

Dilation of regular polygons

Algorithmic aspects (detailed version)

Damien Galant

UMONS - Erasmus Université Paris-Sud

September 24, 2019



Combinatorial explosion

A straightforward way to compute the dilation of regular n -gons is to iterate over all possible triangulations T (up to some symmetries to speed-up the computations).

This was done for instance in Mulzer (2004).

We cannot hope to compute dilations of n -gons even for $n \geq 25$ using this method since the number of triangulations of a n -gon is given by the $(n - 2)$ th Catalan number C_{n-2} , where

$$C_k = \frac{1}{k+1} \cdot \binom{2k}{k}$$

($C_{23} = 343.059.613.650$)

This implies a combinatorial explosion of the number of triangulations as n grows, therefore ruling out purely “bruteforce” approaches to compute the dilation of regular n -gons.

Lower bound: what are we looking for?

We are going to present an algorithm which, given n , returns a *proven* lower bound for the dilation of regular n -gons.
If the found lower bound can be realized by a suitable configuration, then this value will be equal to the dilation of regular n -gons and we are done.

Graphs with cliques

- Given a partial triangulation $P \in \mathcal{P}$, we are interested in *all* triangulations containing P .
- The graph GC_P is obtained by taking all segments between points of S which do not intersect segments of P .
- There is a kind of “duality”: for $T \in \mathcal{T}$, $P \subseteq T \Leftrightarrow T \subseteq GC_P$

Lower bound from a partial configuration

Given a partial triangulation P , a “naive” lower bound on the dilation of *all* triangulations containing P is given by


$$\text{nlb}(P) := \max_{\substack{p, q \in S \\ p \neq q}} \frac{d_{GC_P}(p, q)}{d_{\text{Euclidean}}(p, q)}$$

There is a monotonicity property:

$$P \subseteq P' \Rightarrow \text{nlb}(P) \leq \text{nlb}(P')$$

Futhermore, $\text{nlb}(T) = \text{dil}(T)$ if $T \in \mathcal{T}$ (i.e. the bound is exact for maximal elements of \mathcal{P})

Summary of the “naive” lower bound technique



The lower bound technique

We want to find a better bound, i.e. a value $\text{lb}(P)$ such that

$$\text{nlb}(P) \leq \text{lb}(P) \leq \min_{\substack{T \in \mathcal{T} \\ P \subseteq T}} \text{dil}(T)$$

Since all edges that do not intersect an edge from P will occur in at least one configuration $T \in \mathcal{T}$ such that $P \subseteq T$, we will use the “graph with cliques” GC_P just as for nlb .

The problem can also be seen like this: from the “graph with cliques” GC_P associated to P , find a bound

$$\text{lb}(P) \leq \min_{\substack{T \in \mathcal{T} \\ T \subseteq GC_P}} \text{dil}(T)$$

(this is a kind of “dual problem”).

Indeed, for $T \in \mathcal{T}$ we have $P \subseteq T \Leftrightarrow T \subseteq GC_P$ by “duality”.

Pairs of pairs of points

The dilation of a triangulation $T \in \mathcal{T}$ is given by

$$\max_{\substack{p, q \in S \\ p \neq q}} \frac{d_{\text{Graph of } T}(p, q)}{d_{\text{Euclidean}}(p, q)}$$


The naive lower bound considers a single pair of points $p, q \in S, p \neq q$ and uses the inequality (where $P \in \mathcal{P}, P \subseteq T$)

$$\frac{d_{GC_P}(p, q)}{d_{\text{Euclidean}}(p, q)} \leq \frac{d_{\text{Graph of } T}(p, q)}{d_{\text{Euclidean}}(p, q)}$$

This is not optimal since pairs of points are considered independently. Instead, the new method considers *two pairs of points at once* (this idea was inspired by Dumitrescu and Ghosh (2016)).

Pairs of pairs of points

We begin with a simple observation: if $s_1, s_2, e_1, e_2 \in S$ are distinct points in clockwise order, then the paths from s_1 to e_1 and from s_2 to e_2 must intersect at some point $p \in S$.



Pairs of pairs of points

Since we have no idea of which p is optimal, we take the one which gives the lowest bound for the pair of pairs.

In the end, the bound $\text{lb}(s_1, s_2, e_1, e_2)$ associated to $s_1, s_2, e_1, e_2 \in S$ is

$$\min_{p \in S} \max \left\{ \frac{d_{GC_P}(s_1, p) + d_{GC_P}(p, e_1)}{d_{\text{Euclidean}}(s_1, e_1)}, \frac{d_{GC_P}(s_2, p) + d_{GC_P}(p, e_2)}{d_{\text{Euclidean}}(s_2, e_2)} \right\}$$

We finally obtain

$$\text{lb}(P) = \max_{\substack{s_1, s_2, e_1, e_2 \in S \\ \text{distinct and} \\ \text{in clockwise order}}} \text{lb}(s_1, s_2, e_1, e_2)$$

Global lower bound

The lower bound technique provides a lower bound $\text{lb}(P)$ on the dilation of triangulations which contain P .

