

Topology Preserving Constrained Graph Layout

Tim Dwyer, Kim Marriott, and Michael Wybrow

Clayton School of Information Technology,
Monash University, Clayton, Victoria 3800, Australia,
{Tim.Dwyer, Kim.Marriott, Michael.Wybrow}@infotech.monash.edu.au

Abstract. Constrained graph layout is a recent generalization of force-directed graph layout which allows constraints on node placement. We give a constrained graph layout algorithm that takes an initial feasible layout and improves it while preserving the topology of the initial layout. The algorithm supports poly-line connectors and clusters. During layout the connectors and cluster boundaries act like impervious rubber-bands which try to shrink in length. The intended application for our algorithm is dynamic graph layout, but it can also be used to improve layouts generated by other graph layout techniques.

Keywords: constrained graph layout, polyline routing, clusters

1 Introduction

A core requirement of dynamic graph layout is stability of layout during changes to the graph so as to preserve the user’s mental model of the graph. One natural requirement to achieve this is that, as far as possible, the topology of the current layout is preserved during layout changes. While topology preservation has been used for dynamic layout based on orthogonal graph layout, its use in force-directed approaches to dynamic layout is much less common.

Constrained graph layout is a recent generalization of the force-directed model for graph layout [1–3]. Like force-directed methods, these techniques find a layout minimizing a goal function such as the standard *stress* goal function which tries to place all pairs of nodes their ideal (graph-theoretic) distance apart. However, unlike force directed methods, constrained graph layout algorithms allow the goal to be minimized subject to placement constraints on the nodes. In this paper we detail a constrained graph layout algorithm that preserves the topology of the initial layout. The primary motivation for our development of this algorithm was to support dynamic layout but it can also be used to improve layouts generated by other graph layout techniques such as planarisation techniques [4].

Our algorithm supports network diagrams with poly-line connectors and arbitrary node clusters. It ensures that the nodes do not overlap and that additional constraints on the layout—such as alignment and downward pointing edges—remain satisfied. During layout optimization the *paths*, i.e poly-line connectors and cluster boundaries, act like rubber-bands, trying to shrink in length and

hence, in the case of connectors, straighten. Like physical rubber bands, the paths are impervious and do not allow nodes and other paths to pass through them. Thus, the initial layout topology is preserved. Figure 1 shows example layouts obtained with our algorithm.

Extending constrained graph layout to handle topology preservation is conceptually quite natural since topology preservation can be regarded as a kind of constraint. However, it was not possible to straightforwardly extend existing constrained graph layout algorithms to preserve topology. One issue is that previous algorithms were based on functional majorization whose use relied on particular properties of the stress goal function. The main technical innovations in our new algorithm are fourfold. First, we utilize a new goal function P -stress that encodes the rubber band metaphor, measuring the stretch of paths as well as trying to place objects a minimum distance apart. Importantly, the P -stress is *bend-point invariant* in the sense that merging two consecutive collinear segments in a path does not change the value of the goal function. This is important for convergence since it means that the goal function behaves continuously as the paths change during optimization. Second, we utilize constrained gradient projection rather than functional majorization. This approach is generic in the choice of goal function and so can be used to minimize P -stress. Third, we give a novel algorithm for updating paths in a layout given that nodes are moved in a single dimension. This maintains the relative order of nodes and paths in that dimension and so preserves the initial topology. The final innovation is our uniform treatment of connector routes and cluster boundaries as impervious paths. This allows our algorithm to handle arbitrary clusters.

The algorithm for topology preserving constrained graph layout given here underpins two dynamic graph layout applications we have developed. The first is a network diagram authoring tool, *Dunnart*, which uses the algorithm to provide continuous layout adjustment during user interaction [5]. The second is a network diagram browser which uses the algorithm to update the layout of a detailed view of part of the network as the user changes the focus node or collapses or expands node clusters [6]. The contribution of this paper is to detail the algorithm.





