).

325 Consider a triangle $t \in F_i$ that needs to be removed. To create a cutting tool
 326 move \mathcal{T}_2 so that its only overlap with \mathcal{T}_1 is S_i . Let S'_i denote the area in \mathcal{T}_2
 327 corresponding to S_i and let t' be the projection of t on \mathcal{T}_2 . Our goal will be
 328 to remove $S'_i \setminus t'$ from \mathcal{T}_2 without affecting t' . Note that we can create a cut
 329 where only S'_i overlaps \mathcal{T}_1 in S_i , so the shape of $\mathcal{T}_2 \setminus S'_i$ does not influence the
 330 cut (Figure 8). That means we do not have to cut it away and we do not need
 331 to worry about cutting part of it while creating a cutting tool within S'_i .




Figure 8: A triangle t in S_i is cut out of \mathcal{T}_2 at t' .

332 Consider the halfspace H defined by one of the bounding lines ℓ of t' that
 333 does not contain t' . We can remove $H \cap S'_i$ by rotating \mathcal{T}_1 so that one of the
 334 sides of \mathcal{T}_1 along which S_{i+1} is situated aligns with ℓ and repeatedly snip with
 335 S_{i+1} in a grid-pattern as shown in Figure 9. Because S_{i+1} is large compared
 336 to S'_i we can remove $H \cap S'_i$ in $O(1)$ snips. We then apply the same procedure
 337 for the other two halfspaces that should be removed to obtain the cutting tool

338 for t . This means that, under the assumption that S_{i+1} is large, each triangle
 339 can be removed in $O(1)$ snips. Since there are $O(n)$ triangles in S_i , the linear
 340 bound holds.




Figure 9: If S_{i+1} is large, we can use it to carve out any desired shape in \mathcal{T}_2 with $O(1)$ snips.

341 It remains to consider the case in which S_{i+1} is small. That is, the side
 342 length of S_{i+1} is less than half that of S_i , and potentially much smaller. Al-
 343 though the main idea is the same, we need to remove the triangles in order, and
 344 use portions of Q_i that are still solid to create the cutting tools.

345 Let G_i be the dual graph of T_i . This graph is a tree with at most three
 346 leaves. Two leaves correspond to the unique triangles t_b and t_r that share an
 347 edge with the lower and right boundary of Q_i respectively and the third exists
 348 only if the top-left corner of Q_i is contained in a single triangle t_{ℓ} , that is,
 349 there is at least one segment contained in Q_i that connects the top and left
 350 boundaries; see Figure 7. Finally, we change the coordinate system so that c
 351 is the origin, and S_i is a unit square (note that the vertices of this square are
 352 $(-1, 1)$, $(-1, 0)$, $(0, 1)$, and $c = (0, 0)$).

353 We process the triangles in the following order. We first process the *cross-*
 354 *triangles*, triangles with one endpoint on the left boundary and one on the top
 355 boundary (if any exist), starting from t_{ℓ} following G_i until we find a triangle
 356 that has degree three in G_i which we do not process yet. The remaining *fan-*
 357 *triangles* form a path in G_i which we process from t_b to t_r .

358 **Cross-triangles.** Recall that, by the way in which we nest regions Q_i , there
 359 cannot be vertices to the right or below S_i . In particular, cross-triangles have
 360 all three vertices in the top and left boundaries of Q_i . Hence, while we have
 361 some cross-triangle that has not been processed, the triangle of vertices $(-1, 0)$,
 362 $(0, 1)$ and c must be present in \mathcal{T}_1 . This triangle has half the area of Q_i and
 363 can be used to create cutting pieces in the same way as in the case where we
 364 assumed S_{i+1} is large. Thus, we conclude that any cross-triangle of Q_i can be
 365 removed from \mathcal{T}_1 with $O(1)$ snips.

366 **Fan-triangles.** We now process the fan-triangles in the path from t_b to t_r
 367 in G_i . We treat this sequence in two phases. First consider the triangles that




Figure 10: The triangle used to cut out the fan-triangles. Cut cross-triangles are above the dashed line and cut fan-triangles are below the dotted line.

