

Notes on the Complexity of Exact View Graph Algorithms for Piecewise Smooth Algebraic Surfaces*

J. H. Rieger

ICMSC, Universidade de São Paulo,
 Caixa Postal 668, 13560-970 Sao Carlos – SP, Brazil
 rieger@icmsc.sc.usp.br

Abstract. The view graph of a surface N in 3-space is a graph embedded in the space \mathcal{V} of centers or directions of projection, whose nodes correspond to maximal connected regions of \mathcal{V} which yield equivalent views of N . The size of the view graph of a piecewise smooth algebraic surface N with transverse self-intersection curves and isolated triple-points and cross-caps is $O(n^{K \dim \mathcal{V}} d^{6 \dim \mathcal{V}})$, where n and d denote the number of “component surfaces” of N and their maximal degree, respectively, and where $K = 6$ in general or $K = 3$ for N diffeomorphic to the boundary of a polyhedron. (For surfaces without cross-caps, this bound has been established in [17].) Also, for the special piecewise linear case, where $d = 1$ and $K = 3$, it is known that the size of the view graph is actually $\Theta(n^{3 \dim \mathcal{V}})$.

It is shown that the exact view graphs of such surfaces can be determined in $O(n^{K(2 \dim \mathcal{V} + 1)} \cdot \mathcal{P}(d, L))$ time by a deterministic algorithm and in $O(n^{K \dim \mathcal{V} + \varepsilon} \cdot \mathcal{P}(d, L))$ expected time by a randomized algorithm. Here \mathcal{P} is some polynomial, L is the maximal coefficient size of the defining polynomials of N , and ε is an arbitrarily small positive constant. Note that the randomized algorithm is, in terms of combinatorial complexity (where d and L are assumed to be constants which do not depend on n), nearly optimal—its combinatorial time complexity exceeds the size of the view graph only by ε in the exponent.

1. Introduction

The computation and complexity of view (or aspect) graphs of surfaces in 3-space has been the subject of a fair number of works in the fields of computational geometry and computer vision. The survey article by Bowyer and Dyer [2] gives a good overview of the works on polyhedral surfaces up to 1990 from a computer-vision point of view. Related works in the fields of combinatorics and computational geometry on visibility

* Financial support for this work was given by the research foundation FAPESP (São Paulo, Brazil).

problems and the number of topologically distinct views of “polyhedral terrains” (i.e., special polyhedral surfaces given as function graphs) are summarized in the recent book by Sharir and Agarwal [18] (see also the discussion at the end of this paper). The recent survey [15] concentrates on view graphs of smooth and piecewise smooth algebraic surfaces (the latter include both smooth as well as polyhedral surfaces as special cases); most (algorithmic) works on this subject have appeared after 1990.

In [17] we have studied the complexity of view (or aspect) graphs of piecewise smooth algebraic surfaces M having transverse self-intersection curves (“crease curves”) and isolated triple points and of subsets $M' \subset M$ representing the boundaries of semialgebraic solids. In this work we have also presented exact symbolic algorithms for computing the view graphs of M and M' and have illustrated the results produced by an implementation of these algorithms for simple example surfaces M and $M' \subset M$.

Much remains to be done in terms of analyzing and improving the running times of these algorithms. The algorithm in [17] for the surfaces M has a polynomial running time, but we did not attempt to give a specific bound for the degree of this polynomial (which would have been far from optimal). For bounding surfaces $M' \subset M$ of solids the worst-case running time of the algorithm in [17] is exponential in the number of unions appearing in the defining formula of the solid.

Meanwhile there have been two new developments in the fields of computational geometry and classification of singularities which form the basis for the algorithms described in the present paper. These algorithms are more efficient and (slightly) more general than the ones in [17]. First (and relevant for the efficiency aspect) there is the work by Chazelle et al. [4] and by Sharir and Agarwal [18] in the field of computational geometry and combinatorics which yields algorithms for stratifying semialgebraic sets which are much more efficient, in terms of combinatorial complexity, than previous algorithms. Given n polynomials in D variables, the algorithms in [4] and [18] determine a decomposition of \mathbb{R}^D into $O(n^{2D-3+\epsilon})$ sign-invariant connected cells of “simple shape” for which point-location queries can be answered in $O(\log n)$ time. This decomposition can be determined by a deterministic algorithm in $O(n^{2D+1})$ time and by a randomized algorithm in $O(n^{2D-3+\epsilon})$ (for $D \geq 3$) or $O(n^{2+\epsilon})$ (for $D = 2$) expected time (see Theorems 8.21 and 8.23 in [18] and the introduction in [4]). Second, there is the recent classification by Carter et al. [3] of sequences of projections $\mathbb{R}^4 \rightarrow \mathbb{R}^3 \rightarrow \mathbb{R}^2 \rightarrow \mathbb{R}$ of knotted 2-surfaces in 4-space which, as a by-product, yields the codimension 1 views of singular surfaces in 3-space having not only transverse intersection curves and triple-points but also cross-caps.

These results, together with the rather detailed a priori knowledge about the topology of the view bifurcation set from [17], enable us to obtain—in a fairly straightforward way—exact view graph algorithms for surfaces M and M' (having additional cross-caps) which have a near-optimal combinatorial time complexity.

1.1. *Singular Surfaces and Bounding Surfaces of Solids*

In this paper we are interested in “visibility problems” for the following two classes of curved surfaces:

- (i) Singular algebraic surfaces $M = \bigcup_{i=1}^n M_i$, where M and all M_i have curves of transverse self-intersections and isolated triple-points and cross-caps.

(ii) Bounding surfaces M' of closed semialgebraic solids

$$S := \bigcup_{u=1}^U \bigcap_{i=1}^{I_u} \{p \in \mathbb{R}^3 : h_{u,i}(p) \geq 0\},$$

such that $M = \bigcup_u \bigcup_i h_{u,i}^{-1}(0)$ is an algebraic surface with the kind of singular points described in (i). Note that any M' is a semialgebraic subset of some M .

Note, of course, that piecewise linear surfaces are special cases of the surfaces M and M' above. In terms of space and time complexity of the algorithms below it is, however, convenient to single out a special subclass of the surfaces M' for which sharper complexity bounds can be obtained. Namely, the class (ii)' of bounding surfaces M' of semialgebraic solids which can be deformed—by a nonlinear coordinate change in \mathbb{R}^3 —into the bounding surface of a polyhedron. The surfaces in this subclass have $O(m)$ self-intersection curves, m being the number of faces, and no cross-caps—on the other hand, a general surface M' has $O(m^2)$ self-intersection curves and also has isolated cross-caps.

Also note that, for zero-sets, one expects in the *generic case* that the component surfaces $h_i^{-1}(0)$ (resp. $h_{u,i}^{-1}(0)$) are *nonsingular* and intersect transversely along intersection curves and in isolated triple-points (this is an easy consequence of the Bertini–Sard theorem). Furthermore, such surfaces—without cross-caps—are typically the ones that are relevant in geometric and solid modeling. By contrast, surfaces given as images of *generic mappings* from 2-space into 3-space typically have transverse self-intersection curves, triple-points as well as cross-caps (this is a classical result of Whitney [21]). This more general class of surfaces is considered in this paper.

1.2. Combinatorial and Algebraic Complexity

The *combinatorial complexity* of the surface M is determined by the number n of component surfaces M_i . The combinatorial complexity of M' does a priori depend on the number of unions U and the numbers of intersections I_1, \dots, I_U required to define the solid S . However, the total number $n := \sum_{u=1}^U I_u$ of “basic semialgebraic sets” $\{p \in \mathbb{R}^3 : h_{u,i}(p) \geq 0\}$ will be, for the purpose of this paper, a sufficiently fine measure of the combinatorial complexity of S and its boundary M' .

The *algebraic complexity* of the surfaces M and M' is given by the maximal degree d of the component surfaces M_i and $h_{u,i}^{-1}(0)$, respectively.

The time complexity of the algorithms described below depends on n and d as well as on the coefficient sizes of the defining polynomials h_i of the surfaces. So we set $L = \sup_i |h_i|_1$. Note that a polynomial-time algorithm in (n, d, L) also has a polynomial running time in the *bit-complexity* model.

Finally, the following two (symbolic) constants will appear in the exponents of the complexity estimates below. First, the dimension of the “viewing space” \mathcal{V} which is two in the case of parallel and three in the case of central projection. Second, for any surface M or M' having only $O(n)$ self-intersection curves we set $K = 3$, for general surfaces (having $O(n^2)$ self-intersection curves) we set $K = 6$. Note that, for surfaces

M' which are diffeomorphic to the boundaries of polyhedra, we have that $K = 3$ (recall the discussion in Section 1.1 above).