Fig. 1. Example layouts obtained with the topology preserving constrained graph layout algorithm. In the metabolic pathway, three vertical alignment constraints have been added to improve the layout.

2 Related Work

There has been considerable interest in developing techniques for stable graph layout that preserve the user’s mental model of the graph [7]. These techniques are quite specialized to the underlying layout algorithms. The standard approach for supporting stability in force-directed approaches is to simply add a “stay force” on each node so that it does not move unnecessarily, e.g. [8]. Stable dynamic layout has also been studied for orthogonal graph layout, e.g. [9]. There, stability is preserved by trying to preserve the current bend points and angles. This has the effect of preserving the layout topology. Finally, North and Woodhull [10] have given algorithms for preserving the current horizontal and vertical ordering between nodes in dynamic Sugiyama-style layered layout. Our approach is the first that we are aware of to base stability on topology preservation in a force-directed style layout. It has the advantage over stay forces that the layout is better able to adjust to changes while still preserving the original structure.

Orthogonal graph layout algorithms typically feature a refinement step that attempts to shorten edges while preserving edge crossing topology [11]. However, the approach is very specific to orthogonal drawings and is not applicable to more general layout. Another method, [12], used a force directed approach but only handled abstract graphs with point nodes and straight-line edges. Most closely related is our earlier extension to constrained stress majorization that preserves layout topology while trying to straighten bends in poly-line connectors [13]. This works by introducing dummy nodes in each connector at all possible bend points and adding constraints to ensure a minimum separation between objects and bend-points. Unfortunately, our experience with this algorithm was that straightening bends sometimes meant that connector length was increased and that the algorithm did not scale to moderately sized networks because of the large number of dummy nodes. Even worse it did not always converge because the goal function was not bend-point invariant. The algorithm given here is considerably simpler, convergent and faster.

3 Problem Definition

A graph $G = (V, E, C)$ consists of a set of nodes V , a set of edges $E \subseteq V \times V$, and a set of node clusters $C \subseteq \wp V$. We let $width(v)$ and $height(v)$ give the width and height of the bounding rectangle, r_v , of each node $v \in V$. A 2-D drawing of a graph is specified by a tuple (X, P) where $X = [X_v \equiv (v_x, v_y) | v \in V]$ gives the centre positions for each node and P is a set of *topological boundaries* or paths corresponding to edge routings and cluster boundaries. Each edge routing $P_e \in P$ for edge $e = (s, t) \in E$ is a piecewise linear path through a sequence of points (p_1, \dots, p_n) where point p_1 corresponds to the centre of node s , p_n to the centre of node t and each point $p_k, 1 < k < n$ corresponds to a node v and one of the four corners (labelled TL, TR, BL, BR) of the rectangle r_v . That is, a point in an edge path is a pair $(v \in V, i \in \{Centre, TL, TR, BL, BR\})$. Each cluster boundary $P_c \in P$ for cluster $c \in C$ is similar to an edge routing except that it starts and ends at the same node corner (not a node centre).

Separation constraints are inequality or equality constraints over pairs of position variables in either the horizontal or vertical axes of the drawing, e.g. for a pair of nodes $u, v \in V$ we might define a separation constraint over their x -positions: $u_x + g \leq v_x$ where g specifies a minimum spacing between them.

A *feasible* drawing of a graph (see Fig. 2) is one in which:

- all separation constraints are satisfied;
- no two node rectangles overlap;
- the nodes inside the region defined by the boundary of each cluster c are exactly the nodes in c ;
- every path $p \in P$ is *valid* and *tight*.

