

The Strength of Weak Proximity (Extended Abstract)

Giuseppe Di Battista¹, Giuseppe Liotta², Sue Whitesides³

¹ Dipartimento di Discipline Scientifiche, Sezione Informatica Terza Università di
Roma via Segre 2 00146 Roma

`dibattista@iasi.rm.cnr.it`

² Department of Computer Science, Brown University, 115 Waterman Street,
Providence, RI 02912.

`gl@cs.brown.edu`

³ School of Computer Science, McGill University, 3480 University St. # 318,
Montréal, Québec, H3A 2A7 Canada.

`sue@cs.mcgill.ca`

Abstract. This paper initiates the study of *weak* proximity drawings of graphs and demonstrates their advantages over *strong* proximity drawings in certain cases. Weak proximity drawings are straight line drawings such that if the *proximity region* of two points p and q representing vertices is devoid of other points representing vertices, then segment (p, q) is allowed, but not forced, to appear in the drawing. This differs from the usual, strong, notion of proximity drawing in which such segments must appear in the drawing.

Most previously studied proximity regions are associated with a parameter β , $0 \leq \beta \leq \infty$. For fixed β , weak β -drawability is at least as expressive as strong β -drawability, as a strong β -drawing is also a weak one. We give examples of graph families and β values where the two notions coincide, and a situation in which it is NP-hard to determine weak β -drawability. On the other hand, we give situations where weak proximity significantly increases the expressive power of β -drawability: we show that every graph has, for all sufficiently small β , a weak β -proximity drawing that is computable in linear time, and we show that every tree has, for every β less than 2, a weak β -drawing that is computable in linear time.

1 Introduction and Overview

Given two points u and v of the plane, a *proximity region* of u and v is a portion of the plane, determined by u and v , that contains points relatively close to both of them. A *proximity drawing* of a graph G has been defined in the literature as

* Research supported in part by the National Science Foundation under grant CCR-9423847, by the U.S. Army Research Office under grant 34990-MA-MUR, by FCAR, by Progetto Finalizzato Sistemi Informatici e Calcolo Parallelo of the Italian National Research Council (CNR), and by N.A.T.O.-CNR Advanced Fellowships Programme.

a *straight-line drawing* (vertices of G are mapped to distinct points of the plane, and edges to straight-line segments) such that: (i) for each edge (u, v) of G , the proximity region of the points representing u and v does not contain any other vertex; and (ii) for each pair of non-adjacent vertices u, v of G , the proximity region of the points representing u and v contains at least one other vertex.

Most of the results on proximity drawings take as proximity regions the so-called β -regions [11]. Such regions form an infinite family, each element of the family being identified by a value of the parameter β ($0 \leq \beta \leq \infty$). For example, when $\beta = 1$ the proximity region of u and v is the disk with u and v as antipodal points; when $\beta = 2$ the proximity region is the intersection of two disks with centers at u and v and radius the distance $d(u, v)$ between u and v ; when $\beta = \infty$ the proximity region is the infinite strip perpendicular to the line segment between u and v . A β -drawing is a proximity drawing such that the proximity regions are β -regions. A graph is β -drawable if it has a β -drawing. A brief survey on proximity drawability is in [4]. Besides their theoretical interest, proximity drawings have been studied for their practical characteristics: neighboring graph vertices are clustered in the drawing, and adjacent edges tend to have large angles. Furthermore, proximity drawings are related to minimum spanning tree drawings, minimum weight drawings of triangulations, and Delaunay drawings (e.g., see [6, 10, 2]).

The purpose of this paper is to initiate a study of *weak* proximity drawings, in particular, weak β -drawings. A weak proximity drawing of a graph G is one that ignores requirement (ii) for traditional, or *strong*, drawings. In other words, if (u, v) is *not* an edge of G , then no requirement is placed on the proximity region of u and v in a weak drawing. For example, Fig. 1 shows a weak proximity drawing of a tree. Here, the proximity region of any two points p and q is the disk having p and q as antipodal points. Note that the drawing is not a strong drawing, as no edges between neighbors of the degree six vertex are included.

There are several motivations for studying weak proximity drawings and in particular, weak β -drawings.