1.3. Classifications of Views and the Definition of View Graphs

Let N be either one of the surfaces M or M' , let $S(N)$ denote the singular set of N , and let $\mathcal{V} = \mathbb{P}^2$ (resp. $\mathbb{R}^3 \setminus N$) be the “view space” of all directions (resp. centers) of projection. Given $q \in \mathbb{R}^3$ and $\omega \in \mathcal{V}$, parallel (resp. central) projection maps q to the line through q parallel to ω (resp. to the line joining q and ω). Let $p_\omega: N \rightarrow \mathbb{P}^2$ denote the restriction to N of the parallel (resp. central) projection from \mathbb{R}^3 into the “retinal plane” \mathbb{P}^2 along the direction (resp. from the center) $\omega \in \mathcal{V}$. Define a *view* of N from ω as follows:

$$v_\omega(N) := p_\omega(S(N)) \cup p_\omega(\Sigma_\omega(N)),$$

where Σ_ω denotes the set of critical points of the map p_ω (in which dp_ω fails to have full rank). From “almost all” directions $v_\omega(N)$ is a curve in the retinal plane which is the union of the projected self-intersection curves of N and of the apparent contours of the “faces” of N .

A pair of views of N is *equivalent* if one of them is mapped onto the other by a diffeomorphism of the retinal plane. The actual classifications of views of N use a slightly different equivalence relation: a pair of projections of $N \subset \mathbb{R}^3$ is equivalent if there exist diffeomorphisms of \mathbb{R}^3 , preserving N , and of the retinal plane mapping one projection onto the other. It is known that equivalence of projections implies equivalence of views, but the converse is only true for equivalence classes (of projections and views) of low codimension. In this paper we merely have to know all possible codimension ≤ 1 views (the stable and “minimally unstable” views), and for each of them there is exactly one corresponding type of projection of the same codimension. Views of surfaces N with transverse self-intersection curves and isolated triple points can have six types of *stable* isolated singular points, corresponding to the codimension 0 orbits under the equivalence relation above, and nineteen types of *unstable* codimension 1 singularities, see [17]. For surfaces with additional cross-caps, we get one additional type of stable singular point (a fold of the cross-cap) and three additional codimension 1 singularities, see [3].

The *view bifurcation set* $\mathcal{B} \subset \mathcal{V}$ of the family of all parallel or central projections of N consists of all directions or centers of projection $\omega \in \mathcal{V}$ which yield an unstable view (that is, a view containing at least one singularity of codimension ≥ 1). It has been shown in [17] that the view bifurcation set of any surface with transverse self-intersection curves and isolated triple points is a subset of view space of measure zero—however, surfaces with quadruple points have region-filling view bifurcation sets.

Assuming that the view bifurcation set $\mathcal{B} \subset \mathcal{V}$ of a surface N has measure zero, we can define the *view graph* $G(N)$ of N as follows: $G(N) = (V, E)$ is a graph embedded in \mathcal{V} whose set of vertices V are the connected regions of $\mathcal{V} \setminus \mathcal{B}$ and whose set of edges E are the branches of \mathcal{B} of dimension $\dim \mathcal{V} - 1$ separating adjacent connected regions. Note that all views obtained from within a single connected region of $\mathcal{V} \setminus \mathcal{B}$ are related by a diffeomorphism of the retinal plane. Traversing an edge of the view graph $G(N)$ corresponds to a “catastrophic change” in the view of N .

1.4. Results and Organization of Paper

Section 2 contains some piecewise smooth algebraic example surfaces and their view graphs. (The examples have been computed with a slightly modified implementation of the algorithms in [17]. The modifications take care of cross-caps and other singularities on the component surfaces of M . The more substantial modifications described in Section 4 below have not yet been implemented.)

In Section 3 we present a table of necessary and sufficient conditions for a direction or center of projection ω to lie in the view bifurcation set \mathcal{B} (or \mathcal{B}') of M or M' . In other words, the conditions in this table “recognize” all degenerate views of M (or M'). The conditions are derived from the recent classification of codimension 1 views of such surfaces in [3].

In Section 4 we study the time complexity of algorithms for determining the exact view graphs $G(\cdot)$ of piecewise smooth algebraic surfaces M and of bounding surfaces M' of semialgebraic solids. We show that both $G(M)$ and $G(M')$ can be determined in $O(n^{K(2\dim \mathcal{V}+1)} \cdot \mathcal{P}(d, L))$ time by a deterministic algorithm or in $O(n^{K \dim \mathcal{V} + \varepsilon} \cdot \mathcal{P}(d, L))$ expected time by a randomized algorithm. The combinatorial complexity of the randomized algorithm is only by an arbitrarily small positive constant ε greater than the size of the view graph $G(\cdot)$ which is of order $O(n^{K \dim \mathcal{V}} d^{6 \dim \mathcal{V}})$. It is known that, in the special case of polyhedra where $d = 1$ and $K = 3$, the size of the view graph is actually $\Theta(n^{3 \dim \mathcal{V}})$, see [16]. So, from a combinatorial point of view, the randomized algorithm is nearly optimal.

The overall combinatorial structure of the algorithm for M is the same as in [17], except that, in one substep, the (by now classical) cylindrical algebraic decomposition algorithm by Collins is replaced by the semicylindrical decomposition algorithm described in [4]. Additional modifications of the algebraic part of the algorithm are also necessary, due to the presence of singular component surfaces of M .

For bounding surfaces M' of semialgebraic solids we describe a new polynomial-time algorithm—the algorithm for M' in [17], by contrast, has a running time which is exponential in the number of unions U in the defining formula of the solid.

In Section 5 we present bounds for the degree of the view bifurcation set and for the number of nodes in the view graph of surfaces having cross-caps, as well as curves of transverse self-intersections and isolated triple-points. The complexity of view graphs of surfaces without such cross-caps have been studied previously by Rieger [17] and Petitjean [14]. Using the necessary and sufficient conditions for unstable views presented in Section 3 we show that the asymptotic complexity of the size of view graphs of surfaces with cross-caps is the same as that of piecewise smooth surfaces without cross-caps, namely $O(n^{K \dim \mathcal{V}} d^{6 \dim \mathcal{V}})$.

Finally, in Section 6, we discuss some recent works on visibility problems and view graphs, relate them to the results in the present paper and mention some open problems.

2. A Few Examples

Perhaps it is best to begin with some examples of piecewise smooth algebraic surfaces and their view graphs (for parallel projection). For parallel projection, the view bifurcation

set $\mathcal{B} \subset \mathbb{P}^2$ of a surface N is the dual of the view graph $G(N)$. If $(a : b : c)$ are homogeneous coordinates in \mathbb{P}^2 , then the computation of $G(N)$ can be carried out in the affine chart $(1, b, c)$ with the understanding that pairs of antipodal regions r_+, r_- , i.e., nodes of $G(N)$, are to be identified, unless their closures contain components C of \mathcal{B} at infinity. In this case the nodes r_+ and r_- are either connected by an edge, if C is a subarc of the line at infinity $a = 0$, or not if $\dim C = 0$. All three example surfaces below have certain symmetries which, in turn, induce symmetries of \mathcal{B} and $G(N)$ in the plane (b, c) . Below, we use the notational convention that regions in the complement of \mathcal{B} , i.e., nodes of $G(N)$, which are related by some reflection in the (b, c) -plane are denoted by the same numbers.

Example 1. The zero-set M of $h = x^2y^2 + x^2z^2 + y^2z^2 - xyz$ is known as Steiner's Roman surface (see Chapter VII of [10] and Chapter 6, Section 46 of [9]). The surface M has six cross-caps at $(\pm\frac{1}{2}, 0, 0)$, $(0, \pm\frac{1}{2}, 0)$, $(0, 0, \pm\frac{1}{2})$. There are three lines of double-points (the coordinate axes) connecting pairs of cross-caps and intersecting at the origin in a triple-point. Strictly speaking, Steiner's Roman surface is the semialgebraic surface (having the same view graph as M) given by

$$M \setminus (\{x^2 > \frac{1}{4}, y = z = 0\} \cup \{y^2 > \frac{1}{4}, x = z = 0\} \cup \{z > \frac{1}{4}, x = y = 0\}),$$

i.e., only the intervals $[-\frac{1}{2}, \frac{1}{2}]$ of the coordinate axes belong to the surface. See Fig. 1 for a picture of Steiner's surface. The view graph $G(M)$ of M , shown in Fig. 2, has 88 nodes. For Steiner's surface \mathcal{B} and $G(M)$ are point-symmetric about the origin. Figure 3 shows a view of M associated to node number 36 (for space reasons, we do not show the views associated with the other nodes of $G(M)$).




Fig. 1. Steiner's Roman surface.




Fig. 2. The parallel projection view graph of Steiner's Roman surface: the dashed circle indicates the line at infinity $a = 0$ and edges cutting $a = 0$ connect equally numbered antipodal nodes.

Example 2. The boundary M' of the union of two solid half-spheres defined by

$$(\{2 - (x + 1)^2 - y^2 - z^2 \geq 0\} \cup \{2 - (x - 1)^2 - y^2 - z^2 \geq 0\}) \cap \{z \geq 0\}$$

has two triple points lying on three circles of double-points (see Fig. 4). The view graph $G(M')$ of M' has 26 nodes and is equal to the view graph of the surface $M \supset M'$ which is given by the zero-set of

$$(2 - (x + 1)^2 - y^2 - z^2)(2 - (x - 1)^2 - y^2 - z^2)z.$$

These view graphs are symmetric under reflections in the lines $b = 0$ and $c = 0$ and are shown in Fig. 5. Figure 6 shows the views of M and M' associated to node number 1 in the view graph.