368 have at least one vertex on the left edge of S_i (that is, we process triangles up
 369 to and including the triangle that has degree three in G_i if it exists). Consider
 370 the triangle t of vertices $(0, 1)$, $(0, 3/4)$, and $(-1/4, 3/4)$ (see Figure 10). This
 371 triangle has $1/32$ of the total area of S_i . It is also still fully part of \mathcal{T}_1 until
 372 we cut out the triangle of degree 3. That is, every cross-triangle that was cut
 373 is above the diagonal from $(-1, 0)$ to $(0, 1)$ and any fan-triangle that has at least
 374 one vertex on the left edge of S_i and has degree two in G_i is below the line
 375 from $(-1, 1)$ to $(0, 1/2)$ (technically, below the line from $(-1, 1)$ to the top-right
 376 corner of S_{i+1} , but the higher line suffices for our purposes). So we can use this
 377 triangle t as a cutting tool to create the desired triangle in \mathcal{T}_2 to cut out any
 378 undesired fan-triangles up to and including the triangle of degree 3.

379 The remaining triangles have their vertices in the upper edge of S_i and on
 380 the upper or left edge of S_{i+1} . In this case we must be more careful as we cannot
 381 guarantee the existence of a large square in \mathcal{T}_1 . However, we do not have to
 382 clear the entire space S'_i any longer. Instead it suffices to clear a much smaller
 383 area.

384 Let t denote the next triangle to be removed and let B denote the bounding
 385 box of t and c (see Figure 11). As before consider moving \mathcal{T}_2 so that the only
 386 overlap with \mathcal{T}_1 is B , let B' denote this area in \mathcal{T}_2 and t' the projection of t onto
 387 B' . To create a cutting tool we need only remove the area $B' \setminus t'$.

388 As before, we look for a region in \mathcal{T}_1 that has roughly the area of B to use
 389 for carving the desired shape in \mathcal{T}_2 . Let w be the width of B . Also, let h' be the
 390 height of S_{i+1} . Note that the height of B is 1, and since S_{i+1} is small, we have
 391 $h' < 1/2$. By construction of the bounding box, one of the vertices of t will have
 392 x -coordinate equal to $-w$; let q denote this vertex. The y -coordinate y_q of q is
 393 either 1 or h' as it must be on the upper edge of S_i or on the upper boundary
 394 of S_{i+1} —if t has vertices on the left boundary of S_{i+1} , then there is a vertex on
 395 the upper boundary of S_i with lower x -coordinate. Now consider the triangle
 396 with vertices $(0, 1)$, $(0, h')$, q . This triangle has height at least $1 - h' > 1/2$ and
 397 width w , and thus its area is at least $1/4$ of the area of B . As in the previous




Figure 11: The solid areas (grey) and bounding box B when cutting fan-triangles with no vertices on the left boundary of S_i .

398 cases, we use this triangle to create a cutting tool from \mathcal{T}_2 to remove triangle t
 399 from \mathcal{T}_1 .

400 Thus, it follows that all free-space triangles can be removed with a cutting
 401 tool that is constructed from \mathcal{T}_2 in $O(1)$ snip and reset operations, hence we can
 402 clear Q_i of free-space in total $O(n)$ operations. \square

403 Because there are at most n distinct distances, we repeat this procedure at
 404 most n times, giving us the desired result.

405 **Theorem 8.** *We can cut one of the tools into any target polygonal domain P_1*
 406 *of n vertices using $O(n^2)$ snip and reset operations in the disconnected model.*

407 5 Adding the undo operation

408 We now consider a more powerful model in which we can *undo* the k latest
 409 operations performed on either of the tools. More formally, each snip or reset
 410 operation will change the current shape of one of the two tools (if a snip or reset
 411 operation does not change the shape of either tool, we can ignore it). Given a
 412 sequence of such operations, consider the subsequence o_1, \dots, o_m of operations
 413 that have changed the shape of the first tool. Also, let $P_1^{(i)}$ be the shape of
 414 the first tool after o_i has been executed. The k -undo operation on the first tool
 415 replaces the current shape with $P_1^{(m-k)}$. The k -undo operation on the second
 416 tool is defined analogously.

417 In this section we show that the k -undo operation is very powerful, and allows
 418 us to do much more than we can do without it. In particular, we can transform
 419 two tools into any two target polygonal domains in both the connected and
 420 disconnected model. This statement holds even if we force k to be equal to 1.

421 **5.1 Connected Model**

422 We first consider the connected model. The general idea in this case is that we
 423 first construct the target shape in one of the two tools. In order to construct
 424 the target shape into the second tool, we repeatedly create a needle in the first
 425 tool, cut a part of the second tool, and perform an undo operation to return the
 426 first tool to its target shape.