- Strong proximity drawability is very restrictive, perhaps too much so. By relaxing (ii), a graph G can no longer be reconstructed from the locations of its vertices in a weak drawing; however, many graphs that do not admit strong drawings can be drawn weakly. For example, a tree that has a vertex of degree greater than five has no strong β -drawing for any β [3]. Thus the drawing in Fig. 1 illustrates a graph that is *weak* but not *strong* drawable for the circular disk proximity region defined by antipodal points. Also, characterizations and algorithms for strong β -drawability have been devised only for a few classes of graphs.
- Weak and strong visibility drawings (e.g., see [15]) can be considered as a particular class of proximity drawings. In the field of visibility drawing, the coordinated study of both strong and weak types of drawings led to deep and practical results.
- Weak proximity can be considered as an “edge-vertex resolution rule” in the sense that a vertex cannot enter the region of influence of an edge. Thus, the




Fig. 1. A weak proximity drawing of a tree.

study of weak proximity can contribute to the body of drawing strategies that adopt a resolution rule (e.g., see [5, 9]).

- The weak proximity model may well be sufficient for many drawing applications, particularly ones that do not require recovery of the graph solely from the positions of its vertices.

The main results presented in this paper are as follows.

General graphs. We show that any graph G is weak β -drawable for all β in the range 0 to some upper bound that is a function either of the number of vertices or of the maximum vertex degree of G . (Section 3.)

Planar graphs. First, we show how to extend existing strong proximity drawability results on outerplanar graphs to weak drawability results. Second, we show that, in a certain interval for β , strong and weak β -drawings of triangulated planar graphs coincide. Third, we give new insights on the interplay between angular resolution and proximity drawability. Namely, we show how to interpret any straight-line drawing algorithm for planar triangulated graphs as an algorithm for constructing weak proximity drawings. (Section 4.)

Trees. We provide an algorithm to draw any tree as a weak β -drawing for any value of β less than two. Then we show that for $2 \leq \beta < \infty$, the weak and the strong proximity models give rise approximately to the same class of β -drawable trees. Finally, we show the NP-hardness of deciding whether a tree has a weak proximity drawing for $\beta = \infty$, where the region of influence is an open strip. (Section 5.)

All our algorithms admit a linear time implementation in real RAM. The

above results represent, in many cases, substantial improvements over the known algorithms and characterizations for strong proximity drawability.

2 Preliminaries

We recall some basic definitions concerning proximity drawings.

Given a pair x, y of points in the plane, the *open β -region of influence of x and y* , and the *closed β -region of influence of x and y* , denoted by $R(x, y, \beta)$ and $R[x, y, \beta]$ respectively, are defined as follows:

1. For $0 < \beta < 1$, $R(x, y, \beta)$ is the intersection of the two open disks of radius $d(x, y)/(2\beta)$ passing through both x and y . $R[x, y, \beta]$ is the intersection of the two corresponding closed disks.
2. For $1 \leq \beta < \infty$, $R(x, y, \beta)$ is the intersection of the two open disks of radius $\beta d(x, y)/2$, centered at the points $(1-\beta/2)x + (\beta/2)y$ and $(\beta/2)x + (1-\beta/2)y$. $R[x, y, \beta]$ is the intersection of the two corresponding closed disks.
3. $R(x, y, \infty)$ is the open infinite strip perpendicular to the line segment \overline{xy} and $R[x, y, \infty]$ is the closed infinite strip perpendicular to the line segment \overline{xy} .
4. $R(x, y, 0)$ is the empty set; $R[x, y, 0]$ is the line segment connecting x and y .

Let G be a graph. A *weak (strong) (β) -drawing* of G is a weak (strong) proximity drawing of G such that for each pair of points x, y the proximity region is $R(x, y, \beta)$. Weak and strong (β) -drawings are called w - (β) -drawings and s - (β) -drawings, respectively. Analogously, a *weak (strong) $[\beta]$ -drawing* of G is a weak (strong) proximity drawing of G such that for each pair of points x, y the proximity region is $R[x, y, \beta]$. Weak and strong $[\beta]$ -drawings are called w - $[\beta]$ -drawings and s - $[\beta]$ -drawings, respectively.

