

STRATISFIMAL LAYOUT: A modular optimization model for laying out layered node-link network visualizations

Sara Di Bartolomeo , Mirek Riedewald , Wolfgang Gatterbauer , and Cody Dunne 

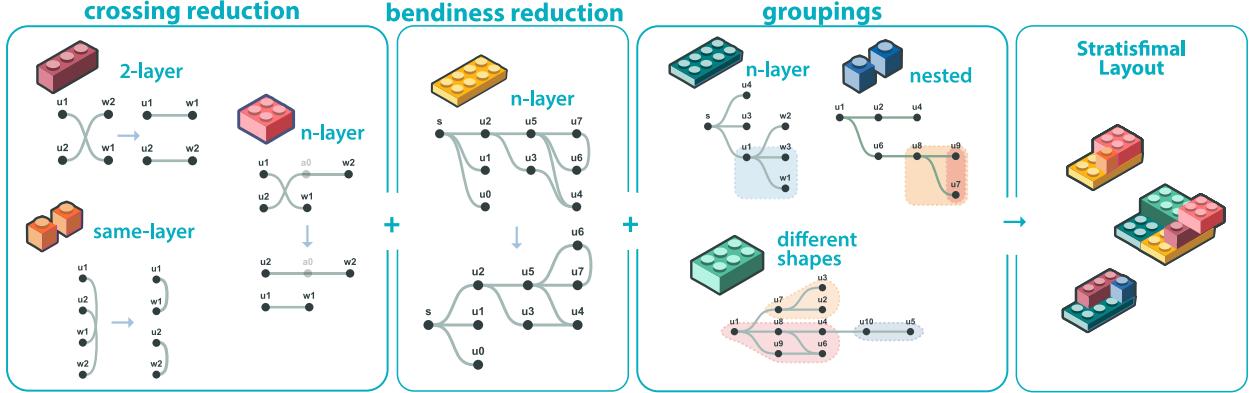


Fig. 1: Our goal is to make complicated networks easier to read. We achieve this by formulating readability criteria as a constrained optimization problem. Our optimization model for laying out layered node-link visualizations includes several modular, customizable components. Each addresses a different readability criteria or network feature — hence the building blocks metaphor. This modularity allows the layout to be tailored for diverse use cases. Here we illustrate some of the features of our Integer Linear Programming (ILP) formulation, which we call STRATISFIMAL LAYOUT — combining the words *stratified* (arranged in layers) and *optimal*.

Abstract—Node-link visualizations are a familiar and powerful tool for displaying the relationships in a network. The readability of these visualizations highly depends on the spatial layout used for the nodes. In this paper, we focus on computing *layered* layouts, in which nodes are aligned on a set of parallel axes to better expose hierarchical or sequential relationships. Heuristic-based layouts are widely used as they scale well to larger networks and usually create readable, albeit sub-optimal, visualizations. We instead use a *layout optimization model* that prioritizes *optimality* — as compared to *scalability* — because an optimal solution not only represents the best attainable result, but can also serve as a baseline to evaluate the effectiveness of layout heuristics. We take an important step towards powerful and flexible network visualization by proposing STRATISFIMAL LAYOUT, a *modular integer-linear-programming formulation* that can consider several important readability criteria *simultaneously* — crossing reduction, edge bendiness, and nested and multi-layer groups. The layout can be adapted to diverse use cases through its modularity. Individual features can be enabled and customized depending on the application. We provide open-source and documented implementations of the layout, both for web-based and desktop visualizations. As a proof-of-concept, we apply it to the problem of visualizing complicated SQL queries, which have features that we believe cannot be addressed by existing layout optimization models. We also include a benchmark network generator and the results of an empirical evaluation to assess the performance trade-offs of our design choices. A full version of this paper with all appendices, data, and source code is available at osf.io/qdyt9 with live examples at <https://visdunneright.github.io/stratisfimal/>.

Index Terms—Layered node-link visualization, integer linear programming, crossing reduction, bendiness reduction, nested groups.

1 INTRODUCTION

Networks are widely used across many disciplines to model entities and the relationships between them [37]. Specifically, a network consists of a finite set of nodes and a finite set of edges, each connecting two nodes. **Node-link visualizations** are a familiar and powerful tool for exposing the network topology (e.g., paths between nodes and highly-connected clusters) by drawing nodes as point marks and edges as connecting line marks [58]. To better show the topology, nodes can be assigned spatial coordinates based on the relationship structure.

We may wish to expose hierarchical or sequential relationships in a network. These can be displayed in a **layered node-link visualization**,

where each node is assigned to exactly one layer (rank) [73]. This layer assignment can be given a priori or obtained algorithmically. The visualization arranges the nodes in a layer along a linear axis and the axes in parallel. Here we use vertical lines for axes, but our methods work horizontally as well. The edges then fall into three categories: **same-layer edges** between nodes that share a layer, **2-layer edges** that connect nodes in adjacent layers, and **3+layer edges** which cross at least one intermediate layer between the incident nodes. In this context, we define **groups** as sets of nodes that must be placed adjacent to each other, both within the same layer and across consecutive layers. Each pair of groups must be either disjoint or one is a subset of the other.

For both general and layered node-link visualizations, assigning spatial coordinates to nodes is key for creating a readable visualization [4, 22, 73]. Optimally creating this assignment, called a visualization **layout**, based on readability criteria remains an open challenge. Just determining the minimal number of edge crossings is NP-complete [29]. The field of *graph drawing* focuses on this and other network visualization problems. Researchers have proposed many layout approaches — see overviews in [3, 4, 21, 31, 73] — but most are **heuristics** that prioritize **scalability** over **optimality**.

• The authors are all with Northeastern University.

E-mail: sdibartolomeo.s@northeastern.edu | m.riedewald@northeastern.edu | w.gatterbauer@northeastern.edu | c.dunne@northeastern.edu

Manuscript received 21 Mar. 2021; revised 13 June 2021; accepted 8 Aug. 2021.

Date of publication 1 Oct. 2021; date of current version 22 Dec. 2021.

8] [] h U ^ C V ^ Y Wh = X Y b h J U T S b c " % \$ " % % heuristics that prioritize scalability over optimality.

There are two main reasons, though, to search for an **optimal solution**: (1) An optimal layout can be used as a *baseline* against which heuristics can be compared, making it possible to quantify the trade-off between increased scalability and worse readability. (2) There are numerous “small” networks in practice where maximal readability is paramount. For such applications, the extra compute time to create an optimal layout to improve readability is worth the wait.

One such example is a metro map, which has few nodes and edges, but needs to be easily and quickly readable. The up-front layout computation time is negligible compared to the possible human and capital expense to deploy a new map.

Another example, which we use as a motivating case study, is helping users interpret database queries. SQL has been the standard language for writing database queries for decades [11], but it has been repeatedly criticized for its complexity and for the time it takes to read a query. Indeed, a query can have a number of nested subqueries, involve many tables in the database, or contain complex joins — all of which make keeping a mental map of the intermediate results of the operations contained in it a very difficult task [9]. Query visualizations, such as the recently-proposed QueryVis diagrams [17, 48], expose the logical structure of SQL queries using layered node-link visualizations. The size of the network QueryVis produces to model a complicated query is small, generally consisting of fewer than 30 nodes. Since one of the uses of QueryVis is to help students learn SQL, unnecessary readability issues should be eliminated as much as possible. Moreover, as a complicated query can take several minutes to understand, the layout computation time is negligible if it stays within a few seconds.

This paper focuses on the use of **Integer Linear Programming (ILP)** to create optimal layouts for layered node-link visualizations. ILP is a method to achieve the optimal solution for a mathematical model within the boundaries established by linear constraints. An ILP formulation is comprised by an objective function, which describes the goal, and a set of constraints. The fundamental challenge of solving a problem through ILP is figuring out a smart way to define these two components, which, together, describe a problem. While ILP is NP-complete, solvers have made tremendous advances over the years to a degree where many practical problems can be solved efficiently. Recent progress on solvers, together with a natural problem-size limit established by human ability to “consume” a network visualization, have renewed interest in finding exact criteria for optimal layout algorithms.

ILP formulations have already been explored for creating optimal layouts for node-link visualizations based on readability criteria [22, 71], including minimizing edge crossings [34, 35, 60, 81], minimizing edge bendiness [28], and contiguously grouping nodes within a single layer [81]. However, existing approaches do not handle same-layer edges, groups that span many layers, or nested groups — nor can these criteria be solved simultaneously.

Our goal is to provide a set of more expressive and modular formulas for readability optimizations that can be “cherry-picked” and flexibly composed like building blocks, depending on the target application and the reader’s needs. Figure 1 illustrates several of our modules and how they can be composed into an overall layout optimization model.

Contributions and Supplemental Materials

In this paper, we contribute:

1. **STRATISFIMAL LAYOUT, an optimization model for laying out layered node-link visualizations.** The model translates important readability criteria into a modular ILP formulation and layout algorithm. It includes:

- (a) **Edge crossing reduction** for 2-layer and 3⁺-layer edges as per Zarate et al. [81] and, for the first time, for same-layer edges.
- (b) **Edge bendiness reduction**, extending Gansner et al.’s formulation [28] to support modular composition with other criteria.
- (c) **Contiguous grouping of nodes** enclosed by concave or convex shapes. We improve efficiency vs. Zarate et al. [81] for groups within a single layer and, novelly, support groups that span many layers and nested groups.

Our **modular, customizable ILP formulation** enables tailoring the layout to diverse use cases. Each module is presented with an in-depth explanation and application example to illustrate the reasoning and trade-offs for including each feature.

2. **Open-source and documented layout implementations**, both as a web-based JavaScript library and for desktop computation via Gurobi [33]. Our model can serve as a layout algorithm directly or an optimal baseline for evaluating fast layout heuristics.
3. **A case study demonstrating the utility of the layout** for improving the readability of logical diagrams of SQL queries.
4. **A benchmark network generator** with parameters to control crucial network properties, including groups of nodes, and **the results of an empirical evaluation** on these networks that illustrate the performance trade-offs of our design choices.