A *valid* path is one in which no segment passes through a node rectangle, except the first and last segments in a path corresponding to an edge which must terminate at the centre of rectangles as specified above. A *tight* path is one where every bend (described by three consecutive points a, b, c in the path) is wrapped around the rectangle r_v associated with the bend point $b = (v, i)$. That is, if $(a_x, a_y), (b_x, b_y), (c_x, c_y)$ are the coordinates of the points a, b, c and (v_x, v_y) is the centre of the node v around whose rectangle v_r the bend is routed, then the points a, b, c in order must constitute a turn in the same direction as the points a, b, v in order, and the points b, c, v must also constitute a turn in the same direction.

A common strategy for finding aesthetically pleasing drawings of graphs is to define a cost function over the positions of nodes X and then to attempt to minimise this cost function by adjusting these positions. In our case we are also interested in the lengths of paths. For each path $p \in P$ we define an ideal length for the path d_p .

The ideal length d_e of an edge $e \equiv (u, v)$ is proportional to the graph theoretic distance D_{uv} between nodes u and v while the desired boundary length of cluster c is $2\sqrt{\pi \sum_{v \in c} \text{width}(v)\text{height}(v)}$ (i.e. the ideal length is proportional to the perimeter of the circle of the same area as that of the constituent nodes). Therefore, we use a cost function *P-stress* which also takes the paths P of the layout into consideration:

$$\sum_{i < j} w_{ij} ((D_{ij} - \|X_i, X_j\|)^+)^2 + \sum_{p \in P} w_p ((\|p\| - d_p)^+)^2$$

where $w_p = \frac{1}{d_p^2}$, $w_{ij} = \frac{1}{D_{ij}^2}$, z^+ is z if $z \geq 0$ and 0 otherwise.

The first component of *P-stress* is a modification of the *stress* function used in the stress majorization [14] and Kamada and Kawai [15] layout methods.




Fig. 2. Example of incorrect (a,b) and correct (c) paths.

However, unlike the stress function, nodes that are more than their ideal distance apart are not penalized, thus eliminating long range attraction since this can cause issues in highly constrained problems. The second component of *P-stress* tries to make the length of each path p in the network, no more than its ideal length d_p . Note that this second component is purely attractive, otherwise minimising *P-stress* could potentially *increase* bends.

Note that *P-stress* is *bend-point invariant* in the sense that merging two consecutive collinear segments in a path does not change the *P-stress* of layout since the overall path length does not change. This is important for convergence of the layout algorithm.

4 Minimising P-stress using gradient projection

Our layout problem is, therefore, given a feasible layout for a graph to find a new layout that is feasible, has the same topology as the original layout, and which locally minimizes *P-stress*. In this section we give an algorithm to do this. An example of its operation is shown in Figure 3.

However, before giving the algorithm we make precise what we mean by topology preservation. We minimise *P-stress* by alternately adjusting the horizontal and vertical positions of all nodes. Considering just the horizontal case since the vertical is symmetrical, we say that a horizontal adjustment of the nodes from feasible layout $L = (X, P)$ to feasible $L' = (X', P')$ is *topology preserving* if no node or line segment moves through another node or line segment. More exactly, let M and M' be the layouts obtained from L and L' , respectively, by infinitesimally reducing the height of each node rectangle and appropriately modifying the paths. This means that node rectangles whose top and bottom were aligned in the original layout now have a

infinitesimal vertical separation between them. Then for any height h we must have that scanning left to right along the horizontal line $y = h$ encounters exactly the same sequence of edges, clusters and nodes in both M and M' where an edge is encountered whenever the line intersects a path segment for the edge, a cluster is encountered whenever the line intersects a path segment for its boundary and a node is encountered when the the line intersects the rectangle corresponding to the node.