A graph is *w- (β) -drawable* (*s- (β) -drawable*) if it has a w - (β) -drawing (s - (β) -drawing). Analogously, a graph is *w- $[\beta]$ -drawable* (*s- $[\beta]$ -drawable*) if it has a w - $[\beta]$ -drawing (s - $[\beta]$ -drawing). When it is clear from the context or when it is not necessary to distinguish between open and closed proximity regions, we simplify the notation by talking about β -drawings and β -drawable graphs. A class of graphs is *w- β -drawable* (resp. *s- β -drawable*) if all its graphs are w - β -drawable (resp. s - β -drawable). A class of graphs is *not w- β -drawable* (resp. *not s- β -drawable*) if it contains at least one graph that is not w - β -drawable (resp. s - β -drawable).

The following properties extend many results on strong drawability to weak drawability.

Property 1. An s - β -drawable graph is also w - β -drawable.

Property 2. Let Γ be a w - β -drawing of a graph G and let P be the set of points of Γ representing the vertices of G . Let Γ' be an s - β -drawing such that Γ' uses the same set of points P to represent the vertices of G' . Then $G \subseteq G'$.

Attractive drawability inclusion properties follow from weak proximity.

Property 3. If G is a w - β -drawable graph, then $G' \subset G$ is a w - β -drawable graph.

Property 4. If G is a w - $\bar{\beta}$ -drawable graph, then G is w - β -drawable for any β such that $0 \leq \beta < \bar{\beta}$.

To analyse w - β -drawings we will frequently use two angles $\alpha(\beta)$ and $\gamma(\beta)$, defined as follows.

1. $\alpha(\beta) = \inf\{\angle xzy \mid z \in R[x, y, \beta]\}$.
2. $\gamma(\beta)$ is only defined for $\beta \geq 2$, and $\gamma(2) = \frac{\pi}{3}$. For $\beta > 2$, let $z \neq y$ be a point on the boundary of $R[x, y, \beta]$ such that $d(x, y) = d(x, z)$. Then $\gamma(\beta) = \angle xzy$.

Property 5. [3]

1. $\beta = \sin \alpha$ for $0 \leq \beta < 1$.
2. $\beta = \frac{1}{1 - \cos \alpha}$ for $1 \leq \beta \leq \infty$.
3. $\beta = \frac{1}{\cos \gamma}$ for $2 \leq \beta \leq \infty$.

3 General Graphs

In this section, we give a simple, fast method for producing w -(β)- and w - $[\beta]$ -drawings of arbitrary graphs on n vertices for certain β in the range $0 \leq \beta < 1$.

Theorem 6. Any graph on n vertices is w - $[\beta]$ -drawable for all values of β such that $0 \leq \beta < \sin(2\pi/n)$; any graph on n vertices is w -($\sin(2\pi/n)$)-drawable.

Sketch of proof. For the trivial case $\beta = 0$, it suffices to place the points on a circular arc of measure $< \pi$. For $\beta > 0$, place n points equally spaced around a circle C of arbitrary radius R . Recall that for $0 < \beta < 1$, the radius r for the circular arcs bounding a proximity region $R(x, y, \beta)$ is given by $d(x, y)/(2\beta)$. For β sufficiently close to 0, the region of influence is a slight widening of the line segment joining x and y . It lies entirely within C and hence contains none of the n points distinct from x and y . Radius r decreases with increasing β . When β increases to the extent that $r \leq R$, then points that are not consecutive on the circle cannot be joined by an edge in a w - $[\beta]$ -drawing (similarly for $r < R$ and w -(β)-drawings). This critical value of β is thus determined by $R = r = d(p, q)/2\beta$, i.e., $\beta = d(p, q)/2R$, where p, q is the closest pair of points that are non-consecutive on C . For n equally spaced points, $\sin(2\pi/n) = (d(p, q)/2)/R$, so $d(p, q) = 2R\sin(2\pi/n)$. Hence, independent of R , the critical value of β is $\sin(2\pi/n)$ (approximately $2\pi/n$ for large n).

The statement of Theorem 6 can be strengthened in certain cases by using a method of [16]. Consider a coloring of G by χ colors, where χ is the chromatic number of G . Divide a circle C of arbitrary radius R into χ arcs of equal length.

