

RE2L: An efficient output-sensitive algorithm for computing Boolean operations on circular-arc polygons and its applications[☆]



Zhi-Jie Wang^{a,b}, Xiao Lin^c, Mei-E Fang^{d,*}, Bin Yao^a, Yong Peng^d, Haibing Guan^a, Minyi Guo^a

^a Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China

^b Department of Computing, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong

^c Department of Science and Engineering, University of Shanghai for Science and Technology, Shanghai, China

^d School of Computer Science, Hangzhou Dianzi University, Hangzhou, China

ARTICLE INFO

Article history:

Received 30 November 2015

Accepted 9 July 2016

Keywords:

Boolean operations
Circular-arc polygons
Related edges
Sequence lists
Appendix points

ABSTRACT

The boundaries of *conic polygons* consist of conic segments or second degree curves. The conic polygon has two degenerate or special cases: the linear polygon and the circular-arc polygon. The natural problem – Boolean operations on linear polygons – has been extensively studied. Surprisingly, (almost) no article focuses on the problem of *Boolean operations on circular-arc polygons*, yet this potentially has many applications. This implies that if there is a targeted solution for Boolean operations on circular-arc polygons, it should be favourable for potential users. In this article, we close the gap by devising a concise data structure and then developing a targeted algorithm called RE2L. Our method is simple, easy-to-implement, but without loss of efficiency. Given two circular-arc polygons with m and n edges respectively, our method runs in $O(m + n + (l + k) \log l)$ time, using $O(m + n + k)$ space, where k is the number of intersections, and l is the number of related edges (defined in Section 5). Our algorithm has the power to approximate to linear complexity when k and l are small. The superiority of the proposed algorithm is also validated through empirical study. In particular, our method is of independent interest and we show that it can be easily extended to compute Boolean operations of other types of polygons.

© 2016 Elsevier Ltd. All rights reserved.

1. Introduction

Boolean operation on polygons is one of the oldest and best-known problems in computer graphics, and it has attracted much attention, due to its simple formulation and broad applications in various disciplines such as computational geometry, CAD, GIS, visual computing, motion planning [1–8]. When the polygons to be processed are *conic polygons* (whose boundaries consist of conic segments or second degree curves), researchers have made some efforts, see e.g., [9–11]. The conic polygon has several special or degenerate cases: (i) the linear polygon (known as traditional polygon), whose boundaries consist of only linear

curves, i.e., straight line segments; and (ii) the circular-arc polygon, whose boundaries consist of circular-arcs and/or straight line segments. The natural problem – Boolean operations on traditional polygons – has been extensively investigated, see e.g., [12–18, 19–22]. However, in existing literature (almost) no article focuses on another natural problem – Boolean operations on circular-arc polygons. Yet, Boolean operation on circular-arc polygons also has many applications. For instance, deploying sensors to ensure wireless coverage is an important problem [23,24]. The sensing range of a single sensor is a circle. With polygonal obstacles, its sensing range is cut off, shaping a circular-arc polygon. When we verify the wireless coverage range of sensors, Boolean operation on circular-arc polygons is needed. As another example, assume there are a group of free-rotating cameras used to monitor a supermarket. The visual range of a single camera can be regarded as a circle, as the camera is to be freely rotated. Various obstacles such as goods shelves, usually impede the visions of cameras; here Boolean operation can be used to check the blind angles. Last but not least, consider a group of free-moving robots used to guide visitors in a museum. Since the energy of a single robot is limited,

[☆] This paper has been recommended for acceptance by Tamal Dey.

* Corresponding author.

E-mail addresses: cszhijwang@comp.polyu.edu.hk (Z.-J. Wang), lin6008@126.com (X. Lin), fme@hdu.edu.cn (M.-E. Fang), yaobin@cs.sjtu.edu.cn (B. Yao), standy.peng@gmail.com (Y. Peng), hbguan@cs.sjtu.edu.cn (H. Guan), guo-my@cs.sjtu.edu.cn (M. Guo).

its movable region is restricted to a circle. With the impact of various obstacles such as exhibits, the original movable region is restricted by obstacles. When we verify if every place in the museum can be served by at least one robot, Boolean operation on circular-arc polygons is also needed.

Although the solution used to handle conic (or more general) polygons can also work for special cases, these cases, however, usually have their unique properties; directly executing the algorithm used to handle conic (or more general) polygons is usually not efficient enough. It is just like when the applications only involve traditional polygons, we usually incline to use the solutions targeted for traditional polygons, rather than those for conic (or more general) polygons. Using the similar argument, when the applications only involve circular-arc polygons, a targeted solution for circular-arc polygons should be favourable for potential users.

Motivated by this, this paper focuses on the problem of *Boolean operations on circular-arc polygons*. In particular, we are interested in developing algorithms with the following features: (i) easy-to-implement for deployment in practice, and (ii) having nice theoretical guarantees. To this end, first of all a concise and easy-to-operate data structure is naturally developed (Section 4.1). Based on this concise structure, we then propose an algorithm dubbed as RE2L that consists of three main steps.

The first step is the kernel (or core) of RE2L, yielding two *special* sequence lists. Specifically, the kernel integrates three simple yet efficient strategies: (i) it introduces the concept of *related edges*, which is used to avoid irrelevant computation as much as possible; (ii) it employs two *special* sequence lists. Each one is a compound structure with three domains. They are used to let the *decomposed arcs*, intersections and *processed related edges* be well organized, and thus immensely simplify the subsequent computation; and (iii) it assigns two labels to each processed related edge before the edge is placed into a balanced tree; this contributes to avoiding the “false” intersections being reported, and speeding up the process of inserting the reported intersections into their corresponding edges (Section 5). The second step produces two *new* linked lists in which the intersections, appendix points and original vertices have been arranged, and the decomposed arcs have been merged. To obtain these two new linked lists, two important but easy-to-ignore issues, ‘inserting new appendix points’ and ‘merging the decomposed arcs’, are addressed (Section 6). The third step is to obtain the resultant (or output) polygon by traversing these two new linked lists. In order to correctly traverse them, the *entry–exit* properties are naturally adopted, and three traversing rules are developed (Section 7).

Viewed from a macro perspective, similar to many methods (see e.g., [25,15,12,22,13,16]) in the literature, our solution also partially inherits two well-known proposals: the Bentley–Ottmann Plane Sweep algorithm [26] and the Weiler–Atherton Clipping algorithm [6], whereas we also advance existing results from various aspects. To summarize, we make the following main contributions.

1. We highlight the circular-arc polygon is one of special cases of the conic polygon, and Boolean operation on circular-arc polygons also has many applications.
2. We devise a concise and easy-to-operate data structure, and develop a targeted algorithm for Boolean operations on circular-arc polygons.
3. While this paper focuses on Boolean operations of circular-arc polygons, we show that our techniques can be easily extended to compute Boolean operations of other types of polygons (Section 9).
4. We provide a rigorous and detailed theoretical analysis for our algorithm. In brief, given two circular-arc polygons with m and n edges respectively, our algorithm runs in $O(m+n+(l+k) \log l)$ time, using $O(m+n+k)$ space, where k is the number of intersections, and l is the number of related edges (Section 8).

5. We conduct extensive experiments to demonstrate the efficiency and effectiveness of our solution (Section 10).

The novelty of our work is threefold: to the best of our knowledge (i) it is the first comprehensive study on Boolean operations of circular-arc polygons; (ii) it is the first time to employ the idea ‘utilizing related edges’ for Boolean operations of polygons, and this technique is simple enough to be of practical value; and (iii) it is the first output-sensitive algorithm that has the potential to approximate to linear complexity for Boolean operations of polygons.

Next, we review previous works most related to ours, and then present our algorithm including rigorous theoretical analysis and extensive empirical study.

2. Related work

We first clarify several technical terms for ease of presentation. It is well known that there are three typical Boolean operations: intersection, union, and difference. Note that *polygon clipping* mentioned in many papers actually computes the *difference* of two polygons [22]. Given two polygons, the one to be clipped is called the *subject polygon*; the other is usually called the *clip polygon* or *clip window* [15,25,22]. Given a polygon, if there is a pair of non-adjacent edges intersecting with each other, this polygon is usually called the *self-intersection* polygon [12,13,21]. Throughout this paper, the *traditional polygon* refers to the polygon whose boundaries consist of *only* straight line segments, while the *circular-arc polygon* refers to the polygon whose boundaries consist of circular arcs, or both straight line segments and circular arcs. We are now ready to review the previous works most related to ours.

2.1. Boolean operations on traditional polygons

In existing literature, there are many papers that study Boolean operations of traditional polygons. For example, Sutherland–Hodgeman [25] proposed an elegant algorithm dealing with the case when the *clip polygon* is convex. Liang et al. [15] carried out an elaborate analysis on the case when the *clip polygon* is rectangular. Andreev [12] presented an algorithm dealing with the case when the *subject polygon* has holes and self-intersections. Vatti [22] and Greiner–Hormann [13] proposed general algorithms that can handle concave polygons with holes and self-intersections for both the clip and the subject polygons. Later, Liu et al. [16] further optimized Greiner–Hormann’s algorithm. Rivero–Feito [21] achieved Boolean operations of polygons based on the concept of *simplicial chains*. Peng et al. [19] also adopted simplicial chains and improved Rivero–Feito’s algorithm. Recently, Martinez et al. [5] proposed to subdivide the edges at the intersection. These works lay a solid foundation for future research. Compared to these studies, this paper focuses on Boolean operations of circular-arc polygons, and therefore is different.

2.2. Boolean operations on conic/general polygons

Researchers have also made some efforts on Boolean operations of conic polygons. For example, Berberich et al. [9] proposed to decompose non- x -monotone curves and compute the arrangement of segments using the plane sweep method, and then compute the *overlap* of two polygons using the results of the arrangement, in order to achieve Boolean operations. Gong et al. [11] achieved Boolean operations of conic polygons using the topological relation between two conic polygons. This method does not require x -monotone conic arc segments. Both algorithms can support Boolean operations of circular-arc polygons, as the conic polygon is the general case of the circular-arc polygon. Moreover, the computational geometry algorithms library (CGAL) [27] can also support

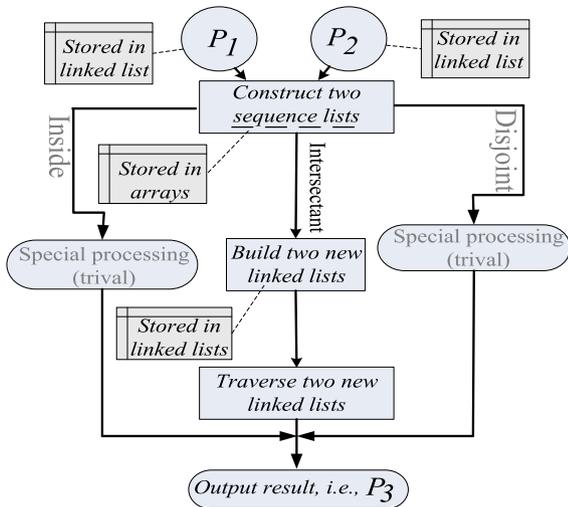


Fig. 1. Framework of our solution. The term ‘Inside’ refers to the case: $\mathcal{P}_1 \subseteq \mathcal{P}_2$ or $\mathcal{P}_2 \subseteq \mathcal{P}_1$. The term ‘Disjoint’ refers to the case: $\mathcal{P}_1 \cap \mathcal{P}_2 = \emptyset$. The term ‘Intersectant’ refers to the case: $\mathcal{P}_1 \cap \mathcal{P}_2 \neq \emptyset$.

