

University of Edinburgh



Department of Computer Science

The Polygon Package

by

E.E. Barton & I. Buchanan

Internal Report

CSR-44-79

**James Clerk Maxwell Building,
The King's Buildings,
Mayfield Road,
Edinburgh,
EH9 3JZ.**

October, 1979

THE POLYGON PACKAGE

E.E. Barton* & I. Buchanan

Department of Computer Science

University of Edinburgh

ABSTRACT

The polygon package is a set of procedures which manipulate geometric objects in the 2D plane. The operations that can be performed on these objects, which may be regarded as polygons, include intersection, union, asymmetric difference, inflation and deflation. Polygons may have both straight and circular arc edges and may include holes within their boundaries. The use of a binary tree of bounding boxes as a sorting aid in the centrally important procedure for calculating edge intersections results in a fast, efficient package.

* now at the Department of Computer Science, California Institute of Technology

1. INTRODUCTION

A polygon package may be used in any program that requires operations such as union, intersection, difference, inflation, and deflation. Examples of these operations are shown in Figure 1-1.

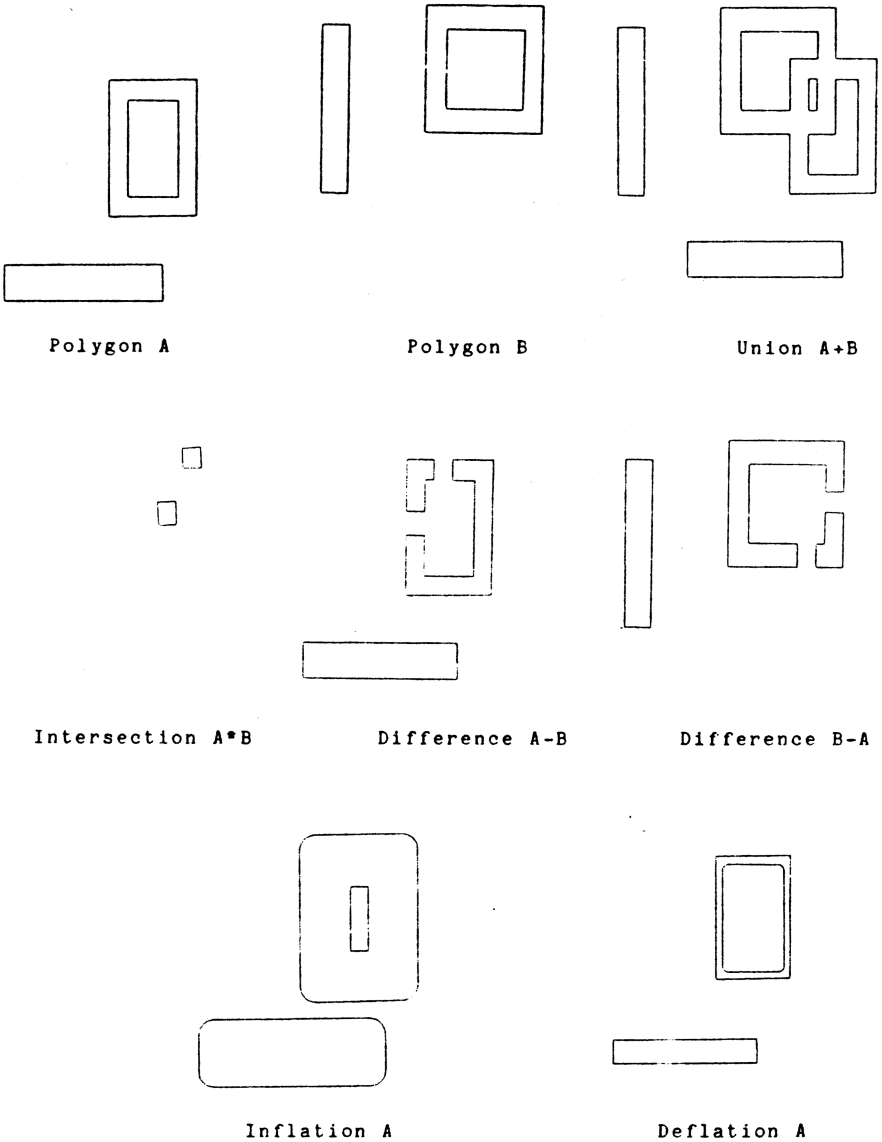


Figure 1-1

Many applications of such a package exist in the fields of, for example, computer graphics, architecture, engineering and

geographical data processing. The following examples illustrate two of the areas in which the package may be used.

