On the Limitations of Combinatorial Visibilities

Y. Disser* D. Bilò* M. Mihalák* S. Suri† E. Vicari* P. Widmayer*

Abstract

We consider *combinatorial visibilities* of vertices of a polygon $\mathcal{P}$, a *local* view of $\mathcal{P}$ from its vertices, and ask how much *global* combinatorial information they provide about $\mathcal{P}$. We study three related questions about the connection of visibility graphs and combinatorial visibilities, and show that in no case the combinatorial visibilities carry enough information.

1 Introduction, Model, and Contribution

In microrobotics, there is a strong tendency to deploy lots of cheap and simple microrobots (instead of few complex, expensive ones) to perform robotic tasks such as the exploration of an unknown environment. This gives rise to studies about the sensory and motor capabilities that microrobots need for carrying out such tasks [4, 1]. One particular model [3] assumes that robots (modeled as points) in a polygonal environment have no notion of (or no way of measuring) coordinates or distances. Instead, these robots are limited to seeing vertices of the polygon that are visible from their location when polygon edges are treated as opaque walls, and seeing whether visible vertices are connected with a polygon edge. This concept is usually referred to as "combinatorial visibility". In a convex polygon, for instance, a robot on any of the vertices "sees" all other vertices and edges (and this is not true in any other polygon). We add that in this model the robots can only move from a vertex to any visible vertex. Since this combinatorial visibility appears to be quite powerful [3], there was hope that the combinatorial visibilities of all vertices together might be enough to derive all the combinatorial (i.e., non-geometric) information about the polygon – the "map". The map (also called "visibility graph" [2]) of the polygon is defined as follows: The map (visibility graph) consists of a vertex for each polygon vertex, and has an edge between any two mutually visible vertices. For a convex polygon, for instance, the visibility graph is the complete graph and can easily be inferred from all combinatorial visibilities together.

In this paper, we show that all combinatorial visibilities of a simple polygon are in fact not enough to infer the map (Theorem 1). For certain robotic tasks, such as the meeting of two robots in an unknown polygon, a map might not be needed, for instance if the simple polygon looks nonperiodic to the robot(s). For periodic-looking polygons, there was hope that the vertices seen from a given vertex and those seen from a "periodic partner" of the given vertex (details follow) would themselves be periodic partners; this property would have allowed a variety of tasks to be solved. We show that this is, unfortunately, not the case (Theorem 2). For map computation, the question arises what to add to the abilities of microrobots to make it possible. We show that adding the "local" ability for a robot to measure the polygon angle at its current vertex is not enough (Theorem 3). The common theme behind these results is the relation between combinatorial visibilities and the well-studied visibility graph concept.

In this work we only consider simple polygons. We denote the n vertices of a simple polygon $\mathcal{P}$ by $V=\{v_0, v_1, \cdots, v_{n-1}\}$, ordered along the boundary in counterclockwise (ccw) order. The polygon has n edges $E = \{e_0, e_1, \ldots, e_{n-1}\}$, where $e_i = (v_i, v_{i+1})$, $i = 0, \ldots, n-1$.[1] The *combinatorial visibility* of a vertex v is given by a binary vector who's j-th component encodes whether the j-th visible vertex and the $(j+1)$-th visible vertex form an edge of $\mathcal{P}$ or not, we call this a *combinatorial visibility vector* cvv(v). The following definitions capture this more formally. Consult Fig. 1 along with the definitions.

Definition 1 *Two vertices $v_i, v_j \in V$ form a* visible pair *in $\mathcal{P}$, if the line segment $\overline{v_iv_j}$ lies entirely within $\mathcal{P}$ (in particular, v_i forms a visible pair with itself for any i). We say v_i and v_j* see each other *and write $v_i \leftrightarrow_{\mathcal{P}} v_j$. We drop the index $\mathcal{P}$ and simply write '$\leftrightarrow$', if the corresponding polygon $\mathcal{P}$ is clear from the context.*

Definition 2 *We define* view(v_i) *of vertex v_i in $\mathcal{P}$, the view of vertex v_i, to be the set of vertices that v_i sees in $\mathcal{P}$. Formally,*

$$\text{view}(v_i) := \{\, v_j \in V \,|\, v_i \leftrightarrow v_j \,\}\,.$$

We write $\text{view}_j(v_i)$ *to denote the j-th vertex, $j \geq 0$, that v_i sees in ccw order, starting at v_i itself, both* $\text{view}_0(v_i)$ *and* $\text{view}_{|\text{view}(v_i)|}(v_i)$ *denoting v_i.*

*Institute of Theoretical Computer Science, ETH Zurich, {ydisser|dbilo|mmihalak|vicariel|widmayer}@inf.ethz.ch

†Department of Computer Science, University of California, Santa Barbara, suri@cs.ucsb.edu

[1] We consider all operations on indices modulo n.