Boolean operations of circular-arc polygons. Inspecting the source codes of CGAL, we realize that its idea is to directly invoke the algorithm of Boolean operations on *general polygons*, defined as `GeneralPolygon_2` in CGAL.¹ To some extent the general polygon can be considered as the most general case, as its edges can be line segments, circular arcs, conic curves, cubic curves, or even more complicated curves. Although the essence of the algorithm in CGAL is basically similar to that in [9] (using the *plane sweep method* to compute the intersections, and the DCEL structure to represent the polygons), CGAL is a very powerful and useful library collecting many classical ideas. For example, Emiris et al. [28] developed a kernel for curved objects and related operations that was targeted for inclusion in CGAL. Note that the CGAL project itself also yields many nice papers in which Boolean operations on polygons with curves are mentioned; see for example [10,29], to mention just a few. These excellent works are the cornerstones of our study, giving us a lot of inspiration.

Compared to these works, our work is different from theirs in the following aspects at least. First, this paper focuses on one of special cases of the conic polygon. We give insights into its unique properties, design a concise data structure customized for this special case, and develop a targeted algorithm, in which the central idea ‘utilizing related edges’ (accompanied by a set of well-designed strategies) is proposed. To our knowledge, it is the first time to employ this technique for Boolean operations on polygons. Moreover, we give the rigorous theoretical analysis for our algorithm. This algorithm runs in $O(m + n + (l + k) \log l)$ time, and approximates to linear complexity when k and l are small (notice: the best-known result for polygon Boolean operation runs in $O((m + n + k) \log(m + n))$ time, which is no better than *linearithmic time*² even if k is small); its superiority is also verified by extensive experiments.

3. Solution overview

This section describes our algorithm at a high level. Let \mathcal{P}_1 and \mathcal{P}_2 be the circular-arc polygons to be processed, and E_1 (resp., E_2) be the set of all edges of \mathcal{P}_1 (resp., \mathcal{P}_2), and \mathcal{P}_3 be the resultant polygon

(i.e., the output of our algorithm). The overall framework of our solution is illustrated in Fig. 1. It contains three main steps: (1) construct two sequence lists; (2) build two new linked lists, and (3) traverse two new linked lists. Note that, the first step is the kernel of our algorithm, as the central idea ‘utilizing related edges’ (accompanied by several other helpful strategies) is included in this step.

Before explaining the construction process of two *sequence lists* (i.e., arrays), we first understand their internal structures. In brief, each item in the two sequence lists is a compound structure consisting of multiple domains; these domains are used to store various information such as edges and intersections (Section 5.2). The essence of constructing two sequence lists is to choose a set E'_1 of edges from E_1 and another set E'_2 of edges from E_2 (Section 5.1), and compute the intersections of these edges based on a *modified plane sweep method* (Section 5.3). Note that, some arcs (i.e., edges) in E'_1 and E'_2 may need to be decomposed before computing the intersections. These edges, intersections, and many other pieces of information are stored in the two sequence lists orderly, completing Step 1. It is noteworthy that in the construction phase, we can determine the geometry relation between \mathcal{P}_1 and \mathcal{P}_2 , based on some available information. It is trivial to obtain the resultant polygon when \mathcal{P}_1 and \mathcal{P}_2 are disjoint (or when one is inside another). In the subsequent discussion, we focus our attention on the case when \mathcal{P}_1 and \mathcal{P}_2 intersect with each other, for ease of exposition.

We proceed to explain Step 2 ‘build two new linked lists’. We need to mention that the input polygon \mathcal{P}_1 (\mathcal{P}_2) is also stored using the linked list-like data structure. Unless stated otherwise, the terms ‘input polygons’ and ‘original linked lists’ are used interchangeably in the rest of the paper. Note that, although the original linked lists and the new linked lists use the same data structure, the information stored in them is different. To understand Step 2, consider that we have two sequence lists (obtained in the previous step) and two original linked lists (i.e., input polygons). The task is to efficiently merge part of information in the original linked lists and the information in the two sequence lists, and store the ‘merged information’ using two new linked lists. In particular, for the two new linked lists we are asked to maintain some properties: for example, the intersections and original vertices should be well arranged and the decomposed arcs should be merged (Section 6). Note that although Step 2 partially inherits the well-known Weiler–Atherton Clipping method, the problem considered here is much more challenging, and thus needs more treatments. On the other hand, as we generate two sequence lists in the previous step, the detailed steps of building the two linked lists are basically different from Weiler–Atherton’s, and also other existing methods (see e.g., [25, 15, 12, 22, 13, 16]).

Step 3 essentially adapts existing traversing methods that were used to compute Boolean operations of traditional polygons. In brief, it first assigns *entry–exit* properties to intersections (Section 7.1), and then traverses the two new linked lists based on various traversing rules (Section 7.2). Our paper discusses various Boolean operations (intersection, union and difference), and investigates various circular-arc polygons: with or without holes, self-intersection or non self-intersection (Sections 7 and 9). It is non-trivial and also meaningful to provide a comprehensive investigation, and we believe it would be helpful for interested readers.

In the following sections, we describe in detail our solution including the data structure, main ideas, algorithms, and rigorous theoretical analysis.

¹ For more information please refer to the site: <http://www.cgal.org>.

² Simply speaking, linearithmic time in Big O notation refers to $O(N \log N)$, provided that the input is $O(N)$ size.

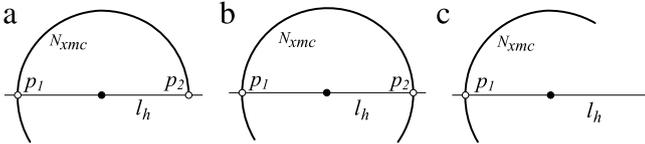


Fig. 2. Illustrations of Lemma 4.1.

4. Preliminary

4.1. Data structure

It is well known that the traditional polygon can be represented by a series of vertices. This method however is invalid for polygons containing circular arcs, as two vertices cannot exactly determine a circular arc segment (note: it may be a major or minor arc). Even so, this ambiguity can be easily eliminated by adding an appendix point, where the appendix point can be an arbitrary point that is located on the arc but it is not the endpoints of the arc. For clarity, a traditional vertex is denoted by v_i , and an appendix point is denoted by \tilde{v}_j . For example, $\{v_1, \tilde{v}_2, v_3, v_4, v_5\}$ determine a circular-arc polygon with four edges (including one circular arc segment $\widehat{v_1\tilde{v}_2v_3}$ and three straight line segments $\overline{v_3v_4}$, $\overline{v_4v_5}$, $\overline{v_5v_1}$). Unless stated otherwise, in the rest of the paper we always use $\overline{\quad}$ and $\widehat{\quad}$ to denote the line segment and the arc segment, respectively. In order to efficiently operate circular-arc polygons, we devise a data structure called APDLL (appendix point based doubly linked list). Specifically, each node in the list consists of several domains below.

- Data: (x, y) , the coordinates of a point.
- Tag: Boolean type; it indicates whether this point is a traditional vertex or an appendix point.
- Crossing: Boolean type; it indicates whether this point is an intersection.
- EE: Boolean type; it indicates what property (entry or exit) an intersection has.
- Prev: Node pointer; it points to the previous node.
- Next: Node pointer; it points to the next node.

4.2. Observation

In this subsection, we introduce a simple yet important observation that will be used frequently later. To explain, we need some preliminaries.

Definition 4.1 (Non-x-monotone Circular Arc). Given any circular arc, it is a non-x-monotone circular arc such that there is at least one vertical line that intersects with the circular arc at two points.

Definition 4.2 (X-monotone Circular Arc). A circular arc is an x-monotone circular arc such that there is at most one intersection with any vertical line.

Lemma 4.1 below formalizes our observation, which can be viewed as a unique property of circular-arc polygons (compared to other types of polygons).

Lemma 4.1. Let N_{xmc} be an arbitrary non-x-monotone circular arc, and C be its corresponding circle. Assume that l_h is a horizontal line passing through the center of C . We have that l_h can decompose N_{xmc} into at least two and at most three x-monotone arcs.

Proof. It is immediate by analytic geometry. See Fig. 2 for some illustrations. □

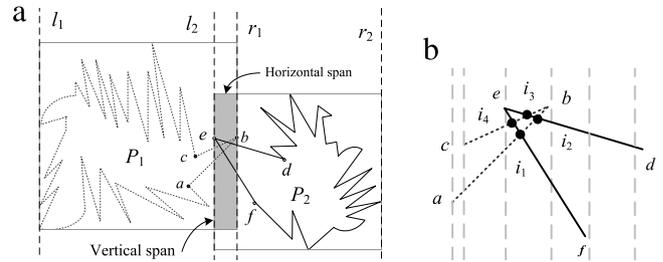


Fig. 3. Example of related edges. (a) Two big rectangles denote the MBRs; the grey rectangle denotes the intersection set of two MBRs, and the dashed vertical lines denote the extended boundary lines. (b) Partial enlarged drawing.

5. The kernel of RE2L

In this section, we detail Step 1 of our solution. Specifically, we first expatiate the main ideas integrated in Step 1 (Sections 5.1–5.3), and then present the detailed algorithms (Section 5.4).

5.1. Utilizing related edges

One of our strategies is to choose related edges (defined later) before doing others. The purpose of choosing related edges is to avoid operations that are irrelevant to obtaining the final result as much as possible. To define related edges formally, we need two notions.

Definition 5.1 (Extended Boundary Lines). Given a circular-arc polygon, w.l.o.g. (without loss of generality), assume the coordinates of the left-bottom corner of its MBR (minimum bounding rectangle) are (x_1, y_1) , those of the right-top corner are (x_2, y_2) . Then, the following four lines, $X = x_1$, $X = x_2$, $Y = y_1$, $Y = y_2$ are respectively the left, right, bottom and top extended boundary lines of this circular-arc polygon.

Definition 5.2 (Effective Axis). Let I_{mm} be the intersection set of two circular-arc polygons' MBRs. If the horizontal span of I_{mm} is larger or equal to its vertical span, then the y-axis is the effective axis. Otherwise, the x-axis is the effective axis.

We now provide the formal definition and inspect more properties of related edges.

Definition 5.3 (Related Edges). Let $l_1(l_2)$ and $r_1(r_2)$ be the left and right extended boundary lines of the circular-arc polygon $\mathcal{P}_1(\mathcal{P}_2)$, respectively. W.l.o.g., assume the effective axis is the x-axis and $l_1 < l_2 < r_1 < r_2$, where $l_1 < l_2$ denotes that l_1 is on the left of l_2 . Then, the following edges are related edges: (i) edges located between l_2 and r_1 ; or (2) edges intersected with l_2 or r_1 .