As we have indicated, our algorithm works by alternately adjusting horizontal and vertical positions of all nodes to incrementally reduces *P-stress*. Constrained stress majorization [3] also uses a similar approach to reduce stress. However,




Fig. 3. Example of how our layout algorithm improves the network layout by reducing *P-stress* by shortening edge routings while preserving the topology of the initial layout.

the useful Cauchy-Schwartz based expansion of the stress function into horizontal and vertical quadratic forms which strictly (upper-)bound the goal function, is no longer easily derived for *P-stress*. Instead, at each iteration k we use a quadratic approximation over the horizontal position variables $x \subset X^k$ from the second order Taylor series expansion of *P-stress* at X^k and compute a descent vector $-g$ and step size α from the first and second derivatives of this quadratic to compute a new position d for the horizontal position variables. We then use the function *project* to project d onto the horizontal constraints necessary to avoid overlap and to preserve topology and any other user specified separation constraints on the horizontal variables C_x . Next we perform an analogous operation to compute a new position for the vertical position variables $y \subset X^k$ and vertical constraints to obtain X^{k+1} . The high-level algorithm is thus:

```

procedure gradient-projection( $dim, X, P, C$ )
   $g \leftarrow \nabla_{dim} P\text{-stress}(X, P)$ 
   $H \leftarrow \nabla_{dim}^2 P\text{-stress}(X, P)$ 
   $\alpha \leftarrow \frac{g^T g}{g^T H g}$ 
   $d \leftarrow X[dim] - \alpha g$ 
return project( $X, P, dim, d, C$ )

procedure improve( $X, P, C_x, C_y$ )
   $k \leftarrow 0$ 
   $(X^k, P^k) \leftarrow (X, P)$ 
  repeat
     $(X', P') \leftarrow \text{gradient-projection}(x, X^k, P^k, C_x)$ 
     $k \leftarrow k + 1$ 
     $(X^k, P^k) \leftarrow \text{gradient-projection}(y, X', P', C_y)$ 
  until  $|P\text{-stress}(X^{k-1}, P^{k-1}) - P\text{-stress}(X, P)|$  sufficiently small
return  $(X^k, P^k)$ 

```

5 Topology preserving projection

The heart of the layout algorithm is the procedure *project*(X, P, dim, d, C) which performs a projection operation in the specified axis dim . More precisely, it returns a new layout (X', P') which is feasible and preserves the topology of (X, P) while ensuring the position of the nodes in dimension dim is as close as possible to the desired positions d and leaving their position in the other dimension unchanged. It has three main steps:

- (1) Generate separation constraints C^{no} to ensure non-overlap of nodes and topology constraints TC to ensure topology preservation.
- (2) Project d on to $SC = C \cup C^{no}$ giving \bar{x} . This is achieved by solving the quadratic program:

$$\min_x \sum_{i \leq n} (x_i - d_i)^2 \text{ subject to } SC \quad (1)$$

- (3) Update the path routing P to give P' by moving the nodes smoothly from x to \bar{x} appropriately adjusting the paths as the nodes move in order to satisfy the topology constraints TC .

Non-overlap and topological constraints are generated as follows. We assume that nodes are moving horizontally: the case for a vertical move is symmetric. A bounding rectangle r_u is associated with each node u . The non-overlap and topological constraints are generated using a top to bottom scan of the drawing. At each step we keep the list of currently open rectangles and line segments. To do so we process the vertical opening and closings of each rectangle OR, CR and line segment OS, CS of the given routing in order from top to bottom and, when two such events occur at the same vertical position, then with precedence:

- OS before CS so that horizontal segments are handled properly
- CR before OR to avoid unnecessary non-overlap constraints (assuming no zero height rectangles)
- CS before OR, CR before OS, OS before OR, CR before CS to ensure all possible segment/rectangle interactions are considered.

For each rectangle opening (i.e. the top of each rectangle) we add to C^{no} a separation constraint between the rectangle and its immediate left and right neighbours in the list of open rectangles at that y -position (the scan position). Each separation constraint has the form $u_x + s \leq v_x$ over the x positions of nodes u and v and preserves the relative horizontal ordering of u and v and prevents the nodes from overlapping, where $s = (\text{width}(r_u) + \text{width}(r_v))/2$.