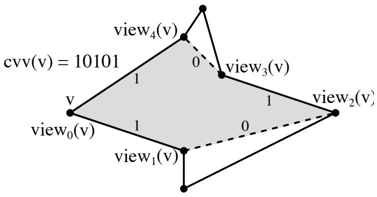

Figure 1: Illustration of a combinatorial visibility vector.

Definition 3 *The* combinatorial visibility vector $\mathrm{cvv}(v_i) \in \{0,1\}^{|\mathrm{view}(v_i)|}$ *of vertex* $v_i \in V$ *is a binary vector with the* j*-th element,* $j \geq 0$*, given by*

$$\mathrm{cvv}_j(v_i) = \begin{cases} 1, & \text{if } (\mathrm{view}_j(v_i), \mathrm{view}_{j+1}(v_i)) \in E, \\ 0, & \text{else.} \end{cases}$$

Note that $\mathrm{view}_1(v_i) = v_{i+1}$ and $\mathrm{view}_{|\mathrm{view}(v_i)|-1}(v_i) = v_{i-1}$ as every vertex sees its neighboring vertices on the polygon boundary. Therefore $\mathrm{cvv}_0(v_i) = \mathrm{cvv}_{|\mathrm{view}(v_i)|-1}(v_i) = 1$ for all $v_i \in V$.

Definition 4 *The* combinatorial visibility sequence cvs *of* $\mathcal{P}$ *lists all combinatorial visibility vectors of the individual vertices of* $\mathcal{P}$ *in ccw order:*

$$\mathrm{cvs} := (\mathrm{cvv}(v_0), \ldots, \mathrm{cvv}(v_{n-1}))\,.$$

In the next section we first show that the cvs of a polygon $\mathcal{P}$ is not enough to determine the map of $\mathcal{P}$. Based on this, we then show a similar result about visibility in "periodical" polygons. Finally, we show that additionally knowing the angles at the vertices of $\mathcal{P}$ does not determine the visibility graph either.

2 Combinatorial Visibility and Visibility Graphs

Theorem 1 *The cvs of a polygon* $\mathcal{P}$ *does not uniquely define its map.*

Proof. In Fig. 2 we present two polygons $\mathcal{P}^A$ and $\mathcal{P}^B$ that share the same cvs, yet they have different maps. The proof is by an inspection of the polygons assisted by a list of the cvv's and view sequences of the needed vertices.

The idea behind the construction of the polygons is to use multiple copies of a "pocket" of vertices (cf. Fig. 2 for an illustration). Each pocket forms a convex curve, but the vertices connecting the pockets form reflex angles, resulting in a non-convex polygon $\mathcal{P}$. The vertices inside a pocket thus do not see all vertices of $\mathcal{P}$, they see (apart from their own pocket) only parts of exactly two pockets. We use the fact that the vertices have no way to distinguish what pockets they are "looking into" and we modify the polygon $\mathcal{P}^A$ by shifting the vertex c (cf. Fig. 2) so that in $\mathcal{P}^B$ the shifted vertex $\tilde{c}$ looks into different pockets, while not changing the cvv of any vertex. □

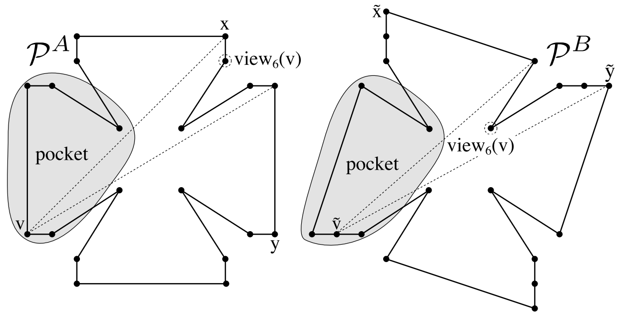

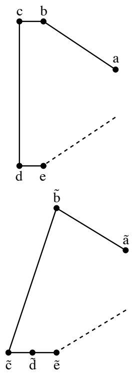

vert	cvv	vision sequence
a	1111101111011111	$abcdeadeabcabcde$
$\tilde{a}$		$\tilde{a}\tilde{b}\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{b}\tilde{a}\tilde{b}\tilde{c}\tilde{d}\tilde{e}$
b	11110111101	$bcdeadeabca$
$\tilde{b}$		$\tilde{b}\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{b}\tilde{a}$
c	11101111011	$cdeadeabcab$
$\tilde{c}$		$\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{b}\tilde{a}\tilde{b}$
d	11011110111	$deadeabcabc$
$\tilde{d}$		$\tilde{d}\tilde{e}\tilde{a}\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{b}\tilde{a}\tilde{b}\tilde{c}$
e	10111101111	$eadeabcabcd$
$\tilde{e}$		$\tilde{e}\tilde{a}\tilde{c}\tilde{d}\tilde{e}\tilde{a}\tilde{b}\tilde{a}\tilde{b}\tilde{c}\tilde{d}$