See Fig. 3(a) for an example. Edges \overline{ab} and \overline{bc} are related edges as they intersect with l_2 . Similarly, edges \overline{de} and \overline{ef} are also related edges. We remark that in Definition 5.3 there are actually other cases, e.g., ' $l_1 < l_2 < r_2 < r_1$ ' or the effective axis is the y-axis; these cases are similar to the listed case, omitted for saving space.

Definition 5.4 (Processed Related Edges). Given a number of related edges, we decompose them if there are non-x-monotone arcs. We call all the edges (after decomposing) the processed related edges.

By Lemma 4.1 and Definition 5.4, we have the following corollary (which will be used later).

Corollary 5.1. Given l related edges, if there is no non-x-monotone arc amongst them, the number of processed related edges is l . Otherwise, the number of processed related edges is larger than l and no more than $3l$.

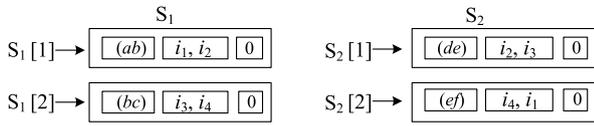


Fig. 4. Example of sequence lists.

Up to now, we have discussed the properties of related edges, and briefly explained how to choose related edges from two circular-arc polygons (remark: more explanations will be given in Algorithm 1 and in the proof of Lemma 5.1). We next show how to use two sequence lists to manage the processed related edges and other important components.

5.2. Managing important components

The main purpose of the two sequence lists (i.e., arrays) is to let the processed related edges, intersections and decomposed arcs be well organized, which can facilitate the subsequent operations. Specifically, each item in the two sequence lists is a compound structure consisting of three domains: (i) the processed related edge; (ii) the intersections (if they exist) on this edge; and (iii) a tri-value switch. For ease of discussion, we denote by S_1 and S_2 the two sequence lists, by $S_i[j]$ the j th item in S_i ($i \in \{1, 2\}$), and by $S_i[j].a$, $S_i[j].b$ and $S_i[j].c$ the three domains of $S_i[j]$, respectively.

The processed related edges in each sequence list are stored in counter-clockwise direction with regard to the original circular-arc polygon. For example, regarding circular-arc polygons in Fig. 3, we construct two sequence lists as shown in Fig. 4. Note that when there are multiple intersections on an edge, we should keep these intersections in order. See $S_1[1].b$ of Fig. 4 for an example; the point i_1 is ahead of point i_2 . Regarding the third domain $S_i[j].c$, it is assigned to either 0, 1, or 2. The assignment rules are as follows. When the edge is not a decomposed arc, we assign ‘0’ to $S_i[j].c$. In this example, for any $1 \leq j \leq |S_i|$ (where $|\cdot|$ denotes the cardinality of S_i), $S_i[j].c$ is set to 0, as there is no decomposed arc. Otherwise, we assign ‘1’ or ‘2’ to $S_i[j].c$. The readers may be curious why we use two different values. The purpose is to differentiate the decomposed arcs which are from different non- x -monotone arcs. This can help us efficiently merge them in the future. (The specific steps on how to merge them will be discussed in Section 6.) Given a series of decomposed arcs, we assign ‘1’ to each decomposed arc that is from the odd (1st, 3rd, ...) non- x -monotone arc, and assign ‘2’ to each decomposed arc that is from the even (2nd, 4th, ...) non- x -monotone arc.

See Fig. 5(a) for an example, there are five related edges in \mathcal{P}_1 . Furthermore, Fig. 5(b) illustrates eight processed related edges (after we decompose them based on Lemma 4.1), implying that $|S_1| = 8$. Based on the assignment rules, the values of the third domains should be ‘0, 1, 1, 2, 2, 1, 1, 0’, respectively.

So far, we have shown how to use two sequence lists to manage the processed related edges and intersections. Note that, in order to obtain the intersections, a standard technique is the plane sweep method [26,30]. In this paper, we do not directly use this algorithm. Instead, we modify it by adding two labels to avoid ‘false’ intersections being reported, and to speed up the process of inserting the reported intersections into their corresponding edges. (Remark: here the false intersections refer to the vertices of polygons.) We next give a brief summary of the plane sweep algorithm, and then show how the two labels work.

Plane sweep method. Let Ω be a priority queue, \mathcal{R} be a balanced tree,³ and l_v be a vertical sweep line. The basic idea of the plane

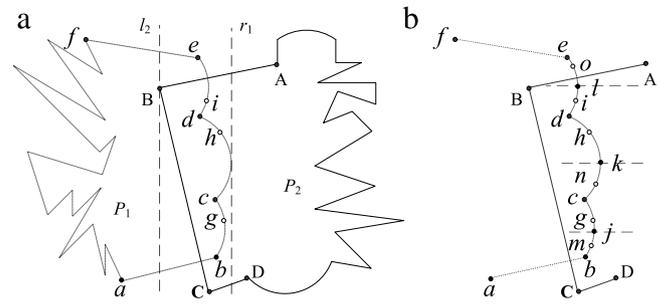


Fig. 5. Example of consecutive non- x -monotone circular arcs. (a) Edges \overline{ab} , \widehat{bgc} , \widehat{chd} , \widehat{die} , \overline{ef} are related edges of \mathcal{P}_1 . (b) Edges \overline{ab} , \widehat{bmj} , \widehat{jgc} , \widehat{cnk} , \widehat{khd} , \widehat{dil} , \widehat{loe} , \overline{ef} are processed related edges of \mathcal{P}_1 ; three dashed lines are the auxiliary lines.

sweep method is as follows. First, it sorts the endpoints of all segments according to their x -coordinates, and puts them into Ω . Next, it sweeps the plane (from left to right) using l_v . At each endpoint during this sweep, if an endpoint is the left endpoint of a segment, it inserts this segment into \mathcal{R} . In contrast, if it is the right endpoint of a segment, it deletes this segment from \mathcal{R} . Note that all the segments that intersect with l_v are stored (in order from bottom to top) in \mathcal{R} . In particular, when l_v moves from one endpoint to another endpoint, it always checks whether or not newly adjacent segments intersect with each other. If so, it computes the intersection. In this way, all intersections can be obtained finally.⁴

5.3. Avoiding false intersections and speeding up lookups

We can easily see that the plane sweep method directly inserts a segment into the balanced tree \mathcal{R} , if the point p ($\in \Omega$) is the left endpoint of the segment. Instead, we assign two labels to the segment before it is inserted into \mathcal{R} . Note that the segment discussed here refers to the processed related edge. For clearness, we denote by lb_1 and lb_2 the two labels, respectively.

lb_1 is the Boolean type, identifying that a segment is from which one of the two circular-arc polygons. Specifically, if the segment is from \mathcal{P}_1 , we assign *true* to lb_1 ; otherwise, we assign *false* to lb_1 . Recall that the plane sweep method always checks whether or not two segments intersect with each other when they are adjacent. Our proposed method does not need to check them, regardless of whether or not they intersect, if the first labels of two adjacent edges have the same value. This can avoid unnecessary tests and ‘false’ intersections.

lb_2 is an integer type denoting a serial number, which corresponds to the ‘id’ of an item stored in the sequence list (note: the ‘id’ information of each item is implied, as we store the items using the sequence list, i.e., array). When we detect an intersection, this label can help us quickly find the item in the sequence list, and then insert the intersection into this item. See Fig. 3(b) for an example. lb_1 and lb_2 of edge \overline{ab} are assigned to *true* and 1, respectively. When we detect the intersection i_1 , we then quickly know that we should insert the intersection into $S_1[1]$ (i.e., the first item of S_1). Otherwise, we have to scan the sequence list in order to insert the intersection into an appropriate item. This way is inefficient, especially when $|S_1|$ (or $|S_2|$) is large.

⁴ Note that, in some cases the segments may be vertical line segments, or they may be tangent, or many segments intersect possibly at one point. For these degenerated cases, please refer to the papers (e.g., [26,30,13,16,20]) for more details. Unless stated otherwise, degenerated cases are processed using existing techniques and/or a straightforward adaptation from existing techniques. We no longer expatiate them separately in order to save space (as they are tedious, and are not the focus of the paper).

³ It is not mandatory to use a priority queue and a balanced tree, whereas they are usually being recommended, for the sake of efficiency [26]. Moreover, both of them are abstract concepts; the priority queue, for example, can be implemented with a heap or other methods.

Algorithm 1 *ConstructSequenceLists*

Input: Circular-arc polygons \mathcal{P}_1 and \mathcal{P}_2
Output: Sequence lists S_1 and S_2 , related edge sets R_1 and R_2

- 1: Find the MBRs, effective axis and extended boundary lines
- 2: **for** each $i \in \{1, 2\}$ **do**
- 3: $R_i \leftarrow$ related edges from \mathcal{P}_i
- 4: Create two empty sequence lists S_1 and S_2
- 5: **for** each $i \in \{1, 2\}$ **do**
- 6: InitializeSequenceList (R_i, S_i) // cf., Algorithm 2
- 7: Sort the endpoints of the segments (from S_1, S_2), and put them into the priority queue \mathcal{Q}
- 8: Initialize the empty balanced tree \mathcal{R}
- 9: **for** each point $p \in \mathcal{Q}$ **do**
- 10: Let s be the segment containing the point p , and t be the segment immediately above or below s
- 11: **if** (p is the left endpoint of segment s)
- 12: Assign two 'labels' to s , and insert s into \mathcal{R}
- 13: **if** ($s.lb_1 \neq t.lb_1$)
- 14: **if** (s intersects with t)
- 15: Insert the intersection into \mathcal{Q} , and also insert it into S_1 and S_2
- 16: **else if** (p is the right endpoint of segment s)
- 17: **if** ($s.lb_1 \neq t.lb_1$)
- 18: **if** (s intersects with t and this intersection $\notin \mathcal{Q}$)
- 19: Insert this intersection into \mathcal{Q} , and also insert it into S_1 and S_2 , respectively; delete s from \mathcal{R}
- 20: **else** // p is an intersection of two segments, say s and t
- 21: Swap the position of s and t // assume s is above t
- 22: Let t_1 be the segment above s , and t_2 be the segment below t
- 23: **if** ($s.lb_1 \neq t_1.lb_1$ or $t.lb_1 \neq t_2.lb_1$)
- 24: **if** (s intersects with t_1 , or t intersects with t_2)
- 25: Insert this intersection into \mathcal{Q} , and also insert it into S_1 and S_2 , respectively
- 26: **return** S_1 and S_2, R_1 and R_2

5.4. The algorithm

Let R_1 and R_2 be the related edges from \mathcal{P}_1 and \mathcal{P}_2 , respectively. Given a segment s , we use $s.lb_1$ and $s.lb_2$ to denote the two labels of segment s . Algorithm 1 illustrates the pseudo-codes of constructing the two sequence lists.

We first choose the *related edges* based on the extended boundary lines (Lines 1–3). Next, we construct two empty sequence lists and initialize them (Lines 4–6). After this, we compute the intersections (Lines 7–25). In particular, when we compute the intersections, two *labels* are assigned to the segment before it is inserted into the balanced tree (Line 12), and we use the two sequence lists to store the intersections (Lines 15, 19 and 25). Note that, the pseudo-codes of initializing the two sequence lists are listed in Algorithm 2. This algorithm decomposes non- x -monotone arcs, puts the *processed related edges* into the sequence lists in an orderly manner, and assigns appropriate values to the tri-value switches.