Cluster points receiving color i about the center of arc i . The closest pair p, q of points that are non-consecutive on C but that are joined by an edge of G must lie in different arcs. Hence given any $t < 2R \sin(\pi/\chi)$, by placing points sufficiently close to the centers of their arcs, we can be sure that $d(p, q) \geq t$. Hence, for $0 \leq \beta < \sin(\pi/\chi)$, w - β -drawings exist for G . Of course, the chromatic number χ is in general hard to compute. Instead of breaking C into χ arcs, one can instead break C into $d + 1$ arcs, where d is the maximum degree of G . A coloring for G by $d + 1$ or fewer colors can be obtained in linear time by a well-known greedy coloring algorithm. Hence for β between 0 and $\sin(2\pi/(d + 1))$, a w -(β)- or w - $\lfloor \beta \rfloor$ -drawing for G can be obtained in linear time. Hence we have:

Theorem 7. *A graph with chromatic number χ is w - β -drawable for all β such that $0 \leq \beta < \sin(\pi/\chi)$. A w - β -drawing for G can be obtained in linear time for any β such that $0 \leq \beta < \sin(\pi/(d + 1))$, where d is the maximum degree of G .*

4 Planar Graphs

Several interesting results on planar graphs can be “imported” from strong proximity results by using Property 1. Other results can be obtained with the same property and little more work. For example, from Lubiw and Sleumer [12] we obtain:

Theorem 8. *Biconnected outerplanar graphs are w - β -drawable for all values of β such that $0 \leq \beta < 2$. Furthermore, such graphs are w -(2)-drawable.*

Sketch of proof. In [12] it is shown that any biconnected outerplanar graph is s -(2)-drawable. The conclusion follows immediately by Properties 1 and 4.

Another example concerns *planar triangulated* graphs, which admit planar drawings with all faces triangular, including the external face.

Theorem 9. *Let G be a planar triangulated graph. For each value of β such that $1 < \beta \leq \infty$, G is w - β -drawable if and only if it is s - β -drawable. Furthermore, G is w - $\lfloor \beta \rfloor$ -drawable if and only if it is s - $\lfloor \beta \rfloor$ -drawable.*

Sketch of proof. The if-part of both statements is trivial by Property 1. For the only-if-part of the first statement (the argument for the second is analogous), consider an allowed value of β and suppose G is w - β -drawable with w - β -drawing Γ . We show that Γ is also an s - β -drawing. Consider the graph G' that has an s - β -drawing with the same set of points for the vertices as Γ . From Property 2 we have that $G \subseteq G'$. From [3] we have that in the above interval of β values, s - β -drawable graphs are planar; hence, G' is planar. The conclusion follows from the maximality of G .

Further consideration of planar triangulated graphs reveals a connection between weak proximity drawability and angular resolution.

Lemma 10. *Let Γ be a straight-line drawing of a planar triangulated graph G such that the angle between two adjacent edges is at most $\alpha(\beta)$, where $0 \leq \beta \leq 1$. Then Γ is a $w(\beta)$ -drawing of G .*

The *angular resolution* of a straight-line drawing is the size of the minimum angle formed by adjacent edges. The *angular resolution of a straight-line drawing algorithm \mathcal{A}* is a number \mathcal{F} that is the infimum of all the angular resolutions of all the drawings that \mathcal{A} could construct. When inputs to \mathcal{A} are restricted to various classes of graphs, examples in the literature show a dependency of \mathcal{F} on the maximum degree (e.g., see [13]) or on the number of vertices of graphs (e.g., see [14]). Lemma 10 is the geometric foundation of the following theorem.

Theorem 11. *Let \mathcal{A} be a straight-line drawing algorithm whose inputs are planar triangulated graphs. Let \mathcal{F} be its angular resolution. We have that any drawing produced by \mathcal{A} is a $w\beta$ -drawing for all values of β such that $0 \leq \beta < \sin(\pi - 2\mathcal{F})$. Furthermore, such a drawing is a $w(\sin(\pi - 2\mathcal{F}))$ -drawing.*

5 Trees

We denote by \mathcal{T}_k a class of trees having maximum vertex degree at most k ; \mathcal{T}_∞ is the class of all trees. We also denote by $\mathcal{T}(\beta)$ the class of $w(\beta)$ -drawable trees. Similarly, $\mathcal{T}[\beta]$ is the class of $w[\beta]$ -drawable trees.

First we prove that every tree is $w(\beta)$ -drawable for $\beta \leq 2$. We begin by constructing a $w(2)$ -drawing for an arbitrary tree. The construction can be formulated as a linear time algorithm for real RAM. In the drawing, each point p representing a tree node has the following construction devices associated with it: an open disk $D(p)$ centered at p ; an open superwedge $W^+(p)$ with vertex at the parent of p (this wedge is left undefined if p is the root); a closed subwedge $W(p)$ with vertex at p .