427 **Theorem 9.** *We can cut two tools \mathcal{T}_1 and \mathcal{T}_2 into any two target polygonal*
 428 *domains P_1 and P_2 of n and m vertices respectively using $O(n + m)$ snip, reset*
 429 *and 1-undo operations in the connected model.*

430 *Proof.* Let e_1 be the longest edge of P_1 not on the boundary of the unit square
 431 and e_2 be the longest edge of P_2 not on the boundary of the unit square. Without
 432 loss of generality, we assume that e_1 is longer than e_2 . We apply Theorem 6
 433 to cut \mathcal{T}_1 into P_1 . To create P_2 we will use a needle to cut along edges as in
 434 Theorem 6. Each needle will be cut along e_1 using a small construction along
 435 e_2 . We will ensure the needle can have varying sizes, so we can cut along each
 436 edge in $O(1)$ cuts. We also guarantee that the needle can be created from P_1 in
 437 a single cut, so we can easily undo the operation.




Figure 12: We can use e_1 , e_2 and a small added edge on e_2 to create a needle in \mathcal{T}_1 that can be used to create P_2 in \mathcal{T}_2 . The needle is indicated in purple.

438 We first explain how to create the needle, as also illustrated in Figure 12.
 439 We create the needle from a segment e of P_1 , which is a subsegment of e_1 that is
 440 half the length of e_1 but centered at its center. The cutting tool will consist of a
 441 subsegment ℓ of e_2 and an edge perpendicular to it, creating a 90° angle in the
 442 freespace. The segment ℓ is also half the length of e_2 and centered at its center.
 443 This is to ensure that there is a constant size rectangle above and below e and
 444 ℓ that does not contain edges or vertices of P_1 or P_2 . Now to cut a needle from
 445 along e , assume that e is horizontal with freespace above it and that the edge
 446 perpendicular to ℓ is on its left endpoint oriented upward. Now place the right
 447 endpoint of ℓ on the right endpoint of e and rotate ℓ counterclockwise around
 448 the right endpoint by an arbitrarily small angle so that the right angle is in the

449 interior of P_1 , just below e . This cut will disconnect a needle from the rest of
 450 P_1 with a length proportional to ℓ . By moving ℓ higher before cutting we can
 451 create shorter needles.

452 For this cutting process to work, the triangle created by ℓ and the perpen-
 453 dicular edge must be empty. So this will be the first piece we remove from \mathcal{T}_2
 454 in the process of creating P_2 . How to do this is illustrated in Figure 13 and
 455 described next. We first reset \mathcal{T}_1 and cut a long narrow rectangle out of the
 456 top left corner of \mathcal{T}_2 . This gives us a long vertical edge and a shorter horizon-
 457 tal edge perpendicular to it. We use this structure to create a narrow triangle
 458 along e_1 as described above. This needle is then aligned with e_2 and cuts out a
 459 narrow triangle above e_2 so that an edge perpendicular to e_2 is created that is
 460 sufficiently far from the endpoint of e_2 .

461 The remainder of the process follows that of Theorem 6 where we use needles
 462 of a specific length to cut edges proportional to that length. The one exception
 463 is e_2 , which is cut last. Note that unlike in Theorem 6, the order in which we cut
 464 the edges is no longer relevant, since we can cut the needle to the size required
 465 for the current edge, cut that edge, and then return the needle to its original
 466 length using a 1-undo operation. This guarantees that the perpendicular edge
 467 required stays attached to the main shape and is removed only when no more
 468 needles need to be created. \square

469 5.2 Disconnected Model

470 Finally, we focus our attention on the disconnected model with undo operations.
 471 We show that allowing undo operations reduces the upper bound on the number
 472 of operations required to cut one target shape out of one tool. In fact, we can
 473 cut any two target shapes out of the two tools, but the number of operations
 474 needed for this depends on the size of the undo stack.