Fig. 3. A view of Steiner's Roman surface from node number 36.




Fig. 4. The union of two solid half-spheres.




Fig. 5. The parallel projection view graph of the bounding surface of two solid half-spheres: the dashed circle indicates the line at infinity $a = 0$ and edges cutting $a = 0$ connect equally numbered antipodal nodes.




Fig. 6. A view of the bounding surface M' of two solid half-spheres from node number 1: crossed-out curves belong to view of $M \setminus M'$.

Remark. The above example is somewhat untypical. In general, the graph $G(M')$ is smaller than the graph $G(M)$, for $M' \subset M$ (see Fig. 2 of [17] for a more typical example).

3. Defining Conditions of the View Bifurcation Set

The view bifurcation set of a surface consists of all directions or centers of projection giving rise to some degenerate view of this surface. The classification of projection-maps, for singular surfaces with transverse self-intersection curves (“crease curves”) and isolated triple-points, and cross-caps, yields 7 equivalence classes of isolated stable singularities (of codimension 0) and 22 minimally unstable singularities of codimension 1. In this classification, two projection-maps are considered to be equivalent if one is mapped onto the other by a diffeomorphism of \mathbb{R}^3 preserving the projected surface and a diffeomorphism of the retinal plane. The “view types” in column 2 of Table 1 refer to the normal forms of the 22 equivalence classes of codimension 1: the view types for $j = 1$ to 19 constitute a complete list of codimension 1 views for surfaces without cross-caps and can be found in Proposition 4.1 of [17]. The three remaining view types ($j = 20, 21, 22$) involve cross-caps: $j = 20$ corresponds to the second normal form

Table 1. Minimally unstable view types.

$j =$	View type (“name”)	Intersection multiplicity	Additional conditions
1	8 (lip/beak)	([3])	$D_{\text{reg}} = 0 *$
2	7 (swallowtail)	([4])	—
3	5+5+5 (fold triple-crossing)	([2], [2], [2])	—
4	6+5 (cusp fold crossing)	([3], [2])	—
5	5++5 (tacnodal fold crossing)	([2], [2])	$N^{\parallel} = 0$
6	III ₂ (semilip/beak)	([2, 1])	$C\Sigma^{\parallel} = 0 *$
7	IV (semicusp)	([3, 1])	—
8	VII ₁ (crease-cusp)	([2, 2])	—
9	1+1+1 (crease triple-crossing)	([1, 1], [1, 1], [1, 1])	—
10	1++1 (tacnodal crease crossing)	([1, 1], [1, 1])	$\pi_C^{\parallel} = 0$
11	6+1 (cusp crease crossing)	([3], [1, 1])	—
12	5++1 (tacnodal fold crease crossing)	([2], [1, 1])	$\pi_{\Sigma C}^{\parallel} = 0$
13	5+5+1 (fold fold crease crossing)	([2], [2], [1, 1])	—
14	5+1+1 (fold crease crease crossing)	([2], [1, 1], [1, 1])	—
15	5+II (fold semifold crossing)	([2], [2, 1])	—
16	II+1 (semifold crease crossing)	([2, 1], [1, 1])	—
17	S(5+Y) (fold vertex crossing)	([2], [1, 1, 1])	—
18	S(1+Y) (crease vertex crossing)	([1, 1], [1, 1, 1])	—
19	S(Y _E) (semifold-vertex)	([2, 1, 1])	—
20	W _{$\Sigma^{1,1}$} (cross-cap-cusp)	([3])	W = 0
21	W+5 (cross-cap fold crossing)	([2], [2])	W = 0
22	W+1 (cross-cap crease crossing)	([2], [1, 1])	W = 0

in Theorem 8.6.1 in the thesis of West [20], and the normal forms for $j = 21$ and 22 are given in Section 4.1(ii) of [3]. We do not list these normal forms here (they are not very enlightening and require some extra notation for multigerms of maps) and refer the reader to the original classifications, but in Section 3.1 we discuss the geometry of the additional view types involving cross-caps.

The unstable views can be “recognized” by certain conditions that involve the intersection type of a ray of projection with a surface at a set of points and, in some cases, the “local geometry” of the surface at these intersection points. More precisely, a view contains one of the 22 codimension 1 singularities (or some codimension ≥ 2 singularity in their closure) if and only if at least one pair of conditions in columns 3 and 4 of Table 1 is satisfied.

The notation for the intersection multiplicities in column 3 is as follows. Let p_1, \dots, p_k be a set of points and let b_i denote the number of surface branches at p_i . Consider a ray of projection $l(t) = p + t \cdot L$, where $L = \omega$ (for $\mathcal{V} = \mathbb{P}^2$) or $L = \omega - p$ (for $\mathcal{V} = \mathbb{R}^3 \setminus M$). The intersection multiplicities of the ray $l(t)$ with each surface branch at a k -tuple of points are then denoted by

$$([i_{11}, \dots, i_{1b_1}], \dots, [i_{k1}, \dots, i_{kb_k}]).$$

Strictly speaking, we should replace all intersection multiplicities i_{ab} in column 3 by $\geq i_{ab}$ (equality holds for the codimension 1 views, greater intersection multiplicities yield codimension ≥ 2 views). The equations in column 4 vanish if certain geometric conditions hold (for example, the point is a parabolic point or a cross-cap of the surface) and will be explained later.

We now have the following: $\omega \in \mathcal{V}$ belongs to \mathcal{B} (resp. \mathcal{B}') if, for some k -tuple of points in M (resp. M'), at least one of the 22 pairs of conditions in columns 3 and 4 of Table 1 are satisfied.

The conditions in Table 1 are described in more detail in the remainder of this section. The major differences between the conditions for unions of component surfaces with self-intersection curves, triple-points, and cross-caps and the corresponding conditions in [17] for unions of regular component surfaces are as follows:

- The cross-caps give rise to additional view singularities that are described in Section 3.1.
- At the singular points of M one has to consider the *intersection multiplicity* of a ray of projection and the variety $M := \bigcup h_i^{-1}(0)$ rather than the contact order—at non-singular points both concepts coincide. Furthermore, there are now, for example, three types of triple-points of M : the three “local surface branches” can globally lie on one, two, or three different component surfaces. The conditions for the intersection multiplicities are slightly different in the three cases, see Section 3.2.
- The “geometric conditions” in [17] (which are briefly summarized in the Appendix) have to be modified substantially in some cases, otherwise they would vanish identically on the singular sets of the component surfaces of M . Also, some “geometric conditions” fail to define complete intersections—this is in contrast to the “generic case” of unions of regular component surfaces studied in [17]. The necessary modifications of the “geometric conditions” are described in Section 3.3.
- Finally, certain types of codimension 1 views are detected by several of the 22

conditions in Table 1, see Section 3.4 below. This is also essentially due to the fact that we now allow nongeneric zero-sets M .

To save indices, we write $M := \bigcup_i h_i^{-1}(0)$ in the present section, rather than $M := \bigcup_{u,i} h_{u,i}^{-1}(0)$.

3.1. The Additional View Types Involving Cross-Caps

The surface parametrized by the map

$$X: \mathbb{R}^2 \rightarrow \mathbb{R}^3, \quad (u, v) \mapsto (u, uv, v^2)$$

has a cross-cap at the origin, and the image $X(\mathbb{R}^2)$ is given by the semialgebraic set

$$\{y^2 - x^2z = 0\} \setminus \{x = y = 0, z < 0\}.$$

Take a line $l(t) := (at, bt, ct)$, $(a : b : t) \in \mathbb{P}^2$, through the origin and compose it with the implicit defining equation $h := y^2 - x^2z$ of the cross-cap:

$$h \circ l(t) = b^2t^2 - a^2ct^3.$$

We see that the intersection multiplicity of the line and the cross-cap is 2 for $b \neq 0$, 3 for $b = 0$ and $ac \neq 0$, and infinite for the special lines $a = b = 0$ and $a = c = 0$. The lines of intersection multiplicity ≥ 3 , corresponding to $b = 0$, are contained in the tangent cone of the cross-cap, which is given by the lowest order homogeneous term of h , i.e., by $y^2 = 0$.

The projections of the cross-cap along some ray not contained in the tangent cone yield folds of the cross-cap, which are locally stable (there is an open set of such rays given by $b \neq 0$). For “most” rays in the tangent cone (the ones for which $b = 0$ but $ac \neq 0$) we obtain cusp views of the cross-cap (of type $W_{\Sigma^{1,1}}$), see entry $j = 20$ in Table 1 and Fig. 7. The last two entries in Table 1 are bilocal view singularities: in type $W+5$ a fold of a cross-cap and an ordinary fold line are superimposed, and in type $W+1$ a fold of a cross-cap and a projected crease line are superimposed, see Fig. 8.