The scan also generates *topology* constraints between nodes and paths which ensure that the paths remain tight and valid. Each topology constraint is a linear inequality over three variables such that the rectangle r_w associated with node w must be to the right or left of a line between the corners of two nodes u and v , e.g. $w_x + g \leq u_x + p(v_x - u_x)$ requires w to be to the left of the line, g is an offset computed based on the widths of the nodes u, v and w (as in Fig. 5) and p is the ratio along the line of the intersection of the scan (y -)position (e.g. the top of r_w). There are two types of topology constraints: *straight* constraints—between a node w and a path segment uv which ensures that the path remains valid, i.e. the node does not overlap the path segment and *bend* constraints associated with a bend point between two line segments which ensures that the path remains tight around the bend point. For straight constraints $0 < p \leq 1$ while for bend constraints $p > 1$. The procedures for creating each type of constraint is given in Fig. 5.




Fig. 4. Constraints generated during a vertical scan. There is a separation constraint $u_x + \frac{1}{2}(\text{width}(u) + \text{width}(v)) \leq v_x$ to prevent overlap, 3 bend constraints (the construction for the constraint $w_x + g \geq u_x + p(v_x - u_x)$ is shown) and 3 straight constraints at the places where the segments s and t may potentially bend.

If $|P|$ denotes the number of path segments, the worst case complexity of Step 1 of *project* is $O(|V|(|P| + \log |V|))$ and up to $O(|V|)$ non-overlap constraints and $O(|P||V|)$ topological constraints can be generated.

In Step 2 of *project* we solve the quadratic program of Eqn. 1 using the incremental active-set procedure *solveQPSC* given in [3]. This steps through the vertices in the constraint boundary in a similar manner to the Simplex algorithm. As such it is difficult to prove polynomial running time, however in practice it is very fast as indicated by our experimental results.

Step 3 of *project* is performed by procedure *move* (Figure 6). This updates the paths by moving the nodes horizontally from the initial feasible solution x for which the routing is correct towards \bar{x} detecting violated topology constraints as they move. A violated *bend* constraint indicates that consecutive segments have become aligned and can be replaced with a single segment. A violated *straight* constraint indicates that a single segment needs to bend (be split into two new segments).

The maximum move that can be made along the line from a to b without violating topology constraint t is determined by solving the linear equation associated with the constraint. For example, if t is the constraint $w + g \leq u + p(v - u)$ then the maximum safe move is:

$$\frac{\alpha}{\beta} = \frac{w_a - g - u_a + p(u_a - v_a)}{u_b - u_a + p(u_a - u_b + v_b - v_a) + w_a - w_b}$$

where u_a, v_a, w_a and u_b, v_b, w_b are, respectively, the position of the nodes u, v and w at a and at b . The iterative process of finding the next such constraint and updating the paths P is accomplished in the *move* procedure, Fig. 6.

Note that the *satisfy* procedure shown in Fig. 6, which satisfies a topology constraint by either merging or splitting segments, must transfer or replace other bend and straight constraints associated with the affected segments. The detail is not shown, but the intuition for the difficult edge case of a horizontal path segment is shown in Fig. 7.

The *move* procedure used for updating the paths to preserve validity and topology can also be thought of as a kind of active-set process, and as such it is difficult to prove that it is polynomial. Again, however, please see our results section for actual running times which indicate that running times scale fairly well with the number of topology constraints generated. Note that the number of bend constraints is exactly the number of bend points in P , and the number of straight constraints—while the worst case is $O(|P||V|)$ —is limited by only generating constraints for segments which are *visible* in the axis of movement from a given rectangle open/close.

