The Combinatorics of Overlapping Convex Polygons in Contact

David Parsons*
EECS department, UC Berkeley, Berkeley CA
Carme Torras†
Institut de Cibernètica, UPC–CSIC, Barcelona, Spain

Abstract

We consider the following problem: given any two convex polygons which are free to translate and rotate arbitrarily in the plane and which have m and n vertices respectively, what is the maximum number of ways they may contact each other such that three independent boundary contacts are made? This problem is similar to several others which have been investigated in the field of combinatorial geometry, such as the polygon containment problem, but has the significant difference that in our context the polygons are allowed to overlap. We present the results of an empirical study of the number of such configurations, and we prove that an upper bound on the total is $O(m^3n^2 + m^2n^3)$.

1 Introduction

This paper is concerned with a problem in combinatorial geometry in the plane. Roughly speaking, we are interested in the number of ways in which two convex polygons may touch each other such that three independent boundary contacts are made. In other words, the polygons are considered to be free to translate and rotate arbitrarily, including in overlapping configurations, and we wish to establish an upper bound, in terms of the complexity of the polygons, on the number of relative configurations satisfying the triple contact condition. We present the best-known bound to date, as well as some empirical evidence that perhaps a better bound is possible.

Given two convex polygons P and Q, we denote the set of vertices of P (respectively Q) by V_P (resp. V_Q), and the set of edges of P (Q) by E_P (E_Q). Then a *simple contact* is defined as an ordered pair of one of the following two types:

- type I: (v_P, e_Q), where $v_P \in V_P$ and $e_Q \in E_Q$
- type II: (e_P, v_Q), where $e_P \in E_P$ and $v_Q \in V_Q$

*This paper was written while the first author was on leave at the Institut de Cibernètica.

†This research was supported by the ESPRIT II Basic Research Actions Program of the EEC under contract No. 3075 (project ALCOM).

The *contact manifold* associated with a simple contact (v_P, e_Q) (respectively, (e_P, v_Q)) is defined as the set of all relative configurations of P and Q in which v_P lies on e_Q (resp., v_Q lies on e_P). Now, a *critical configuration* of the two polygons is defined as a relative configuration in which three independent simple contacts are made simultaneously. Such a configuration is the intersection of three independent contact manifolds. The term *independent* in this definition ensures that no relative configuration between P and Q will be counted more than once. For example, when a vertex of P lies on a vertex of Q, it can be seen that four simple contacts are made; however since there remains one degree of freedom only two of them are independent.

Define $K(P, Q)$ as the number of critical configurations for a particular pair of polygons P and Q. The problem, then, is to find an upper bound on

$$K(m, n) \equiv \max\{K(P, Q) \mid P \text{ has } m \text{ vertices and } Q \text{ has } n\}.$$

Various similar combinatorial problems have been discussed in the robotics and computational geometry literature. The aspect of our problem which makes it distinct is that we allow the polygons to overlap. In most of the robotics applications, for example, one of the polygons is taken as a model of a mobile robot, while the other models a workspace, and in such a context any overlap represents a non-free configuration and hence does not count. For example, [LS87] gives an upper bound of $O(mn\lambda_s(mn))$ on the number of independent contact triples (called "critical contacts") for the case of a convex robot P and an arbitrary polygonal environment Q, which may be non-convex and even disconnected, where s is a constant less than or equal to 6 and $\lambda_s(q)$ is an almost linear function known to be in $O(q \log^* q)$ for any constant s (see [Sze74]). Our problem is also related to the "polygon containment" problem. There the goal is to determine whether a given polygon will fit inside another. [Cha83] proves that this is true if and only if there exists a relative configuration in which three simple contacts are made, and in addition gives an algorithm to compute all such configurations in $O(m^3n^3(m+n)\log(mn))$ time for arbitrary polygons of size m and n. In this problem too the polygons are not allowed to overlap.