Fig. 7. A cusp view of the cross-cap (center) and nearby stable views (left and right).




Fig. 8. A (fold or crease) line passing through a fold view of the cross-cap (center) and nearby stable views (left and right).

3.2. Singular Component Surfaces: Intersection Multiplicities

The intersection multiplicity of $l(t)$ and a *single* component surface $M_i = h_i^{-1}(0)$ at p_m is given by the order of the function $K(t) := h_i \circ l(t)$ at $t_m = l^{-1}(p_m)$. For K to have order a at t_m , it is necessary that $d^j K(t)/dt^j$, $0 \leq j \leq a - 1$, vanish at $t = t_m$.

For surfaces M , whose component surfaces M_i are singular, there are, for example, several distinct possibilities for interpreting the intersection multiplicity $[a, b, c]$ at a triple point p_m of M . Note that p_m can (globally) lie on

- (i) three distinct component surfaces M_1, M_2, M_3 of M whose intersection multiplicities with $l(t)$ are a, b , and c , respectively;
- (ii) on a pair of component surfaces M_1, M_2 , one of them having a self-intersection $S(M_1)$ at p_m , and intersection multiplicities $a + b$ and c (or $a + c, b$ or $a, b + c$);
or
- (iii) on a single component surface M_1 , having a triple point $T(M_1)$ at p_m and intersection multiplicity $a + b + c$.

In cases (ii) and (iii) we sometimes write $[a + b, c]$ and $[a + b + c]$, rather than $[a, b, c]$, for the required intersection multiplicity at p_m ; and in case (ii) we also require that $p_m \in S(M_1)$ and in case (iii) that $p_m \in T(M_1)$ (the conditions for a self-intersection and a triple point of a component surface are described below).

3.3. Singular Component Surfaces: Geometric Conditions

We begin with the conditions for the singular points of M . Let $S(M)$ denote the closure of the self-intersections of M , and let $T(M)$ and $W(M)$ denote the sets of triple-points and of cross-caps (Whitney umbrellas). The conditions for a point p to lie in one of these sets are as follows—recall that $M_i = h_i^{-1}(0)$ and note that the conditions on the right-hand side must hold for some $1 \leq i \leq n$ or for some k -tuple of distinct indices $1 \leq i_m \leq n$ ($1 \leq m \leq k$):

- $p \in S(M) \Leftrightarrow$ (i) $p \in M_{i_1} \cap M_{i_2}$ or (ii) $p \in S(M_{i_1})$;
- $p \in T(M) \Leftrightarrow$ (i) $p \in M_{i_1} \cap M_{i_2} \cap M_{i_3}$ or (ii) $p \in S(M_{i_1}) \cap M_{i_2}$ or (iii) $p \in T(M_{i_1})$;
- $p \in W(M) \Leftrightarrow p \in W(M_{i_1})$.

Furthermore, the conditions for p to lie in the self-intersection locus or in the triple-point or cross-cap set of a single component-surface M_i are as follows:

- $p \in S(M_i) \Leftrightarrow h_i(p) = dh_i(p) = 0$;
- $p \in T(M_i) \Leftrightarrow p \in S(M_i)$ and $\text{rank}(d^2h_i)_p = 0$;
- $p \in W(M_i) \Leftrightarrow p \in S(M_i)$ and $\text{rank}(d^2h_i)_p = 1$.

The triple-points of M_i are therefore given by the vanishing of h_i , of the components of dh_i , and of the entries of the Hessian d^2h_i . The conditions for the cross-caps are slightly more complicated: if the entries of the Hessian are simply replaced by the 2 by 2 minors, then the resulting conditions detect cross-caps as well as triple-points. The vanishing ideal of the cross-caps is the union of (h_i, dh_i) and the ideal quotient $I : J$, where I is generated by the 2 by 2 minors and J by the entries of d^2h_i . The system of generators of this vanishing ideal is denoted by \mathbf{W} in Table 1.

Next, we consider the remaining “geometric” conditions in column 4 of Table 1. For $j = 1, 5, 6, 10, 12$ and regular component surfaces M_i , these conditions are given in each case by the vanishing of a single polynomial equation supplementing the conditions for the intersection multiplicities in column 3 (see the Appendix and Section 3.2 of [17] for explicit expressions). For singular component surfaces, certain modifications of these conditions are necessary which, except for $j = 1$ and 6, are very minor.

The modifications for $j = 1$ and 6 are as follows. The condition $D_i = 0$ for a parabolic point of a regular zero-set $M_i = h_i^{-1}(0)$ (see the Appendix) vanishes not only on the parabolic set $P(M_i)$ but also on the singular set $S(M_i)$ of a singular component surface M_i . What makes matters worse, the (unwanted) component $S(M_i)$ creates solutions of the recognition equations having higher dimension than the view bifurcation set that we want to compute. We therefore replace the conditions $h_i = D_i = 0$ by a set of generators, denoted by D_{reg} , of the following ideal quotient:

$$I(D_{\text{reg}}) := (h_i, D_i) : J_{h_i},$$

where J_{h_i} denotes the Jacobian ideal of h_i (which is generated by the components of dh_i). Note that $I_2 := (h_i) + J_{h_i}$ is the ideal of functions vanishing on $S(M_i)$ and that such functions also belong to the ideal $I_1 := (h_i, D_i)$. The quotient $I_1 : I_2$ is contained in the vanishing ideal of the closure of the set $(P(M_i) \cup S(M_i)) \setminus S(M_i)$ (and, for $I_1 = \sqrt{I_1}$, is equal to it) and is equal to $I(D_{\text{reg}})$ (the latter equality follows from $I_1 : I_2 = \bigcap_{f \in I_2} I_1 : (f)$ and $I_1 : (h_i) = k[p]$, where k denotes the ground field).

For $j = 6$ there are two cases: the double-curve (i) is the intersection of distinct component surfaces and (ii) is the self-intersection of a single component surface. In case (i) we use the condition given in the Appendix, in the latter case (ii) we replace the expression for the tangent lines of the double-curves of M , which vanishes identically on $S(M_i)$, by the following. Observe that d^2h_i has corank 1 along $S(M_i)$ except at isolated points: at cross-caps and triple-points of M_i the corank is 2 and 3, respectively. Let α be a vector along the kernel direction of the gradient map $dh_i: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ and set $m := \langle \alpha, d^2h_i(L) \rangle$. In the absence of cross-caps and triple-points on $S(M_i)$, m could be taken as the new condition for $C\Sigma^{\parallel}$, but any such singular point would create an (unwanted) excess component in the bifurcation set. Hence, in general, we again have to compute the generators of an ideal quotient $I_1 : I_2$, where I_1 is the ideal generated

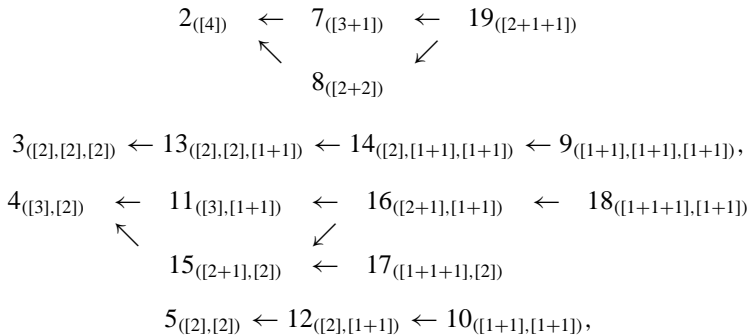
by m, h_i, J_{h_i} , and $d^2h_i(L, L)$ and where I_2 is generated by the vanishing of the 2 by 2 minors of d^2h_i .

3.4. *Singular Component Surfaces: Hierarchy of Dependencies between Defining Conditions*

For surfaces M which are unions of nonsingular zero-sets M_i (the “generic case”) each condition in Table 1 will “recognize” exactly one type of codimension 1 view of M , in this sense the conditions are “independent.” For singular component surfaces M_i this is no longer true. In this case there exist strictly antisymmetric dependency relations between certain pairs of conditions A and B , denoted by $A \rightarrow B$. That is, A holds automatically if B does, but not vice versa. In principal, B could be modified to recover independence from A : if $I(A)$ and $I(B)$ are the ideals defined by the conditions of A and B , then we had to replace B by the generators of the ideal $I(B) : I(A)$. However, the (costly) computation of these generators can be avoided, knowing the dependencies is sufficient.

In describing the dependencies we denote each condition by its number j in Table 1 and a subscript indicating its intersection multiplicity. Recall that k -local conditions involve k -tuples of intersection multiplicities $([i_1], \dots, [i_k])$, where $[i_l] = [a, b]$ or $[a, b, c]$ at a double- or triple-point of M . We now have to distinguish multiple points of M which arise from the same or from distinct component surfaces: for example, $[a, b, c]$, $[a + b, c]$, and $[a + b + c]$ denote the intersection multiplicities at triple-points of M cut out by 3, 2, and 1 component surface.