A full version of this paper with all appendices, data, and source code is available at osf.io/qdyt9 with live examples at <https://visdunneright.github.io/stratisfimal/>.

2 MOTIVATING CASE STUDY: INTERPRETING SQL QUERIES

We introduce STRATISFIMAL LAYOUT in the context of a motivating case study on visualizing the logical structure of SQL queries. The goal is to improve the layout of QueryVis diagrams [48], which incorporate layered node-link visualizations and hierarchical grouping of nodes within and across layers. This application was our original motivation, as we know of no existing heuristic or optimization-based layout algorithm that can address all the necessary constraints simultaneously.

SQL (Structured Query Language) has been the standard query language for relational databases for decades [11]. Formulating or interpreting a non-trivial query can be challenging and time consuming, even for experts. SQL queries can be verbose, deeply nested, and involve complex logical constructs. Several attempts have been made to improve query readability, either by mapping SQL to natural language [43] or through visual query languages [10, 30, 38]. QueryVis [17, 48], the subject of our case study, focuses on displaying the underlying logic behind an SQL query. A controlled experiment found that existing SQL users were faster at correctly interpreting queries with QueryVis than with SQL alone, even after only a brief exposure to the visual language.

QueryVis models an SQL query as a network (see Figure 2, *Middle*). Nodes represent attributes, and edges represent relationships between them (e.g., joins). Each node belongs to a group for its database table; these groups are disjoint sets. Additional groups are used to show any hierarchical nesting of subqueries. This enables modeling query nesting and logical quantifiers (e.g., NOT EXISTS). This network is then displayed using a layered node-link visualization.

Optimal layouts for readability. The visualizations produced by QueryVis may still be challenging to interpret, even for well-trained users, and despite being comparatively easier to read than the original SQL queries. This is a consequence of the logic behind SQL, which may take long to understand for complicated queries (e.g., see the unique beer taste query in Figure 2, *Left*). In many cases, running a misinterpreted query can incur hours or even days of delay for the user. Making interpretation of a visualization as easy as possible is imperative in this situation.

For that reason, we would like to find an optimal layout (subject to the constraints imposed by the query logic). QueryVis originally used a heuristic-based layout: GraphViz’s *dot* [27], which is based on Gansner et al.’s edge crossing reduction algorithm for layered networks [28]. The latter is fast but does not guarantee minimal edge crossings, a widely accepted readability criteria. (See, e.g., Figure 5 and our appendix at osf.io/qdyt9). It also does not support groups, a key aspect of QueryVis, but *dot* can lay out groups by treating them as subproblems. This further compromises readability, especially as it does not reduce crossings of edges within a layer. The network for a complicated SQL query generally consists of fewer than 30 nodes. In this case, spending the time to compute an optimal layout for the visualization is clearly warranted.

Another case study based on StoryLines [32, 50] is in our appendix at osf.io/qdyt9.

```

SELECT
  FROM
    Likes L1
    WHERE
      NOT EXISTS
        (SELECT *
          FROM Likes L2
          WHERE
            L1.drinker <> L2.drinker
            NOT EXISTS
              (SELECT *
                FROM Likes L3
                WHERE L3.drinker = L2.drinker
                AND NOT EXISTS
                  (SELECT *
                    FROM Likes L4
                    WHERE L4.drinker = L1.drinker
                    L4.beer = L3.beer))
        AND NOT EXISTS
          (SELECT *
            FROM Likes L5
            WHERE L5.drinker = L1.drinker
            NOT EXISTS
              (SELECT *
                FROM Likes L6
                WHERE L6.drinker = L2.drinker
                L6.beer = L5.beer)))
  AND

```

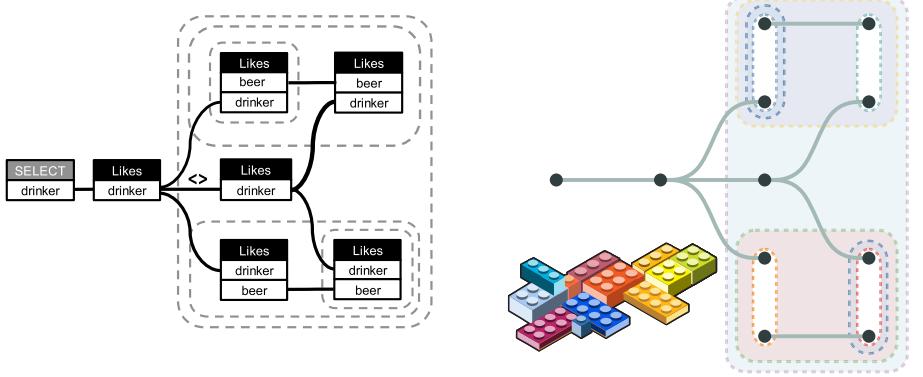


Fig. 2: Visualizing an SQL query (adapted from QueryVis, Fig. 1 [48]). *Left*: The SQL query for “find drinkers that like a unique set of beers.” *Middle*: The same query shown using QueryVis. Each table used is represented by a rectangle, subdivided vertically to show the title (black background) and relevant attributes (white background). The SELECT statement is similarly displayed, just with a gray background for the title. Curves show relationships between attributes (e.g., joins) and may be labeled with comparison operators (e.g., $<>$). NOT EXISTS statements are shown as dashed lines surrounding hierarchically-nested groups for subqueries. Layers are arranged left-to-right to show depth from the SELECT statement, and the contained nodes are aligned to a vertical line. *Right*: An abstract representation of the QueryVis network, simplified by omitting any layout-irrelevant information (e.g., title nodes). The layout is optimized for readability using our method, with circles for nodes/attributes and nested convex hulls for groups. We illustrate which building blocks are used in the layout formulation (their arrangement is arbitrary).

3 BACKGROUND AND RELATED WORK

Node-link visualization of networks. Networks are used to represent data in an innumerable amount of contexts, among which linguistics [12], cartography [2, 54, 79], computer networks [5], and social network analysis [16, 56, 70] are just a few examples [37]. Effectively visualizing a network can help a reader better understand relationships between elements in the data, find insights [36], and discover patterns and outliers [63]. Research on how to best represent a network as a node-link visualization started as early as 1934 with Moreno [56]. The discipline of representing networks visually is called *graph drawing* [4], though in this context we use the terminology of networks, node-link visualizations, and layouts of them.

Scientists have developed several readability criteria for visualizations [22, 66], and algorithms to implement those criteria in practice [22, 45]. Computational layout approaches began appearing in the 1960s [41, 75, 76]. Though most layout algorithm research has focused on heuristics (relevant examples: [28, 72]), there have been efforts at using neural networks [14, 46, 53] and Integer Linear Programming (ILP). The problem is still open and the general agreement is that there is no universal layout solution that embraces all the possible applications and use cases [6, 31].

Integer Linear Programming (ILP). Our STRATISFIMAL LAYOUT approach uses ILP to find an optimal layout based on readability criteria, rather than layout heuristics. ILP [18] is a generic optimization technique that can solve a wide range of combinatorial optimization problems. In order to use ILP for laying out a node-link visualization, a composition of a set of readability criteria needs to be translated into a quantitative objective function, such as minimizing the number of edge crossings, minimizing edge bendiness, and creating uniform edge length. Additional constraints may be specified in the formulation. The formulation can then be fed to an ILP solver, e.g., GLPK.js¹ or Gurobi [33], in order to find the optimal solution.

We integrate a discussion of the literature most relevant to our optimization model formulation as part of Section 5, but here mention a closely-related problem: metro map layouts. Nöllenburg and Wolff [62] used a mixed-integer linear program to lay out a metro map, optimizing for uniform edge length and avoiding bends. They proposed constraints that enforce that the solution keeps the original topology intact, so that readers can easily trace back the network topology to the actual geography of the represented region. Metro maps have been a popular subject for this kind of approach, e.g., [2, 54, 79].

Crossing reduction: Minimizing edge crossings is one of the best known and important optimization criteria for node-link visualization

layout algorithms, as crossings clearly affect network readability for path-following tasks [66]. Buchheim et al. first proposed a branch-and-cut approach [8] to reduce the number of variables involved in an integer-linear-programming problem for crossing-restricted networks (networks that allow at most one intersection between two edges) by dividing the network into smaller subproblems. A similar approach is also used in [40]. Chimani et al. later worked on improving the efficiency of a similar method based on column generation [13]. The crossing minimization problem has also been reduced to a planarization problem [59], at first with two layers [60], then on k-level networks [34, 35]. Gange et al. [26] use ILP to obtain crossing reduction and minimal edge deletion for planarization including an optimization based on $K_{2,2}$ graphs [44]. Alternatively, we consider the entire graph for the layout as well as edge bendiness and groups. An extensive survey on crossing minimization in node-link visualizations can be found in [73].

The most relevant approach to discuss is Zarate et al.’s optimization model for layered node-link visualizations [81], which can minimize crossings of 2-layer and 3⁺-layer edges. Similar to Zarate et al. and without loss of generality, we assume that all edges are same-layer or 2-layer — we show in Section 5.1.6 how to “break up” 3⁺-layer edges using dummy anchor nodes to satisfy this condition. We based our initial formulation for minimizing edge crossings on Zarate et al.’s, but further extend it to support same-layer edges.