The original motivation for considering this problem was that it relates to the problem of computing the five-dimensional "contact space" between two convex polyhedra. (This contact space is equivalent to the boundary of the configuration space obstacle which would arise in the problem of motion planning for one convex polyhedron in the presence of a convex polyhedral obstacle.) Contact space can be used, for example, for planning relative motions between two contact configurations during which the polyhedra remain in contact throughout. In [TT89], it was shown that the graph of the set of 5 and 4 dimensional strata of the contact space along with their adjacencies is sufficient for solving this contact motion planning problem, and furthermore that this graph has size $\Theta(mn)$ and can be computed in (optimal) time $O(mn)$, where m and n in this case are the numbers of edges in the polyhedra. This time bound is linear in the number of constraints. However, in that paper the question of the complexity of the lower-dimensional strata of contact space was left open. It turns out that the key question concerning the numbers of lower-dimensional strata is the complexity of the sub-graph induced by a planar contact between a pair of the faces of the polyhedra. The size of this sub-graph depends linearly on the number of 0-dimensional strata, which correspond to the combination of a planar face-to-face contact and 3 independent simple contacts between those faces. The number of these latter is exactly the question addressed in the present paper.

Given any fixed triple of independent simple contacts, it is known that there cannot exist more than 4 relative configurations of P and Q which satisfy all three of them (see e.g. [SS83]; in fact, [LS87] conjecture that in fact there can be at most 2 such configurations). Therefore, a trivial upper bound on $K(m,n)$ is $4\binom{2nm}{3} \in O(m^3n^3)$. However this bound, which is cubic in the number of simple contacts, holds also for non-convex or even disjoint polygons, and it is intuitively clear that convexity should substantially reduce the number of possible critical configurations.

The next section describes some empirical evidence that indeed for specific sequences of pairs of real convex polygons the number of critical configurations grows considerably slower than the cubic bound. This evidence has motivated the search for a rigorous theoretical upper bound, which is the subject of section 3.

2 An algorithm for computing critical configurations

In ([PT91]), we described an algorithm for computing all the critical configurations which may arise for a specific pair of input polygons. That algorithm uses what is essentially a brute-force generate-and-test approach, in that in some sense it considers all possible triples of single contacts and determines whether they are simultaneously satisfiable. It uses pair-wise distance constraints between the contacts to prune from the tree of possibilities those triples which need not be considered; each triple which passes the distance constraints is then tested by solving explicitly for a rigid transformation which would bring one polygon into contact with the other such that each of the three contacts is made. The details of the algorithm and of its implementation are described in [PT91]; here we simply wish to quote the main empirical results from that report.

Critical configurations may be classified into different types, depending on the types of the simple contacts which they satisfy. First of all, we distinguish among the following three configuration types:

type A: 1 vertex–vertex contact, 1 simple contact

type B: 1 edge–edge contact, 1 simple contact

type C: any 3 simple contacts NOT satisfying the definition of type A or type B

Note that types A and B are indeed configurations satisfying 3 simple contacts, but which have an additional constraint concerning which sets of contacts may constitute such a configuration. For example, to constitute a type A configuration, two of the contacts must involve the same vertex, and the edges of those two contacts must share a common endpoint. We have been primarily interested in establishing non-trivial upper bounds for the number of type C critical configurations which may occur. This is because the numbers of the other two types are easier to predict; indeed, an immediate trivial upper bound for both types is $O(m^2n^2)$. It is also useful to classify critical configurations according to the types of the simple contacts they contain. There are four possibilities: of the three contacts, either 0, 1, 2, or 3 may be of type II (edge–vertex), the others being of type I (vertex–edge). We denote these possibilities as case 0, case 1, case 2, and case 3, respectively. Note that due to the intrinsic symmetry of the problem, the complexity of the case 0 configurations is the same as that of the case 3 configurations, but with the roles of P and Q reversed; the same symmetry exists between cases 1 and 2.

We now turn to describing some results obtained from an implementation of this critical configuration generating algorithm. The algorithm has been implemented in the C language on a Sun workstation. The program takes as input a pair (P, Q) of convex polygons and proceeds to compute the type C critical configurations of those polygons. Because of the symmetry noted above between case 0 (resp. case 1) critical configurations and case 3 (resp. case 2) critical configurations, we have not bothered to implement the latter two cases. The output consists of a list of pairs, each containing a simple contact triple and a transformation; each transformation when applied to polygon P yields a configuration in which P contacts Q in exactly the ways specified by the corresponding triple.