One of the computationally most expensive operations in computer graphics is the removal of hidden lines or surfaces in three dimensional images. The polygon package may be used to remove the hidden regions provided that the surfaces do not intersect [1,2]. The relevant surfaces are projected onto the viewplane where they may be treated as polygons. By sorting these polygons on their distance from the view point and using the difference operator the hidden regions may be removed.

In VLSI circuit fabrication, the masks used to expose selective portions of resist on the silicon substrate are subject to certain dimensional design rules. These rules state that certain regions in the mask may not be closer than a defined distance, while others may not be less than a minimum size. Manual checking for violations of these rules is a time consuming, expensive and unreliable procedure. If the mask, or set of masks, is passed to the polygon package as a collection of polygons, the proximity of a set of regions to each other may be quickly and accurately tested by inflating those regions by half the test distance and calculating the intersection. If such intersecting areas are found then the design rule has been violated at precisely those locations.

Similar techniques are applicable in architecture and various engineering disciplines where structures may be subject to minimum and maximum size requirements.

2. DEFINITIONS

The generality and correctness of the definition of a polygon is the foundation on which the package builds the set of manipulation procedures. The set of polygons, corresponding to the definition below, must be closed under the operations union, intersection, difference, inflation and deflation. Briefly, a polygon consists of a number of sheets, each of which may have a positive or negative sense [3]. A sheet is a sequence of directed edges, which may be straight lines or circular arcs, connecting at vertices. Every point in the plane has an associated wrap number which is a measure of its level of enclosure by a set of sheets (Figure 2-1).




Figure 2-1

Detailed definitions follow:

An edge is a straight line or a circular arc connecting two vertices. The edge possesses direction from its tail vertex to its head vertex. The tail vertex is omitted from the set of points on the edge. Circular arc edges are represented by their centre point and radius. The radius is positive if the arc is described anti-clockwise from tail to head and negative if it is described clockwise.

A vertex is a point on the 2D coordinate plane at which two edges connect. It is the head vertex of its preceding edge and the tail vertex of its succeeding edge.

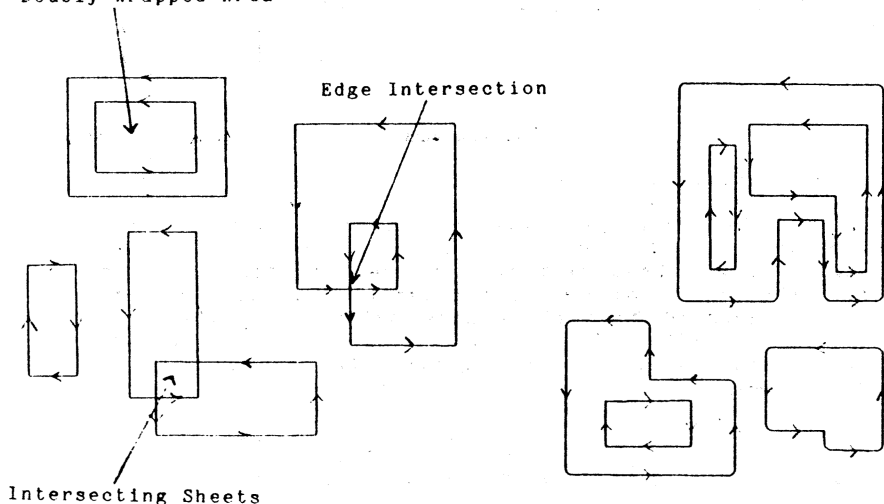
A sheet is a set of points enclosed by an ordered ring of alternating vertices and edges. The edges may not exhibit any self intersections. The sense of the sheet is either clockwise (negative) or anti-clockwise (positive) depending on the sense of the cyclic order of the edges and vertices.

The wrap number of a point is a measure of the number of times the point is encircled by a set of sheets. This number is positive if the point is enclosed by a net number of anti-clockwise sheets. It is negative if the point is enclosed by a net number of clockwise sheets. If the point lies on the edge of a negative sheet then it is considered to be outside the sheet and the sheet makes no contribution to the wrap number of the point. If the point lies on the edge of a positive sheet then the point is considered to be inside the sheet and the wrap number is incremented.