Bendiness reduction: *Edge bendiness* is another important layout optimization criteria for node-link visualizations. That is to say, having predominantly straight edges versus edges that have bends, curves, or have to be drawn diagonally. Low edge bendiness is important for users performing path-following tasks [66]. Consistent with previous work on layered networks with vertically aligned layer axes, we define an edge to be “bent” if and only if it is not horizontal, i.e., its ends have different y coordinates. The first optimization model solution by Sugiyama et al. [72] formalized the optimization objective using the squared difference in vertical coordinates between the two end nodes of an edge: $(y_{u_1} - y_{w_1})^2$. This quadratic formulation makes the optimization expensive. Gansner et al. [28] simplified the optimization objective into a linear function, though they also added a new variable and two constraints for each edge. Both Sugiyama et al. and Gansner et al. counterbalance the costliness of their solutions by additionally providing a heuristic-based layout algorithm as an alternative. There are also several other heuristic-based approaches for optimizing edge bendiness [7, 24, 69]. Our formulation here extends Gansner et al.’s [28] approach to support modular composition with other readability criteria. Optimizing for horizontal edges requires modifications to the objective function and introduction of variables that capture the vertical position (i.e., the y coordinate) of a node and the resulting bendiness of an edge.

¹<https://github.com/hgourvest/glpk.js>

This paper generalizes several previous optimization model approaches by defining a modular mathematical framework that allows for combining and optimizing several readability criteria *simultaneously* — minimizing edge crossings, minimizing edge bendiness, and contiguously grouping nodes. These criteria can be mixed and matched to adapt to different use cases. Moreover, we contribute extensions to previous ILP-based approaches [28, 81] to novelly support minimizing edge crossings for same-layer edges, add support for groups that span multiple layers, and add support for nested groups.

4 READABILITY OBJECTIVES AND CONSTRAINTS

Based on an analysis of previous work (e.g., [22, 66, 67, 71]) and our motivating case study on QueryVis (Section 2), we identified the following objectives and constraints necessary for creating an optimal, readable layout for a layered node-link visualization.

Optimization objectives: These are the metrics we aim to optimize.

- **Minimize edge crossings** [22, 66]: In a 1997 study [66], Purchase found that crossing edges are the main characteristic negatively affecting network readability. Placing nodes so that the number of crossings is minimized is thus given the highest priority
- **Minimize edge bendiness** [22, 66]: Edge bendiness is considered a criterion negatively affecting readability as well. In a layered node-link visualization with parallel layers, we define an edge to be “bent” if its points do not lie on a line perpendicular to the layers.

Hard constraints: These conditions must be valid in any context.

- **No node-node or node-edge overlaps** [22]: Nodes cannot be drawn on top of each other and no edge can lie over a node.
- **Groups cannot enclose non-member nodes**: Only nodes that are members of a group can be drawn within the area enclosed by the group’s mark.

Soft constraints: Nice-to-have features, but not strictly necessary.

- **Rectangular groups** [77, 80]: Groups can be constrained to fit in a rectangular area.
- **Minimize group area**: The area enclosed by a group should be as small as possible, given any necessary whitespace padding.

5 THE STRATISFIMAL LAYOUT OPTIMIZATION MODEL

We now discuss how to formalize all these varied optimization objectives and constraints for laying out a layered node-link visualization. Our optimization model formulation uses Integer Linear Programming (ILP) and is modular — depending on the desired visualization properties, the reader can combine the corresponding features. An extensive demonstration of our approach on real networks can be found at <https://visdunneright.github.io/stratisfimal/proofs.html>.

How to read this section: Each subsection will detail a feature using a modular ILP formulation, which the reader can simply add to the overall layout optimization model. The actual formulation is marked by an image of a building block. Relevant information about the formulation is shown in an accompanying gray box. The *Intuition* behind the constraints explains the idea more plainly (but imprecisely). The *Number of constraints generated* provides some indication of how computationally expensive it would be to add the feature to the optimization model. Finally, for cases where constraints build atop each other, all the *Prerequisites* from other subsections are shown. Please refer to Figure 3 for an explanation of the notation we use.

Preprocessing: Our method assumes an already-established and immutable layer assignment for the nodes of the input network. Layer assignment may be given a priori, such as the query depth in our motivating case study (Section 2). Alternatively, layers can be assigned algorithmically based on network topology, e.g., using approach detailed by Gansner et al. [28], Rüegg et al.’s extension of it [68], or the ILP formulation by Tang & Hu [74].

5.1 Crossing reduction

5.1.1 Objective function for crossing reduction

Crossings can only occur between edges in the set E_k . The following formula minimizes the number edge crossings:

Definitions:	
$G = \{N, E\}$	The network (graph) consists of a set of nodes N and edges E .
N_k	The nodes in layer k .
$E_k^=$	Same-layer edges with both ends in layer k .
$E_k^<$	2-layer edges with one end in layer k and the other in layer $k+1$.
$E_k = E_k^= \cup E_k^<$	2-layer edges (layer k to $k+1$) and same-layer edges (only layer k).
$L = \{1, 2, \dots, \ell\}$	The set of ℓ layers in G .
$u_1 w_1$	An edge between nodes u_1 and w_1 .
Γ	The set of groups in G .
g_1	A group; $g_1 \in \Gamma$.
L_{g_1}	Set of layers in which g_1 has at least one node.

Decision variables:	
x_{u_1, u_2}	The relative vertical order of nodes. Boolean equal to 1 if u_1 is above u_2 , 0 otherwise.
$c_{u_1 w_1, u_2 w_2}$	Indicates if edges $u_1 w_1$ and $u_2 w_2$ cross. Boolean equal to 1 if they cross, 0 otherwise.
y_{u_1}	Vertical coordinate of node u_1 .
$b_{u_1 w_1}$	Bendiness of edge $u_1 w_1$, defined as $ y_{u_1} - y_{w_1} $.
$y_{g_1}^T$	Topmost boundary of group g_1
$y_{g_1}^B$	Bottommost boundary of group g_1

Parameters:	
m	Greatest allowed absolute vertical distance between the topmost and bottommost nodes in a layer. Default: 50, but depends on implementation details.
γ_1, γ_2	Weights for the two parts of the objective function. Default: γ_1 (crossings) = 10, γ_2 (bendiness) = 1.

Fig. 3: The notation used in this paper.

$$\text{Minimize} : \sum_{k \in L} \sum_{\substack{u_1 w_1, u_2 w_2 \in E_k \\ u_1 w_1 \neq u_2 w_2}} c_{u_1 w_1, u_2 w_2} \quad (1)$$

5.1.2 Transitivity constraints

Laying out nodes in a layer establishes a total order of nodes within it. We enforce that the ILP solution is consistent with such per-layer orders by defining transitivity constraints for each triplet of nodes in a layer. The idea is that if node u_1 is above node u_2 and node u_2 is above node u_3 , then node u_1 must be above node u_3 . The opposite must also be true. We can write this as $x_{u_1, u_2} \wedge x_{u_2, u_3} \Rightarrow x_{u_1, u_3}$ and $\neg x_{u_1, u_2} \wedge \neg x_{u_2, u_3} \Rightarrow \neg x_{u_1, u_3}$. These formulas, translated to ILP constraints, are written thus:

$$\begin{aligned} x_{u_1, u_2} + x_{u_2, u_3} - x_{u_1, u_3} &\geq 0 \\ -x_{u_1, u_2} - x_{u_2, u_3} + x_{u_1, u_3} &\geq -1 \end{aligned} \quad (2)$$

$$(\forall k \in L : \forall u_1, u_2, u_3 \in N_k, \text{ where } u_1 \neq u_2 \neq u_3 \neq u_1)$$

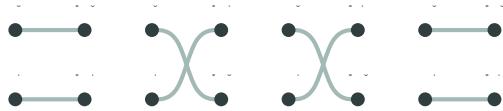
Intuition: If u_1 is above u_2 , and u_2 is above u_3 , then u_1 must be above u_3 .

Number of constraints generated: $O(|N|^3)$

Note that, in any one of the constraints, x_{u_1, u_2} can always be written as $1 - x_{u_2, u_1}$. This reduces the number of variables needed in a model.

5.1.3 2-layer edge crossings

Crossing indicators c are computed by comparing position indicators x for each pair of 2-layer edges in a layer. There are 4 possible cases, captured by the visualization and formulas below.



The visualization shows that there is a crossing if the two ends of the edges are in opposite order on the first and on the second layer. This can be written as $x_{u_1, u_2} \wedge \neg x_{w_1, w_2} \Rightarrow c_{u_1 w_1, u_2 w_2}$ and $\neg x_{u_1, u_2} \wedge x_{w_1, w_2} \Rightarrow c_{u_1 w_1, u_2 w_2}$, which in turn translates to:

$$\begin{aligned} x_{u_1, u_2} - x_{w_1, w_2} - c_{u_1 w_1, u_2 w_2} &\leq 1 \\ -x_{u_1, u_2} + x_{w_1, w_2} - c_{u_1 w_1, u_2 w_2} &\leq 1 \end{aligned} \quad (3)$$

$$(\forall k \in L : \forall u_1 w_1, u_2 w_2 \in E_k^<, \text{ where } u_1 w_1 \neq u_2 w_2)$$

Intuition: There is a crossing ($c_{u_1 w_1, u_2 w_2} = 1$) if the ends of a 2-layer edge have opposite positions in the two adjacent layers: either u_1 is above u_2 but w_1 is below w_2 or vice versa.

Number of constraints generated: $O(|E|^2)$

Prerequisites: 5.1.2 Transitivity constraints

We provide a more detailed explanation of this constraint, which can be found in our appendix at osf.io/qdyt9.

5.1.4 Same-layer edge crossings

We need additional constraints to detect crossings of same-layer edges. There are $4! = 24$ permutations of 4 nodes in a layer; 8 of these permutations can produce a crossing between two same-layer edges. The 8 formulas of this module, illustrated and detailed below, each encode one such permutation:



$$\begin{aligned} c_{u_1 w_1, u_2 w_2} + x_{u_2, u_1} + x_{w_1, u_2} + x_{w_2, w_1} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{w_2, u_1} + x_{w_1, w_2} + x_{u_2, w_1} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{u_1, u_2} + x_{w_2, u_1} + x_{w_1, w_2} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{w_1, u_2} + x_{w_2, w_1} + x_{u_1, w_2} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{u_2, w_1} + x_{u_1, u_2} + x_{w_2, u_1} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{w_2, w_1} + x_{u_1, w_2} + x_{u_2, u_1} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{w_1, w_2} + x_{u_2, w_1} + x_{u_1, u_2} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{u_1, w_2} + x_{u_2, u_1} + x_{w_1, u_2} &\geq 1 \end{aligned} \quad (4)$$

$$(\forall k \in L : \forall u_1 w_1, u_2 w_2 \in E_k^=, \text{ where } u_1 w_1 \neq u_2 w_2)$$

Intuition: There is a crossing between two same-layer edges if nodes are in one of 8 specific orderings. An example: there is a crossing if u_1 is above u_2 , u_2 is above w_1 , and w_1 is above w_2 .