Figure 2: Top: Two polygons P^A and P^B with $n = 20$, identical cvs and different visibility graphs. Bottom: cvv and vision sequence of every vertex within a pocket.

Definition 5 *We say that a* cvs C *is* periodical *with a period* $p \geq 2$*, if* $C_i = C_{i+k\cdot\frac{n}{p}}$ *for all* $0 \leq i < n$ *and for all* $1 \leq k < p$*. For each* $0 \leq i < n$ *we say* $\left\{ v_{i+k\cdot\frac{n}{p}} \,\middle|\, 0 \leq k < p \right\}$ *are* periodical partners.

Theorem 1 shows that the cvs is not enough to determine the visibility graph of a polygon. In the following we are interested in a similar question:[2] Assume vertex v_i sees a vertex v_{i+j} as its k-th visible vertex (in ccw order), does vertex $v_{(i+n/2)}$ see vertex $v_{(i+n/2)+j}$ as its k-th visible vertex? The following theorem implies that this is not the case.

Theorem 2 *There is a polygon* $\mathcal{P}$ *with a periodical* cvs *of period* $p \geq 2$ *for which we have*

$$\exists v_i \in V\ \exists j\ :\ \mathrm{cvv}(\mathrm{view}_j(v_i)) \neq \mathrm{cvv}\Big(\mathrm{view}_j\Big(v_{i+\frac{n}{p}}\Big)\Big)\,.$$

Proof. We construct a polygon $\mathcal{P}$ with the aforementioned property from the two polygons $\mathcal{P}^A$ and $\mathcal{P}^B$ in Fig. 2. The polygon will have period $p = 2$. At the end of the proof we show how to generalize the construction to arbitrary periods.

The idea of the construction is to "glue" $\mathcal{P}^A$ and $\mathcal{P}^B$ together at vertices v and $\tilde{v}$ of $\mathcal{P}^A$ and $\mathcal{P}^B$, respectively, where v and $\tilde{v}$ are as depicted in Fig. 2. We

[2] A positive answer to this question would have an impact on various interesting problems in the field of simple robots; for example, on the weak version of the *rendezvous* problem in symmetric polygons.

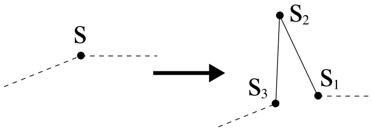

Figure 3: The concept of inserting spikes at vertices.

want to glue the polygons such that every two corresponding vertices w and $\tilde{w}$ of the two polygons form periodical partners in $\mathcal{P}$. Thus, we need to glue the polygons such that the cvv's of corresponding vertices w and $\tilde{w}$ are the same. We can then use the result of Theorem 1 which guarantees the existence of vertices w and $\tilde{w}$ with the same cvv but different views. Formally, if w from $\mathcal{P}^A$ is a vertex v_i in $\mathcal{P}$ and $\tilde{w}$ from $\mathcal{P}^B$ is a vertex $v_{i+n/2}$ in $\mathcal{P}$ (where n is the number of vertices of $\mathcal{P}$), there is a position j in their visions such that $\text{view}_j(v_i) = v_k$ and $\text{view}_j(v_{i+\frac{n}{2}}) = v_l \neq v_{k+\frac{n}{2}}$. Because of the structure of the two polygons, we will see that $\text{cvv}(v_k) \neq \text{cvv}(v_l)$ which proves the theorem.

We glue both polygons together at a vertex $v \in V_{\mathcal{P}^A}$ and the corresponding vertex $\tilde{v} \in V_{\mathcal{P}^B}$ with the same cvv by splitting these vertices into v, v' and $\tilde{v}, \tilde{v}'$ respectively and merging v with $\tilde{v}'$ and v' with $\tilde{v}$.[3] Because we perform the splitting on both v and $\tilde{v}$, the altered cvv's of the new vertices will still be pairwise equal. Similarly we do not change the cvv of vertices which did not see the split vertices. The problem however is the change in the cvv's of all vertices that see v or $\tilde{v}$ in $\mathcal{P}^A$ or $\mathcal{P}^B$ respectively. Such a vertex now, instead of seeing the original vertex, sees two unconnected vertices (zero in the cvv). Thus, vertices would be able to distinguish between split and non-split vertices in their fields of vision and therefore they would be able to distinguish into which pocket they are "looking".