A polygon is a set of associated sheets such that none of the sheets exhibit edge intersections with any of the others and, furthermore, relative to this set of sheets all points in the plane are either inside the polygon and have wrap number +1 or are outside the polygon and have wrap number 0 ie no sheet can be enclosed by another of the same sense and no negative sheet can exist outside a positive sheet.

All well-formed polygons conform to the above definition. Polygons which violate any of the rules are termed malformed and cannot be reliably manipulated by the union, intersection, difference, inflation and deflation procedures unless they are first processed by the package to remove the irregularities. Examples of well-formed and malformed polygons are shown in Figure 2-2.

Doubly Wrapped Area



Malformed Polygon

Figure 2-2

Well-formed Polygon

3. PRIMITIVE OPERATIONS

The high level polygon manipulation procedures of union, intersection, difference, inflation and deflation depend on several lower level operations. The most basic of these is the edge intersection routine. On this may be built the calculation of the wrap number which is the basis for deducing the topology of any polygonal interaction. The use of a sorting aid increases the speed of these primitive operations by reducing the number of calculations to be performed. The technique used in this package is a bounding box based binary tree [4].

3.1 BOUNDING BOXES

Manipulating polygons ultimately reduces to calculations on pairs of edges to discover the existence and locations of any intersections. Testing each edge in one polygon against every edge in another yields an algorithm of $O(n^2)$ complexity. For large polygons with few intersections this is clearly unacceptable. It is necessary to develop some method of disregarding unproductive pairs of edges. The technique used in this package is based on the paraxial bounding box (Figure 3-1), a simply described rectangle which encloses all the points on an edge and within a surrounding rectangular region.



Figure 3-1

To achieve an algorithm with the desired performance it is necessary to merge the primitive bounding boxes so that sets of edges may be compared (Figure 3-2).

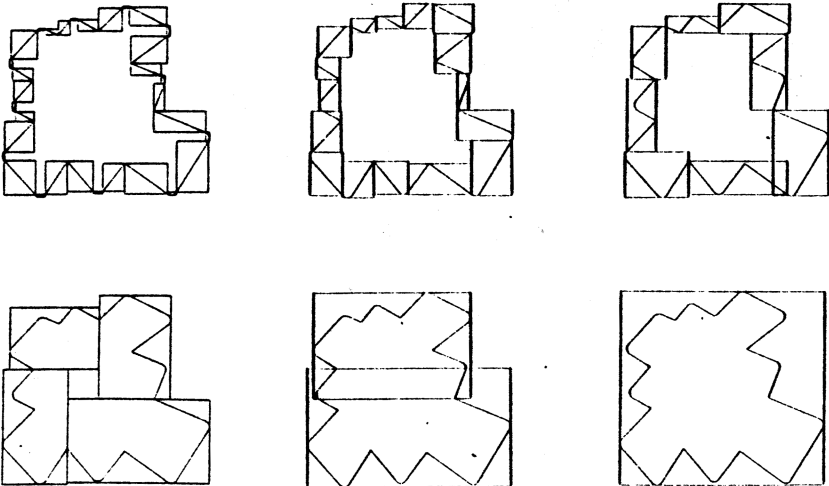


Figure 3-2

The necessary hierarchy is realized in a binary tree structure in which the root is the bounding box of the complete polygon and the leaves are the bounding boxes of the individual edges. The tree is built up by merging the primitive bounding boxes according to the cyclic order of the edges in the sheets. The sheets are sorted by proximity and merged accordingly until the root of the tree, the polygon, is reached and is enclosed in a single bounding box. Below the sheet level the tree building is performed by an algorithm of $O(n)$ complexity, in fact in $2n-1$ operations. Above that level the sorting of the sheets imposes an $O(n \log n)$ complexity.