Number of constraints generated: $O(|E|^2)$

Prerequisites: 5.1.2 Transitivity constraints

5.1.5 Crossings between same-layer and 2-layer edges

Detecting crossings between same-layer and 2-layer edges requires yet another type of constraint:



$$\begin{aligned} c_{u_1 w_1, u_2 w_2} + x_{u_1, u_2} + x_{w_2, w_1} &\geq 1 \\ c_{u_1 w_1, u_2 w_2} + x_{u_1, w_2} + x_{u_2, u_1} &\geq 1 \end{aligned} \quad (5)$$

$$(\forall k \in L : \forall u_1 w_1 \in E_k^= : \forall u_2 w_2 \in E_k^<)$$

Intuition: There is a crossing between a same-layer and a 2-layer edge if an end of the 2-layer edge is in-between the two ends of the same-layer edge.

Number of constraints generated: $O(|E|^2)$

Prerequisites: 5.1.2 Transitivity constraints

5.1.6 3⁺-layer edges

Similar to Chimani et al. [13] and Zarate et al. [81], we split any 3⁺-layer edge, i.e., an edge whose endpoints are in layers i and $i+n-1$, for $n > 2$, into a path consisting of $n+1$ 2-layer edges between adjacent layers. This is done during a pre-processing step, so that the ILP formulation only has to deal with 1- and 2-layer edges, eliminating the need for more complex constraint types for 3⁺-layer edges. (The pseudocode for this step can be found in our appendix at osf.io/qdyt9.) The added nodes in the intermediate layers are called *anchors*, and they will not appear in the visualization. Breaking up 3⁺-layer edges increases the cardinality of E and N , thus increasing the number of constraints accordingly. The picture below shows an example effect of anchors used in combination with crossing reduction. Anchors are represented with semi-transparent nodes that are added here only for illustration.



5.2 Bendiness reduction

5.2.1 Objective function for bendiness reduction

For bendiness reduction, we add a term to the objective function to account for bent edges. Parameters γ_1 and γ_2 are set by the user to weight the relative importance of minimizing edge crossings vs. bendiness. Note that in our formulation, the bendiness $b_{u_1 w_1}$ of an edge $u_1 w_1$ is defined as the absolute difference $|y_{u_1} - y_{w_1}|$ between the vertical coordinates of its endpoints. Other formulations with more complex distance measures (e.g., squared difference) are possible, but may fall outside the capabilities of ILP and hence are left for future work.

$$\text{Minimize} : \gamma_1 \sum_{k \in L} \sum_{u_1 w_1, u_2 w_2 \in E_k^=, u_1 w_1 \neq u_2 w_2} c_{u_1 w_1, u_2 w_2} + \gamma_2 \sum_{u_1 w_1 \in E} b_{u_1 w_1}$$

5.2.2 Vertical position constraints

For nodes in the same layer, we must ensure that their vertical positions (y variables) are consistent with the ordering established by the x variables: if $x_{u_1, u_2} = 1$ then $y_{u_1} < y_{u_2}$; and if $x_{u_1, u_2} = 0$ then $y_{u_1} > y_{u_2}$. These implications can be expressed as $y_{u_1} > y_{u_2} x_{u_1, u_2}$ and $y_{u_2} > y_{u_1} x_{u_1, u_2}$, respectively. (Note that $x_{u_1, u_2} = 1 - x_{u_2, u_1}$.) Since ILP does not allow non-linear constraints — like those involving the product of 2 variables — we need to convert them to equivalent linear constraints. For the specific structure of our problem, we identified an effective conversion detailed by Coelho [15]. The conversion introduces an auxiliary variable z_i for each node u_i and requires a predefined constant, m , that limits the maximal vertical span (i.e., the greatest absolute difference between any pair of y coordinates) of the layer:



$$\begin{aligned}
 z_i - m \cdot x_{u_2, u_1} &\leq 0 \\
 z_i - y_{u_2} + m \cdot x_{u_2, u_1} &\geq -m \\
 y_{u_1} - z_i - 1 \cdot x_{u_2, u_1} &\geq 0 \quad (\forall k \in L : \forall u_1, u_2 \in N_k) \\
 z_i &\leq y_{u_2} \\
 z_i &\geq 0
 \end{aligned} \tag{6}$$

Intuition: Vertical node coordinates must be consistent with their vertical position ordering.

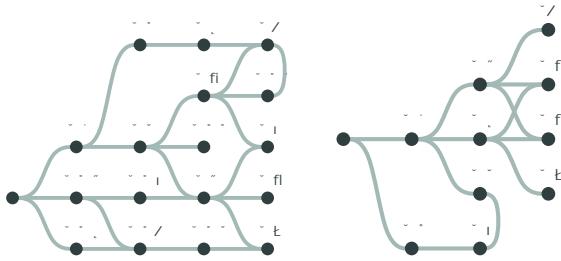
Number of constraints generated: $O(|N|^2)$

Prerequisites: 5.1.2 Transitivity constraints,

Equation (6) contains a term $-1 \cdot x_{u_2, u_1}$, which defines the distance between two nodes and can be adjusted by increasing the constant factor.

5.2.3 Computing edge bendiness

We define the bendiness of an edge as the absolute difference between the vertical coordinates of its endpoints. Unfortunately the absolute function is not allowed in ILP constraints as the resulting function would be non-linear. We therefore convert each constraint $b_{u_1 w_1} \geq |y_{u_1} - y_{w_1}|$ to the following equivalent linear constraints:





$$\begin{aligned}
 y_{u_1} - y_{w_1} &\leq b_{u_1 w_1} \quad (\forall u_1 w_1 \in E) \\
 y_{w_1} - y_{u_1} &\leq b_{u_1 w_1}
 \end{aligned} \tag{7}$$

Intuition: The bendiness of every edge is the difference between the y variables of the two ends of the edge.

Number of constraints generated: $O(|E|)$

Prerequisites: 5.2.2 Vertical position constraints

5.3 Grouping

We now discuss how to lay out enclosed and contiguous groups of nodes and the associated constraints we developed.

Showing contiguous groups of nodes within a single layer is useful in many contexts. E.g., the proportionally sized groups in each layer in Sankey diagrams and Di Bartolomeo et al.'s Sequence Braiding visualization for temporal event sequence data — both showing an overview of flow patterns. In our motivating case study on QueryVis (Section 2, Figure 2), groups are used to display tables of attributes contiguously within a layer, akin to an entity relationship (ER) diagram. Thus, Zarate et al.'s ILP approach for laying out Sankey diagrams [81] includes a simple constraint to support such single-layer groups.

However, showing any hierarchical subqueries in an SQL query additionally requires support for groups that (1) **span multiple layers** and (2) **can be nested**. To the best of our knowledge, there are no optimization models for laying out layered node-link visualizations that have multi-level or nested groups. Providing this support in STRATISFIMAL LAYOUT is one of the main technical innovations of this work.

For ease of presentation and without loss of generality, the discussion here focuses on rectangular group shapes. Our framework supports more general shapes by removing the corresponding “rectangularity” constraints below, as illustrated near the end of Section 5.3.4.

5.3.1 Single-layer groups

We are not aware of any prior framework that could seamlessly deal with groups, especially in combination with crossings and bendiness. Importantly, as the experiments will confirm, in our framework groups do *not* fundamentally increase the asymptotic complexity of the optimization problem.

To enforce adjacent placement of group members in a single layer, it is sufficient to specify that if two nodes u_1 and u_2 are in the same group, then no node u_3 not belonging to that group can be placed between them. The constraint below equivalently expresses that u_3 has to be either above both u_1 and u_2 , or below both.

$$x_{u_1, u_3} - x_{u_2, u_3} = 0 \tag{8}$$

($\forall g \in \Gamma$, where Γ is a single-layer group : $\forall u_1, u_2 \in g : \forall u_3 \notin g$)

Intuition: If nodes u_1 and u_2 are in group g_1 , and node u_3 is not in g_1 , then either u_3 is above both u_1 and u_2 , or below both of them.

Number of constraints generated: Reduces the number of constraints. For every single-layer group constraint we add, 2 transitivity constraints are removed.

Prerequisites: 5.1.2 Transitivity constraints

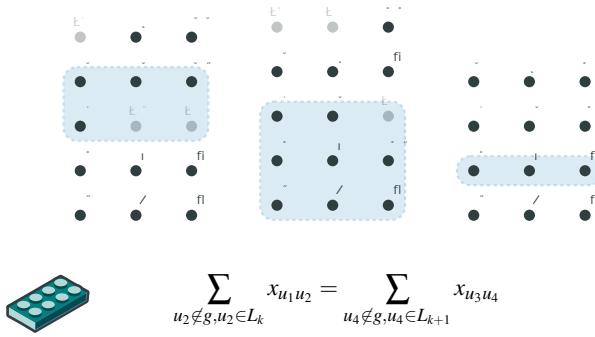
Notice that Equation (8) implies the 2 transitivity constraints in Equation (2), meaning that after adding Equation (8), *both implied constraints can be removed without changing the solution*. We show the implication for $x_{u_1, u_2} + x_{u_2, u_3} - x_{u_1, u_3} \geq 0$; the other one is analogous: From $x_{u_1, u_3} - x_{u_2, u_3} = 0$ follows $x_{u_2, u_3} - x_{u_1, u_3} = 0$, which together with $x_{u_1, u_2} \geq 0$ yields the desired $x_{u_1, u_2} + x_{u_2, u_3} - x_{u_1, u_3} \geq 0$.