To *generate the children* of a point p means to compute for each child q_i its superwedge, its coordinates, its disk, and its subwedge. The construction continues in breadth-first fashion from the root. Each time the children q_i of some point p are generated, the following invariants are maintained.

1. Each superwedge $W^+(q_i)$ belongs to $W(p)$, and the superwedges of distinct children of p are disjoint.
2. Each disk $D(q_i)$ lies inside the superwedge $W^+(q_i)$, tangent to its sides.
3. Subwedge $W(q_i)$ with vertex at q_i lies inside the superwedge $W^+(q_i)$ and has sides parallel to those of the superwedge $W^+(q_i)$. (Its purpose is to contain q_i and all its descendants.)

The root of the tree is placed at the origin. If the root has $k \geq 2$ children q_1, \dots, q_k , these are generated by dividing the plane into k equal angle superwedges $W^+(q_i)$ with vertex at the origin. Then each q_i is placed distance 1 from the origin on the bisector of its superwedge. This determines the disks and the subwedges of the q_i . If the root has only one child q_1 , then the superwedge

$W^+(q_1)$ is given vertex angle $\pi/2$ and q_1 is placed at unit distance from the origin on the bisector of its superwedge; its disk $D(q_1)$ has unit radius. Once the coordinates, disks and wedges have been determined for all points at depth 0 and 1, the construction continues in a breadth-first manner. Suppose the tree has depth d , where the root has depth 0. For each depth $i = 1$ to $d - 1$, the children of each point p at depth i are generated from p and its disk and wedges as follows (see Fig. 2). Suppose p has $k \geq 1$ children q_i , $1 \leq i \leq k$. Equally subdivide $W(p)$ into k superwedges $W^+(q_1) \dots W^+(q_k)$. Place q_i at the intersection of the bisector of its superwedge with the boundary of $D(p)$. This determines the subwedge of p , since its sides are parallel to those of the superwedge of p , and the disk of q_i , since it is tangent to the superwedge.

Theorem 12. $\mathcal{T}(2) = \mathcal{T}_\infty$. Furthermore, given a tree $T \in \mathcal{T}_\infty$, a $w(2)$ -drawing of T can be computed in time proportional to the size of T in the real RAM.




Fig. 2. Construction for trees.

We can now exploit Property 4 to extend the above result to infinitely many values of β .

Corollary 13. $\mathcal{T}(\beta) = \mathcal{T}_\infty$ and $\mathcal{T}[\beta] = \mathcal{T}_\infty$ for any $0 \leq \beta < 2$.

Surprisingly, it turns out that if $\beta = 2$ and the region of influence is a closed set, then the class of $w[2]$ -drawable trees does not contain trees with arbitrarily large vertex degree.

Lemma 14. $\mathcal{T}[2] = \mathcal{T}_5$.

The previous lemma can be generalized to values of β such that $2 < \beta \leq \infty$.

Lemma 15.

1. Let Γ be a w -(β)-drawing of a tree for $2 < \beta \leq \infty$. Then the angle between two adjacent edges is at least $\gamma(\beta)$.
2. Let Γ be a w - $[\beta]$ -drawing of a tree for $2 < \beta \leq \infty$. Then the angle between two adjacent edges is greater than $\gamma(\beta)$.

From Lemma 15, an upper bound can be deduced for the maximum vertex degree of w - β -drawable trees for $2 < \beta \leq \infty$. It is worth noticing that the same lemma, and hence the same upper bounds, hold for s - β -drawable trees [3]. Thus, we can use Lemma 15 and Property 1 to import the results of [3] on the s - β -drawability of trees.

	β	$\mathcal{T}(\beta)$	$\mathcal{T}[\beta]$
1	$2 < \beta < \frac{1}{\cos(2\pi/5)}$	$T_4 \subset \mathcal{T}(\beta) \subseteq T_5$	$T_4 \subset \mathcal{T}[\beta] \subseteq T_5$
2	$\beta = \frac{1}{\cos(2\pi/5)}$	$T_4 \subset \mathcal{T}(\beta) \subset T_5$	$\mathcal{T}[\beta] = T_4$
3	$\frac{1}{\cos(2\pi/5)} < \beta < \infty$	$\mathcal{T}(\beta) = T_4$	$\mathcal{T}[\beta] = T_4$
4	$\beta = \infty$	$T_3 \subset \mathcal{T}(\beta) \subset T_4$	$\mathcal{T}[\beta] = T_3$

Table 1. Table for Theorem 16

Theorem 16. Table 1 describes w - β -drawable trees for all values of β such that $2 < \beta \leq \infty$.