To facilitate the input and output to the program, a graphical interface has been constructed. The user can input the pair of polygons by drawing them on a graphics window on the screen, and then can browse through the resulting list of critical configurations. Each configuration is displayed by redrawing the polygon P superimposed on Q, so that it is possible to verify that the contacts are being made. In addition, for each run of the program some statistics are printed. An example of these is as follows:

```
----------------
|P| = 5, |Q| = 7
----------------

Number of possible simple contacts: 70

<applying distance constraints>

Of 54740 possible triples of contacts:
    906 are feasible and 53834 are infeasible.

<solving for critical configurations>

93 type-C critical configurations found:
    21 case-0, 72 case-1
```

In this example, P was a convex polygon with 5 vertices and Q had 7. Thus the number of simple contacts was $2 \times 5 \times 7 = 70$. The number of a priori possible triples was therefore $\binom{70}{3}$, which is 54740. By applying the distance constraints (aluded to above and described in detail in [PT91]), the number of these which needed to be examined further was 906, or about 1.7%. Of these, only 93 turned out to yield transformations which really placed P on Q in such a way that the three contact constraints were met.

Experience with the output of the program reveals that the largest number of critical configurations tends to be produced when the two input polygons are approximately regular, and have the same radius and number of vertices. To give an idea of how many critical configurations this situation may yield, the following table has been culled from program output:

SUMMARY of approximately-regular polygon data

n	K(P,Q) case-0	case-1	total	K/n^2	K/n^3	K/n^4
4	4	26	30	1.88	0.469	0.117
5	15	90	105	4.20	0.840	0.168
6	85	232	317	8.81	1.468	0.245
7	185	297	482	9.84	1.405	0.201
8	287	829	1116	17.44	2.180	0.272
9	430	852	1282	15.83	1.759	0.195
10	662	1281	1943	19.43	1.943	0.194
11	956	2220	3176	26.25	2.386	0.217
12	1567	3329	4896	34.00	2.833	0.236
16	2850	6225	9075	35.45	2.216	0.138
20	7979	19814	27793	69.48	3.474	0.174

Each row of the table corresponds to a separate run of the program. The quantity n refers to the number of edges in each of the polygons; the next three columns give the number of critical configurations which were computed for that run. Each row for $n < 16$ is the data from the run which, out of three or four runs for that value of n, yielded the highest number of critical configurations. (The runs for $n = 16$ and $n = 20$ were performed only once, since they each required over an hour of computer time.) The rightmost three columns compare the critical configuration totals with the second, third, and fourth powers of n, in order to give an idea of the growth rate of the critical configuration count. Although it is of course impossible to make any definite conclusions based on such sparse data, it certainly seems plausible that the number of critical configurations grows *faster* than n^3, and possibly as fast as n^4. Similarly, it seems unlikely that the growth rate is any larger than n^4. Thus we conjecture that $K(n,n) \in O(n^4)$, which is quadratic in the number of constraints.

This (admittedly weak) empirical evidence that $K(m,n)$ grows considerably slower than the trivial cubic bound $O(m^3n^3)$ has led us to try to establish an upper bound closer to the quadratic one. The results of this effort will be described next.

3 An upper bound on $K(m, n)$

In this section we present a proof of an upper bound on $K(m, n)$ which is somewhat better than the cubic trivial bound, although not as good as the quadratic bound conjectured in section 2 above. Specifically, we show that

$$K(m, n) \in O(m^2n^3 + m^3n^2).$$

For technical reasons, we will assume throughout this section that neither of the two polygons P and Q contains a pair of parallel edges. In particular, this assumption implies that both polygons are "strictly" convex, in the sense that no 3 consecutive vertices are collinear.

The proof of the upper bound relies on a geometric analysis of the set of contact manifolds. Given any two convex polygons P and Q, a fixed vertex $v \in V_P$, and a fixed edge $e \in E_Q$, consider the set of relative configurations satisfying the simple contact (v, e). This is a two-dimensional set, isomorphic to $S_1 \times (0, 1)$, which can be parametrized as follows. Let $q \in V_Q$ be one of the two endpoints of the edge e, and let $\vec{u}$ be any direction fixed relative to P. A configuration c on the contact manifold of (v, e) may now be expressed by a pair $(d, \alpha) \in \Re^+ \times S^1$, where d is the distance from q to v and α is the angle between $\vec{u}$ and the edge e, considered as oriented from q to the other endpoint of e.