475 **Theorem 10.** *We can cut one of the tools into any target polygonal domain P_1*
 476 *of n vertices using $O(n)$ snip, reset and 1-undo operations in the disconnected*
 477 *model.*

478 *Proof.* We first triangulate the free-space $\mathcal{T}_1 \setminus P_1$. Then, we remove each triangle
 479 t by making a congruent triangle t' in \mathcal{T}_2 . Each time we create a triangle t' in
 480 \mathcal{T}_2 we first reset \mathcal{T}_1 and \mathcal{T}_2 . Then, we can remove $\mathcal{T}_2 \setminus t'$ using \mathcal{T}_1 with a constant
 481 number of snips. Since we only apply one operation on \mathcal{T}_1 , we can use an undo
 482 operation to restore \mathcal{T}_1 to its previous shape, which is the partially constructed
 483 shape towards the target shape P_1 . Next, we can cut the triangle t in \mathcal{T}_1 using
 484 the congruent triangle t' in \mathcal{T}_2 . Thus, we use $O(1)$ snip, reset and 1-undo
 485 operations. We apply this process for each triangle in the free-space. Hence,
 486 since the triangulation has linear complexity, we can remove the free-space with
 487 $O(n)$ operations in total. \square

488 Next, we show that we can cut the two tools into any two target shapes
 489 using only snip, reset and 1-undo operations.




Figure 13: Steps illustrating the creation of a freespace triangle above e_2 that can be used to create needles along e_1 . Small squares indicate the shape that is on the undo stack (omitted when not used later).

490 **Theorem 11.** *We can cut two tools \mathcal{T}_1 and \mathcal{T}_2 into any two target polygonal*
 491 *domains P_1 and P_2 using a finite number of snip, reset and 1-undo operations*
 492 *in the disconnected model.*

493 *Proof.* We apply Theorem 10 to cut \mathcal{T}_1 into P_1 . Then, the idea is that we can
 494 shape P_1 into a very narrow triangle, a *needle*, by using one snip operation, and
 495 use the needle to cut all the free-space $\mathcal{T}_2 \setminus P_2$. After we get P_2 , we can perform
 496 a 1-undo operation to restore \mathcal{T}_1 to P_1 .

497 Let α be the smallest angle between any two adjacent edges of P_2 , d be the
 498 length of the shortest edge of P_2 , and h be the shortest distance between any
 499 vertex and a non-adjacent edge of P_2 . These values will define how small the
 500 needle we create needs to be. Let ℓ_1 be the vertical line touching the leftmost
 501 vertex of P_1 . Since there may be multiple such vertices, let p be the bottommost
 502 vertex of P_1 on ℓ_1 . Let ℓ_2 be the vertical line touching the first vertex on the
 503 right side of ℓ_1 in P_1 . We first reset \mathcal{T}_2 to a unit square. We align the left edge
 504 L_2 of \mathcal{T}_2 with ℓ_1 such that \mathcal{T}_2 fully covers P_1 . Then, we shift \mathcal{T}_2 a little bit
 505 to the right such that L_2 is between ℓ_1 and the bisector of ℓ_1 and ℓ_2 , and the
 506 length of the bottommost edge of P_1 between ℓ_1 and L_2 is less than $d/2$. We
 507 cut P_1 with \mathcal{T}_2 so that we have a set T of triangles (or quadrangles) left in \mathcal{T}_1
 508 (see Figure 14).

509 Let e be the bottommost edge of T and let t be the bottommost object of
 510 T . Let R be the function that rotates the input by 180° around the midpoint
 511 of e , i.e., $R(T)$ is the set of triangles (or quadrangles) obtained by rotating T
 512 180° around the midpoint of e , and $R(t)$ is the triangle obtained by rotating t
 513 in the same manner. Notice that the intersection of $R(T)$ and T is only e . Let
 514 R_ϵ be the function that rotates the input by 180° around the midpoint of e and
 515 then rotates it by a small angle ϵ counterclockwise around p of T . We pick a
 516 small $\epsilon < \alpha/2$ such that no triangle in $R_\epsilon(T)$ crosses ℓ_2 , only $R_\epsilon(t)$ intersects
 517 with t , and the distance between $R_\epsilon(p)$ and e is less than $h/2$. We shift \mathcal{T}_2 back
 518 to the left such that L_2 is on ℓ_1 . Then, we perform the rotation R_ϵ on T and
 519 cut \mathcal{T}_2 with $R_\epsilon(T)$. After this cut, we perform an undo operation to restore \mathcal{T}_1
 520 to P_1 and rotate P_1 back to its starting orientation. Finally, we cut P_1 with \mathcal{T}_2
 521 to obtain the needle (see Figure 15).