We conclude this section by proving that it is NP-hard to determine whether a tree of maximum degree four is w -(∞)-drawable. The proof follows the NAE-3SAT paradigm introduced by Bhatt and Cosmadakis [1] and exploited for geometric graph realizability questions by Eades and Whitesides [6], [7]. First we introduce some terminology for w -(∞)-drawings. These are also called *weak open strip* drawings, as the region of influence of two points x, y is the open strip perpendicular to the line segment between them.

A weak open strip drawing is *orthogonal* if each of its edges is parallel to one of two given orthogonal direction vectors. These vectors may be assumed to be horizontal and vertical. A *normalized* orthogonal drawing is obtained from an orthogonal drawing as follows. Order the vertices in the drawing from left to right, with vertices having the same x -coordinate assigned the same place in the order. Order the vertices from bottom to top similarly. Assign each vertex a new x -coordinate equal to its order in the left-to-right order; assign each vertex a new y -coordinate equal to its order in the bottom-to-top order. The drawing resulting from this coordinatization is a normalized orthogonal drawing. Normalized drawings are weak open strip drawings if the original weak open strip drawing was orthogonal. This is *not* necessarily true in general.

A graph G is *orthogonally unique* if a) any weak open strip drawing of it must be orthogonal, and b) the set of points and edges of any normalized orthogonal drawing of G is unique up to rigid motions (i.e., combinations of reflections, rotations and translations) of the plane.

Observation 17. Suppose H is a subgraph of some graph G that has a weak open strip drawing Γ . When the part of Γ that does not represent H is discarded, what remains is a weak open strip drawing of H .

It follows from this observation that if H is orthogonally unique, then any weak open strip drawing of a graph G containing H must respect the constraints on the drawing of H .

Observation 18. Any angle formed by edges with a common endpoint in a weak open strip drawing must be $\geq \pi/2$.




Fig. 3. Orthogonally unique graphs.

Lemma 19. Each of the graphs a) through e) of Figure 3 is orthogonally unique and has a normalized orthogonal drawing as shown.

Now consider an instance of NAE-3SAT consisting of m clauses $C_1 \dots C_m$ and n variables and their complements, $X_1, \dots, X_n, X'_1, \dots, X'_n$. Each clause contains

three literals, but no clause contains both a literal X and its complement X' . The instance is a “yes” instance if and only if a truth assignment can be found such that each clause contains at least one true literal and at least one false literal. We encode the instance in polynomial time in its length by designing a graph G having the form shown in part e) of Figure 3. In particular, G should have m internal vertical columns, n internal rows lying above the horizontal chain containing the crosses, and n internal rows lying below this horizontal chain.

By Lemma 19, G is orthogonally unique. Note, however, that the vertical columns can be flipped individually around the central horizontal chain without changing the drawing. Similarly, short 2-edge horizontal paths joining degree 1 vertices can be flipped around their points of attachment to long paths without changing the drawing.

Theorem 20. *To determine whether a given tree of maximum degree four has a weak open strip drawing is NP-hard.*

Sketch of proof. We modify graph G to encode a particular instance of NAE-3SAT. This is done following the Bhatt Cosmadakis paradigm. Applications of this paradigm are familiar from [6] and [7], and the application of the paradigm in this case is straight-forward once Lemma 19 is known. Hence we sketch only briefly how this is done, referring the reader to the original paper by Bhatt and Cosmadakis [1] for more details.

The first pair of rows above and below the horizontal chain of crosses represents clause C_1 , the second pair represents C_2 , and so on. Similarly, the left-most internal column represents variable X_1 and its complement X'_1 . For the j^{th} internal column, label one half the column with X_j and the other half with the corresponding X'_j . Now add some extra edges to G as follows. If C_i fails to contain X_j (or X'_j), add an extra edge to the short horizontal path in one of the two rows corresponding to C_i , on the half-column corresponding to X_j (or X'_j). Clearly the resulting graph has a weak open strip drawing if and only if the two columns and the short paths can be flipped around so that there is at least one position in each row to which no extra edge is attached. But such flips can be found if and only if the NAE-3SAT instance is a “yes” instance. This is because half-columns appearing above the horizontal chain can be interpreted as “true”, and half-columns appearing below the horizontal chain can be interpreted as “false”. Hence a missing edge in each of a pair of corresponding rows above and below the horizontal chain means that the clause associated with this pair of rows contains a true literal and a false literal.