Lemma 5.1. *Given two circular-arc polygons with m and n edges, respectively, and assume there are l related edges between the two polygons, we have that constructing the two sequence lists can be completed in $O(m + n + l + (l + k) \log l)$ time, where k is the number of intersections.*

Proof. To obtain the *related edges*, we first need to find the MBRs of two polygons, which takes linear time. We next determine the *effective axis* by comparing the horizontal and vertical spans of the intersection set of two MBRs, which can be finished in constant time. Furthermore, the *extended boundary lines* can be obtained in constant time once we obtain the effective axis. Based on two extended boundary lines, we finally obtain the *related edges* by comparing the geometrical relationship between each edge and extended boundary lines, which also takes linear time. Thus, Lines 1–3 take $O(m + n)$ time.

Creating two empty sequence lists takes constant time. In addition, in order to initialize the two sequence lists, we need

to decompose each non- x -monotone arc. Decomposing a single arc can be finished in constant time. In the worst case, all the *related edges* are non- x -monotone arcs. Even so, there are no more than $3l$ items in the two sequence lists, according to Lemma 4.1. Hence initializing two sequence lists takes $O(l)$ time. Sorting all the endpoints of segments in the priority queue \mathcal{Q} takes $O(l \log l)$ time, and initializing the balanced tree \mathcal{R} takes constant time. Thus, Lines 4–8 take $O(l + l \log l)$ time.

As there are no more than $3l$ segments in S_1 and S_2 , the number of endpoints thus is no more than $6l$. So we know that the number of executions of the **for** loop (Line 9) is no more than $6l + k$. Within the **for** loop, each operation (e.g., insert, delete, swap, find the above/below segment) on \mathcal{R} can be finished in $O(\log l)$ time, as the number of segments in \mathcal{R} never exceeds $3l$. Additionally, each of other operations (e.g., assign labels to the segment, determine if two segments intersect with each other) can be finished in constant time. Thus, Lines 9–25 take $O((6l + k) \log l)$ time, i.e., $O((l + k) \log l)$ time. Putting it together, this completes the proof. \square

Algorithm 2 *InitializeSequenceList*

Input: R_i, S_i
Output: S_i

- 1: $temp \leftarrow 1$ // the $temp$ is used to set the tri-value switch
- 2: **for** each related edge $r \in R_i$ **do**
- 3: **if** (r is a non- x -monotone circular arc)
- 4: Decompose it and put the decomposed arcs into S_i , and set the value of each tri-value switch to ' $temp$ '
- 5: **if** ($temp=1$)
- 6: $temp \leftarrow 2$;
- 7: **else** // $temp=2$
- 8: $temp \leftarrow 1$;
- 9: **else** // r is not a non- x -monotone circular arc
- 10: Put it into S_i , set the value of tri-value switch to '0'

We have shown how to construct two sequence lists. Clearly, we cannot get the resultant polygon based on *only* the information stored in the two sequence lists. Next, we are ready to merge the information in them and part of information in the original linked lists, and store the 'merged information' using two *new linked lists*. For ease of exposition, we call this step 'building two new linked lists'. Note that the two new linked lists will be significantly used in Section 7, as we need to traverse them to obtain the resultant polygon.

6. Building two new linked lists

To construct two new linked lists, on the whole we first initialize two (empty) new linked lists, and then copy the information from the *original linked lists* to the *new linked lists*, while we replace those *related edges* using the information stored in the two sequence lists. Note that there are two important yet easy-to-ignore issues that need to be handled when we construct new linked lists. We next check these issues, and then present the algorithm of constructing new linked lists.

6.1. Eliminating the ambiguity

Recalling Section 4.1, we always add an *appendix point* between two vertices if an edge is a circular arc. When we replace *related edges* with the information stored in sequence lists, we also have to ensure this property. It is easy to see that, when the intersections appear on a circular arc, this arc will be decomposed by these intersections. We thus have to add the new appendix point for each sub-segment, in order to eliminate the ambiguity.

Lemma 6.1. *Suppose there are k intersections on a circular arc; then, we need to insert at least k and at most $k + 1$ new appendix points.*

Proof. Since k intersections can subdivide a complete circular arc into $k + 1$ small circular arcs, and for each small circular arc one *appendix point* is needed, and is enough to eliminate the ambiguity. Clearly, $k + 1$ *appendix points* are needed for $k + 1$ small circular arcs. Note that, there is an *appendix point* beforehand. Therefore, when there is no intersection (among all these intersections) that *coincides with this existing appendix point*, only k new *appendix points* are needed. Otherwise, we need $k + 1$ new *appendix points*. \square

Besides the above issue, another issue is on how to handle the decomposed arcs. We decomposed non-x-monotone arcs into x-monotone arcs ever. We thus need to merge them. The natural solution is to compare each pair of adjacent edges of the resultant polygon, checking if they can be merged into a single arc. This way, however, is inefficient because (i) most of edges of the resultant polygon may not need to be merged; and (ii) given two adjacent arcs, let C_1 and C_2 respectively denote their corresponding circles. Checking if the two arcs can be merged into a single arc requires to compute the centres of C_1 and C_2 . This will use trigonometric functions, which could have been avoided. We next show how to efficiently merge them, with the help of the *tri-value switch* (recall Section 5.2).

6.2. Efficiently merging decomposed arcs

We merge the decomposed arcs when constructing new linked lists, rather than merge them after obtaining the resultant polygon. In particular here we utilize the information stored in the tri-value switch to improve the efficiency. Specifically, given an item $S_i[j]$, if $S_i[j].c = 1$ (or 2), we continue to fetch its next item $S_i[j + 1]$ from the sequence list if $S_i[j].c = S_i[j + 1].c$. In this way, a group of consecutive items are fetched from the sequence list. W.l.o.g, assume that we have fetched λ consecutive items, $S_i[j], \dots, S_i[j + \lambda - 1]$. Then, we do as follows.

- If $S_i[j].b = S_i[j + 1].b = \dots = S_i[j + \lambda - 1].b = \emptyset$, we discard the fetched items instead of merging them. This is because there is no intersection on these decomposed arcs and so the merged result should be the same as the edge in the original linked list.
- Otherwise, we insert new appendix points, merge decomposed arcs, and replace the edge in the original linked list.

Let us revisit Fig. 5(b). Recall that there are eight items in S_1 , and the values in their tri-value switches are '0, 1, 1, 2, 2, 1, 1, 0', respectively. Although $S_1[2].c = S_1[3].c = 1$, we discard the two items instead of merging them, as $S_1[2].b = S_1[3].b = \emptyset$. Similarly, we also discard the items $S_1[4]$ and $S_1[5]$. Note that, for the 6th and 7th items, $S_1[6].c = S_1[7].c = 1$ and $S_1[7].b \neq \emptyset$; thus, we insert a new appendix point, merge the two decomposed arcs, and use the merged result to replace the edge in the original linked list.

Note that the consecutive items mentioned earlier are actually the decomposed arcs generated from a single non-x-monotone circular arc. According to Lemma 4.1, we can easily obtain the following corollary.

Corollary 6.1. *Let λ be the number of consecutive items, we have that $\lambda \leq 3$ and $\lambda \geq 2$.*

6.3. The algorithm

Let \mathcal{P}_1^* and \mathcal{P}_2^* be the two new linked lists, respectively. Algorithm 3 depicts the pseudo-codes of constructing two new linked lists. For each edge e in the *original linked list*, we check whether it is a *related edge*. If so, we further check whether $S_i[j]$ is a decomposed arc. Lines 7–13 are used to process the case when it is not a decomposed arc. In contrast, Lines 14–22 are

used to handle the opposite case. In this case, we first fetch all the consecutive decomposed arcs (Lines 15–18), and then check if there are intersections on these decomposed arcs. If not, we put the edge e into \mathcal{P}_i^* (Lines 19–20). Otherwise, we insert new appendix points, merge decomposed arcs, and put the merged result (instead of e) into \mathcal{P}_i^* (Lines 21–22).

Algorithm 3 BuildNewLinkedLists

Input: \mathcal{P}_1 and \mathcal{P}_2 , S_1 and S_2 , R_1 and R_2

Output: \mathcal{P}_1^* and \mathcal{P}_2^*

```

1: Set  $\mathcal{P}_1^* = \mathcal{P}_2^* = \emptyset$ , and  $j \leftarrow 1$ 
2: for each  $i \in \{1, 2\}$  do
3:   for each edge  $e \in \mathcal{P}_i$  do
4:     if ( $e \notin R_i$ )
5:       Add  $e$  to  $\mathcal{P}_i^*$ 
6:     else //  $e$  is a related edge
7:       if ( $S_i[j].c = 0$ ) // not a decomposed arc
8:         if ( $S_i[j].b = \emptyset$ ) // no intersection
9:            $j \leftarrow j + 1$ , and add  $e$  to  $\mathcal{P}_i^*$ 
10:        else //  $S_i[j].b \neq \emptyset$ 
11:          if ( $S_i[j]$  is a circular arc)
12:            Insert new appendix points
13:            Put the information from  $S_i[j]$  into  $\mathcal{P}_i^*$ , and set  $j \leftarrow j + 1$ 
14:          else //  $S_i[j].c = 1$  (or 2)
15:            Set  $tri = S_i[j].c$ , and  $\lambda \leftarrow 0$ 
16:            do // copy the consecutive decomposed arcs
17:               $\lambda \leftarrow \lambda + 1$ ,  $temp[\lambda] \leftarrow S_i[j]$ ,  $j \leftarrow j + 1$ 
18:            while  $S_i[j].c = tri$ 
19:            if ( $temp[1].b = \dots = temp[\lambda].b = \emptyset$ )
20:              Put  $e$  into  $\mathcal{P}_i^*$ 
21:            else
22:              Insert new appendix points, merge decomposed arcs, and
                put the merged result into  $\mathcal{P}_i^*$ 
23: return  $\mathcal{P}_1^*$  and  $\mathcal{P}_2^*$ 

```

Lemma 6.2. *Suppose we have obtained the two sequence lists S_1 and S_2 . Then, constructing the two new linked lists \mathcal{P}_1^* and \mathcal{P}_2^* takes $O(m + n + l + k)$ time.*

Proof. Inserting a single *appendix point* takes constant time. In the worst case, all the intersections are located on arcs rather than on line segments. Even so, there are no more than $2k$ new *appendix points* according to Lemma 6.1. Hence inserting *appendix points* takes $O(k)$ time. Merging λ consecutive decomposed arcs takes constant time, as $\lambda \leq 3$ (cf., Corollary 6.1). In the worst case, all the related edges are non-x-monotone arcs; hence merging all consecutive decomposed arcs takes $O(l)$ time.