Our goal is to find a good (unique) global lower bound glb with

$$\text{glb} \leq \min_{T \in \mathcal{T}} \text{dil}(T)$$

with an inequality as sharp as possible (with equality if possible, since this would mean that glb is the exact value of the dilation).

Our algorithm will take

$$\text{glb} = \min_{P \in \mathcal{C}} \text{lb}(P)$$

where $\mathcal{C} \subseteq \mathcal{P}$ is a set of partial configurations such that

$$\forall T \in \mathcal{T}, \exists P \in \mathcal{C}, P \subseteq T$$

Note that the exhaustive method corresponds exactly to the case $\mathcal{C} = \mathcal{T}$!

Global lower bound: proof of correctness

glb satisfies

$$\text{glb} \leq \min_{T \in \mathcal{T}} \text{dil}(T)$$

Indeed if $T \in \mathcal{T}$ then, by hypothesis on \mathcal{C} , there exists $P_T \in \mathcal{C}$ such that $P_T \subseteq T$. We then have

$$\text{glb} = \min_{P \in \mathcal{C}} \text{lb}(P) \leq \text{lb}(P_T) \leq \text{lb}(T) = \text{dil}(T)$$

Global lower bound: which configurations should we consider?

What remains to be done is to explain how the algorithm chooses \mathcal{C} , the set of partial configurations to be considered.

The key point is to find a good tradeoff between having \mathcal{C} small, with a fast algorithm but a possibly poor bound, and \mathcal{C} large, with a slower method but a better bound (maybe optimal, if \mathcal{C} is large enough)

The search tree

In the end, we have an abstract “search tree” of partial configurations (the lattice of sets ¹ of \mathcal{P} , ordered by inclusion) and for each $P \in \mathcal{P}$, we have a bound $\text{lb}(P)$.

Recall that, if $P_0 \subseteq P_1 \subseteq \dots \subseteq P_n = T \in \mathcal{T}$, then

$$\text{lb}(P_0) \leq \text{lb}(P_1) \leq \dots \leq \text{lb}(P_k) = \text{dil}(T)$$

¹Which will be abusively called the “search tree”, even though one might argue that it is not really a tree.

Pruning the search tree

A very common and efficient technique to deal with optimisation problems on search trees is *pruning*, i.e. stop exploring branches of the tree which do not lead to optimal solutions.

We therefore provide a “target value” to our lower bound algorithm:

Lower bound, with a “target value” c

Given a constant

$$c \geq \min_{T \in \mathcal{T}} \text{dil}(T)$$

return a *proven* lower bound

$$\text{glb} \leq \min_{T \in \mathcal{T}} \text{dil}(T)$$

Typically, we will use $c = \text{dil}(T_{\text{candidate}}) \in \mathcal{T}$, the dilation of a possibly optimal triangulation (or at least a very good triangulation).

What is c useful for?

c is only used for pruning purposes: if a partial triangulation $P \in \mathcal{P}$ is considered and

$$\text{lb}(P) \geq c$$

then we know that

$$\forall T \in \mathcal{T} : \left((P \subseteq T) \Rightarrow c \leq \text{lb}(P) \leq \text{lb}(T) = \text{dil}(T) \right)$$

and we can “cut” that branch of the search tree.

A peculiar feature of the lower bound method is that c , given as input to the lower bound algorithm, *does not change the result returned by the algorithm* (!)

The speed of the proposed method depends *crucially* on the “quality” of c . It is therefore important to find good configurations beforehand.


The hope is that the lower bound algorithm will prove that c is in fact equal to the dilation, i.e.

$$\text{glb} = c = \text{dil}(T_{\text{candidate}})$$

Important edges first

- The last (important!) fact we did not mention is the order in which partial configurations are considered.
- Indeed, it is important to first put some edges that will likely cause $lb(P)$ to be big and hopefully to cut this branch of the search tree.
- In practice, our program puts the *edges of the triangle which contains the center* first.
- It then puts (or tries to put) three smaller triangles on the 3 zones delimited by the central triangle.

Central triangle



A possible central triangle in a 10-gon.

Putting it all together

To summarize, our algorithm works as follows:

Lower bound algorithm

- 1 Take a positive integer n and a “target value” c as input.
- 2 Go through the search tree of partial triangulations, considering important edges first (adding triangles gradually).
- 3 Prune while going through the search tree.
- 4 Stop at a specified depth (or try to do some processing at the leaves of the search tree, which correspond to configurations which were not pruned).
- 5 Return the global lower bound glb .