3.2 EDGE INTERSECTIONS

The bounding box hierarchy is used to concentrate on only those areas where intersections are likely to occur. At each call of the edge intersection procedure two bounding boxes are referenced and compared for a possible intersection. Initially these are the bounding boxes of the two polygons. No intersection between bounding boxes implies no intersection between any of the edges they contain and therefore the search down these branches of the two trees can be terminated. If an intersection does occur then one of the bounding boxes, that containing the greater number of edges, is divided into its two successors. The algorithm is then applied recursively to each new pair of bounding boxes. Once two primitive bounding boxes have been reached an edge intersection is possible and the appropriate calculation is performed. Figure 3-3 shows a pair of intersecting polygons and the bounding boxes examined by the searching algorithm.




Figure 3-3

The number of bounding boxes tested by the algorithm is greatest near the edge intersection points, thus demonstrating the efficiency of the search method.

The edge intersection procedure must distinguish between three types of interactions viz straight edge and straight edge, straight edge and circular arc, circular arc and circular arc. Calculations involving circular arc edges are more complex than their straight edge counterparts but, since edge intersection calculations occur only when strictly necessary, this makes no appreciable difference to the overall performance of the package while considerably enhancing its functionality.

3.3 WRAP NUMBER

The wrap number of a point determines its level of nesting within a set of sheets. A line is constructed from the test point rightwards for an infinite distance. The number of intersections between this line and the edges of the polygon are counted and the sense of the intersection determines whether the count is incremented or decremented. Again, the bounding box hierarchy produces great savings in computation by restricting the number of edge intersection calculations to be performed. The procedure calls itself recursively and at each level it references a bounding box and the test point. Bounding boxes are trivially rejected if they lie completely to the left, above or below the constructed test line. A bounding box which is intersected by the line is divided into its two successors and the algorithm is applied to both, recursively, until a primitive bounding box is reached (Figure 3-4).




Figure 3-4

At this stage the normal edge intersection calculation is performed and several alternative cases are recognised:

- (1) The test point lies to the left of the bounding box.

The test line crosses the edge, if and only if one of its vertices lies on or above the line and the other lies below the line (Figure 3-5). Circular arcs therefore make no contribution to the wrap number if both their vertices lie to one side of the line, as their crossings are in opposite directions and cancel each other out (Figure 3-6). The wrap number is incremented for edges crossing in an anti-clockwise direction ie the tail vertex lies below the head vertex, and decremented for edges crossing in a clockwise direction.




Figure 3-5




Figure 3-6

(ii) The test point lies within the bounding box.

(a) The edge within the bounding box is a straight line.

The position of the point relative to this edge is tested. If the point lies to the left of the edge, the wrap number is incremented or decremented as in (i) (Figure 3-7). If the point lies to the right of the edge, no crossing of the test line has occurred and no contribution to the wrap number is made (Figure 3-8).




Figure 3-7




Figure 3-8

If the point lies on the edge then either the edge points upwards or it points downwards; horizontal edges are classed as pointing upwards if they point to the right and downwards if they point to the left. The wrap number is incremented for all points on upward pointing edges, excluding the tail and head vertices. A point on the head of an upward pointing edge increments the wrap number if the next edge points upwards or to the left of the edge (Figure 3-9). The wrap number is not affected by downward pointing edges unless the point is on the head vertex and the next edge points upwards and to the right of the edge (Figure 3-10).




Figure 3-9




Figure 3-10

(b) The edge within the bounding box is a circular arc.

Tests are made to discover whether or not the point lies on the arc and also how many times the test line crosses the arc. If the arc is crossed twice no net change occurs to the wrap number as the directions of the crossings are opposite (Figure 3-11). If the line crosses the arc only once then the wrap number is incremented or decremented according to the direction of the arc at the crossing point (Figure 3-12).




Figure 3-11




Figure 3-12

Once again, particular attention must be paid to points on the edge and its two endpoints. If the test point lies on the arc but not at a vertex then the wrap number is incremented only for arcs of positive radius (Figure 3-13). If the test point lies on the tail vertex (Figure 3-14) then no change is made to the wrap number for arcs of positive or negative radius unless the test line intersects with another portion of the same arc.






Figure 3-13




Figure 3-14

If the test point lies on the head vertex then a similar series of further tests is required as for head vertices of

straight edges. If the tangent to the arc at the head vertex points upwards and the next edge points upwards or to the left of the arc (Figure 3-15), or the tangent to the arc at the head vertex points downwards and the next edge points upwards and to the right of the arc (Figure 3-16) then the wrap number is incremented.