The analysis proceeds by considering the form of the intersections of the remaining contact manifolds with the manifold of (v, e). Each such intersection is a one-dimensional set represented by a small number of curve segments in $d - \alpha$ space. Any critical configuration which satisfies the contact (v, e), along with two other contacts, will therefore be represented by the intersection of two such curve segments. In the following sequence of claims, we give an upper bound on the number of such pairwise intersections. From this the overall bound follows by summing over all contacts (v, e).

In each of the following claims, let v, e, d, and α be defined as above.

Claim 1 *The set S of relative configurations satisfying both the simple contact (v, e) and any additional simple contact (of either type) simultaneously is a one-dimensional set consisting of at most 4 connected components.*

Proof: That S is one-dimensional is obvious, since with two contacts satisfied only one degree of freedom remains from the original three.

Now suppose that the second contact is also of type I, i.e. that it has the form (v', e'). Let l be the line segment joining the two vertices of P involved in the contacts, and let c be a relative configuration which satisfies both contacts such that the vertices v and v' lie on the interiors of the respective edges e and e', and let θ be the orientation of l at c. Because we have assumed that Q has no pair of parallel edges, as the orientation of l varies within a sufficiently small neighborhood of θ there exist, by continuity, configurations also satisfying both contacts. Therefore the extremal points of S can only occur in double-contact configurations in which one of the vertices touches one of the endpoints of its corresponding edge. For each of the four possible vertex–endpoint pairs, at most two such configurations may exist (corresponding to the possible intersections of one edge with the circle of diameter length(l) centered at the endpoint of the other edge). Thus there are at most 8 extremal points possible, which implies that S has at most 4 connected components, as claimed.

In the case that the second contact is of type II, i.e. has the form (e', v'), a similar continuity argument shows that again extremal points of S occur only when a vertex touches an edge endpoint; again, this yields at most 8 extremal points. □

Considered as a subset of the contact manifold of (v, e) parametrized by d and α, each connected component of an additional simple contact appears as a continuous curve segment. In addition, we have

Claim 2 *Let c be a simple contact NOT satisfying either:*

- $c = (v', e)$, *or*
- $c = (e', v')$ *and* $v \in \delta(e')$ *and* $v' \in \delta(e)$,

where $\delta(e)$ *denotes the set of two endpoints of the edge e. Then each of the curve segments representing a component of the additional simple contact c may be expressed as a continuous function of the form* $d = d(\alpha)$.

Proof: This is equivalent to the assertion that, for any fixed value $\alpha = \alpha_0$, there exists at most one configuration satisfying both (v, e) and c at that angle. This follows directly from our assumption that the polygons have no parallel edges. □

The exclusion of the two degenerate cases in claim 2 will not affect our final bound. The second case can only yield critical configurations of type B, for which we already have a trivial bound of $O(m^2n^2)$, and a bound of $O(m^3n^2)$ for configurations arising from the first degenerate case is easy to obtain, since both v and v' are restricted to lie on the same edge e.

Claim 3 *Any two contact curve segments may intersect in at most 4 points in* $d - \alpha$ *space.*

Proof: An intersection of two additional curve segments with the manifold of (v, e) represents a critical configuration satisfying the simple contact (v, e) and the two simple contacts corresponding to the curve segments. As mentioned above in the introduction, it is known ([SS83]) that for any fixed triple of contacts at most four configurations exist which satisfy them, which establishes the claim. □

Now consider a set of all the contacts which involve a particular fixed vertex, i.e. a set of the form $\{(v', e') | e' \in E_Q\}$ (type I contacts) or $\{(e', v') | e' \in E_P\}$ (type II contacts). The intersection of each such "vertex–polygon" contact set with the manifold of (v, e) is a set of curve segments in $d - \alpha$ space. We will refer to such a set in the sequel as a *chain*. Because of claim 1, the number of connected components in each chain is at most 4 times the number of simple contacts in the "vertex–polygon" contact set; thus a chain of type I simple contacts has $O(n)$ connected components and a chain of type II simple contacts has $O(m)$.