If $glb = c$, we managed to compute the dilation!

Upper bound: what are we looking for?

As we saw before, we need a good target constant $c = \text{dil}(T_{\text{good}})$ if we want our lower bound algorithm to run fast enough, and we can only conclude if

$$c = \min_{T \in \mathcal{T}} \text{dil}(T)$$

Discussion of such techniques

These techniques have two main advantages:

- The number of considered configurations is polynomial in n , allowing to find bounds even for large values of n .
- Finding the best configuration among the class of considered configurations is doable either with a computer (since the state space has a polynomial size) or even “by hand” due to the specific structure of the considered triangulations.

However, we argue that these approaches also have intrinsic issues:

- Often, there is no formal justification regarding why these classes are considered, but rather heuristic motivations: the class contains the optimal examples for small n , has enough parameters, . . .
- (!) *There is no control on the sharpness of the inequality*

$$\min_{T \in \mathcal{T}} \text{dil}(T) \leq \text{dil}(T_{\text{good}})$$

Discussion of such techniques

- The second issue is really due to the nature of the methods, which consist in living in a (small and better understood) subset

$$S \subseteq \mathcal{T}$$

and forgetting about the rest of \mathcal{T} .

- The whole point of the lower bound algorithm is to respond to the second issue.
- As we want to investigate what happens for $n \geq 25$ and do not have convincing candidates of polynomial-size subclasses of \mathcal{T} which should contain optimal triangulations, we will use *metaheuristics* instead to find good configurations, among which we hope to find optimal ones.

Metaheuristics

- We want to explore the search space \mathcal{T} and find good configurations.
- Finding maxima ² of real functions defined on large discrete search spaces is an intensively studied subject, and generic methods to solve these kinds of optimization problems are often called metaheuristics.
- We used one of the most classical approaches: *hill climbing*.

²Here we want to find a global minimum but we will stick with the common terminology used for maximization problems to explain briefly what are metaheuristics.

Hill climbing

Given “neighbourhood operations” on the search space, the following algorithm can be used:

Hill climbing


- ① Start from some initial state s_0 in the configuration space.
- ② Consider all neighbours of s_0 .
- ③ Go to the neighbour which corresponds to the highest value.
- ④ When all neighbours produce a lower value, stop the algorithm and return the current state and the current value.

From local maxima to candidates of global maxima

- It is clear that the mentioned algorithm only gives rise to local maxima.
- Generally, the problem of finding the global maximum of a function on a search space is hard without further hypotheses.
- In practice, we used a “randomized multistart” strategy: we run the hill climbing algorithm from a lot of triangulations chosen uniformly at random from \mathcal{T} .
- By doing so, we explore different regions of the search space and are less likely to miss the global maxima.

An example of neighbourhood operation

Take a quadrilateral from T and replace its diagonal by the other diagonal.



To avoid being stuck in local minima, we can use this neighbourhood operation with depth 2.

Example of 42-gons

- In less than 10s, the upper bound program gives a value of 1.4222217969.
- The lower bound program then proves that this bound is correct in less than 3s.

Known values for the dilation before our work

n	$\text{dil}(S_n)$	n	$\text{dil}(S_n)$	n	$\text{dil}(S_n)$
4	1.4142	12	1.3836	20	1.4142
5	1.2360	13	1.3912	21	1.4161
6	1.3660	14	1.4053	22	1.4047
7	1.3351	15	1.4089	23	1.4308
8	1.4142	16	1.4092	24	1.4013
9	1.3472	17	1.4084	25	< 1.4296
10	1.3968	18	1.3816	26	< 1.4202
11	1.3770	19	1.4098		

The values of $\text{dil}(S_n)$ for $n = 4, \dots, 26$, from Dumitrescu and Ghosh (2016)

Maxmial dilation of a convex polygon

- Our lower program shows (after approximately 30min) that the dilation of 53-gons is at least 1.4336430827. We do not know the exact value of the dilation of 53-gons however.
- We thereby improve the bound of $\text{dil}(23\text{-gons}) \approx 1.4308$ obtained in Dumitrescu and Ghosh (2016) for the “worst-case dilation of plane spanners”, i.e. the maximal dilation of a set of points:

$$\sup_{\substack{S \subset \mathbb{R}^2 \\ S \text{ finite}}} \text{dil}(S)$$

Further goals

- One of the main goals is to use the hindsight given by small cases to study the asymptotic case, i.e. the dilation of the circle.
- An interesting observation would be to find “small” classes containing optimal configurations, therefore going back (but with more support) to classical upper bound methods.
- Our program allows to have finer information about small configurations: find all good configurations, their symmetries, ...
- We might move towards a “real branch-and-bound” instead of our “two-steps” method.