Using this notation, the hierarchy of dependencies amongst the defining conditions is given by the following sequences:



and

$$21_{(2),[2]} \leftarrow 22_{(2),[1+1]}.$$

All the codimension ≥ 1 views recognized by the conditions of some given sequence are simultaneously recognized by the corresponding “terminating condition” on the left. Each sequence could therefore be represented by its terminating condition. In practice, however, it is often better first to determine the components of the bifurcation set defined by the rightmost conditions in a sequence and then work from right to left. In this way, we can factor out the components already recognized by other conditions in the sequence.

4. The Time Complexity of Computing View Graphs

Applying the recognition conditions in Table 1 to algebraic surfaces M yields the defining equations of algebraic sets $\tilde{\mathcal{B}}_{j,r} \subset \mathcal{V} \times \mathbb{R}^{k_j+2}$, where $k_j = 1, 2,$ or 3 for local, bilocal, or trilocal singularities of the projection (there are no ≥ 4 -local codimension 1 views, and the views of higher codimension are all “recognized” by the conditions in Table 1). The first index j of $\tilde{\mathcal{B}}_{j,r}$ ranges over the 22 view types and the second index r enumerates the possible tuples of component surfaces M_i involved in a view of type j . We can think of the union $\tilde{\mathcal{B}}$ of these algebraic sets as being embedded in $\mathcal{V} \times \mathbb{R}^5$, note that \mathbb{R}^{k_j+2} has at most dimension 5. The restriction of the projection $\pi: \mathcal{V} \times \mathbb{R}^5 \rightarrow \mathcal{V}$ to $\tilde{\mathcal{B}}$ yields the view bifurcation set \mathcal{B} of M , which is a closed semialgebraic subset of \mathcal{V} . The view bifurcation set \mathcal{B}' of $M' \subset M$ is a semialgebraic subset of \mathcal{B} —the “branches” of $\mathcal{B} \setminus \mathcal{B}'$ consist of centers or directions of projection $\omega \in \mathcal{V}$ which yield unstable views of “pieces” of M which do not belong to the boundary M' . On the other hand, \mathcal{B} is a subset of the closed real algebraic set $\hat{\mathcal{B}}$ which is the projection of the complexification of $\tilde{\mathcal{B}}$ into \mathcal{V} . This set-up is summarized in the following diagram:

$$\begin{array}{c} \tilde{\mathcal{B}} \subset \mathcal{V} \times \mathbb{R}^5 \\ \downarrow \pi \\ \mathcal{B}' \subset \mathcal{B} \subset \hat{\mathcal{B}} \subset \mathcal{V} \end{array}$$

The following properties of the view bifurcation sets \mathcal{B} and \mathcal{B}' of surfaces M and M' are essential in the algorithms described below. For surfaces with transverse self-intersection curves and isolated triple points, these properties have been established in [17]. Looking at the proofs in [17] we see that these properties still hold for surfaces with additional cross-caps. Let \mathcal{P}^∞ (resp. \mathcal{P}^{alg}) denote the class of infinitely differentiable (resp. algebraic) surfaces with transverse self-intersection curves and isolated triple-points and cross-caps (where, of course, $\mathcal{P}^{\text{alg}} \subset \mathcal{P}^\infty$), then we have the following:

1. \mathcal{B} (and hence \mathcal{B}') has positive codimension in \mathcal{V} for any $M \in \mathcal{P}^\infty$ (see Proposition 2.1 of [17]).
2. \mathcal{B} and \mathcal{B}' have no “free boundaries” in \mathcal{V} of codimension 2 for any $M, M' \in \mathcal{P}^{\text{alg}}$. In particular, the boundary of any “component” $\mathcal{B}_{j,r}$ (or $\mathcal{B}'_{j,r}$) of \mathcal{B} (or \mathcal{B}') always lies in some other component $\mathcal{B}_{j',r'}$ (or $\mathcal{B}'_{j',r'}$) (see Propositions 2.3, 5.2, and 5.3 of [17]).
3. There are $O(n^K)$, $K = 6$ or 3 , components $\mathcal{B}_{j,r}$. However, for any $M \in \mathcal{P}^{\text{alg}}$, the boundary of $\mathcal{B}_{j,r}$ will lie in $O(1)$ other components which are known a priori (see case (i), Section 5(d)(i) of [17]).
4. Likewise, there are $O(n^K)$, $K = 6$ or 3 , components $\mathcal{B}'_{j,r}$. However, for any $M' \in \mathcal{P}^{\text{alg}}$, the boundary of $\mathcal{B}'_{j,r}$ will lie in $O(n)$ other components which are known a priori (see case (ii), Section 5(d)(i) of [17]).

Property 1 ensures that the view graph of M (and hence of M') always exists. Property 2 implies that we can “recover” the semialgebraic sets \mathcal{B} and \mathcal{B}' from the algebraic set $\hat{\mathcal{B}}$ (recall the diagram above) by computing certain cell-decompositions in \mathcal{V} —by contrast, for general semialgebraic sets \mathcal{B} arising as projections of some algebraic set $\tilde{\mathcal{B}} \subset \mathcal{V} \times \mathbb{R}^5$ one has to compute cell-decompositions of the total space $\mathcal{V} \times \mathbb{R}^5$. Finally, properties 3

(and 4) yield relatively coarse “stratifications” of $\hat{\mathcal{B}}$ (and \mathcal{B})—consisting of $O(n^K)$ (or of $O(n^{K+\dim \mathcal{V}})$), as opposed to $O(n^{K \dim \mathcal{V}})$, “branches”—from which \mathcal{B} (and \mathcal{B}') can be constructed by deleting those branches that lie in $\hat{\mathcal{B}} \setminus \mathcal{B}$ (or in $\mathcal{B} \setminus \mathcal{B}'$). Note that there are $O(n^K)$ components $\hat{\mathcal{B}}_{j,r}$ and $\mathcal{B}_{j,r}$: cutting each $\hat{\mathcal{B}}_{j,r}$ with at most $O(1)$ other components yields $O(n^K)$ branches of $\hat{\mathcal{B}}$. On the other hand, cutting a component $\mathcal{B}_{j,r} \subset \mathcal{V}$ with at most $O(n)$ other components yields $O(n^{\dim \mathcal{V}})$ branches of $\mathcal{B}_{j,r}$ and hence a total number of $O(n^{K+\dim \mathcal{V}})$ branches of \mathcal{B} . Note that the estimate of $O(n^{\dim \mathcal{V}})$ for the number of branches of $\mathcal{B}_{j,r}$ is very conservative: a bound of Milnor [12] says that the sum of the Betti numbers of the union of $O(n)$ components is $O(n^{\dim \mathcal{V}})$, but we are only interested in the number of branches on a single component.

We can now give a rough outline of our algorithm, which consists of a preprocessing step and three (resp. four) main steps for M (resp. M'). Compared with the algorithm in [17], the modified step 2 now has the highest combinatorial complexity and the new step 4 has a polynomial running time (by contrast, step 4 in [17] has an exponential running time). The exposition below concentrates on the combinatorial aspects of the algorithm, the algebraic parts are the same as in our earlier algorithm and have been described in [17] (note, in particular, that for a fixed number n of input surfaces $M = \bigcup_{i=1}^n M_i$ the running time of both the old and the new algorithm is polynomial in the maximal degree and bit-length of the defining polynomials of the surfaces).

First, we need the following notation. If $\bigcup_{u=1}^U \bigcap_{i=1}^{I_u} \{p \in \mathbb{R}^3 : h_{u,i}(p) \geq 0\}$ is the defining formula of the solid with boundary M' , then we set

$$H := \{h_{1,1}, \dots, h_{U,I_U}\} \subset \mathbb{Q}[p].$$

Recall that a connected subset Y of \mathbb{R}^3 is said to be H -invariant if all polynomials in H are sign-invariant on Y (i.e., are either strictly positive, negative, or vanish).

We first apply the semicylindrical stratification algorithm of Chazelle et al. [4] to the set of defining polynomials $H \subset \mathbb{Q}[p]$ of the input solid, in order to prune away the sets $\tilde{\mathcal{B}}_{j,r}$ that correspond to (empty!) components of the bifurcation set arising from surface-tuples that do not intersect (over \mathbb{R}). The time required for this preprocessing step is $O(n^7)$ (using the deterministic algorithm in [4]) or $O(n^{3+\epsilon})$ (using the randomized algorithm in [4]). This completes the “preprocessing” for the view graph algorithm for the zero-sets M .