5.3.2 Multi-layer groups

When a group spans multiple layers, Equation (8) is not sufficient. In addition to adjacent node placement within each layer, the vertical positions across layers needs to be taken into account. To this end, we could follow an approach similar to edge bendiness by penalizing differences in vertical coordinates between group members in adjacent layers. Since that is straightforward, we instead demonstrate our framework's flexibility by enforcing **rectangular multi-layer groups**. A rectangle is a simple easy-to-read convex hull [57], popular for enabling easy identification of group boundaries and group membership via enclosure [77, 80].

To achieve rectangular groups, we use a trick: for a given group, we determine the maximum number of group members in each layer and then add “filler nodes” to ensure the same number of group members across all layers it spans. We apply this procedure to all groups, then analogously ensure that all layers have the same total number of nodes. The pseudocode for this step is in our appendix at osf.io/qdyt9.

With the filler nodes in place, we only need to ensure that the top-most node of the group in each layer has the same y coordinate. Our ILP framework does this by “counting” the nodes above a group through the x variables. Consider a group g_1 with nodes u_1 and u_2 in some layer. A non-group node u_3 is above the group if and only if $x_{u_1, u_3} = 1$, or, equivalently, $x_{u_2, u_3} = 1$. In general, since Equation (8) enforces that all group members in a layer must be adjacent to each other, we only need to select one group member node per layer to count the number of non-group nodes above it in that layer. This count must be consistent across all layers the group spans:



In the above formula, u_1 is an arbitrary member of group g in layer k and u_3 is an arbitrary member of group g in the next layer $k+1$, i.e., $u_1 \in g \wedge u_1 \in L_k \wedge u_3 \in g \wedge u_3 \in L_{k+1}$.

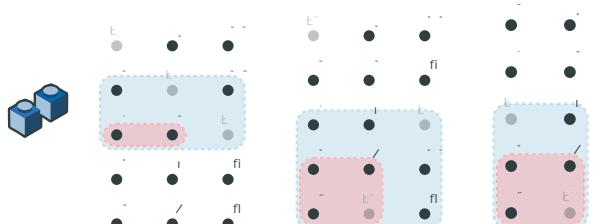
Intuition: The number of non-group nodes above the group must be the same across all layers the group spans.

Number of constraints generated: $O(\ell \cdot |\Gamma|) = O(\ell \cdot |N|)$

Prerequisites: 5.1.2 Transitivity constraints, 5.3.1 Single-layer groups

5.3.3 Nested multi-layer groups

Nested groups play an important role in hierarchical networks, including our motivating case study of visualizing SQL queries. There, the relative scopes of tables are nested according to the nesting structure of the query (see Section 2). Interestingly, the ILP constraints introduced so far can, by design, already handle nested groups. We simply apply the group constraints discussed above (Equations (8) and (9)) recursively such that the subgroup (red box in the diagram below) takes on the role of the group and the enclosing group (blue box) takes on the role of the entire network. This recursive construction is applied by starting with the innermost group. The pseudocode for this procedure is described in our appendix at osf.io/qdty9.



5.3.4 Multi-layer groups with vertical coordinates

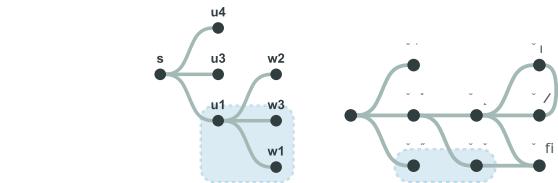
Notice that the ILP constraints for multi-layer groups in Equations (8) and (9) can be used regardless of whether the problem otherwise requires y coordinates (for bendiness reduction). We now show an alternative formulation that relies on the y coordinates and has the major benefits that no filler nodes are needed and that additional requirements, e.g., regarding the shape and area of non-rectangular groups, can be expressed easily. To this end, we introduce two new variables, y_g^T and y_g^B , for every group $g \in \Gamma$ to represent the top and bottom boundary of the group, respectively. For a rectangular group, all group members must be within those boundaries:

$$\begin{aligned} y_{u_1} &\geq y_g^T & (\forall g \in \Gamma : \forall u_1 \in g) \\ y_{u_1} &\leq y_g^B \end{aligned} \quad (10)$$

Intuition: All nodes in a group must fall between top and bottom boundary of the group.

Number of constraints generated: $O(|N|)$

Prerequisites: 5.2.2 Vertical position constraints, 5.3.1 Single-layer groups



Similarly, all nodes not in the group must fall outside the group boundaries in all layers spanned by the group. The direct encoding — enforcing for a node u_2 outside the group that either $y_{u_2} < y_g^T$ or $y_{u_2} > y_g^B$ — requires disjunction (“OR”), which is not supported by ILP. We therefore convert the disjunction to:

$$y_{u_2} - m \cdot x_{u_1, u_2} < y_g^T \quad (11)$$

$$-y_{u_2} + m \cdot x_{u_1, u_2} < m - y_g^B \quad (12)$$

$$(\forall g \in \Gamma : \forall k \in L_g : \forall u_2 \in L_k, \text{ where } u_2 \notin g)$$

Recall that m denotes the largest vertical coordinate difference allowed within a layer. Node u_1 is an arbitrary node in layer k that belongs to group g . (Since group nodes in the same layer are adjacent to each other due to Equation (8), any one of them being above u_2 implies that all of them are.)

Intuition: In a layer spanned by the group, nodes outside the group should be either above or below the group.

Number of constraints generated: $O(|N| \cdot |\Gamma|)$

Prerequisites: 5.2.2 Vertical position constraints, 5.3.1 Single-layer groups

To prevent the solver from setting all top and bottom group boundaries to 0 (the lowest possible y value) and m (the highest possible y value), respectively, we need to add constraints that force the group boundaries to not overlap. This again introduces disjunctions (to not overlap, in each layer the top boundary of one group must be below the bottom boundary of the other, or vice versa), which we transform to ILP-supported linear constraints:

$$\begin{aligned} y_{g_2}^B - m \cdot x_{u_2, u_1} - y_{g_1}^T &< 0 \\ -y_{g_2}^T + m \cdot x_{u_2, u_1} + y_{g_1}^B &< m \end{aligned} \quad (13)$$

$$(\forall g_1, g_2 \in \Gamma, \text{ where } g_1 \neq g_2 \wedge L_{g_1} \cap L_{g_2} \neq \emptyset)$$

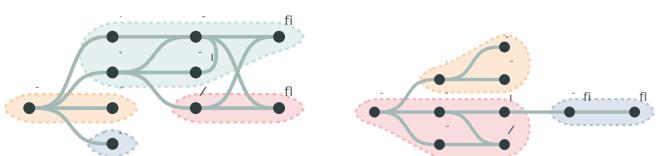
Notice that the constraints are only introduced for groups that contain nodes in at least one common layer. Here $u_1 \in g_1$ and $u_2 \in g_2$ refer to arbitrary group members from one of the common layers.

Intuition: For each two groups that have in common at least one layer, either the top boundary of one is below the bottom boundary of the other, or vice versa.

Number of constraints generated: $O(|\Gamma|^2)$

Prerequisites: 5.2.2 Vertical position constraints, 5.3.1 Single-layer groups

Non-rectangular groups. The y -coordinate variables enable more flexible group-shape constraints. For non-rectangular groups, one can replace variables y_g^T and y_g^B with per-layer boundaries like y_{gk}^T for each layer k spanned by group g . Or the user may decide for a group to have a common bottom boundary y_g^B , but flexible per-layer top boundaries y_{gk}^T . In particular cases like the ones below, in which all the nodes in the network are enclosed in completely distinct groups, flexible group shapes can simply be attained by leaving out constraints 11 and 12.



Constraining the area of a group. So far, none of the constraints forces a group to be “compact,” e.g., there could be a large gap between adjacent group members in the same layer. The introduction of variables for vertical coordinates offers an easy fix for this issue through constraints on the difference between top and bottom boundary of a group. For example, let $groupsize$ be a user-controlled parameter, then we can constrain group area as follows:



$$y_g^B - y_g^T \leq groupsize \quad (\forall g \in \Gamma) \quad (14)$$

Intuition: For each group, the difference between top and bottom boundary must not exceed a given limit.

Number of constraints generated: $O(|\Gamma|)$

Prerequisites: 5.2.2 Vertical position constraints, 5.3.4 Multi-layer groups with vertical coordinates

Note that one could similarly specify a different limit for each group. In all cases, the limit must be at least as high as the number of group members in a single layer. If it is set too low, the ILP solver will not find a solution and needs to be re-started with a higher limit.

5.4 Optimizing the number of variables and constraints

Complement variables. We cut the number of order-indicator variables in half by removing complement variables: instead of using both x_{u_1, u_2} and x_{u_2, u_1} , we use only x_{u_1, u_2} and replace x_{u_2, u_1} with $1 - x_{u_1, u_2}$.

Collapsing groups. We propose to exploit groups to reduce the *effective* network size, and thus optimization cost, for all types of constraints by essentially collapsing groups into single nodes. A special case of this idea is in Section 5.3.1, where each group-related constraint enabled the removal of 2 transitivity constraints. For the general approach, consider Figure 4. The “inclusion network” on the right encodes the hierarchy of groups and subgroups from the network on the left, making each subgroup a child of the directly enclosing group.

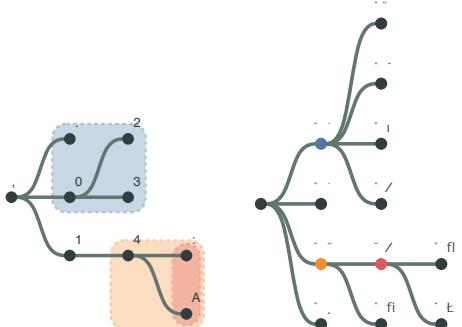


Fig. 4: A network with nested groups and its corresponding inclusion network. Colored nodes on the right represent groups on the left.