Theorem 1. *Let (X, P) be a feasible layout with respect to separation constraints C in the dimension corresponding to \dim . Then $\text{project}(X, P, \dim, d, C)$ will return a new layout (X', P') that is the layout that minimizes the distance between $X'[\dim]$ and d subject to only changing the positions $X[\dim]$ of X , remaining feasible and being topology preserving with respect to (X, P) .*


```

procedure createStraightConstraint( $s, w, y, TC$ )
  % for segment  $s = uv$  and node  $w$  at scan pos  $y$ 
   $p \leftarrow (y - u_y)/(v_y - u_y)$ 
   $x_p \leftarrow u_x + p(v_x - u_x)$ 
   $leftOf \leftarrow xCentre(r_w) < x_p$ 
   $corner \leftarrow$  if  $y < yCentre(r_w)$  then if  $leftOf$  then BR else BL
                    else if  $leftOf$  then TR else TL

   $o \leftarrow width(r_w)/2$  (-ve if  $leftOf$ )
   $g \leftarrow offset(u) + p(offset(v) - offset(u)) - o$ 
   $TC \leftarrow TC \cup \{TopologyConstraint(straight, u, v, w, p, g)\}$ 

procedure createBendConstraint( $b, TC$ )
  % for bend point  $b = (v, i)$ , between segments  $ab$  and  $bc$ 
  if  $i$  is the centre of  $v$  then return
  if existing bend constraint  $t$  on  $b$  then remove  $t$ 
   $leftOf \leftarrow i \in \{TR, BR\}$ 
  if  $|a_y - b_y| > |b_y - c_y|$  then
     $p \leftarrow (c_y - a_y)/(b_y - a_y)$ 
     $g \leftarrow offset(a) + p(offset(b) - offset(a)) - offset(c)$ 
     $t \leftarrow TopologyConstraint(bend, a, b, c, p, g, leftOf)$ 
  else
     $p \leftarrow (a_y - c_y)/(b_y - c_y)$ 
     $g \leftarrow offset(c) + p(offset(b) - offset(c)) - offset(a)$ 
     $t \leftarrow TopologyConstraint(bend, c, b, a, p, g, leftOf)$ 
   $TC \leftarrow TC \cup \{t\}$ 

```

Fig. 5. The procedures for creating straight constraints and bend constraints are used in both the initial scan to set up topology constraints and by the procedure *satisfy* (Fig. 6).

Proof. (Sketch) Any feasible and topology preserving layout must satisfy $SC = C \cup C^{no}$. By construction the new node position X' changes only the positions $X[dim]$ of X and since $X'[dim]$ is the projection of d onto SC it is the closest node position that satisfies SC . Furthermore, one can prove by induction that the *satisfy* procedure returns updated paths P' that are topology preserving, tight and valid.

6 Finding a feasible topology

We can apply our topology preserving layout adjustment to layout obtained by any graph drawing algorithm, assuming the generated layout is feasible as defined in §3. Although not the primary focus of this paper we have developed an algorithm to find an initial feasible layout. This has two main steps:

- (1) Perform standard stress majorization to find an initial position X^0 for the nodes. A position for the nodes satisfying the constraints is found by projecting X^0 on to the user specified separation constraints and then using a greedy heuristic to satisfy the non-overlap constraints and cluster containment constraints. We use the approach detailed in [3].

```

procedure satisfy( $t, TC, P$ )
   $TC \leftarrow TC \setminus \{t\}$ 
  if  $t$  is a bend constraint over points  $a, b, c$  then
    %  $b = (v, i)$  is the bend point of  $t$ 
    replace segments  $ab$  and  $bc$  in  $P$  with new segment  $ac$ 
    createStraightConstraint( $ac, v, b_y$ )
  else %  $t$  is a straight constraint over  $u, v, w$ 
    replace segment  $uw$  in  $P$  with segments  $uv$  and  $vw$ 
    transfer straight constraints on  $uw$  to either  $uv$  or  $vw$ 
    createBendConstraint( $u$ )
    createBendConstraint( $v$ )
procedure move( $x, \bar{x}, TC, P$ )
  repeat
     $\alpha \leftarrow \beta \leftarrow 1$ 
     $t^* \leftarrow \text{None}$ 
    for  $t \in TC$ 
      %  $t$  is a Topology Constraint over  $u, v, w$  with constants  $p, g$ 
       $a \leftarrow w_a - g - u_a + p(u_a - v_a)$ 
       $b \leftarrow u_b - u_a + p(u_a - u_b + v_b - v_a) + w_a - w_b$ 
      if  $a\beta < \alpha b$  then
         $\alpha \leftarrow a, \beta \leftarrow b, t^* \leftarrow t$ 
       $x \leftarrow x + \frac{\alpha}{\beta}(\bar{x} - x)$ 
    if  $t^* \neq \text{None}$  then satisfy( $t^*, TC, P$ )
  until  $\frac{\alpha}{\beta} = 1$ 