For the bounding surfaces $M' \subset M$ we also prune away those $\tilde{\mathcal{B}}_{j,r}$ that correspond to surfaces, surface-pairs, and -triples that belong to $M \setminus M'$. Such surface-tuples correspond to empty components $\mathcal{B}'_{j,r}$ of the view bifurcation set of M' . This extra pruning can be done within the same asymptotic time, and the criteria for being part of the boundary M' are the same as in step 4 of the algorithm below. This leaves

$$n^K \sim O(n^3 + n^2 D + n D^2 + D^3 + n T + D T)$$

components $\tilde{\mathcal{B}}_{j,r}$ to be considered further, where $n = \sum I_u$ and where D and T are the numbers of surface-pairs and -triples that intersect over \mathbb{R} (passing the first pruning stage) or lie in the boundary M' of the solid (second pruning). For surfaces of general type, $D \sim O(n^2)$ and $T \sim O(n^3)$ so that $K = 6$. For surfaces diffeomorphic to a polyhedral surface having $O(n)$ edges and vertices, we have $K = 3$.

The view graph algorithms for M (for M') in the present work, and the ones in [17], then consist of the following very high-level steps 1–3 (1–4):

1. Compute, for the remaining $\tilde{\mathcal{B}}_{j,r}$, the (radicals of the) elimination ideals $I(\tilde{\mathcal{B}}_{j,r}) \cap \mathbb{Q}[\omega]$, for $1 \leq j \leq 22$ and $1 \leq r \leq c(j) \sim O(n^K)$. Result: the defining polynomials of real algebraic sets $\hat{\mathcal{B}}_{j,r}$.
2. Determine the connected components of $\mathcal{V} \setminus \hat{\mathcal{B}}$, where $\hat{\mathcal{B}} = \bigcup_{j,r} \hat{\mathcal{B}}_{j,r}$.
3. Decompose $\hat{\mathcal{B}}$ into $O(n^K)$ branches and remove the branches that lie in $\hat{\mathcal{B}} \setminus \mathcal{B}$. Result: the view graph $G(M)$.
4. Decompose \mathcal{B} into $O(n^{K+\dim \mathcal{V}})$ branches and remove the branches that lie in $\mathcal{B} \setminus \mathcal{B}'$. Result: the view graph $G(M')$.

Remark. The “branches” in steps 3 and 4 are, in general, smaller than the components $\mathcal{B}_{j,r}$ and $\mathcal{B}'_{j,r}$ of \mathcal{B} and \mathcal{B}' , respectively. We have already pruned away those components of \mathcal{B} and \mathcal{B}' corresponding to surface-tuples that don’t intersect over the reals or do not lie in the boundary M' . Steps 3 and 4 remove certain branches of the remaining components of \mathcal{B} and \mathcal{B}' .

We now consider some of these steps in more detail in order to determine their time complexity.

The elimination **step 1** uses standard tools from computational algebra and clearly requires $O(n^K) \cdot \mathcal{P}(d, L)$ time (note that the defining polynomials of $\mathcal{B}_{j,r}$ have a constant number of variables and their degree and coefficient size is $O(d)$ and $O(L)$, respectively).

Step 2 is based on a sign-invariant decomposition of \mathcal{V} with respect to the defining polynomials of the $O(n^K)$ algebraic sets $\hat{\mathcal{B}}_{j,r}$. In [17] step 2 is based on one (for $\mathcal{V} = \mathbb{P}^2$) or several (for $\mathcal{V} = \mathbb{R}^3$) cylindrical algebraic decompositions of \mathbb{R}^2 . Replacing the cylindrical algebraic decomposition (which is the computational bottleneck in step 2) by a semicylindrical decomposition (as defined by Chazelle et al. [4]) of \mathbb{R}^2 or \mathbb{R}^3 still allows us to determine the regions of $\mathcal{V} \setminus \hat{\mathcal{B}}$ in the same way as in [17]. Using the deterministic algorithm described in [4], step 2 requires $O(n^{K(2\dim \mathcal{V}+1)}) \cdot \mathcal{P}(d, L)$ time, and with the randomized algorithm from [4] the expected running time of step 2 becomes $O(n^{K \dim \mathcal{V} + \varepsilon}) \cdot \mathcal{P}(d, L)$.

In **step 3** we pick one sample point $\omega' \in S$ in each of the $O(n^K)$ branches of $\hat{\mathcal{B}}$. For $\omega' \in \hat{\mathcal{B}}_{j,r}$ we check whether the specialization of the ideal $I(\tilde{\mathcal{B}}_{j,r})$ to $\omega = \omega'$ has a positive number of real roots, if not we remove S . Step 3 can be carried out in $O(n^K) \cdot \mathcal{P}(d, L)$ time.

Remark. The sample points ω' in steps 3 and 4 can be chosen such that the specialization of $I(\tilde{\mathcal{B}}_{j,r})$ to $\omega = \omega'$ has isolated roots, see [17]. The definitions of the “branches” (with the appropriate extensions for the additional view types $j = 20, 21, 22$) and the determination of “good” sample points in steps 3 and 4 are the same as in our earlier algorithm, and we refer the reader to Section 5(d)(ii) of [17] for the (algebraic) details.

Likewise, in **step 4** we pick one sample point $\omega' \in S$ in each of the $O(n^{K+\dim \mathcal{V}})$ branches of \mathcal{B} . Below, we describe a new polynomial-time decision procedure for checking whether ω' belongs to \mathcal{B}' ; if not, we remove the branch S . (Note that the “old” decision procedure described in [17] has to count the number of real roots of up to $n^{3(U-1)}$ sys-

tems of polynomial equations and inequalities.) The new procedure for removing from \mathcal{B} those branches S that do not belong to \mathcal{B}' is as follows:

```

begin {Step 4}
  Decompose  $\mathbb{R}^3$  into  $H$ -invariant cells  $C$  and label the cells  $C \subset M'$  as
  boundary-cells;
  for each branch  $S \subset \mathcal{B}$  do
    Pick some algebraic representative  $\omega' \in S$ ;
     $b := FALSE$ ;
    for each root  $(p^m; \lambda_1^m, \dots, \lambda_{k_j-1}^m)$ ,  $1 \leq m \leq \rho$ , “over”  $\omega'$  do
      Determine the real algebraic number coordinates of the  $k_j$ -tuple
      of points  $p^m, l(\lambda_1^m), \dots, l(\lambda_{k_j-1}^m)$ ;
      Determine the cells  $C_1, \dots, C_{k_j}$  containing the points above;
      if all these cells are boundary-cells then  $b := TRUE$ 
    od
  if  $b = FALSE$  then delete  $S$ 
od
end

```

Using the deterministic (resp. randomized) semicylindrical stratification algorithm of Chazelle et al. [4] yields $O(n^{3+\varepsilon}) \cdot \mathcal{P}(d)$ H -invariant cells in $O(n^7) \cdot \mathcal{P}(d, L)$ time (resp. $O(n^{3+\varepsilon}) \cdot \mathcal{P}(d, L)$ expected time). (Recall that this stratification is already known from the preprocessing step.) For each H -invariant cell C of dimension ≤ 2 we pick a sample point p (with algebraic number coordinates) and check whether $p \in M'$, as follows. Let $\bar{h}_1, \dots, \bar{h}_t \in H$ denote the polynomials vanishing on the cell C and define $H_u := \{h_{u,1}, \dots, h_{u,l_u}\} \subset H$, $1 \leq u \leq U$. Also set $\bar{H} := \{\bar{h}_1, \dots, \bar{h}_t\}$,

$$H^+ := \bigcup_{H_u \cap \bar{H} \neq \emptyset} (H_u \setminus H_u \cap \bar{H})$$

and denote the H_u whose intersection with \bar{H} is empty by H_u^- . We label C as a boundary-cell if (i) for all $h \in H^+$ the algebraic number $h(p)$ has positive sign and (ii) for each H_u^- there exists at least one $h' \in H_u^-$ for which $h'(p)$ has negative sign. Checking whether C belongs to the boundary M' therefore requires $O(n) \cdot \mathcal{P}(d, L)$ time for each cell (recall that $|H| = n$ and note that the signs of these algebraic numbers can be determined in polynomial time). The first line of the above procedure therefore requires $O(n^7) \cdot \mathcal{P}(d, L)$ deterministic or $O(n^{4+\varepsilon}) \cdot \mathcal{P}(d, L)$ expected time.

Next, note that the outer loop will be executed at most $O(n^{K+\dim \mathcal{V}}) \cdot \mathcal{P}(d)$ times and the inner loop $\mathcal{P}(d)$ times. Recall that the points $l(\lambda_i)$ are either given by $p + \lambda_i \cdot \omega'$ (for $\mathcal{V} = \mathbb{P}^2$) or otherwise by $p + \lambda_i \cdot (\omega' - p)$, where λ_i and the coordinates of p and ω' are algebraic numbers. We can compute the minimal polynomials of the coordinates of the $l(\lambda_i)$ from the minimal polynomials of these algebraic numbers in $\mathcal{P}(d, L)$ time (using the usual algorithms for adding and multiplying algebraic numbers encoded by their minimal polynomials, see, e.g., the article by Loos [11] or Chapter 8.5 of the book by Mishra [13]). Finally note that, again using the algorithms in [4], we can determine the H -invariant cell containing some given point p or $l(\lambda_i)$ in $O(\log n) \cdot \mathcal{P}(d, L)$ time (note that the algebraic number coordinates of these points and the H -invariant cells are defined by polynomials of size $\mathcal{P}(d, L)$).