In the inclusion network, comparable elements are children of the same node. It is easy to show that every constraint involving inclusion-network nodes that are not siblings can be removed without affecting the ILP solution. E.g., in Figure 4, we can replace x_{u_6, u_2} and x_{u_6, u_3} with x_{u_6, g_1} , as node u_6 can only be above or below both nodes in g_1 . For networks with large nested groups, this can vastly reduce the number of variables and constraints, resulting in lower optimization cost.

6 IMPLEMENTATION AND EVALUATION

We implemented two versions of STRATISFIMAL LAYOUT: (1) a web-based JavaScript library using GLPK.js² as solver and (2) a formulation for desktop computation via Gurobi [33]. Our implementations, website, and supplemental material can be found at osf.io/qdty9 with a live demo also available online at <https://visdunneright.github.io/stratisfimal/>.

The JavaScript implementation allows for easy distribution of web-based visualizations built upon STRATISFIMAL LAYOUT, requiring no server-side computation or local installation. The formulation of the problem, the solving process, and the visualization can all be done seamlessly client-side in the browser. The downside of using a JavaScript solver is reduced performance. But most of the networks in our target size (a few dozen nodes) can still be solved quickly.

The website also contains a network generator with parameters that can be tweaked and processed on the spot, a page³ containing more extensive explanations and examples of the application of the modules to the network, comparisons of the application of our method compared to other layout algorithms, and several examples and benchmarks.

Applying the method to QueryVis on the web: We provide initial validation of our approach using of our motivating case study on QueryVis (Section 2). We conducted an experimental evaluation which compared STRATISFIMAL LAYOUT with Gansner et al.’s [28] heuristic-based layout based on computation time and the number of crossings. Bendiness reduction was omitted other than a very simple postprocessing step which did not impact our measurements. We used the constraints from sections 5.1.2–5.1.5 and 5.3.1. At first, a network is randomly generated (Figure 5a), then the results from the application of Gansner et al.’s method (Figure 5b) and our method (Figure 5c). While the original, randomly-generated network had 30 crossings, Gansner is able to reduce the number of crossings to 8 in 76 ms, while our method reduces the number of crossings to 2 in 191 ms. Comparisons on a number of different networks and network sizes are available on our website and appendices. As expected, our method takes more time but finds an optimal solution — demonstrating the utility of STRATISFIMAL LAYOUT for improving QueryVis visualizations.

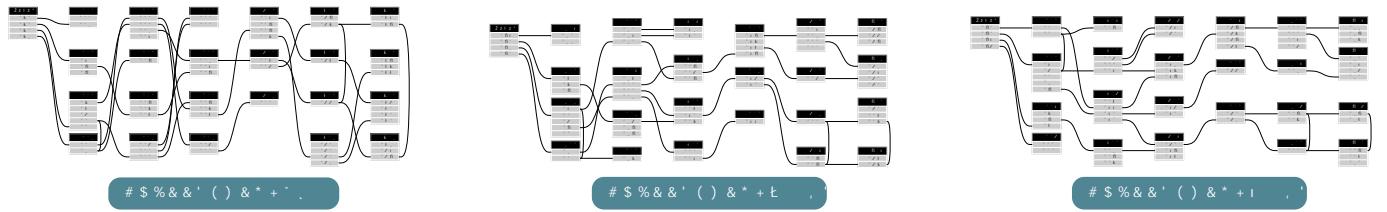
The values reported in Figure 5 are obtained by running the comparison in Firefox 86, on a 2020 MacBook with 32 GB RAM and an Intel Core i5-1038NG7 CPU. We wrote both layouts in JavaScript and run entirely on the client, including the ILP solver GLPK.js.

Rome-Lib benchmark on the desktop: We also experimented with a commercial desktop-based solver, Gurobi [33] (version 9.1). In Figure 6, we report results obtained using the Rome-Lib benchmark dataset [20] — comparing the application of just crossing reduction (CM) to both crossing and bendiness reduction (CM + BR). All the tests are run using the same subset of problems from Rome-Lib — the first 26 networks with a number of nodes that is a multiple of 5, starting from 10. The problems were all run on a computer with an Intel i7-7700K processor, 32 GB RAM, and running Ubuntu 20.04. Note that Rome-Lib contains undirected, non-layered networks with no groups. In order to obtain these results, we preprocessed the Rome-Lib dataset, performing a layer-assignment step (in appendix at osf.io/qdty9) to obtain layered networks. We used Navlakha et al.’s [61] network summarization algorithm to obtain groups of nodes. In Figure 6, we report the time spent solving problems of increasing size. The results are split in three figures: the first compares CM vs. CM + BR on Rome-Lib with no groups, the second includes groups but does not apply group collapsing, the third has groups and group collapsing. The visualizations show the median time to solve problems with increasing number of nodes, with boundaries at the 25% and 75% quartiles. We set a timeout of 1000 seconds for each problem, and the lines in the visualization are cut wherever the success rate of solving problems within the timeout became less than 75%, which would affect the shown data. A more extensive report can be found at osf.io/qdty9.

A word on scalability. While our experiments indicate that STRATISFIMAL LAYOUT could easily scale to networks with 50 to 100 nodes, it generally is difficult to predict ILP optimization time from network size alone. Other factors such as the connection pattern and the selection of constraints play an important role as well. In theory, the complexity of an ILP problem is determined by the number of variables, which define the space of possible combinations. Here Integer variables are more expensive than Boolean variables. However, in practice, an ILP solver can exploit special properties of a given problem to find the optimal solution much more rapidly. In some cases, adding

²<https://github.com/hgourvest/glpk.js>

³<https://visdunneright.github.io/stratisfimal/proofs.html>



(a) A randomly-generated QueryVis [48] visualization.

(b) Gansner et al. [28] applied to the network.

(c) STRATISFIMAL LAYOUT applied to the network.

Fig. 5: Comparison of different layouts applied to a randomly generated QueryVis example.

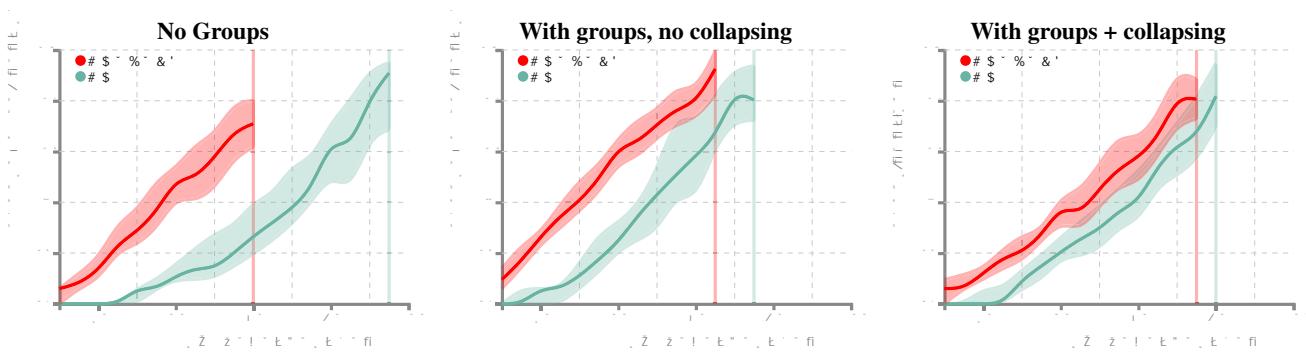


Fig. 6: The median time spent solving a problem from the Rome-Lib [20] network benchmark. The interval between the first and fourth percentiles is shown in lighter colors. Each line represents the execution on the first 26 networks in Rome-Lib with n nodes, where n is a multiple of 5 between 10 and 100, for a total of 468 networks per line. The timeout was set to 1000 seconds per individual network, and the lines are cut whenever less than 75% of problems with the same number of nodes could be completed within the timeout. Crossing reduction (CM) + bendiness reduction (BR) is always several times costlier than CM only, and groups improve performance and the number of problems solvable within the timeout. Some oscillations in the lines are only due to varying difficulty in the set of problems.

redundant constraints that are already implied by existing ones can improve the solver’s ability to home in on the feasibility region [47, 55]. Another approach is the introduction of variance in performance via *erraticism* [25, 51], due to imperfect tie-breaking by the ILP solver’s search algorithm [42] and the choice of initial conditions for the search algorithm, such as the seed for the random-number generator [47]. Lastly, we have to take into account how performance changes with different solvers and machines [51].

For our problem, the number of variables grows more rapidly with the number of nodes in a layer than with the number of layers, because the number of $x_{u,w}$ order indicator variables is quadratic in the number of nodes in a layer. Our group-collapsing optimization (Section 5.4) ultimately reduces constraint redundancy, yet it also succeeds in reducing the number of variables. It thus effectively lowers problem complexity — in the best case, halving the number of variables. Thus, an ILP solver will perform better on a network with pre-established groups than on a network of a similar size without groups. Specific examples for this effect can be found in our appendix at osf.io/qdyt9.

Heuristics vs. ILP. Our method complements work on heuristic solutions by establishing a baseline against which the solutions found by the faster techniques can be compared, as done for network layout algorithms and other problems in [23, 26, 39, 52, 64]. Furthermore, ILP solvers, together with other optimization tools, have been improving at a staggering rate in terms of optimization speed and features supported [1]. Both Gurobi⁴ and CPLEX’s⁵ most recent technical reports claim substantial improvements in processing times and the number of benchmark problems solved, including a $3\times$ improvement over the last five years in processing time. This means that STRATISFIMAL LAYOUT will become faster with each improvement in ILP solver technology, allowing it to scale to larger problems even as new visualization constraints may be added.