Since the number of edges in \mathcal{P}_1 and \mathcal{P}_2 is $m + n$, the number of executions of the second **for** loop (Line 3) is $m + n$. Specifically, the number of executions of Line 4 is $m + n - l$, and that of Line 6 is l . Even if all related edges are non-x-monotone arcs, the number of executions of Line 17 is no more than $3l$. Furthermore, within the **for** loop, each operation takes constant time (note: here we no longer consider the time for inserting new appendix points and merging decomposed arcs, as we have analysed them in the previous paragraph). Therefore, Lines 4–5 and Lines 7–22 take $O(m + n - l)$ and $O(3l)$ time, respectively. Putting it all together, this completes the proof. \square

7. Traversing

In the previous section, we obtained \mathcal{P}_1^* and \mathcal{P}_2^* . This section shows in detail how to obtain the resultant polygon by traversing them. In order to correctly traverse \mathcal{P}_1^* and \mathcal{P}_2^* , we need to assign *entry-exit* properties to intersections.

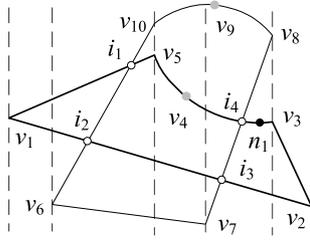


Fig. 6. Example of assigning entry–exit properties. $\mathcal{P}_1 = \{v_1, v_2, v_3, \tilde{v}_4, v_5\}$ and $\mathcal{P}_2 = \{v_6, v_7, v_8, \tilde{v}_9, v_{10}\}$.

7.1. Entry–exit property

The entry–exit property is an important symbol that was used in many papers focusing on Boolean operations of traditional polygons (see e.g., [13,16]). The following shows that this technique can be used to the case of our concern as well. Specifically, we assign the intersections with the *entry* or *exit* property *alternately*. Note that the *entry–exit* property for the first intersection in \mathcal{P}_1^* (\mathcal{P}_2^*) is determined as follows. W.l.o.g, assume the first intersection is i (i') in \mathcal{P}_1^* (\mathcal{P}_2^*), and the previous node of i (i') is $i.prev$ ($i'.prev$). We check if $i.pre$ ($i'.pre$) is outside the input polygon \mathcal{P}_2 (\mathcal{P}_1). If so, we assign the *entry* (*exit*) property to i (i'). See Fig. 6 for an example. \mathcal{P}_1 and \mathcal{P}_2 are two input polygons. We can easily obtain the two new linked lists \mathcal{P}_1^* and \mathcal{P}_2^* using algorithms discussed before. Here $\mathcal{P}_1^* = \{v_1, i_2, i_3, v_2, v_3, \tilde{n}_1, i_4, \tilde{v}_4, v_5, i_1\}$, and $\mathcal{P}_2^* = \{v_6, v_7, i_3, i_4, v_8, \tilde{v}_9, v_{10}, i_1, i_2\}$. Regarding \mathcal{P}_1^* , since the previous node of i_2 is v_1 which is outside \mathcal{P}_2 , we assign ‘*entry, exit, entry, exit*’ to ‘ i_2, i_3, i_4, i_1 ’, respectively. The *entry–exit* properties of ‘ i_2, i_3, i_4, i_1 ’ in \mathcal{P}_2^* are obtained similarly. See Fig. 7(a).

Once the *entry–exit* properties are assigned to intersections, we then obtain the resultant polygon based on the traversing rules below.

7.2. Traversing rules

Let i_s be an intersection (point) of \mathcal{P}_1^* such that i_s has the *entry* property. Let v_s be a vertex of \mathcal{P}_1^* such that v_s is not located in \mathcal{P}_2^* . There are three typical Boolean operations: intersection, union and difference. Note that in the rest of the discussion, the default traversing direction is counter-clockwise, unless stated otherwise.

Intersection. We start to traverse \mathcal{P}_1^* using i_s as the starting point. Once we meet an intersection with the *exit* property, we shift to \mathcal{P}_2^* , and traverse it. Similarly, if we meet an intersection with the *entry* property in \mathcal{P}_2^* , we shift back to \mathcal{P}_1^* . In this way, a circuit will be produced. After this, we check if there is another intersection of \mathcal{P}_1^* such that (i) it has the *entry* property; and (ii) it is not a vertex of the

produced circuit. If there is no such an intersection, we terminate the traversal, and this circuit is the intersection between \mathcal{P}_1 and \mathcal{P}_2 . Otherwise, we use this intersection as a new starting point, and traverse the two new linked lists (using the same method discussed just now), until no such an intersection exists. In the end, we get multiple circuits, which are the intersection between \mathcal{P}_1 and \mathcal{P}_2 . See Fig. 7(b) for an example. We first choose i_2 in \mathcal{P}_1^* as the starting point, and then traverse the two linked lists, obtaining a circuit (see the dashed lines). Moreover, there is no other intersection that satisfies the two conditions outlined before. Therefore, the intersection is $\{i_2, i_3, i_4, \tilde{v}_4, v_5, i_1\}$.

Union. For this case, we traverse \mathcal{P}_1^* using v_s (instead of i_s) as the starting point. Once we meet an intersection with the *entry* property, we shift to \mathcal{P}_2^* , and traverse it. Similarly, if we meet an intersection with the *exit* property in \mathcal{P}_2^* , we shift back to \mathcal{P}_1^* . In this way, a circuit is produced; it is just the union between \mathcal{P}_1 and \mathcal{P}_2 . See Fig. 7(c) for an example. We first choose v_1 as the starting point, and then traverse the two linked lists, until we are back to the starting point v_1 . Therefore, $\{v_1, i_2, v_6, v_7, i_3, v_2, v_3, \tilde{n}_1, i_4, v_8, \tilde{v}_9, v_{10}, i_1\}$ is the union between \mathcal{P}_1 and \mathcal{P}_2 .

Difference. The first few steps are the same as those in the *union* operation, *but* we traverse \mathcal{P}_2^* in a clockwise direction. Similarly, if we meet an intersection with the *exit* property in \mathcal{P}_2^* , we shift back to \mathcal{P}_1^* . In this way, a circuit will be produced. Furthermore, we check if there is another vertex of \mathcal{P}_1^* such that (i) it is not a vertex of any produced circuit; and (ii) it is not located in \mathcal{P}_2^* . If no such a vertex exists, we terminate the traversal, and this circuit is the difference between \mathcal{P}_1 and \mathcal{P}_2 . Otherwise, we use the vertex as a new starting point, and traverse the two new linked lists (using the same method discussed previously), until no such a vertex exists. In the end, we obtain multiple circuits that are the difference between \mathcal{P}_1 and \mathcal{P}_2 . See Fig. 7(d) for an example. The difference between \mathcal{P}_1 and \mathcal{P}_2 consists of two parts, i.e., $\{v_1, i_2, i_1\}$ and $\{v_2, v_3, \tilde{n}_1, i_4, i_3\}$.

7.3. The algorithm

In some cases the result consists of multiple circuits; we denote by l_j the linked list used to store the j th circuit. Let n_d be a node of \mathcal{P}_i^* where $i \in \{1, 2\}$. We denote by c_n the current node when we traverse \mathcal{P}_i^* , and denote by $c_n.next$ the next node of c_n . Furthermore, we denote by i'_s the intersection of \mathcal{P}_1^* such that i'_s has the *entry* property and it is not a vertex of any produced circuit, and we use $\exists(i'_s) = true$ to denote that there exists such a point. Algorithm 4 illustrates the pseudo-codes of obtaining the intersection result (remark: the pseudo-codes of the other two operations can be constructed similarly by traversing rules).

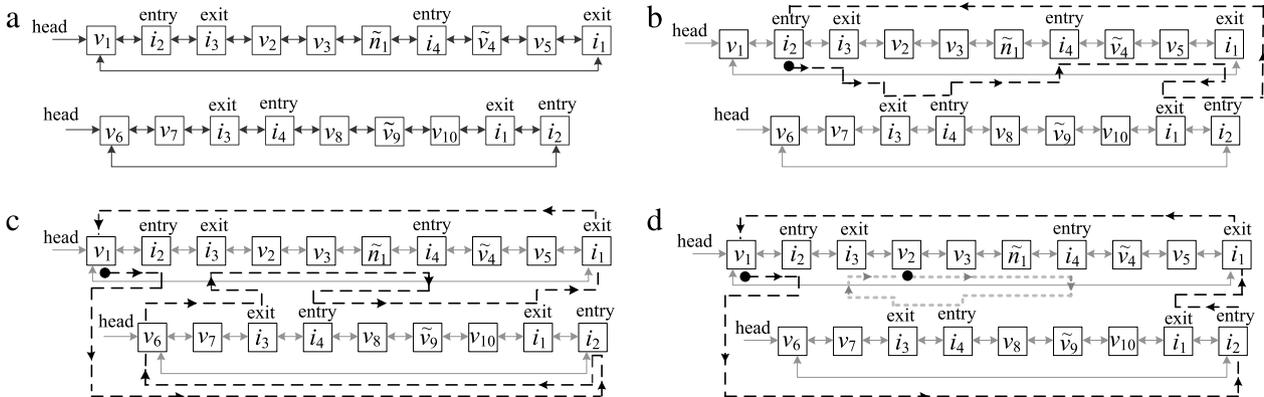


Fig. 7. Illustrations of entry–exit properties and traversing rules. (a) Entry–exit properties of \mathcal{P}_1^* and \mathcal{P}_2^* . (b) Intersection operation. (c) Union operation. (d) Difference operation.

Lemma 7.1. Given the two new linked lists \mathcal{P}_1^* and \mathcal{P}_2^* , to obtain the resultant polygon takes $O(k + m + n + l)$ time.

Proof. Assigning the entry–exit property to each intersection takes constant time, and there are k intersections on each new linked list. Thus, assigning entry–exit properties to intersections takes $O(k)$ time.

\mathcal{P}_1^* and \mathcal{P}_2^* are used to generate the resultant polygon; they store the vertices, appendix points and intersections. The number of vertices is $2(m + n)$. In the worst case, all edges of two input polygons are circular arc segments, implying that the number of appendix points in the input polygons is $m + n$. In this case, all the k intersections are located on arcs, we need to insert new appendix points, and the number of new appendix points is no more than $k + 1$, according to Lemma 6.1. So the number of all appendix points in \mathcal{P}_1^* and \mathcal{P}_2^* is no more than $m + n + k + 1$. Therefore, the total number of nodes in \mathcal{P}_1^* and \mathcal{P}_2^* is no more than $3(m + n) + 2k + 1$. Furthermore, each operation on a node (e.g., determine the type of a node, insert a node into the resultant polygon) takes constant time. Thus, the traversal takes $O(3(m + n) + 2k + 1)$ time. To summarize, we have that obtaining the resultant polygon takes $O(m + n + k + l)$ time when \mathcal{P}_1^* and \mathcal{P}_2^* are given beforehand.⁵ □

Algorithm 4 TraverseLinkedLists

Input: \mathcal{P}_1^* and \mathcal{P}_2^*
Output: \mathcal{P}_3

- 1: Set $j = 1$
- 2: **for** each $i \in \{1, 2\}$ **do**
- 3: Assign entry–exit property to \mathcal{P}_i^*
- 4: **do**
- 5: **if** ($j=1$)
- 6: Choose a starting point i_s from \mathcal{P}_1^*
- 7: **else**
- 8: Let $i_s \leftarrow i'_s$ // i.e., let i'_s be a new starting point
- 9: Set $c_n \leftarrow i_s$, and $l_j = \emptyset$
- 10: **do**
- 11: Put c_n into l_j
- 12: **if** ($c_n.next$ is not an intersection point)
- 13: Let $c_n \leftarrow c_n.next$, and put c_n into l_j
- 14: **else**
- 15: Shift to \mathcal{P}_2^* , choose the node n_d such that $n_d = c_n.next$, let $c_n \leftarrow n_d$, and put c_n into l_j
- 16: **if** ($c_n.next$ is not an intersection point)
- 17: $c_n \leftarrow c_n.next$, and put c_n into l_j
- 18: **else**
- 19: Shift to \mathcal{P}_1^* , choose the node n_d such that $n_d = c_n.next$, and let $c_n \leftarrow n_d$
- 20: **while** ($c_n \neq i_s$)
- 21: Let $\mathcal{P}_3 \leftarrow \mathcal{P}_3 \cup l_j$, and set $j \leftarrow j + 1$
- 22: **while** ($\exists(i'_s) = true$)
- 23: **return** \mathcal{P}_3

Up to now, we have addressed all the main steps of our algorithm—RE2L. We next analyse its complexity.