Our next goal is to establish an upper bound on the number of intersections between any given pair of chains. For this purpose it would be convenient if each chain were also a (not necessarily continuous) function $d = d(\alpha)$. Unfortunately, as can be seen in figure 1, for the same relative orientation of the two polygons it is possible that double contact configurations exist which satisfy both the contact (v, e) and a contact of the chain at two distinct values of d, d_1 and d_2. In these cases, the chain cannot be described as a function. Nevertheless, we have

Claim 4 *For any fixed value α_0 and any particular chain C, the line $\alpha = \alpha_0$ may intersect C in at most 2 points in $d - \alpha$ space.*

Proof: Recall that a chain is composed of a set of segments, each representing a connected component of the contact of the vertex corresponding to the chain with one edge. From claim 2, the number of intersections of any single contact segment with the line $\alpha = \alpha_0$ is at most one; thus we just need to determine the number of segments which the line may cross. If C is a chain of type I segments, let $v' \in V_P$ denote the vertex of P involved in the contacts of C, and let l be the line swept by v' as P moves with the fixed orientation α_0 with v sliding along e (see figure 2). Ignoring possible contacts of v' with vertices of Q (which correspond to type A critical configurations), we can see that l can intersect at most 2 edges of Q, because of our assumption that Q is strictly convex. This establishes the claim for chains of type I segments. Similarly, if C is a chain of type II segments, the strict convexity of P implies that at most two of its segments may cross any line of fixed α. □

It can be seen that, given the condition described in claim 4, any connected component of a chain may consist of at most 3 maximal functional sections. Thus we may artificially split each component into these sections; the splitting points, if any, will be those of minimum and maximum α along the component. Once this is done, the resulting set of connected components of the chain may be partitioned into two sets in such a way that the curves in each set do indeed form a function of α (this also follows from claim 4). To avoid the proliferation of unnecessary terminology, in the sequel we will continue to use the term "chain" to refer to a set of curve segments in one of these partitions of the original chains.

In summary, within the contact manifold of a fixed type I contact (v, e) parametrized by α and d, for each additional vertex of P (respectively Q) we obtain two chains, each of which is a function of α composed of at most $O(n)$ (respectively $O(m)$) connected components. We now define a ***breakpoint*** of a chain as either an extremal point of one of the connected components of the chain or a point at which two simple contact curve segments of the chain meet. Clearly the number of breakpoints of a type I (respectively type II) chain is also in $O(n)$ (respectively $O(m)$).

Claim 5 *Given a chain $C_1(\alpha)$ which has k breakpoints and another chain $C_2(\alpha)$ with l breakpoints, the intersection of C_1, C_2, and the contact manifold of (v, e) is a finite point set of size $O(k + l)$.*

Proof: Let $\{s_i\}, i = 1, \ldots, k$ denote the sequence of values of α corresponding to the breakpoints of C_1, and let $\{t_j\}, j = 1, \ldots, l$ be the sequence of α values of the breakpoints of C_2. Since each chain is a function of α, we may assume that each of these breakpoint sequences is sorted in order of increasing α. (Strictly speaking, this is incorrect, since the range of α is circular. However, by breaking the circle at some arbitrary point, e.g. at $\alpha = 0$, and splitting any chain components which cross that point into two, we can impose a strict ordering among angles and this increases the number of chains by at most two.) By definition, between any two adjacent breakpoints $C_1(s_i)$ and $C_1(s_{i+1})$, the chain C_1 either has a portion of a single simple contact curve segment or is empty. (The same is true of C_2 for its breakpoints.)

Now consider the sorted sequence of values of α obtained by merging the sequences $\{s_i\}$ and $\{t_j\}$. Within the interval defined by any two consecutive values in this merged sequence, the number of intersections between the chains is at most 4 by claim 3. Since the number of such intervals is $k + l - 1$, the result follows. □

We are now ready to establish our

Main Result $K(m,n) \in O(m^2n^3 + m^3n^2)$.