7 CONCLUSION AND FUTURE WORK

Algorithms that lay out node-link network visualizations based on human factors can help users understand complicated networks faster. Motivated by the real-world application of helping users understand complicated SQL queries faster, we devised a modular layout optimization framework that supports a combination of existing and novel readability criteria and features that together had not been addressed by any previous optimization model. Novel features include crossing minimization for same-layer edges and multi-layer and nested groups. Our approach finds optimal layouts for networks with 30 variables and complicated groupings within a second, which is a reasonable time, given that a diagram can take minutes to understand, and then be reused multiple times. The framework is versatile and allows the user to cherry-pick modules according to the features they require for their visualizations. Moreover, we provided our solution as open-source code together with a benchmark to measure optimization time for problems of various sizes and features.

Future work. Many more network features could be integrated into our modular formulation, e.g., to optimize edge-crossing angles [78], to improve the symmetry of the visualization [49, 66, 67], to bundle groups of edges [19, 65], or to minimize edge length [80]. A feature that is not necessary for our motivating example (but may be relevant for other applications) is to handle groups with overlapping intersections such as in Venn diagrams.

Acknowledgements. This work was supported in part by the National Science Foundation (NSF) under award numbers CAREER IIS-1762268 and IIS-1956096.

REFERENCES

- [1] R. Anand, D. Aggarwal, and V. Kumar. A comparative analysis of optimization solvers. *Journal of Statistics and Management Systems*, 20(4):623–635, 2017. doi: 10.1080/09720510.2017.1395182
- [2] M. Asquith, J. Gudmundsson, and D. Merrick. An ILP for the metro-line crossing problem. In *Proceedings of the Fourteenth Symposium on Computing: The Australasian Theory - Volume 77*, CATS ’08, p. 49–56. Australian Computer Society, Inc., 2008.

⁴<https://www.gurobi.com/wp-content/uploads/2020/02/Performance-Gurobi-9.0-1.pdf>

⁵<https://www.ibm.com/downloads/cas/KEVDB4NZ>

[3] G. D. Battista, P. Eades, R. Tamassia, and I. G. Tollis. Algorithms for drawing graphs: an annotated bibliography. *Computational Geometry*, 4(5):235 – 282, 1994. doi: 10.1016/0925-7721(94)00014-X

[4] G. D. Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall PTR, USA, 1st ed., 1998.

[5] R. Blue, C. Dunne, A. Fuchs, K. King, and A. Schulman. Visualizing real-time network resource usage. In J. R. Goodall, G. Conti, and K.-L. Ma, eds., *Visualization for Computer Security*, pp. 119–135, 2008. doi: 10.1007/978-3-540-85933-8_12

[6] J. Blythe, C. McGrath, and D. Krackhardt. The effect of graph layout on inference from social network data. In F. J. Brandenburg, ed., *Graph Drawing*, pp. 40–51, 1996. doi: 10.1007/BFb0021789

[7] U. Brandes and B. Köpf. Fast and simple horizontal coordinate assignment. In P. Mutzel, M. Jünger, and S. Leipert, eds., *Graph Drawing*, pp. 31–44, 2002. doi: 10.1007/3-540-45848-4_3

[8] C. Buchheim, D. Ebner, M. Jünger, G. W. Klau, P. Mutzel, and R. Weiskircher. Exact crossing minimization. In P. Healy and N. S. Nikolov, eds., *Graph Drawing*, pp. 37–48, 2006. doi: 10.1007/11618058_4

[9] M. Cembalo, A. De Santis, and U. Ferraro Petrillo. SAVI: A new system for advanced SQL visualization. In *Proceedings of the 2011 Conference on Information Technology Education*, SIGITE ’11, p. 165–170, 2011. doi: 10.1145/2047594.2047641

[10] C. Cerullo and M. Porta. A system for database visual querying and query visualization: Complementing text and graphics to increase expressiveness. In *18th International Workshop on Database and Expert Systems Applications (DEXA 2007)*, pp. 109–113, 2007. doi: 10.1109/DEXA.2007.91

[11] D. D. Chamberlin and R. F. Boyce. SEQUEL: A structured english query language. In *Proceedings of the 1974 ACM SIGFIDET (Now SIGMOD) Workshop on Data Description, Access and Control*, SIGFIDET ’74, p. 249–264, 1974. doi: 10.1145/800296.811515

[12] Z. Chen and H. Ji. Graph-based clustering for computational linguistics: A survey. In *Proceedings of the 2010 Workshop on Graph-Based Methods for Natural Language Processing*, TextGraphs-5, p. 1–9. Association for Computational Linguistics, USA, 2010.

[13] M. Chimani, C. Gutwenger, and P. Mutzel. Experiments on exact crossing minimization using column generation. In C. Álvarez and M. Serna, eds., *Experimental Algorithms*, pp. 303–315, 2006. doi: 10.1145/1498698.1564504

[14] A. Cimikowski and P. Shope. A neural-network algorithm for a graph layout problem. *IEEE Transactions on Neural Networks*, 7(2):341–345, 1996. doi: 10.1109/72.485670

[15] L. C. Coelho. Linearization of the product of two variables. <https://www.leandro-coelho.com/linearization-product-variables>. Last accessed on March 31, 2021.

[16] T. Crnovrsanin, C. W. Muelder, R. Faris, D. Felmlee, and K.-L. Ma. Visualization techniques for categorical analysis of social networks with multiple edge sets. *Social Networks*, 37:56–64, 2014. doi: 10.1016/j.socnet.2013.12.002

[17] J. Danaparamita and W. Gatterbauer. QueryViz: Helping users understand SQL queries and their patterns. In *Proceedings of the 14th International Conference on Extending Database Technology (EDBT)*, EDBT ’11, pp. 558–561, 2011. doi: 10.1145/1951365.1951440

[18] G. B. Dantzig and M. N. Thapa. *Linear Programming 1: Introduction*. Springer-Verlag, Berlin, Heidelberg, 1997.

[19] S. Di Bartolomeo, Y. Zhang, F. Sheng, and C. Dunne. Sequence braiding: Visual overviews of temporal event sequences and attributes. *IEEE transactions on visualization and computer graphics*, 27(2):1353–1363, February 2021. doi: 10.1109/tvcg.2020.3030442

[20] G. Di Battista, A. Garg, G. Liotta, R. Tamassia, E. Tassinari, and F. Vargiu. An experimental comparison of four graph drawing algorithms. *Computational Geometry*, 7(5):303 – 325, 1997. 11th ACM Symposium on Computational Geometry. doi: 10.1016/S0925-7721(96)00005-3

[21] W. Didimo, G. Liotta, and F. Montecchiani. A survey on graph drawing beyond planarity. *ACM Computing Surveys*, 52(1), 2019. doi: 10.1145/3301281

[22] C. Dunne, S. I. Ross, B. Shneiderman, and M. Martino. Readability metric feedback for aiding node-link visualization designers. *IBM Journal of Research and Development*, 59(2/3):14:1–14:16, 2015. doi: 10.1147/JRD.2015.2411412

[23] T. Dwyer, N. Henry Riche, K. Marriott, and C. Mears. Edge compression techniques for visualization of dense directed graphs. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2596–2605, 2013. doi: 10.1109/TVCG.2013.151

[24] P. Eades, X. Lin, and R. Tamassia. An algorithm for drawing a hierarchical graph. *International Journal of Computational Geometry & Applications*, 06(02):145–155, 1996. doi: 10.1142/S0218195996000101

[25] M. Fischetti and M. Monaci. Exploiting erraticism in search. *Operations Research*, 62(1):114–122, 2014. doi: 10.1287/opre.2013.1231

[26] G. Gange, P. J. Stuckey, and K. Marriott. Optimal k-level planarization and crossing minimization. In U. Brandes and S. Cornelsen, eds., *Graph Drawing*, pp. 238–249, 2011. doi: 10.1007/978-3-642-18469-7_22

[27] E. R. Gansner, E. Koutsofios, and S. C. North. Drawing graphs with dot, 1996.

[28] E. R. Gansner, E. Koutsofios, S. C. North, and K.-P. Vo. A technique for drawing directed graphs. *IEEE Transactions on Software Engineering*, 19(3):214–230, 1993. doi: 10.1109/32.221135

[29] M. R. Garey and D. S. Johnson. Crossing number is NP-complete. *SIAM Journal on Algebraic Discrete Methods*, 4(3):312–316, 1983. doi: 10.1137/0604033

[30] W. Gatterbauer. Databases will visualize queries too. *PVLDB*, 4(12):1498–1501, 2011. doi: 10.14778/3402755.3402805

[31] H. Gibson, J. Faith, and P. Vickers. A survey of two-dimensional graph layout techniques for information visualisation. *Information Visualization*, 12(3-4):324–357, 2013. doi: 10.1177/1473871612455749

[32] M. Gronemann, M. Jünger, F. Liers, and F. Mambelli. Crossing minimization in Storyline visualization. In Y. Hu and M. Nöllenburg, eds., *Graph Drawing and Network Visualization*, pp. 367–381, 2016. doi: 10.1007/978-3-319-50106-2_29

[33] Gurobi Optimization, LLC. Gurobi optimizer reference manual, 2021.