522 We argue why the final cut indeed leaves only the needle. Since \mathcal{T}_2 almost
 523 covers P_1 except for the missing part $R_\epsilon(T)$, it is essential to show that the
 524 intersection of $R_\epsilon(T)$ and P_1 is the needle. Since $R(T)$ lies between ℓ_1 and the
 525 bisector of ℓ_1 and ℓ_2 , there exists a small ϵ such that $R_\epsilon(T)$ lies between ℓ_1
 526 and ℓ_2 . In addition, e is the bottommost edge of T , so there cannot be any
 527 intersection of $R_\epsilon(T)$ and P_1 below e . The intersection of P_1 and $R(T)$ is e and
 528 all the triangles in $R(T)$ are below e , so we can rotate them by a small angle ϵ
 529 around p so that only one vertex $R_\epsilon(p)$ in $R_\epsilon(T)$ lies above e (see Figure 14).
 530 As one of the endpoints of the edge of P_1 that contains e lies on or to the
 531 right side of ℓ_2 , the intersection of P_1 and $R_\epsilon(t)$ is a triangle. In particular, the
 532 intersection of $R_\epsilon(T)$ and P_1 is a narrow triangle with a base length of at most
 533 $d/2$, height of at most $h/2$ and a small angle of at most $\alpha/2$.

534 After we obtain the needle, we reset \mathcal{T}_2 and use the needle to cut the free-




Figure 14: The figure shows the set T of the triangles and quadrangles (with filled colors) after cutting P_1 with the unit square \mathcal{T}_2 and the set $R_\epsilon(T)$ obtained by rotating T 180° around the midpoint of e and then rotating a small angle ϵ counterclockwise around p .

535 space $\mathcal{T}_2 \setminus P_2$ in a finite number of snip operations, because the free-space is a
 536 compact object. Finally, we perform an undo operation to restore the needle to
 537 P_1 , resulting in the two target polygonal domains. \square

538 Finally, we show that if we are allowed to use a 2-undo operation instead of
 539 a 1-undo, the number of required operations reduces to linear in the complexity
 540 of the two target polygonal domains.

541 **Theorem 12.** *We can cut two tools \mathcal{T}_1 and \mathcal{T}_2 into any two target polygonal*
 542 *domains P_1 and P_2 using $O(n + m)$ snip, reset and 2-undo operations in the*
 543 *disconnected model.*

544 *Proof.* We apply Theorem 10 to cut \mathcal{T}_1 into P_1 . Then, we define a cover of the
 545 free-space $\mathcal{T}_2 \setminus P_2$ with only small right triangles. We remove each right triangle
 546 t by making a congruent triangle t' in \mathcal{T}_1 by performing at most two operations
 547 on P_1 , so we can get the target shape P_2 and restore \mathcal{T}_1 to P_1 .

548 We first explain how to define the cover of the free-space with only right
 549 triangles. We start with any triangulation on the free-space $\mathcal{T}_2 \setminus P_2$. Then, we
 550 subdivide each triangle into a constant number of smaller triangles such that




Figure 15: The figure shows \mathcal{T}_2 after removing $R_\epsilon(T)$ and the part of the boundary of P_1 .

551 each smaller triangle fits in a $\frac{1}{2} \times \frac{1}{2}$ square. For each triangle, we draw a line
 552 segment from the vertex of the largest angle perpendicular to its opposite edge
 553 in order to split the triangle into two right triangles. Hence, there are $O(m)$
 554 right triangles in the cover.

555 Next, we describe how to create the cutting tool in \mathcal{T}_1 (see Figure 16). For
 556 each right triangle t in the free-space, we first reset both \mathcal{T}_1 and \mathcal{T}_2 (P_1 and the
 557 partially constructed P_2 are stored at the top of their stacks). We use the unit
 558 square \mathcal{T}_2 to cut the unit square \mathcal{T}_1 to get a triangle t' congruent to t at a corner
 559 of \mathcal{T}_1 (P_1 is stored at the second element of its stack). Note that there are other
 560 garbage components left in \mathcal{T}_1 . Then, we translate \mathcal{T}_1 in such a way that only
 561 t' overlaps \mathcal{T}_2 , and cut \mathcal{T}_2 to make a square with a triangular hole (the partially
 562 constructed P_2 is at the second element of its stack). We perform an undo
 563 operation to restore \mathcal{T}_1 back to the unit square. The next step is to align the
 564 bounding unit square of \mathcal{T}_1 and \mathcal{T}_2 , and cut \mathcal{T}_1 with \mathcal{T}_2 so that we get only t' in
 565 \mathcal{T}_1 . After we get the cutting tool t' , we perform two undo operations to restore
 566 \mathcal{T}_2 to the partially constructed P_2 , and use t' to remove t from the free-space.
 567 Finally, we perform two undo operations to restore \mathcal{T}_1 to P_1 . Overall, we use
 568 $O(1)$ snip, reset and undo operations to make some progress on \mathcal{T}_2 towards P_2
 569 while maintaining P_1 .