Figure 3-15

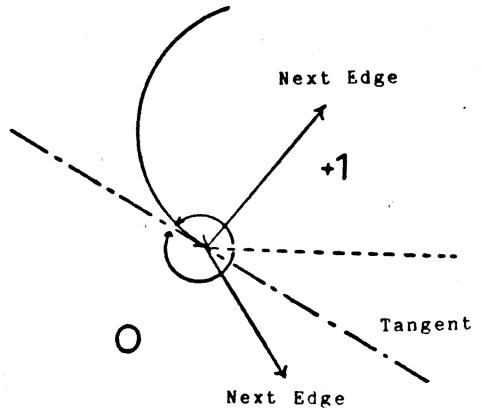


Figure 3-16


This set of tests is consistent for all well-formed polygons. More tests are performed when the test point lies on a vertex because more information is needed to decide whether the point forms an upper or lower extremity and, if so, in a clockwise or anti-clockwise sheet.

4. POLYGON OPERATIONS


The polygon manipulation operations that the package provides fall into two categories; unary (inflation and deflation) and binary (intersection, union and difference). All operations result in a single well-formed polygon. In addition three other unary operations (self-union, overlap and underlap) are supplied to discover and repair malformed polygons. Self union trims off any malformed regions of the polygon, overlap extracts all areas where the polygon intersects itself and underlap finds all the regions that are encircled a net negative number of times by the edges of the polygon.

4.1 BINARY OPERATIONS


Let $wrap(r,R)$ be a function which returns the wrap number of a point r with respect to a polygon R . Then for a given point p in polygon P , $wrap(p,P)=1$. Thus for polygons P & Q , p is in union $P+Q$ if and only if $wrap(p,P)=1$ or $wrap(p,Q)=1$. Similarly p is in intersection $P*Q$ if and only if $wrap(p,P)=1$ and $wrap(p,Q)=1$. The asymmetric difference may be expressed as p is in difference $P-Q$ if and only if $wrap(p,P)=1$ and $wrap(p,Q)=0$. Figure 4-1 shows the results of these operations on an interaction between two polygons, one with two curved sheets and the other with a single straight-edged sheet. In each case the result of the operation has been slightly inflated or deflated to make its position obvious.




Polygon P




Polygon Q




Union P+Q



Intersection P*Q



Difference P-Q
Figure 4-1



Difference Q-P

Clearly it is necessary to consider only points on the edges of sheets when manipulating the polygons within the package. In fact, the points of interest can be limited even further to the edge intersections. Accordingly an algorithm has been developed which constructs a new polygon by tracing all the edges with the appropriate wrap number.

To calculate the result of a set operation on two polygons the algorithm starts by finding all edge intersections. Not all of these intersections will be useful to the tracing routines. The following tests select only those intersections which are relevant and will form decision points during the tracing of the output polygon.

(i) The intersection lies on the tail vertex of one or both of the two edges.

The intersection is always ignored as the tail vertex of an edge is excluded from the edge for the purpose of finding edge intersections since the tail vertex of one edge forms the head vertex of the preceding edge.

(ii) The intersection lies on the head vertices of both the edges.

The intersection is ignored if either the two edges are parallel and their successors are parallel, or if both edges are anti-parallel to the other edge's successor (Figure 4-2).




Figure 4-2

(iii) The intersection occurs on the head of one edge.

If the successor of the edge with the intersecting head is parallel to the other edge then it is ignored (Figure 4-3).




Figure 4-3

In effect, parallel edges, either singly or in groups, will yield a single intersection where they diverge.

The edge intersection routines build the data structure required by the tracing routines. The tracing routines must now examine the valid intersections and decide which paths to follow round the initial polygons to produce the desired binary operation result. The same procedure, parameterised by the appropriate wrap number, is used to produce the union, intersection or difference. The process begins by passing the first valid edge intersection to the tracing routine as the initial vertex of a sheet.

The step that selects which edge to follow makes the decision with reference to a number of conditions. If nothing is known about the current edge intersection *eg* on first entering the procedure, then a test must be made of the wrap number of a pair of points on the edges about to be traced (Figure 4-4).




Figure 4-4

These wrap numbers are compared with the desired wrap number with the following possible outcomes.

(i) Neither wrap number equals the parameter.