Summing up, we see that the above procedure for removing the branches of \mathcal{B} that do not belong to \mathcal{B}' requires $O(n^{K+\dim \mathcal{V}} \log n + n^7)$ time using the deterministic algorithm in [4] for computing an H -invariant stratification of 3-space or $O(n^{K+\dim \mathcal{V}} \log n)$ expected time using the randomized version.

Looking at the above bounds for the running times of the three (resp. four) steps of the view graph algorithm for M (resp. M') we see that, from a combinatorial point of view, step 2 dominates the asymptotic time complexity. The bounds for the running times of the deterministic and randomized view graph algorithms stated in the introduction now follow.

5. The Size of View Graphs

It is shown below that the upper bound in [17] of $O(n^{K \dim \mathcal{V}} d^{6 \dim \mathcal{V}})$ for the number of nodes in the view graph of a piecewise smooth surface, which is the union of n nonsingular surfaces of degree $\leq d$ intersecting in double-curves and triple-points, remains valid for the surfaces M (and hence M') studied in this paper. The double-curves and triple-points of the surfaces studied in [17] are cut-out by pairs and triples of regular component surfaces. On the other hand, the surfaces M are unions of singular component surfaces with double-curves, triple-points, and cross-caps.

Note that this estimate for the number of nodes $|V|$ in the view graph $G = (V, E)$ yields an upper bound for the size of G , because $|E| \sim O(|V|)$. This special property of view graphs—which is in contrast to complete graphs having $O(|V|^2)$ edges—boils down to the fact that the edges in E are top-dimensional branches of the bifurcation set \mathcal{B} and there are at most $O(n^{K \dim \mathcal{V}} d^{6 \dim \mathcal{V}}) = O(|V|)$ such branches.

One checks that the degrees of the view bifurcation sets \mathcal{B} of M and of the surfaces studied in [17] are of the same order, namely $O(n^K d^6)$, which yields the desired bound. First, one observes that the presence of double-curves and triple-points on singular component surfaces of M does not change the degree bounds stated in [17] for the components \mathcal{B}_j , $1 \leq j \leq 19$, of the view bifurcation set. (Essentially this follows from the fact that the degrees of the double-curve and of the triple-point set of a component surface of M are of order d^2 and d^3 , despite the fact that many equations are required to define these sets.)

The result then follows from the following degree bounds for the additional components \mathcal{B}_{20} to \mathcal{B}_{22} of the view bifurcation set of surfaces with cross-caps.

Proposition 5.1. *Let $M = \bigcup_{i=1}^n M_i$, $d = \sup_i \deg M_i$, be a piecewise smooth algebraic surface with transverse double-curves and isolated cross-caps and triple-points. Then the degree orders of the components \mathcal{B}_{20} , \mathcal{B}_{21} , and \mathcal{B}_{22} of the view bifurcation set are nd^3 , n^2d^5 , and n^3d^5 , respectively.*

Proof. The components of the bifurcation set are unions $\mathcal{B}_j = \bigcup_{r=1}^{c(j)} \mathcal{B}_{j,r}$ whose combinatorial complexity is given by

$$c(j) = \binom{n}{m(j)} \sim O(n^{m(j)}),$$

where $m(j)$ is the maximal number of distinct component surfaces M_i involved in a view singularity of type j . Clearly, $m(j) \leq 1, 2$ and 3 for $j = 20, 21$, and 22 , respectively.

The remaining task, then, is to estimate the degree (as a function of d) of a single subset $\mathcal{B}_{j,r}$ ($j = 20, 21, 22$). In doing this, the following lemma will be useful

Lemma 5.2. *Let $T(M_i)$ and $W(M_i)$ denote the sets of triple-points and cross-caps (Whitney umbrellas) of a degree d surface M_i with transverse double-curves and isolated triple-points and cross-caps. Then the following (asymptotically tight) upper bound holds:*

$$|T(M_i)| + |W(M_i)| \sim \Theta(d^3).$$

Proof. Consider the following stratification of M_i : take as zero-dimensional strata the triple-points and cross-caps, as one-dimensional strata the arcs of double-curves in the complement of the zero-strata and as two-dimensional strata the faces of M_i cut out by the closure of the double-point arcs. Denote the number of i -dimensional strata by e_i . It is convenient to distinguish bounded arcs, which contain at least one triple-point or cross-cap, from unbounded ones—denote the number of bounded and unbounded arcs by e_1^b and e_1^u , respectively.

Now $e_1^b \leq 6e_0$, where equality holds in the worst case where each bounded arc has infinite length and terminates in a triple-point. For the unbounded arcs we note that each contains four faces with, in turn, either one or two unbounded arcs in their closure (note, any face with ≥ 3 arcs in its closure must be cut out by bounded arcs), hence $e_1^u \leq e_2/2$. Therefore:

$$e_1 = e_1^b + e_1^u \sim O(e_0 + e_2).$$

On the other hand, we have that

$$e_0 - e_1 + e_2 = \chi(M_i) = \sum_j (-1)^j b_j(M_i) \leq \sum_j b_j(M_i),$$

and the sum of the Betti numbers b_j is, by a result of Milnor [12], at most $O(d^3)$. Hence, in particular,

$$e_0 = |T(M_i)| + |W(M_i)| \sim O(d^3).$$

Finally, the zero-set of $\prod_{j=1}^3 \prod_{k=1}^d (x_j - k)$ in \mathbb{R}^3 has degree $3d$ and d^3 triple-points. This example shows that the degree of our bound is exact. \square

Proof of Proposition (Conclusion). Let $l(t) = p + t \cdot L$, where $L = \omega$ (for $\mathcal{V} = \mathbb{P}^2$) or $L = \omega - p$ (for $\mathcal{V} = \mathbb{R}^3 \setminus M$), be a ray through a cross-cap $p \in M_i = h_i^{-1}(0)$. The intersection multiplicity of l and M_i at p is given by the order of $h_i \circ l(t)$ at $t = 0$. One checks that, at a cross-cap p , this order is at least two.

So, for ω to lie in $\mathcal{B}_{20,i}$, p has to be a cross-cap of M_i and l an asymptotic line. The latter means that $h_i \circ l(t)$ has order ≥ 3 at $t = 0$ which—at a cross-cap p where the order is automatically ≥ 2 —corresponds to the single condition

$$\left. \frac{d^2(h_i \circ l(t))}{dt^2} \right|_{t=0} = d^2 h_i|_p(L, L) = 0.$$

Hence

$$\mathcal{B}_{20,i} = \bigcup_{p \in W(M_i)} \{\omega \in \mathcal{V} : d^2 h_i|_p(L, L) = 0\}$$

is the union of $|W(M_i)| \sim O(d^3)$ quadrics in \mathcal{V} , which implies that $\deg \mathcal{B}_{20,i} \sim O(d^3)$.

For ω to lie in $\mathcal{B}_{21,r}$ there must exist—not necessarily distinct—surfaces $M_i, M_j \in \{M_1, \dots, M_n\}$ such that $p \in M_i$ is a cross-cap and the ray $l(t) = p + t \cdot L$ has at least 2-point contact with M_j at $l(\lambda), \lambda \neq 0$. This boils down to the following (where $\text{cl}\{\cdot\}$ denotes the closure):

$$\mathcal{B}_{21,r} = \bigcup_{p \in W(M_i)} \text{cl}\{\omega \in \mathcal{V} : \exists \lambda \neq 0 : h_j(p + \lambda \cdot L) = dh_j|_{p+\lambda \cdot L}(L) = 0\}.$$

The degrees of $h_j, dh_j(L) \in k[\lambda; \omega]$ are $O(d)$ —hence the union $\mathcal{B}_{21,r}$ of these $O(d^3)$ algebraic sets has degree at most $O(d^5)$.