6 Open Problems

Several remaining open problems make weak proximity drawability an attractive direction of research. One class of problems concerns the use of the weak model for proximity regions other than β -regions. For example weak rectangle of influence drawings [8] could be tackled. It is easy to see that any planar graph that admits an st -orientation without transitive edges has a weak rectangle of

influence drawing. Another problem area is to consider weak proximity models that do not allow *edges*, as opposed to vertices, to enter the proximity regions of other edges. This seems too restrictive to consider in a strong proximity model.

References

1. S. Bhatt and S. Cosmadakis. The Complexity of Minimizing Wire Lengths in VLSI Layouts. *Information Processing Letters*, **25**, 1987, pp. 263-267.
2. P. Bose, W. Lenhart, and G. Liotta. Characterizing Proximity Trees. *Algorithmica*, Special Issue on Graph Drawing (to appear).
3. P. Bose, G. Di Battista, W. Lenhart, and G. Liotta. Proximity Constraints and Representable Trees. *Graph Drawing, Proc. of the DIMACS International Workshop GD94, New Jersey, USA, October 1994*, LNCS **894**, R. Tamassia and I.G. Tollis eds., Springer-Verlag, 1995, pp. 340-351.
4. G. Di Battista, W. Lenhart, and G. Liotta. Proximity Drawability: a Survey. *Graph Drawing, Proc. of the DIMACS International Workshop GD94, New Jersey, USA, October 1994*, LNCS **894**, R. Tamassia and I.G. Tollis eds., Springer-Verlag, 1995, pp. 328-339.
5. G. Di Battista, R. Tamassia, and I. G. Tollis. Area Requirement and Symmetry Display of Planar Upward Drawings. *Discr. and Comp. Geometry*, **7**, 1992, pp. 381-401.
6. P. Eades and S. Whitesides. The Realization Problem for Euclidean Minimum Spanning Trees is NP-hard. *Proc. 10th ACM Symposium on Computational Geometry*, 1994, pp. 49-56.
7. P. Eades and S. Whitesides. Nearest Neighbor Graph Realizability is NP-hard. *Proc. Latin'95, Valparaiso, Chile, 1995*, LNCS **911**, R. Baeza-Yates, E. Goles, P. Poblete eds., Springer-Verlag, 1995, pp. 245-256.
8. M. Ichino and J. Sklansky. The Relative Neighborhood Graph for Mixed Feature Variables. *Pattern Recognition*, **18**, 1985, pp. 161-167.
9. G. Kant, G. Liotta, R. Tamassia, and I.G. Tollis. Area Requirement of Visibility Representations of Trees. *Proc. 5th CCCG, Waterloo, 1993*, pp. 192-197.
10. J.M. Keil. Computing a Subgraph of the Minimum Weight Triangulation. *Comp. Geom.: Theory and Appl.*, **4**, 1994, pp. 13-26.
11. D. G. Kirkpatrick and J. D. Radke. A Framework for Computational Morphology. *Computational Geometry*, ed. G. T. Toussaint, Elsevier, Amsterdam, 1985, pp. 217-248.
12. A. Lubiw and N. Sleumer. All Maximal Outerplanar Graphs are Relative Neighborhood Graphs. *Proc. 5th CCCG, Waterloo, 1993*, pp. 198-203.
13. S. Malitz and A. Papakostas. On the Angular Resolution of Planar Graphs. *Proc. STOC*, 1993, pp. 431-437.
14. W. Schnyder. Embedding Planar Graphs on the Grid. *Proc. SODA*, 1990, pp. 41-51.
15. R. Tamassia and I. G. Tollis. A Unified Approach to Visibility Representations of Planar Graphs. *Discr. and Comp. Geometry*, **1**, 1986, pp. 321-341.
16. M. Formann, T. Hagerup, J. Haralambides, M. Kaufmann, F. T. Leighton, A. Simvonis, E. Welzl, G. Woeginger. Drawing Graphs in the Plane with High Resolution. *Proc. FOCS*, 1990, pp. 86-95.