Proof: We derive a bound for each type of critical configuration separately. First, consider the critical configurations consisting of three type I contacts, i.e. in the language of section 2, those of case 0. There are mn contacts of type I. We have seen that for each such contact, the number of functional "chains" representing "vertex–polygon" contacts is linear in the number of such contacts, which is $O(m)$ for the case of additional type I contacts; also, all critical configurations are represented by intersections between contact functions. Furthermore, the number of intersection points between any fixed pair of type I contact functions is $O(n)$, because that is a bound on the number of breakpoints in each and by claim 5. Summing over all type I contacts and over all possible pairs of chains from each, we get a total of $O(mn\binom{m}{2}n) = O(m^3n^2)$ critical configurations of case 0.

Similarly, for case 2 critical configurations we have mn type I contacts, each of which has $O(n)$ chains of type II, and each chain has $O(m)$ breakpoints; this gives $O(mn\binom{n}{2}m) = O(m^2n^3)$ critical configurations of case 2.

Cases 1 and 3 are symmetric to cases 2 and 0, respectively, with the roles of P and Q reversed. Thus we can simply exchange m and n in the above analysis, and this yields the bounds $O(m^3n^2)$ for case 1 critical configurations and $O(m^2n^3)$ for those of case 3. Hence the overall bound for critical configurations of all types is $O(m^2n^3 + m^3n^2)$. □

4 Conclusions

We have presented the best asymptotic bound known to us on the number of ways in which a convex polygon with m vertices may contact another with n vertices in such a way that three independent simple contacts are made between the two. The proof technique relies on considering a two-dimensional subset of relative configurations in which one contact is satisfied, and using simple geometry arguments to derive bounds on the number of ways pairs of additional contacts may intersect simultaneously with such a subset. The key idea lies in claim 4, which is where we exploit the convexity of the polygons.

It is an open question whether the bound presented is tight. The empirical evidence given in section 2 suggests that a better bound should be possible. It is interesting to note that the distance constraints used by the algorithm briefly described in that section are very efficient in rejecting most triples of contacts from further consideration. Accordingly it would be interesting to devise an upper bound whose proof in some way exploits these same constraints. Although we have made some effort in this direction, as of this writing we have not been able to obtain a non-trivial bound in this way.

Another direction for future work is in the area of algorithms. Following the lines of the proof in section 3, one approach would be to explicitly construct the contact chains for each simple contact and to use a sweepline algorithm to find their intersections; this could probably be done in time $O((m^3n^2 + m^2n^3)\log(mn))$. However this is little better than the trivial cubic bound attained by the brute-force algorithm, and it is possible that a better upper bound on the number of critical configurations would yield a better algorithmic approach as well. A more general question is that of algorithms for constructing the set of strata of all dimensions, either for just a pair of convex polygons or for the stratification of the contact space of two convex polyhedra. Again, some kind of sweeping plane approach suggests itself, where the event queue would be the orientations at which contact space vertices (i.e. critical configurations) occur.

References

[Cha83] Bernard Chazelle. The polygon containment problem. *Advances in Computing Research*, 1:1–33, 1983.

[LS87] Daniel Leven and Micha Sharir. On the number of critical free contacts of a convex polygonal object moving in two-dimensional polygonal space. *Discrete and Computational Geometry*, 2:255–270, 1987.

[PT91] David Parsons and Carme Torras. On the complexity of the contact space between convex polyhedra. Technical report, Institut de Cibernètica, Universitat Politècnica de Catalunya–CSIC, Av. Diagonal 647 planta 2, 08028 Barcelona, Spain, 1991.

[SS83] Jacob T. Schwartz and Micha Sharir. On the piano movers' problem: I. the case of a two-dimensional rigid polygonal body moving amidst polygonal barriers. *Communications on Pure and Applied Mathematics*, 36:345–398, 1983.

[Sze74] E. Szemeredi. On a problem by Davenport and Schinzel. *Acta Arithmetica*, 25:213–224, 1974.

[TT89] Carme Torras and Federico Thomas. Motion of two convex polyhedra in contact. Technical Report IC-DT-1989.04, Institut de Cibernètica, Universitat Politècnica de Catalunya–CSIC, Av. Diagonal 647 planta 2, 08028 Barcelona, Spain, July 1989.