```

Fig. 6. The procedure $satisfy(t, TC, P)$ satisfies a topology constraint $t \in TC$ that is at equality, by modifying P with a *valid* and *tight* system of segments. Procedure $move(x, \bar{x}, TC, P)$ updates the path P by moving nodes in one dimension from position x to \bar{x} to satisfy the topology constraints TC .

- (2) Edge routing is performed using the incremental poly-line connector routing library `libavoid` [16] to compute poly-line routes for each edge, which minimise edge length and amount of bend. An initial cluster boundary is obtained by taking the convex hull of the nodes in the cluster.

We note that the edge routing library has been extended to handle clusters and finds routes for edges that do not unnecessarily pass through clusters.

7 Experimental results

Table 1 gives some indicative run-times for finding an initial layout then using the topology-preserving constrained graph layout algorithm to find an optimal layout on various size graphs. The topology-preserving constrained graph layout algorithm is quite fast with less than two seconds required to layout networks of around 350 nodes. We have found that the main cost for each iteration is computation of the descent vector and step size. We also note that our experience with the algorithm in interactive applications is that it provides real-time updating of layout for graphs with up to 100 nodes.




Fig. 7. The result of each iteration of *move* is shown for a path with a horizontal segment. The iterations progress from left to right. The node v is required to move to the right relative to the other nodes. The four central nodes are shown slightly separated for clarity, but we assume that the boundaries of these nodes are actually touching—hence creating, initially, a horizontal segment. The small circles represent bend points, while the ‘-’s represent straight topology constraints. Note that, to properly preserve topology as the segments are split to satisfy a straight constraint, the remaining straight constraints must be transferred to the correct sub-segments.

V	E	Feasible layout		Optimize	Total
		Stage 1	Stage 2		
49	51	0.08	0.11	0.06	0.17
93	105	0.22	0.50	0.24	0.74
128	144	0.51	1.02	0.55	1.57
144	156	0.92	1.31	0.45	1.76
169	195	0.83	1.97	0.82	2.79
199	238	1.31	2.94	1.45	4.39
343	487	2.65	13.94	1.89	15.83

For each graph we give the number of nodes and edges. The number of separation constraints imposing downward edges is $|E|$. We give the time to find an initial feasible layout (Stage 1 and Stage 2) from a random starting configuration; and then to optimize the result using the topology preserving constrained graph layout algorithm. Optimization algorithms were set to terminate when the change in P -stress or stress was $< 10^{-5}$.

Table 1. Indicative running times for layout on an average (1GHz) PC for various size randomly generated directed networks with constraints imposing downward pointing edges. All times are in seconds.

Computing an initial layout is more expensive, and the dominating cost in finding an initial layout is finding the initial connector routing.

8 Conclusion

We have presented a constrained graph layout algorithm that preserves the topology of the initial layout. It supports network diagrams with poly-line connectors and arbitrary node clusters. It ensures that nodes do not overlap and that additional placement constraints on the layout remain satisfied. The algorithm is fast enough for real-time layout of networks with up to 100 nodes in two dynamic graph layout applications we have developed: a network diagram authoring tool and a network diagram browser. While the primary motivation for our development of the algorithm was to support dynamic layout it can also be used to improve layouts generated by other graph layout techniques.