8. Time/space complexity

We analyse the complexity of our algorithm using the intersection operation as a sample (note: the complexity of the other two operations is the same as that of this operation, and can be derived similarly).

⁵ It is simple to determine i'_s (cf., Lines 8 and 22). We just need to collect all the intersections with entry properties in a data structure when we assign entry–exit properties to intersections, and then remove the intersections from this data structure once they have been visited in the process of traversing. After a circuit is formed, we check if this data structure is empty. If not, any intersection stored in this data structure can be taken as i'_s .

Theorem 8.1. Given two circular-arc polygons with m and n edges, respectively, and assume there are l related edges between them. Then, to perform Boolean operation on them takes $O(m + n + (l + k) \log l)$ time, using $O(m + n + k)$ space, where k is the number of intersections.

Proof. Our algorithm consists of three main steps that take time $O(m + n + l + (l + k) \log l)$, $O(m + n + k + l)$, and $O(m + n + k + l)$, respectively (see Lemmas 5.1, 6.2 and 7.1). Putting these results together, the time complexity is $O(m + n + (l + k) \log l)$.

The space used in our algorithm mainly consists of two groups of related edges R_1 and R_2 , two sequence lists S_1 and S_2 , the priority queue \mathcal{Q} , two new linked lists \mathcal{P}_1^* and \mathcal{P}_2^* , and the balanced tree \mathcal{R} . (Remark: the input polygons $\mathcal{P}_1, \mathcal{P}_2$, and the output polygon \mathcal{P}_3 are the input and output data; by the convention,⁶ we do not need to consider them when we analyse the space complexity.)

Specifically, R_1 and R_2 have size $O(l)$, as they are used to store the related edges. S_1 and S_2 are used to store the processed related edges. In the worst case, the number of processed related edges is no more than $3l$, according to Corollary 5.1. So S_1 and S_2 have the size of $O(3l)$. Recall Algorithm 1, the priority queue \mathcal{Q} is used to store the endpoints of processed related edges and the intersections, and so \mathcal{Q} has the size of $O(6l + k)$. Regarding the balanced tree \mathcal{R} , it has the size of $O(l)$ at most, as it is used to store the segments currently intersecting the sweep line. Furthermore, \mathcal{P}_1^* and \mathcal{P}_2^* have the size of $O(3(m + n) + 2k + 1)$, see the proof of Lemma 7.1. Putting it all together, we have that the space complexity of our algorithm is $O(m + n + k + l)$. In addition, the upper bound of the parameter l is $m + n$. This completes the proof. □

We can see that our algorithm consumes roughly linear space when k is small. The running time also approximates to linear complexity when l and k are small. It is noteworthy that a straightforward adaptation from the plane sweep algorithm (see e.g., the ‘Standard’ method described in Section 10.1) requires $O((m + n + k) \log(m + n))$ time. In other words, even if k is small (e.g., $k \ll m + n$), it is also linearithmic time. We remark that, for Boolean operations on circular-arc polygons, the $O((m + n + k) \log(m + n))$ result is actually the state-of-the-art competitor in terms of computational complexity.

Although our algorithm has some advantages to some extent, we have to point out that in the worst case (note: this case is possible although it is not the usual case), i.e., $l = m + n$, the running time deteriorates to $O((m + n + k) \log(m + n))$, which is equal to that of the standard method. In this case, the advantages of the proposed algorithm disappear, viewed from the theoretical perspective. To this step, an interesting issue is: when $l = m + n$, whether its practical efficiency is the same as that of the standard method? We will experimentally evaluate our algorithm as well as the competitors in Section 10, after we introduce some immediate extensions.

9. Extensions

While this paper focuses on Boolean operations of circular-arc polygons, our techniques can be easily extended to compute Boolean operations of traditional polygons. Assuming that there are two traditional polygons, for example, we can also use two lists to represent them. In this case, the Tag domain is not needed as the traditional polygons do not need the appendix points. We can also choose related edges based on the extended boundary lines, and store them using two sequence lists. Note that the third domain of the sequence list is unnecessary, as here all related edges

⁶ As an example, the bubble sort algorithm takes $O(1)$ space for sorting arbitrary n natural numbers.

Table 1

The coordinates of vertices are listed counter-clockwise, and the left-bottom vertex is listed first. The values tagged with ‘_’ denote the coordinates of *appendix points*.

Polygon	Coordinates
\mathcal{P}_1	(10, 10), (40, 10), (40, 30), (32.5, 40), (20, 40), (15, 30), (25, 22.5), (15, 15)
\mathcal{P}_2	(20, 20), (32.5, 25), (45, 20), (55, 30), (47, 38.625), (50, 50), (30, 45)

are straight line segments. When computing the intersections, we can also use *two labels* to speed up the process of inserting the intersections into their corresponding edges, and to avoid *false intersections*. Next, we also construct *two new linked lists*, by using the information stored in two sequence lists to replace the related edges in the original linked lists. In particular, here we do not need to insert new appendix points and merge the decomposed arcs, as the traditional polygons have no such information. We finally obtain the resultant polygon by *traversing* the two new linked lists, as the same as that discussed in Section 7.

Furthermore, the discussions presented in previous sections assumed that the circular-arc polygons to be processed have no holes. If we want to handle the opposite case, this can be easily achieved by a straightforward adaptation of our proposed method. Assume we want to compute the intersection of two circular-arc polygons with holes, we can use multiple lists to represent the circular-arc polygon with holes. One is used to represent the outer boundaries of the circular-arc polygon, while others are respectively used to represent the boundaries of each hole. We can first compute the intersection of the outer boundaries of two polygons, and then use this intersection result to subtract each hole of the two polygons. All of these steps are quite straightforward when our proposed method is given beforehand.

Finally it is also immediate to compute Boolean operations on circular-arc polygons with *self-intersection*. For example, assume that we want to compute the intersection. We just need to make a minor modification to the traversing rule presented in Section 7.2. We also start to traverse \mathcal{P}_1^* using an intersection with the *entry* property as the starting point, and shift to \mathcal{P}_2^* when the number of intersections we meet in \mathcal{P}_1^* is even. The main difference is that the *traversing direction* here is not always constant. To determine whether or not the traversing direction needs to change when we shift from $\mathcal{P}_1^* (\mathcal{P}_2^*)$ to $\mathcal{P}_2^* (\mathcal{P}_1^*)$, a key step is to check if the *entry-exit* property of this intersection in \mathcal{P}_1^* is different from the one in \mathcal{P}_2^* . If so, we need to change the traversing direction. Otherwise, we needn't.

10. Performance evaluation

This section evaluates our algorithm experimentally. Specifically, Section 10.1 describes the baselines. The experimental settings and datasets are introduced in Section 10.2, and the results are reported in Section 10.3. We also investigate our proposed algorithm based on a specific application (Section 10.4).

10.1. Methodologies

We compare our method (i.e., RE2L) with several baselines, which are either the algorithms used to handle more general cases of polygons, or the simpler versions of our proposed method. We introduce them briefly as follows.

CGAL. We directly use the implementation of CGAL. The essence of this method is to directly invoke the algorithm of Boolean operations on *general polygons*, defined as `GeneralPolygon_2` in CGAL. Specifically, the ‘CGAL::Cartesian<Number_type>’ is used as the kernel, in which ‘Number_type’ denotes the exact rational number-type by default (see the header file ‘arr_rational_nt.h’ in CGAL for reference). Based on this kernel,

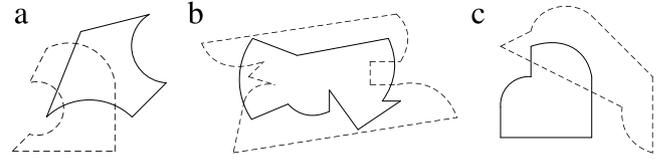


Fig. 8. The use cases for our experiments. The polygon with dashed lines is \mathcal{P}_1 ; \mathcal{P}_2 is another one. (a) For the first set of experiments. (b) For the fourth set of experiments. (c) A sampled example for which all algorithms work well, even if we use floating point type data.

we construct ‘CGAL::Gps_circle_segment_traits_2’ trait class, and the following objects ‘Curve_2, X_monotone_curve_2, General_polygon_2, Point_2’ are used, which inherit the trait class above.

Berberich. We directly use the algorithm in [9], which was initially developed for computing Boolean operations of conic polygons. This method employs the DCEL structure to represent the polygons. It first decomposes non-x-monotone conic curves and then computes the arrangement of segments using the plane sweep method. Next, it uses the results of the arrangement to compute the *overlap* of two polygons in order to achieve Boolean operations. This algorithm is similar to that of CGAL, recall Section 2. (Remark: more information about the DCEL structure and computing the overlap of two polygons can be found in [1].)

Naïve. It is one of simpler versions of our proposed method. This method employs our proposed data structure. However, it computes the intersections by comparing each pair of edges. Specifically, for each edge e of \mathcal{P}_1 , it checks whether e intersects with the edges of \mathcal{P}_2 . If so, it computes the intersections and inserts them into corresponding edges. It repeats these steps, until all edges of \mathcal{P}_1 are processed. The rest of the steps are to assign the *entry-exit* properties and to traverse, which are the same as the ones in our proposed method. Note that Greiner–Hormann’s algorithm [13], which was initially developed for Boolean operations of traditional polygons, also computes the intersections by comparing each pair of edges, and the rest of steps also include traversing. In this regard, the Naïve method can be also looked upon as a generalization of Greiner–Hormann’s method.

Standard. This is also a simpler version of our proposed method, but is different from the previous one. Specifically, it employs the standard plane sweep algorithm to compute the intersections instead of checking each pair of edges. Note that it does not adopt the proposed optimization strategies (e.g., ‘using related edges’, ‘avoiding false intersections’, ‘speeding up the lookups’, ‘speeding up the merging of decomposed arcs’). Other steps are the similar to the RE2L method. (Remark: the idea of the Standard method is roughly the same as that of CGAL and Berberich, but it simply removes the generality and employs a targeted data structure.)