A similar argument shows that the set

$$\mathcal{B}_{22,r} = \bigcup_{p \in W(M_i)} \text{cl}\{\omega \in \mathcal{V} : \exists \lambda \neq 0 : h_j(p + \lambda \cdot L) = h_k(p + \lambda \cdot L) = 0\},$$

corresponding to bilocal projections in which a cross-cap and a self-intersection curve appear superimposed, has degree at most $O(d^5)$. □

6. Additional Remarks and Open Problems

For the special case of *opaque views* of “polyhedral terrains”—that is, for graphs of piecewise linear functions in two variables having $O(n)$ edges, vertices, and faces—there is a bound for the number of nodes in the view graph, namely $O(n^{3 \dim \mathcal{V} - 1 + \epsilon})$, which is by about a factor of n sharper than the one for general polyhedra (see Theorems 8.31 and 8.33 of Sharir and Agarwal [18] and their paper [1]). Note that the size of the view graph for opaque views is less than or equal to the size of the view graph for transparent views (actually, the former can be obtained from the latter by contracting certain edges). De Berg et al. [5] have obtained a lower bound of $\Omega(n^{3 \dim \mathcal{V} - 1} \alpha(n))$ (α being the functional inverse of Ackermann’s function) for the number of distinct opaque views of polyhedral terrains, which means that the above upper bound is almost tight (actually, for parallel projection there is a slightly sharper upper bound of $O(n^5 2^c \sqrt{\log n})$ due to de Berg et al. [5] and Halperin and Sharir [8]). Replacing the piecewise linear functions by piecewise algebraic ones—but maintaining the restriction that the surfaces are function graphs with $O(n)$ edges, vertices, and faces—the combinatorial results in [18] can be combined with our algebraic estimates to obtain an $O(n^{3 \dim \mathcal{V} - 1 + \epsilon} d^{6 \dim \mathcal{V}})$ bound for the nodes in the view graphs of opaque views of such surfaces. In principal, the view graph of opaque views of “polyhedral terrains,” and more generally of semialgebraic function graphs, could be computed by supplementing our algorithm by a postprocessing step which merges those nodes in the view graph which correspond to distinct transparent but not to distinct opaque views. The resulting algorithm would have the same asymptotic

complexity as the algorithm described in this paper and hence would not be optimal for function graphs (terrains), its time complexity would be by about a factor of n too high.

Petitjean [14] has obtained exact formulas, as opposed to asymptotic bounds, for the degrees of the complexified view bifurcation sets of algebraic surfaces with double-curves and triple-points using techniques from enumerative geometry. These techniques also apply to surfaces with additional cross-caps. The transition from the complex to the real case is, however, very difficult: there exist no nontrivial lower bounds for the number of connected regions in the complement of \mathcal{B} in terms of d .

For the case of parallel projections of polyhedra, Gigus et al. [7] have presented an algorithm that determines the view graph $G = (V, E)$, as well as an explicit description of a view for each node, in $O(|V| \log |V| + n^4 \log n)$ time. The design of such an “output sensitive” algorithm, whose running time depends on the actual size $|V|$ of the view graph, for curved (piecewise smooth) algebraic surfaces is highly desirable but seems to be a formidable task. Note that, for (piecewise smooth) algebraic surfaces M of high degree d but small actual size of $|V|$, the real bifurcation set \mathcal{B} of M cuts out only a small number of regions, but the complexification of \mathcal{B} is nevertheless an algebraic variety of very high degree (namely $O(n^k d^6)$). It just happens that the high-degree polynomials at the very end of the view graph computation have relatively few *real* roots—but the polynomials up to that point could have lots of real roots. Obtaining an output-sensitive algorithm in such a situation would be of general interest in computational mathematics and would be a major breakthrough in computational real algebraic geometry.

For polyhedra some recent works have considered the computation of finite-resolution view graphs (see [19]), which take into account that details of a view that are smaller than some size-threshold cannot be detected by cameras of limited spatial resolution. An extension of this work to curved semialgebraic surfaces seems possible but not very enlightening. Contrary to the claim by some authors (see the paper [6] based on a panel discussion on aspect graphs in computer vision) that there exist no mathematical techniques for computing finite-resolution view graphs of curved surfaces, one should note that the finite-resolution partition of view space of (piecewise smooth) algebraic surfaces can be defined by Tarski sentences and hence can be effectively computed. (Recall that Tarski sentences are Boolean formulas with quantifiers and with polynomial (in-)equalities as predicates.) However, the boundaries of the regions of this partition are no longer view bifurcation sets (with well-understood topological properties) but rather unrestricted semialgebraic sets. The exact computation of this finite-resolution partition might therefore be prohibitive, perhaps one has to be satisfied with approximations.

Appendix. The Geometric Conditions in [17]

Below, $M_i = h_i^{-1}(0)$ is a regular surface and $l(t) = p + t \cdot L$ is a ray of projection, where $L = \omega$ (for parallel projection) or $L = \omega - p$ (for central projection):

- (i) M_i has zero Gaussian curvature at p if and only if

$$D_i(p) := \det \begin{pmatrix} d^2 h_i|_p & dh_i|_p \\ (dh_i|_p)' & 0 \end{pmatrix} = 0.$$

- (ii) A pair of surface normals $dh_1|_{p=l(0)}, dh_2|_{l(t)}, t \neq 0$, in L^\perp is parallel if and only if

$$N^\parallel(p, t, \omega) := \langle dh_1|_p \wedge L, dh_2|_{l(t)} \rangle = 0.$$

- (iii) The tangent lines of the crease curve $M_1 \cap M_2$ and of the critical set of M_1 are parallel at p if and only if

$$C\Sigma^\parallel(p, \omega) := \langle dh_1|_p \wedge dh_2|_p, d^2h_1|_p(L) \rangle = 0.$$

- (iv) The tangent lines of the projections of a pair of crease curves $M_1 \cap M_2, M_3 \cap M_4$ are parallel at their point of intersection if and only if

$$\pi_C^\parallel(p, t, \omega) := \langle (dh_1|_p \wedge dh_2|_p) \wedge (dh_3|_{l(t)} \wedge dh_4|_{l(t)}), L \rangle = 0, \quad t \neq 0.$$

- (v) Finally, the tangent lines of a fold of M_1 and of a projected crease curve $M_2 \cap M_3$ are parallel at their point of intersection if and only if

$$\pi_{\Sigma C}^\parallel(p, t, \omega) := \langle dh_1|_p, dh_2|_{l(t)} \wedge dh_3|_{l(t)} \rangle = 0, \quad t \neq 0.$$

References

1. P.K. Agarwal and M. Sharir, On the number of views of polyhedral terrains, *Discrete Comput. Geom.* 12 (1994), 177–182
2. K.W. Bowyer and C.R. Dyer, Aspect graphs: an introduction and survey of recent results, *Internat. J. Imaging Technol.* 2 (1990), 315–328
3. J.S. Carter, J.H. Rieger, and M. Saito, A combinatorial description of knotted surfaces and their isotopies, *Adv. in Math.* 127:1 (1997), 1–51
4. B. Chazelle, H. Edelsbrunner, L. Guibas, and M. Sharir, A singly exponential stratification scheme for real semi-algebraic varieties and its applications, *Theoret. Comput. Sci.* 84 (1991), 77–105
5. M. de Berg, D. Halperin, M. Overmars, and M. van Kreveld, Sparse Arrangements and the Number of Views of Polyhedral Scenes, *Internat. J. Comput. Geom. Appl.* 7 (1997), 175–195.
6. O. Faugeras et al., Panel theme: why aspect graphs are not (yet) practical for computer vision, *Proc. IEEE Workshop on Directions Automated CAD-Based Vision*, 1991, pp. 98–104
7. Z. Gigus, J. Canny, and R. Seidel, Efficiently computing and representing aspect graphs of polyhedral objects, *IEEE Trans. Pattern Anal. Mach. Intell.* 13:6 (1991), 542–551
8. D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visibility in terrains, *Discrete Comput. Geom.* 12 (1994), 313–326
9. D. Hilbert and S. Cohn-Vossen, *Anschauliche Geometrie*, Springer-Verlag, Berlin, 1932.
10. C.M. Jessop, *Quartic Surfaces with Singular Points*, Cambridge University Press, Cambridge, 1916.
11. R. Loos, Computing in algebraic extensions, in *Computer Algebra: Symbolic and Algebraic Computation*, B. Buchberger, G.E. Collins, and R. Loos, eds., Springer-Verlag, Wien, 1982, pp. 173–187
12. J. Milnor, On the Betti numbers of real varieties, *Proc. Amer. Math. Soc.* 15 (1964), 275–280
13. B. Mishra, *Algorithmic Algebra*, Springer-Verlag, New York, 1993.
14. S. Petitjean, The enumerative geometry of projective algebraic surfaces and the complexity of aspect graphs, *Internat. J. Comput. Vision* 19:3 (1996), 261–287
15. S. Petitjean, J.H. Rieger, and D. Forsyth, Recognizing algebraic surfaces from aspects, in *Algebraic Surfaces in Computer Vision*, J. Ponce, ed., Springer-Verlag, Heidelberg, in press.
16. W. Platinga and C. Dyer, Visibility, occlusion and the aspect graph, *Internat. J. Comput. Vision* 5 (1990), 137–160
17. J.H. Rieger, On the complexity and computation of view graphs of piecewise smooth algebraic surfaces, *Philos. Trans. Roy. Soc. London Ser. A* 354 (1996), 1899–1940

18. M. Sharir and P.K. Agarwal, *Davenport–Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, Cambridge, 1995.
19. I. Shimshoni and J. Ponce, Finite-resolution aspect graphs of polyhedral objects, *IEEE Trans. Pattern Anal. Mach. Intell.* 19:4 (1997), 315–327
20. J.M. West, *The Differential Geometry of the Crosscap*, Ph.D. thesis, University of Liverpool, 1995.
21. H. Whitney, The singularities of mappings from smooth n -manifolds into $(2n - 1)$ -space, *Ann. of Math.* 45 (1944), 247–293

Received July 11, 1996, and in revised form July 14, 1997.