The intersection point must lie on some malformed region of one or both of the two polygons. Such a configuration should never arise if the polygons passed to the package are well-formed. Any edges added to the output polygon by the tracing procedure up to this point are disassembled and the point is marked as having been analysed.

(ii) One of the points has a wrap number equal to the parameter.

The edge with the matching wrap number is selected as the path on which tracing will be continued. The intersection is marked as having been analysed.

(iii) Both points have wrap numbers equal to the parameter.

This case occurs when two edges which will appear in the output polygon touch, but do not intersect, one another. The point is marked to indicate the edge on which tracing will be continued so that the next time this point is reached the other edge will be selected immediately without recalculation of the wrap numbers. If the procedure is generating the intersection of two polygons then tracing will be continued on the current path. If the union or difference is being traced

then the path is switched.

The procedure continues by copying edges and vertices to the output sheet from the selected path until another intersection point is reached. If the current and intersecting edges are continuous over the intersection point then there is no need to repeat the above tests and the tracing path is switched to the other polygon (Figure 4-4). The intersection point is marked as having been analysed.

The procedure continues to add edges to the output sheet until the initial intersection is rediscovered. If only a single valid path exists from this point then the output sheet is closed. However, if both paths from the intersection were valid and only one has been traced, then the current sheet may be continued. Tracing is continued for the following reasons.

(i) The intersection is being constructed and the starting point is not approached on the path that was initially selected (Figure 4-5).

(ii) The union or difference is being constructed and the starting point is approached on the initially selected path (Figure 4-6).




Figure 4-5






Figure 4-6

In all other cases the output sheet is closed (Figures 4-7 & 4-8).



Union
Figure 4-7



Intersection
Figure 4-8

When a new sheet has just been closed the next intersection that has either not been analysed or possesses a remaining untraced path is passed to the tracing procedure as an initial vertex. Sheets are added to the output polygon until all intersections have been analysed.

Any sheets in the two initial polygons which did not exhibit edge intersections are now tested. They are included in the output polygon if the wrap number of a point on one of their edges, calculated relative to both polygons, equals the desired wrap number. This completes the construction of the output polygon.

4.2 INFLATION AND DEFLATION

Inflating a polygon expands all the anti-clockwise sheets and contracts all the clockwise sheets by a given distance. Conversely, deflating a polygon contracts all the anti-clockwise sheets and expands the clockwise sheets by a given distance. These two operations are implemented in the same procedure which takes a positive or negative parameter for inflation and deflation respectively (Figure 4-9).




Figure 4-9

A strict definition of the output polygon places the edges of the output polygon exactly the given distance from the edges of the input polygon. However, some algorithms have been content merely to shift the edges of the input polygon outwards or inwards by the desired distance without regard to the discrepancies that can occur at the vertices of the output polygon. This course of action has been followed by packages which do not include edges modelled by circular arcs. Some simplifications in code and data can be obtained by this view of the problem but at a cost in accuracy and functionality.

The expansion of a vertex on an anti-clockwise sheet, whose edges form an internal angle of less than 180 degrees, should produce a circular arc centred on the original vertex with radius the inflation distance (Figure 4-10). An algorithm that acts only on straight lines approximates this by extending the two edges until they intersect. If the edges are almost parallel but in opposite directions then their endpoints will be extended by a distance many times greater than the inflation distance (Figure 4-11). Such inaccuracies are unacceptable in many applications.




Figure 4-10








Figure 4-11

The polygon package uses a single algorithm to produce both inflation and deflation of sheets. Every edge in the input sheet is shifted outwards or inwards by the appropriate distance. Straight edges are translated along their unit normal vectors and circular arcs have their radii increased or decreased by the given distance. Each vertex in the input sheet is the source of a new circular arc in the output sheet centred on that vertex. This process produces a curlicue at each vertex (Figure 4-12). The curlicues may cause malformed regions in the output sheet if the edges become intersecting. Such regions are removed by the repair procedures included in the package. The output includes only those edges in the output sheet which have wrap numbers equal to that of the edges in the input sheet. Sheets which have been transformed so far as to reverse the sense of the sheet are discarded (Figure 4-12).