10.2. Experimental settings and datasets

10.2.1. Settings

All the algorithms are implemented in C++ language. The versions of LEDA, CGAL and Boost library are 6.2, 4.3 and 1.46.1, respectively. The proposed algorithm and its simpler versions do not employ other libraries, except the *standard template library* (STL). The priority queue and the balanced tree mentioned in previous sections are implemented using a heap and a red–black tree, respectively. The experiments are conducted on a computer with 2.16 GHz dual core CPU and 1.86 GB of memory. By convention, we use the execution time to measure the efficiency. In our experiments, we run each algorithm 100 times by default, and then compute the average running time.

Table 2

The coordinates of vertices are listed counter-clockwise, and the left-bottom vertex is listed at first.

Polygon	Coordinates
\mathcal{P}_1	(15, 5), (88.5, 16.5), (80.95, 25.9), (69, 27.5), (60, 27.5), (60, 35), (70, 35), (72.35, 43.05), (68.5, 50.5), (4.5, 40.5), (13.65, 33.65), (25, 35), (20, 30), (28.5, 27.5), (22.4, 26.75), (18.5, 22)
\mathcal{P}_2	(33.5, 21), (39.5, 17.15), (46.5, 18.5), (46.5, 25.5), (56, 12.5), (70, 22), (64, 22), (67.9375, 31.9643), (66, 42.5), (36, 37), (21.5, 42.5), (17.4643, 29.32145), (21, 16)

Table 3

The average running time in the first set of experiments.

Method	CGAL	Berberich	Naive	Standard	RE2L
Time (s)	0.0273	0.0287	0.00239	0.00175	0.00112
Impr. fac.	24.375	25.625	2.13	1.5625	–

10.2.2. Datasets

Experiment 1. We manually produce two circular-arc polygons. Each of them is less than 10 edges, for simplicity and for ease of reproducing the findings. The vertex coordinates of the two polygons (cf., Fig. 8(a)) are listed in Table 1.

Experiment 2. To study the overall performance of these algorithms, we adopt thousands of circular-arc polygons. Specifically, given an integer n , a pair of circular-arc polygons with n edges are generated at random,⁷ and then each algorithm is executed alternately, using the pair of polygons as the input. This is done one thousand times. In each trial, the running time of each algorithm is recorded, and accumulated to previous trials. We compute the average value for estimating the average running time of each algorithm. Furthermore, we vary the value of n , and obtain the running time of each algorithm using the same method mentioned above.

Experiment 3. We use the parameter ‘double’ to replace the parameter ‘Number_type’ in the kernel ‘CGAL::Cartesian<Number_type>’ (the 1st approach), and also the parameter ‘Rational’ in the kernel ‘CGAL::Cartesian<Rational>’ (the 2nd approach), in order to investigate the performance when all these algorithms use the floating point number-type. We randomly generate pairwise circular-arc polygons as the test data.

Experiment 4. To answer the interesting issue mentioned in Section 8, we use polygons satisfying the condition $l = m + n$ as the input. See Fig. 8(b) (it is a sample). The vertex coordinates of these two polygons are listed in Table 2.

Experiment 5. Inspired by the curiosity, we also investigate the scalability of our proposed algorithm using polygons with a larger number of edges. Note that in this set of experiments almost all the polygons generated are self-intersection polygons when the number of edges is equal to or larger than 100. This is because it is pretty difficult to generate polygons without self-intersections when the number of edges is large, using the method mentioned in Footnote 7. Specifically, in this case we do not employ the constraint (ii), see Footnote 7.

10.3. Experimental results

10.3.1. Results of the first experiment

Table 3 lists the results of Experiment 1. The methods are listed in Row one, the average running time of each method

is listed in Row two, and the improvement factors⁸ of our algorithm over the baseline methods are shown in Row 3. From this table, we can see that the proposed method outperforms its simpler versions, demonstrating the effectiveness of our proposed strategies. Interestingly, the simpler versions of the proposed method outperform the former two methods, let alone the proposed method. To some extent this verifies our previous claim—directly executing existing algorithms used to compute Boolean operations of conic and/or general polygons is usually not efficient enough.

Compared to the former two methods, although the latter three ones adopt a different data structure, but we note that the ‘Standard’ method also directly employs the plane sweep method, similar to the former two ones. Viewed from a theoretical perspective, the former two methods should have performance similar to the ‘Standard’ method. To further verify this phenomenon and explain it, we conducted another set of experiments, evaluating the overall performance based on larger datasets.

10.3.2. Larger datasets

Table 4 lists the detailed results of Experiment 2. The results show that the proposed method outperforms the other four ones as well, and is several orders of magnitude faster than the former two. By comparing the differences of these methods, we can easily see that the reason for the larger running time of the former two methods may well be that both methods employ the CGAL library and the DCEL structure.⁹ Even so, it is noteworthy that comparing the former two methods with the latter three ones might be not very fair, since the latter three use *floating point arithmetic* (similar to that in [13,15–18,5,19–22]), whereas the former two ones use *exact algebraic arithmetic*, which is more robust.

To make a more fair comparison, a simple remedy is to let the input data type also be *machine floating point* for the former two methods. In the next paragraph, we report our findings when all these methods use machine floating point type.

10.3.3. Floating point number-type

Specifically, we use double type as the input data type for all these methods, recall Section 10.2.2. Next, we randomly generate a pair of circular-arc polygons and then attempt to call these methods. Unfortunately, the former two methods fail with the message ‘precondition violation’ for many inputs. For this step, we also attempt to generate many other (pairs of) circular-arc polygons, and to test them. As a result, in most cases the runtime exceptions are also reported. We also realize that, for a few test data (i.e., circular-arc polygons generated randomly), all these

⁷ Generally speaking, we first randomly generate two rectangles such that they overlap each other. Then, we randomly generate n points in each rectangle one by one such that they satisfy two constraints: (i) the segment joining the j th point and the $(j - 1)$ th one cannot intersect with any other segment except the segment joining the $(j - 1)$ th point and the $(j - 2)$ th one, where $j \geq 4$; and (ii) the segment joining the 1st point and the n th one cannot intersect with any other segments except its two adjacent segments. These n points will be used as the vertices of the circular-arc polygon. Finally, we import circular-arc segments by inserting a set of *appendix points*.

⁸ Here the improvement factor refers to the ratio of time. Assume that the execution time of the ‘Standard’ method is 0.2 s, that of the proposed method is 0.05 s, for example, the improvement factor is $\frac{0.2}{0.05} = 4$.

⁹ We remark that computing the intersections of two polygons is unavoidable for any clipping algorithm, and it is a dominant step [1,2,13]. Both the former two methods and the ‘Standard’ method adopt the plane sweep method to compute the intersections; their performance differences however are vast. This reminds us that the gap may well be due to the usage of the CGAL library and the DCEL structure (the former might be the major reason).

Table 4

The average running time of each algorithm, where n denotes the number of edges of each polygon.

n	CGAL	Berberich	Naive	Standard	RE2L
5	0.0281	0.0293	0.00234	0.0018	0.00107
10	0.0609	0.0772	0.0062	0.0041	0.00256
20	0.1282	0.1297	0.0125	0.0072	0.00369
30	0.1656	0.1741	0.0328	0.0176	0.00614
40	0.5172	0.5672	0.0391	0.0182	0.00683
50	0.61	0.681	0.0594	0.0213	0.00851

Table 5

The average running time when machine floating type is used for all these methods.

Method	CGAL	Berberich	Naive	Standard	RE2L
Time (s)	0.0112	0.0127	0.00192	0.00151	0.000948
Imp. fac.	11.814	13.396	2.025	1.593	–

Table 6

The average running time when the two polygons satisfy the condition $l = m + n$.

Method	CGAL	Berberich	Naive	Standard	RE2L
Time (s)	0.07844	0.08023	0.00718	0.004652	0.002965
Imp. fac.	26.4553	27.058	2.4215	1.5689	–

methods can work correctly.¹⁰ For example, when the vertex coordinates of two input circular-arc polygons (cf., Fig. 8(c)) are $\{(5, 2), (5.125, 1), (6, 0.5), (6, 3), (4, 5), (2.875, 5.25), (2, 4.5), (1, 4)\}$ and $\{(2, 4), (2, 3), (1.25, 2.75), (1, 2), (1, 1), (4, 1), (4, 3), (3.25, 4)\}$, all these methods can work correctly. We remember these two polygons and run 100 times for each method (using these two polygons as the input). Table 5 depicts their average running time. This table shows us that the proposed method also outperforms its simpler versions. Again, the simpler versions of the proposed method also outperforms the former two methods, which is similar to our previous findings (although the improvement factors decrease in terms of the former two methods). This verifies our original claim in a more justified manner.

10.3.4. The case $l = m + n$

Table 6 reports the results of Experiment 4. Interestingly, the proposed method still outperforms the ‘Standard’ method¹¹ (note: both algorithms have the same time complexity in this case, recall Section 8). This phenomenon reveals that (i) two algorithms with the same time complexity might have different performance results in terms of the execution time¹²; (ii) here other heuristics or optimization strategies (except the heuristic ‘utilizing related edges’) also bring us the benefits; and (iii) the performance gain of other optimization strategies is larger than the cost consumed by the operator ‘choosing related edges’.

10.3.5. Scalability

Fig. 9 reports the results when we vary the number of edges (of polygons) from 5 to 500. The columns denote the numbers of intersections and related edges respectively, whereas the curves

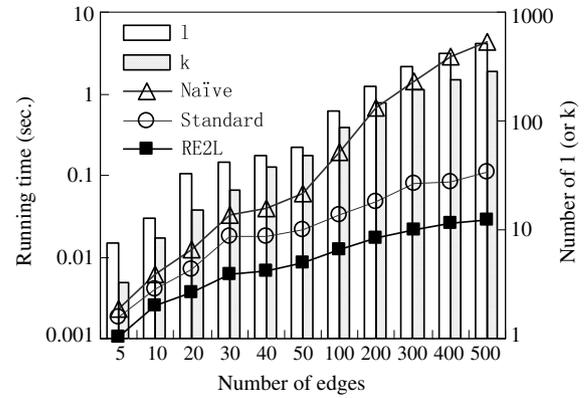


Fig. 9. The experimental results when we use larger data sets.

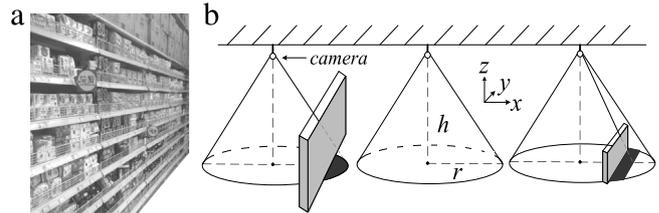


Fig. 10. (a) A corner of the ‘Auchan supermarket’. (b) An illustration of the visual range.

Table 7

The height information. Notice that when the heights of several different obstacles are almost the same, we adopt the average value for clearness.

Obstacle	[Color gradient bar]						
Height	4.50	2.85	2.15	1.8	1.65	1.2	0.5

denote the running time. From this figure, we can see that the RE2L has better scalability compared to its competitors, as the growth speed of the running time is slower than that of other two methods, when the number of edges increases. Compared to the ‘Standard’ method, the better performance of our proposed method is ascribed to those optimization strategies, and the poorer performance of the ‘Naive’ method is due to that computing the intersections in such a way is (somewhat) inefficient. Especially, this deficiency is more obvious when the number of edges of polygons is large.