In order to maintain equal cvs' in both polygons we split all vertices in a similar way. Fig. 3 shows how to split vertices by inserting *spikes*. A spike substitutes a vertex s by three new vertices s_1, s_2, and s_3 as illustrated. It is important that the new vertex s_2 at the tip of the spike does not see any other vertices except for s_1 and s_3. That way, other vertices will have the impression to see two vertices s_1 and s_3 that are not connected by an edge. This is exactly how the spikes resulting from splitting v and $\tilde{v}$ look from outside, so that they are indistinguishable from other spikes. Fig. 4 shows $\mathcal{P}^A$ and $\mathcal{P}^B$ after inserting spikes.[4] Fig. 5 gives a listing of the cvv of each vertex within a pocket. For the sake of brevity we do not give the visibility graph explicitly.

[3] We place the newly created vertices very close together so that if vertices v and x saw each other originally, then after the splitting x sees both v and v'.

[4] Note that we do not assume general position for our construction. However it is easy to modify the spiked versions of our polygons accordingly.

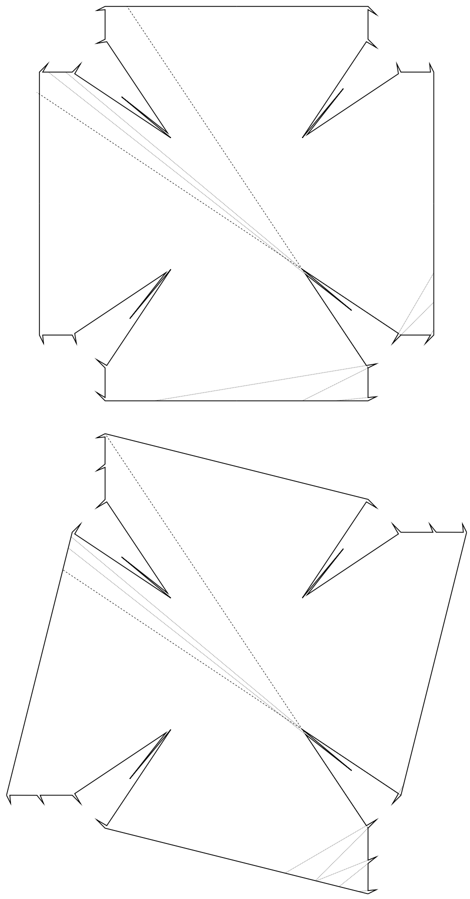

Figure 4: The two polygons from Fig. 2 equipped with spikes and still with identical cvs. The areas visible from different tips are indicated.

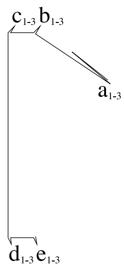

vert	cvv
a_1	11001010101001010101 01
a_2	101
a_3	101010101000101010100101010101 11
b_1	11101010100010101010 01
b_2	101
b_3	10101010001010101001 11
c_1	11101010001010101001 01
c_2	101
c_3	10101000101010100101 11
d_1	11101000101010100101 01
d_2	101
d_3	10100010101010010101 11
e_1	11100010101010010101 01
e_2	101
e_3	10001010101001010101 11

Figure 5: The combinatorial visibilities of each vertex in a pocket of $\mathcal{P}^A$ after adding spikes (the same cvv's arise for $\mathcal{P}^B$). We write v_{1-3} to denote the group of vertices v_1, v_2, v_3.

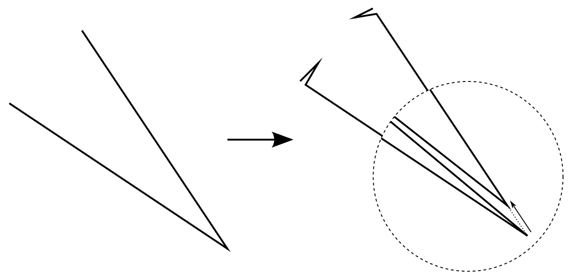

Figure 6: Illustration of how the spikes are inserted at reflex vertices. We chose our modification such that the right neighbor of the spike tip retains the visibility of the original vertex.