Inflation



Deflation

Figure 4-12

While the individual sheets of the output polygon are well-formed the polygon itself may be malformed. This occurs when two adjacent sheets have been inflated to such an extent that their boundaries now cross. A well-formed polygon may be produced in the usual way by the repair procedures but this is not done automatically since the user may require the data in the pre-repair format, possibly so that some check on overlap areas can be performed.

4.3 UNARY OPERATIONS

The definition of the polygon has so far provided a set of rules within which the operations of union, intersection, difference, inflation and deflation have been implemented. These operations must be guaranteed to produce correct output polygons from well-formed input data. Polygons passed to the package could be malformed and it is necessary to provide a checking procedure in order that a user may have confidence in the package. Additionally such checks are of use internally eg following an inflation which may cause self-overlap. Malformations are detected by the presence of edge intersections combined with topological tests. A number of configurations can cause the polygon to be malformed and a number of these are shown in figure 4-13. Any of these conditions will cause the result false to be returned by the checking procedure. The self-union and self-intersection procedures repair the malformed polygons by tracing the singly and doubly wrapped regions respectively. Any holes in free space are discovered by following edges with wrap number 0.




Figure 4-13

The implementation of the malformation detector and the repair procedures rely heavily on the intersection and tracing routines already described. In the binary operations the tracing routine is used to discover all the sheets possessing edges with the required wrap number. Exactly this operation is performed on single polygons in the unary operations. The difference lies in constructing the intersection lists which the tracing procedure uses.

All the edge intersections of the input polygon must be found before tracing can commence. The usual edge intersection procedure is used to find these intersections but it is only efficient when finding the intersections of groups of edges. A second procedure is required to present a series of groups of edges to the intersection routine in such a manner that, eventually, all the edges have been compared for intersections with all other edges. The bounding box trees consist of groups of edges. Every edge in the polygon will be implicitly compared with every other edge if the intersection procedure is called with the two sons of every element in the tree, excluding the leaves. This has been implemented in a recursive procedure which is initialised with the bounding box of the whole polygon. It calls the intersection routine with the two sons of the current bounding box and then recurses on these. The routine returns immediately if the bounding box is primitive. If there are n edges in the polygon then there will be $2n-1$ elements in the tree and this procedure will perform $n-1$ calls of the edge intersection procedure. Thereafter the tracing procedure will be called upon to retrieve the negatively, singly and doubly wrapped regions with a wrap number parameter of 0, 1, & 2 respectively (Figure 4-13). The output polygon is well-formed and can be used in further polygon package operations.

4.4 AREA AND PERIMETER

Two further procedures are provided viz area and perimeter, both of which return real number results. The area of a polygon is the total area of all its sheets, remembering that negative sheets have negative area. A polygon may not however, have a negative area since a negative sheet may not exist outside a positive sheet.

The perimeter of a well-formed polygon is simply the total absolute length of all its edges, ignoring direction and sheet nesting depth.

5. CONCLUSIONS

The polygon package provides a powerful and efficient tool which may be employed in a variety of application areas. Within the stated definitions, the set of polygons is closed under the operations of intersection, union, difference, inflation and deflation supplied by the polygon package. The bounding box hierarchy enables groups of edges to be compared for intersection, and is effective in reducing the average number of edge intersection calculations below the upper bound of $O(n^2)$ complexity. The inclusion of circular arcs considerably extends the application area of the package beyond that available to its straight line counterpart. The additional procedures for detecting and repairing malformed polygons enable input data to be verified so that the polygons resulting from any of the operations supplied by the package will be well-formed. The resulting package is therefore both functionally complete and efficient in performance.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the major contribution to this work made by Ivan Sutherland, Professor of Computer Science at the California Institute of Technology. I. Buchanan was a graduate student at Caltech during 1978. This implementation was undertaken by E. Barton as a final year undergraduate project.

REFERENCES

- [1] Newman W.M. & Sproull R.F.
Principles of Interactive Computer Graphics
Second Edition, McGraw-Hill, 1979.
- [2] Sutherland I.E., Sproull R.F. & Schumacker R.A.
A Characterisation of Ten Hidden Surface Algorithms
Computer Surveys Vol 6, No 1, March 1974.
- [3] Sutherland I.E.
The Polygon Package
California Institute of Technology, 1978.
- [4] Burton W.
Representation of Many-Sided Polygons and Polygonal
Lines For Rapid Processing
CACM Volume 20, Number 3, March 1977.