10.4. Case study

In this subsection, we study our proposed algorithm based on a specific application. We use the ‘Auchan supermarket’ located at 2092 Dongchuan Rd., Shanghai, China, as a sample. We collect the deployment information of the shopping area on the 2nd floor of this supermarket. Fig. 10(a) exposes a corner of this shopping area. The length and width of the area are 100 and 60 meters, respectively. Its plane layout is depicted in Fig. 11 for the sake of intuition. The heights of different shelves or obstacles are listed in Table 7.

To connect our algorithm with the specific application, we consider a set of n_c free-rotating cameras which are to be placed on the ceiling for monitoring the shopping area; later, we shall employ our algorithm to check the blind angles. In the following, we restrict our attention to checking the blind angles on the ground, i.e., the $z = 0$ plane. (Remark: the blind angles for the $z \neq 0$ plane can be achieved similarly.)

In our experiments, we adopt two types of manners to deploy the n_c cameras: (i) they are placed uniformly; and (ii) they are placed randomly. For short, we denote by \mathbf{U} and \mathbf{R} these two

¹⁰ The former two methods are originally designed for the exact algebraic arithmetic. Here we use the machine floating point as the input data type. This could be the reason why most cases cannot work correctly.

¹¹ We remark that the results are similar when input polygons with more edges are used, omitted for saving space.

¹² This argument could be another reason why the performance of the ‘Standard’ method is significantly different from that of the former two methods. That is, it is possible that the CGAL library used in the former two methods pays more attention to robustness and stability, at the expense of performance. The further explanation is beyond the topic of this paper.

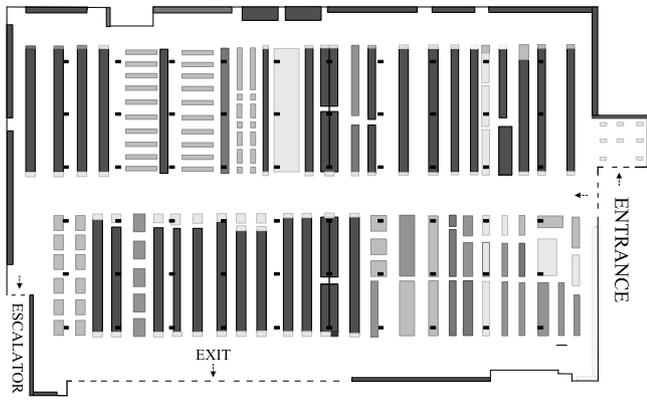


Fig. 11. The plane layout of the shopping area. The geometries filled in different colors are for different heights.

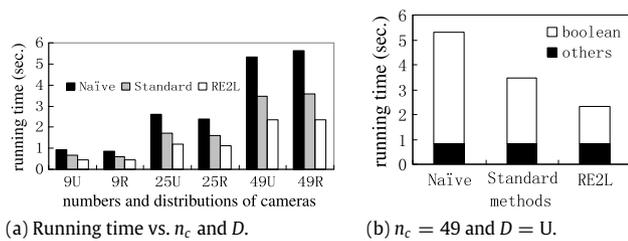


Fig. 12. Performances for computing the blind angles.

distributions, respectively. We regard the visual range of a single camera as a circle with center c and radius r when there is no obstacle near to it, where the x - and y -coordinates of c are the same as those of the camera but $z = 0$. Otherwise, we compute its visual range by considering the effect of obstacles. See Fig. 10(b) for an illustration of computing the visual range.

Let B_a , S_a , O_a , and V_a^i denote the blind angles (notice: here we do not consider the areas occupied by obstacles), the shopping area, the areas occupied by obstacles such as shelves, and the i th camera's visual range, respectively, where $i \in [1, \dots, n_c]$. We compute B_a based on the following equation.

$$B_a = (S_a - O_a) - \bigcup_{i=1}^{n_c} V_a^i.$$

For ease of reference, Table 8 lists the parameters used in our experiments. We use the 'Standard' and the 'Naïve' methods as baselines when we compute the blind angles, as both of them work well for the floating point data type. Fig. 12(a) reports the performance results when we use different settings. We can easily see that the execution time increases for all the methods when we vary the number of cameras from 9 to 49. This is mainly because for those 'new' cameras we also need to compute their visual areas and to remove them by Boolean operations, and so the total number of Boolean operations needed to be executed increases, which consumes more computation time. From this figure, we realize another interesting phenomenon: the distributions of cameras make little impact on the execution time (remark: here the distributions refer to *uniform* and *random*). Besides the above two findings, we can also see that our preferred method, RE2L, always outperforms its competitors for all these settings. This further validates the superiority of our algorithms. Furthermore, Fig. 12(b) reports the time consumed by Boolean operations and that consumed by all other operations. This shows us that improving the efficiency of Boolean operations is feasible and also important for applications like this.

Table 8

The parameters used in this set of experiments.

Para.	Desc.	Value
n_c	Number of cameras	9, 25, 49
r	Radius of the visual range	5
h	Height of cameras being placed	4
D	Distribution of cameras	U, R

11. Conclusions

In this paper we investigated the problem of Boolean operations on circular-arc polygons. By well considering the nature of the problem, the concise data structure and targeted algorithms were proposed. The proposed method runs in time $O(m + n + (l + k) \log l)$, using $O(m + n + k)$ space. We conducted extensive experiments demonstrating the effectiveness and efficiency of the proposed method, and showed that our techniques can be easily extended to compute Boolean operations of other types of polygons. We conclude this paper with two research topics: (i) As we know, multiprocessor and multi-GPU systems are widely used nowadays. It could be interesting to design efficient parallel algorithms for computing Boolean operations of polygons. (ii) As we showed in this paper, our techniques can be easily extended to compute Boolean operations of traditional polygons, and circular-arc polygons with holes and self-intersections; it is still open whether our techniques can be extended to compute Boolean operations on conic polygons or more general cases.

Acknowledgments

This work was supported by the National Basic Research 973 Program of China (No. 2015CB352403), the NSFC (No. 61272032, 61202024, 61202025, 61502220 and U1304616), the EU FP7 CLIMBER project (No. PIRSES-GA-2012-318939), Shanghai Pujiang Program (No. 13PJ1403900), Hong Kong GRF (152201/14E), Singapore NRF (CREATE E2S2), the Program for Changjiang Scholars and Innovative Research Team in University of China (IRT1158, PCSIRT).

References

- [1] de Berg M, Cheong O, van Kreveld M, Overmars M. *Computational geometry: algorithms and applications*. 3rd ed. Berlin: Springer; 2008.
- [2] Foley JD, van Dam A, Feiner SK, Hughes JF. *Computer graphics: principles and practice*. 2nd ed. Massachusetts: Addison-Wesley; 1996.
- [3] Gardan Y, Perrin E. An algorithm reducing 3D Boolean operations to a 2D problem: concepts and results. *Comput-Aided Des (CAD)* 1996;28:277–87.
- [4] Hoffmann CM, Hopcroft JE, Karasick MS. Robust set operations on polyhedral solids. *IEEE Comput Graph Appl* 1989;9:50–9.
- [5] Martinez F, Rueda AJ, Feito FR. A new algorithm for computing Boolean operations on polygons. *Comput Geosci (GANDC)* 2009;35:1177–85.
- [6] Weiler K, Atherton P. Hidden surface removal using polygon area sorting. In: *International conference on computer graphics and interactive techniques*, 1977, p. 214–22.
- [7] Wang ZJ, Wang DH, Yao B, Guo M. Probabilistic range query over uncertain moving objects in constrained two-dimensional space. *IEEE Trans Knowl Data Eng* 2015;27:866–79.
- [8] Wang ZJ, Yao B, Cheng R, Gao X, Zou L, Guan H, Guo M. Sme: explicit & implicit constrained-space probabilistic threshold range queries for moving objects. *Geoinformatica* 2016;20:19–58.
- [9] Berberich E, Eigenwillig A, Hemmer M, Hert S, Mehlhorn K, Schömer E. A computational basis for conic arcs and Boolean operations on conic polygons. In: *European symposium on algorithm*, 2002, p. 174–86.
- [10] Berberich E, Hemmer M, Kerber M. A generic algebraic kernel for non-linear geometric applications. In: *Symposium on computational geometry*, 2011, pp. 179–86.
- [11] Gong YX, Liu Y, Wu L, Xie YB. Boolean operations on conic polygons. *J Comput Sci Technol* 2009;24:568–77.
- [12] Andreev RD. Algorithm for clipping arbitrary polygons. *Comput Graph Forum* 1989;8:183–91.
- [13] Greiner G, Hormann K. Efficient clipping of arbitrary polygons. *ACM Trans Graph* 1998;17:71–83.

- [14] Krishnan D, Patnaik LM. Systolic architecture for Boolean operations on polygons and polyhedra. *Comput Graph Forum* 1987;6:203–10.
- [15] Liang YD, Barsky BA. An analysis and algorithm for polygon clipping. *Commun ACM* 1983;26:868–77.
- [16] Liu YK, Wang XQ, Bao SZ, Gombosi M, Zalik B. An algorithm for polygon clipping, and for determining polygon intersections and unions. *Comput Geosci* 2007;33:589–98.
- [17] Maillot PG. A new, fast method for 2D polygon clipping: Analysis and software implementation. *ACM Trans Graph* 1992;11:276–90.
- [18] Margalit A, Knott GD. An algorithm for computing the union, intersection or difference of two polygons. *Comput Graph* 1989;13:167–83.
- [19] Peng Y, Yong JH, Dong WM, Zhang H, Sun JG. A new algorithm for Boolean operations on general polygons. *Comput Graph* 2005;29:57–70.
- [20] Rappoport A. An efficient algorithm for line and polygon clipping. *Vis Comput* 1991;7:19–28.
- [21] Rivero M, Feito FR. Boolean operations on general planar polygons. *Comput Graph* 2000;24:881–96.
- [22] Vatti BR. A generic solution to polygon clipping. *Commun ACM* 1992;35:56–63.
- [23] Meguerdichian S, Koushanfar F, Potkonjak M, Srivastava MB. Coverage problems in wireless ad-hoc sensor networks. In: *IEEE international conference on computer communications*, 2001, p. 1380–87.
- [24] Wang B. Coverage problems in sensor networks: A surveys. *ACM Comput Surv* 2011;43:1–53.
- [25] Sutherland IE, Hodgman GW. Reentrant polygon clipping. *Commun ACM* 1974;17:32–42.
- [26] Bentley JL, Ottmann T. Algorithms for reporting and counting geometric intersections. *IEEE Trans Comput* 1979;28:643–7.
- [27] Fogel E, Halperin D, Wein R. *CGAL arrangements and their applications: A step-by-step guide. geometry and computing*. Springer; 2012.
- [28] Emiris IZ, Kakargias A, Pion S, Teillaud M, Tsigaridas EP. Towards and open curved kernel. In: *Symposium on computational geometry*, 2004, p. 438–46.
- [29] Wein R. Exact and approximate construction of offset polygons. *Comput-Aided Des* 2007;39:518–27.
- [30] Shamos MI, Hoey D. Geometric intersection problems. In: *IEEE symposium on foundations of computer science*, 1976, p. 208–15.