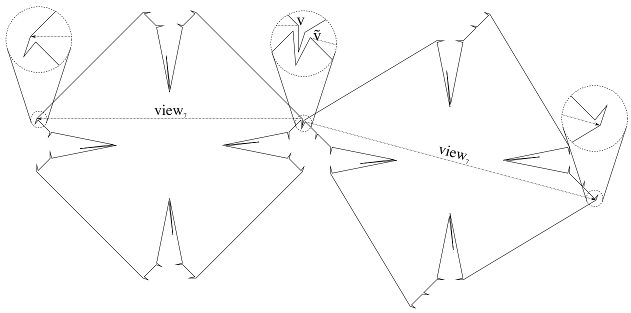

Figure 7: A polygon with $n = 120$ that proves Theorem 2.

For each convex vertex s we are able to insert a spike such that the two new vertices s_1 and s_3 that lie at each side of s_2 have the same vision as s had (apart from seeing s_2 and seeing "gaps" instead of single vertices everywhere else). We do this by splitting the convex vertex along its tangential direction. It is always possible to position the spike tip such that it is seen by its two neighbors only. Because we are able to insert the spike without qualitatively changing the vision of s_1 and s_3, it is obvious that the cvv's of all convex vertices remain equal in $\mathcal{P}^A$ and $\mathcal{P}^B$. For reflex vertices r, however, we cannot introduce a spike such that r_1 and r_3 share the same vision apart from seeing r_2 on different sides. In general we need to split reflex vertices "manually" and make sure that the resulting cvs' of both polygons are equal. Fig. 6 shows how to do this with the four reflex vertices in our case.

Once we have spiked versions of $\mathcal{P}^A$ and $\mathcal{P}^B$, we can glue them together in a straightforward way by simply splitting the spike tip of v and $\tilde{v}$ and attaching the open ends. This modification only affects the cvv of $v_1, \tilde{v}_3$ and $v_3, \tilde{v}_1$ in a similar manner as well as the cvv of the duplicated spike tips $v_2, \tilde{v}_2$ each of which are periodical partners in $\mathcal{P}$. The resulting cvs is still periodical with period 2. Fig. 7 shows the resulting polygon $\mathcal{P}$.

The extension to $p > 2$ is easily made, as we can attach more than two copies of the two spiked polygons around a common center. □

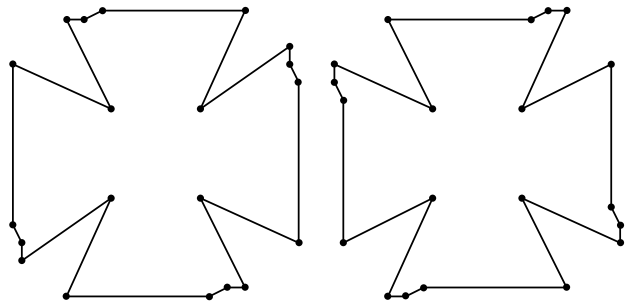

Figure 8: Two polygons with identical cvs and identical interior angles but different G_V. The visibilities are similar to those of $\mathcal{P}^A$ and $\mathcal{P}^B$.

Theorem 1 shows that the cvs is not sufficient to reconstruct the visibility graph of a polygon. A natural question is to determine a "minimal" structural information we need to add to the cvs for reconstructing the visibility graph. In the following we show that knowing the precise inner angle of every vertex of a polygon in addition to its cvs is still not enough for reconstructing G_V. We show the following theorem.

Theorem 3 *The* cvs *and precise inner angles of a polygon $\mathcal{P}$ do not uniquely determine the map of $\mathcal{P}$.*

Proof. Fig. 8 shows a modified version of the polygons $\mathcal{P}^A$ and $\mathcal{P}^B$ of Fig. 2. As one can easily check, these polygons still have the same cvs and different visibility graphs. In addition, they also have the same set of inner angles at the vertices. The existence of such polygons proves the theorem. □

References

[1] A. Ganguli, J. Cortés, and F. Bullo. Distributed deployment of asynchronous guards in art galleries. In *Proceedings of the American Control Conference*, pages 1416–1421, June 2006.

[2] S. K. Ghosh. *Visibility Algorithms in the Plane.* Cambridge University Press, 1st edition, April 2007.

[3] S. Suri, E. Vicari, and P. Widmayer. Simple robots with minimal sensing: From local visibility to global geometry. *International Journal of Robotics Research*, 27(9):1055–1067, September 2008.

[4] A. Yershova, B. Tovar, R. Ghrist, and S. LaValle. Bitbots: Simple robots solving complex tasks. In *Proceedings of the Twentieth National Conference on Artificial Intelligence and the Seventeenth Innovative Applications of Artificial Intelligence Conference*, pages 1336–1341, 2005.