Figure 16: The figure shows how to remove a triangle t in the partially constructed P_2 of \mathcal{T}_2 while maintaining P_1 . Smaller squares indicate which shapes are on the undo stack.

570 We repeat the above process for each right triangle in the free-space, so we
571 use $O(m)$ operations to carve out P_2 . Including the $O(n)$ operations to carve
572 out P_1 , we use $O(n + m)$ operations in total. \square

573 6 Open Problems

574 The natural open problem is to close the gap between our algorithms and the
575 lower bound. Specifically, we are interested in a method that could extend
576 our lower bound approach to the case in which you have the undo operation.
577 We believe that without the undo operation there must exist a shape in the
578 disconnected model that needs $\omega(n)$ operations to carve.

579 Our algorithms focus on worst-case bounds, but we also find the minimiza-
580 tion problem interesting. Specifically, can we design an algorithm that cuts one
581 (or two) target shapes with the fewest possible cuts? Is this problem NP-hard?
582 If so, can we design an approximation algorithm? Although it is not always pos-
583 sible to cut two tools simultaneously into the desired polygonal shapes, it would
584 be interesting to characterize when this is possible. Is the decision problem
585 NP-hard?

586 It would also be interesting to consider the initial shape implemented in the
587 Snipperclips game (instead of the unit squares we used for simplicity), namely,
588 a unit square adjoined with half a unit-diameter disk. This initial shape opens
589 up the possibility of making curved target shapes bounded by line segments and
590 circular arcs of matching curvature. Can all such shapes be made, and if so, by
591 how many cuts?

592 The stack size has a big impact in the capabilities of what we can do and on
593 how fast can we do it. Additional tools can have a similar effect, since they can
594 be used to *store* previous shapes. It would be interesting to explore if additional
595 tools have the same impact as the undo operation or they actually allow more
596 shapes to be constructed faster.

597 Acknowledgments

598 This work was initiated at the 32nd Bellairs Winter Workshop on Computa-
599 tional Geometry held January 2017 in Holetown, Barbados. We thank the other
600 participants of that workshop for providing a fun and stimulating research en-
601 vironment. We also thank Jason Ku for helpful discussions about (and games
602 of) Snipperclips.

603 References

- 604 [1] Glowforge — the 3D laser printer. <https://glowforge.com/>, 2015.
605 [2] D. A. Applegate, G. Calinescu, D. S. Johnson, H. Karloff, K. Ligett, and J. Wang.
606 Compressing rectilinear pictures and minimizing access control lists. In *Proceed-*

- 607 *ings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages
608 1066–1075, New Orleans, Louisiana, 2007.
- 609 [3] E. D. Demaine, M. L. Demaine, and C. S. Kaplan. Polygons cuttable by a circular
610 saw. *Computational Geometry: Theory and Applications*, 20(1–2):69–84, October
611 2001. CCCG 2000.
- 612 [4] E. D. Demaine, M. Korman, A. van Renssen, and M. Roeloffzen. Snipperclips:
613 Cutting tools into desired polygons using themselves. In *Proceedings of the 29th*
614 *Canadian Conference on Computational Geometry (CCCG 2017)*, pages 56–61,
615 Ottawa, Canada, 2017.
- 616 [5] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics:*
617 *Principles and Practice*. Addison-Wesley Professional, 1996.
- 618 [6] GDC. European innovative games showcase: Friendshapes. YouTube video, 2015.
619 <https://youtu.be/WJGooKIoy1Q>.
- 620 [7] J. W. Jaromczyk and M. aw Kowaluk. Sets of lines and cutting out polyhedral
621 objects. *Computational Geometry: Theory and Applications*, 25(1):67–95, 2003.
622 CCCG 2001.
- 623 [8] M. H. Overmars and E. Welzl. The complexity of cutting paper. In *Proceedings*
624 *of the 1st Annual ACM Symposium on Computational Geometry*, pages 316–321,
625 Baltimore, Maryland, June 1985.
- 626 [9] J. Pach and G. Tardos. Cutting glass. *Discrete & Computational Geometry*,
627 24:481–495, 2000.
- 628 [10] Wikipedia. Snipperclips. <https://en.wikipedia.org/wiki/Snipperclips>, 2017.