[34] P. Healy and A. Kuusik. The vertex-exchange graph: A new concept for multi-level crossing minimisation. In J. Kratochvíl, ed., *Graph Drawing*, pp. 205–216, 1999. doi: 10.1007/3-540-46648-7_21

[35] P. Healy and A. Kuusik. Algorithms for multi-level graph planarity testing and layout. *Theoretical Computer Science*, 320(2):331–344, 2004. doi: 10.1016/j.tcs.2004.02.033

[36] N. Henry and J. Fekete. MatrixExplorer: a dual-representation system to explore social networks. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):677–684, 2006. doi: 10.1109/TVCG.2006.160

[37] I. Herman, G. Melancon, and M. S. Marshall. Graph visualization and navigation in information visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):24–43, 2000. doi: 10.1109/2945.841119

[38] H. Jaakkola and B. Thalheim. Visual SQL – high-quality ER-based query treatment. In M. A. Jeusfeld and Ó. Pastor, eds., *Conceptual Modeling for Novel Application Domains*, pp. 129–139, 2003. doi: 10.1007/978-3-540-39597-3_13

[39] M. Jünger, E. K. Lee, P. Mutzel, and T. Odenthal. A polyhedral approach to the multi-layer crossing minimization problem. In G. DiBattista, ed., *Graph Drawing*, pp. 13–24, 1997. doi: 10.1007/3-540-63938-1_46

[40] M. Jünger and P. Mutzel. Exact and heuristic algorithms for 2-layer straightline crossing minimization. In F. J. Brandenburg, ed., *Graph Drawing*, pp. 337–348, 1996. doi: 10.1007/BFb0021817

[41] D. E. Knuth. Computer-drawn flowcharts. *Communications of the ACM*, 6(9):555–563, 1963. doi: 10.1145/367593.367620

[42] T. Koch, T. Achterberg, E. Andersen, O. Bastert, T. Berthold, R. Bixby, E. Danna, G. Gamrath, A. Gleixner, S. Heinz, A. Lodi, H. Mittelmann, T. Ralphs, D. Salvagnin, D. Steffy, and K. Wolter. MIPLIB 2010: Mixed integer programming library version 5. *Mathematical Programming Computation*, 3(2):103–163, 2011. doi: 10.1007/s12532-011-0025-9

[43] G. Koutrika, A. Simitsis, and Y. E. Ioannidis. Explaining structured queries in natural language. In *2010 IEEE 26th International Conference on Data Engineering (ICDE 2010)*, pp. 333–344, 2010. doi: 10.1109/ICDE.2010.5447824

[44] C. Kuratowski. Sur le problème des courbes gauches en topologie. *Fundamenta Mathematicae*, 15(1):271–283, 1930.

[45] O.-H. Kwon, T. Crnovrsanin, and K.-L. Ma. What would a graph look like in this layout? a machine learning approach to large graph visualization. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):478–488, Jan 2018. doi: 10.1109/tvcg.2017.2743858

[46] O. H. Kwon and K. L. Ma. A deep generative model for graph layout. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):665–675, 2020. doi: 10.1109/TVCG.2019.2934396

[47] E. Lalla-Ruiz and S. Voss. Improving solver performance through redun-

dancy. *Journal of Systems Science and Systems Engineering*, 25(3):303–325, 2016. doi: 10.1007/s11518-016-5301-9

[48] A. Leventidis, J. Zhang, C. Dunne, W. Gatterbauer, H. Jagadish, and M. Riedewald. QueryVis: Logic-based diagrams help users understand complicated SQL queries faster. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, SIGMOD '20, p. 2303–2318, 2020. doi: 10.1145/3318464.3389767

[49] R. J. Lipton, S. C. North, and J. S. Sandberg. A method for drawing graphs. In *Proceedings of the First Annual Symposium on Computational Geometry*, SCG '85, p. 153–160, 1985. doi: 10.1145/323233.323254

[50] Y. T. K. Liu Ma. Design considerations for optimizing StoryLine visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2679–2688, 2012. doi: 10.1109/TVCG.2012.212

[51] A. Lodi and A. Tramontani. *Performance Variability in Mixed-Integer Programming*, chap. Chapter 1, pp. 1–12. INFORMS, 2013. doi: 10.1287/educ.2013.0112

[52] R. Martí, G. Reinelt, and A. Duarte. A benchmark library and a comparison of heuristic methods for the linear ordering problem. *Computational Optimization and Applications*, 51(3):1297–1317, Apr. 2012. doi: 10.1007/s10589-010-9384-9

[53] B. Meyer. Self-organizing graphs — a neural network perspective of graph layout. In S. H. Whitesides, ed., *Graph Drawing*, pp. 246–262, 1998. doi: 10.1007/3-540-37623-2_19

[54] T. Milea, O. Schrijvers, K. Buchin, and H. Haverkort. Shortest-paths preserving metro maps. In M. van Kreveld and B. Speckmann, eds., *Graph Drawing*, pp. 445–446, 2012. doi: 10.1007/978-3-642-25878-7_45

[55] A. Miyauchi and N. Sukegawa. Redundant constraints in the standard formulation for the clique partitioning problem. *Optimization Letters*, 9:199–207, 2015. doi: 10.1007/s11590-014-0754-6

[56] J. L. Moreno. *Who shall survive?: A new approach to the problem of human interrelations*. Nervous and mental disease publishing co, 1934.

[57] D. M. Mount. Geometric intersection. In *Handbook of Discrete and Computational Geometry*, chapter 33, pp. 615–632. CRC Press LLC, Boca, 1997.

[58] T. Munzner. *Visualization analysis and design*. CRC press, 2014.

[59] P. Mutzel. An alternative method to crossing minimization on hierarchical graphs. In S. North, ed., *Graph Drawing*, pp. 318–333, 1997. doi: 10.1007/3-540-62495-3_57

[60] P. Mutzel and R. Weiskircher. Two-layer planarization in graph drawing. In K.-Y. Chwa and O. H. Ibarra, eds., *Algorithms and Computation*, pp. 72–79, 1998. doi: 10.1007/3-540-49381-6_9

[61] S. Navlakha, R. Rastogi, and N. Shrivastava. Graph summarization with bounded error. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data*, SIGMOD '08, p. 419–432, 2008. doi: 10.1145/1376616.1376661

[62] M. Nollenburg and A. Wolff. Drawing and labeling high-quality metro maps by mixed-integer programming. *IEEE Transactions on Visualization and Computer Graphics*, 17(5):626–641, 2011. doi: 10.1109/TVCG.2010.81

[63] A. Perer and B. Shneiderman. Integrating statistics and visualization: Case studies of gaining clarity during exploratory data analysis. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, p. 265–274, 2008. doi: 10.1145/1357054.1357101

[64] S. Polyakovskiy, M. R. Bonyadi, M. Wagner, Z. Michalewicz, and F. Neumann. A comprehensive benchmark set and heuristics for the traveling thief problem. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, GECCO '14, p. 477–484, 2014. doi: 10.1145/2576768.2598249

[65] S. Pupyrev, L. Nachmanson, and M. Kaufmann. Improving layered graph layouts with edge bundling. In U. Brandes and S. Cornelsen, eds., *Graph Drawing*, pp. 329–340, 2011. doi: 10.1007/978-3-642-18469-7_30

[66] H. Purchase. Which aesthetic has the greatest effect on human understanding? In G. DiBattista, ed., *Graph Drawing*, pp. 248–261, 1997. doi: 10.1007/3-540-63938-1_67

[67] H. Purchase. Metrics for graph drawing aesthetics. *Journal of Visual Languages & Computing*, 13(5):501 – 516, 2002. doi: 10.1006/jvlc.2002.0232

[68] U. Rüegg, T. Ehlers, M. Spönemann, and R. von Hanxleden. A generalization of the directed graph layering problem. In Y. Hu and M. Nöllenburg, eds., *Graph Drawing and Network Visualization*, pp. 196–208, 2016. doi: 10.1007/978-3-319-50106-2_16

[69] G. Sander. A fast heuristic for hierarchical Manhattan layout. In F. J. Brandenburg, ed., *Graph Drawing*, pp. 447–458, 1996. doi: 10.1007/BFb0021828

[70] M. A. Smith, B. Shneiderman, N. Milic-Frayling, E. Mendes Rodrigues, V. Barash, C. Dunne, T. Capone, A. Perer, and E. Gleave. Analyzing (social media) networks with NodeXL. In *Proceedings of the Fourth International Conference on Communities and Technologies*, C&T '09, p. 255–264, 2009. doi: 10.1145/1556460.1556497

[71] K. Sugiyama. *Graph Drawing and Applications for Software and Knowledge Engineers*. World Scientific, 2002. doi: 10.1142/4902

[72] K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(2):109–125, 1981. doi: 10.1109/TSMC.1981.4308636

[73] R. Tamassia. *Handbook of Graph Drawing and Visualization*. Chapman & Hall/CRC, 1st ed., 2016.

[74] H. Tang and Z. Hu. Network simplex algorithm for DAG layering. In *2013 International Conference on Computational and Information Sciences*, pp. 1525–1528, 2013. doi: 10.1109/ICCIS.2013.401

[75] W. T. Tutte. Convex representations of graphs. *Proceedings of the London Mathematical Society*, s3-10(1):304–320, 1960. doi: 10.1112/plms/s3-10.1.304

[76] W. T. Tutte. How to draw a graph. *Proceedings of the London Mathematical Society*, s3-13(1):743–767, 1963. doi: 10.1112/plms/s3-13.1.743

[77] C. Vehlow, F. Beck, and D. Weiskopf. Visualizing group structures in graphs: A survey. *Computer Graphics Forum*, 36(6):201–225, 2017. doi: 10.1111/cgf.12872

[78] C. Ware, H. Purchase, L. Colpoys, and M. McGill. Cognitive measurements of graph aesthetics. *Information Visualization*, 1(2):103–110, 2002. doi: 10.1057/palgrave.ivs.9500013

[79] H.-Y. Wu, S. Takahashi, D. Hiroto, M. Arikawa, C.-C. Lin, and H.-C. Yen. Spatially efficient design of annotated metro maps. *Computer Graphics Forum*, 32(3pt3):261–270, 2013. doi: 10.1111/cgf.12113

[80] V. Yoghoudjian, T. Dwyer, G. Gange, S. Kieffer, K. Klein, and K. Marriott. High-quality ultra-compact grid layout of grouped networks. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):339–348, 2016. doi: 10.1109/TVCG.2015.2467251

[81] D. C. Zarate, P. L. Bodic, T. Dwyer, G. Gange, and P. Stucky. Optimal Sankey diagrams via integer programming. In *2018 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 135–139, 2018. doi: 10.1109/PacificVis.2018.000025