One of the strengths of the algorithm is that can be straightforwardly modified to work with other goal function, so long as the second derivative is computable and the goal function is bend-point invariant. We plan to explore other

goal functions. We also plan to explore generalizing the algorithm to handle arbitrary linear constraints, not only separation constraints.

References

1. He, W., Marriott, K.: Constrained graph layout. *Constraints* **3** (1998) 289–314
2. Dwyer, T., Koren, Y., Marriott, K.: Drawing directed graphs using quadratic programming. *IEEE Transactions on Visualization and Computer Graphics* **12**(4) (2006) 536–548
3. Dwyer, T., Koren, Y., Marriott, K.: IPSep-CoLa: An incremental procedure for separation constraint layout of graphs. *IEEE Transactions on Visualization and Computer Graphics* **12**(5) (2006) 821–828
4. Gutwenger, C., Mutzel, P., Weiskircher, R.: Inserting an edge into a planar graph. In: *SODA '01: Proc. of the 12th Annual ACM-SIAM Symp. on Discrete Algorithms, Society for Industrial and Applied Mathematics* (2001) 246–255
5. Dwyer, T., Marriott, K., Wybrow, M.: Dunnart: A constraint-based network diagram authoring tool. In: *Proc. 16th Int. Symp. on Graph Drawing (GD'08)*. Submitted.
6. Dwyer, T., Marriott, K., Wybrow, M.: Exploration of networks using over view+detail with constraint-based cooperative layout. *IEEE Transactions on Visualization and Computer Graphics (InfoVis 2008)* To appear.
7. Misue, K., Eades, P., Lai, W., Sugiyama, K.: Layout adjustment and the mental map. *Journal of Visual Languages and Computing* **6**(2) (1995) 183–210
8. Frishman, Y., Tal, A.: Online dynamic graph drawing. In: *Eurographics/IEEE-VGTC Symp. on Visualization, Eurographics Association* (2007)
9. Bridgeman, S.S., Fanto, J., Garg, A., Tamassia, R., Vismara, L.: InteractiveGiotto: An algorithm for interactive orthogonal graph drawing. In: *GD 1997: Revised Papers from the 5th International Symposium on Graph Drawing, Berlin, Springer-Verlag* (1998) 303–308
10. North, S.C., Woodhull, G.: Online hierarchical graph drawing. In: *GD '01: Revised Papers from the 9th International Symposium on Graph Drawing, Berlin, Springer-Verlag* (2002) 232–246
11. Eiglsperger, M., Fekete, S.P., Klau, G.W.: Orthogonal graph drawing. In: *Drawing Graphs: Methods and Models*. Springer, London, UK (2001) 121–171
12. Bertault, F.: A force-directed algorithm that preserves edge crossing properties. In: *Proc. 7th Int. Symp. on Graph Drawing (GD'99)*. Volume 1731 of LNCS., Springer (1999) 351–358
13. Dwyer, T., Marriott, K., Wybrow, M.: Integrating edge routing into force-directed layout. In: *Proc. 14th Intl. Symp. Graph Drawing (GD '06)*. Volume 4372 of *Lecture Notes in Computer Science.*, Springer (2007) 8–19
14. Gansner, E., Koren, Y., North, S.: Graph drawing by stress majorization. In: *Proceedings of 12th Int. Symp. Graph Drawing (GD'04)*. Volume 3383 of *Lecture Notes in Computer Science.*, Springer (2004) 239–250
15. Kamada, T., Kawai, S.: An algorithm for drawing general undirected graphs. *Information Processing Letters* **31** (1989) 7–15
16. Wybrow, M., Marriott, K., Stuckey, P.J.: Incremental connector routing. In: *Proc. 13th Int. Symp. on Graph Drawing (GD'05)*. Volume 3843 of *Lecture Notes in Computer Science.*, Springer (2006) 446–457