WHAT IS...

Treewidth?

David Eppstein

The *treewidth* of a graph, a positive integer defined using a tree of sets of vertices, is central to graph structure theory and the parametrized complexity of algorithms. Its many areas of application include sparse numerical linear algebra, Bayesian inference, control theory, game theory, low-dimensional topology, network science, and algebraic geometry.

Definitions

A common definition of treewidth involves *tree-decompositions*, of which the following is a special case. A graph is *outerplanar* if it can be drawn with its vertices on the boundary of a disk and its edges as noncrossing curves in the disk. Adding artificial edges, we can augment this drawing to a *triangulation*, a subdivision of the disk into triangles (Fig. 1), with the following properties:

1. The edge-to-edge adjacencies of the triangles form a tree.
2. Each graph vertex belongs to triangles forming a connected subtree.
3. Each graph edge belongs to at least one triangle.

A tree-decomposition, then, consists of a tree (like the tree of triangles) whose nodes are associated with sets of graph vertices, called *bags*. In the outerplanar case, each bag contains the three corners of a triangle. More generally, tree-decompositions may have larger bags, but must satisfy properties (2) and (3): each graph vertex belongs to bags forming a connected subtree, and each graph edge has both endpoints in at least one bag. The *width* of a tree-decomposition is the size of the largest bag, minus one. The treewidth of any finite undirected graph is the smallest width of any of its tree-decompositions. The "minus one" adjustment ensures that trees have width one: root any tree arbitrarily, then use bags containing each vertex and its parent. The tree of triangles of an outerplanar graph has width two.

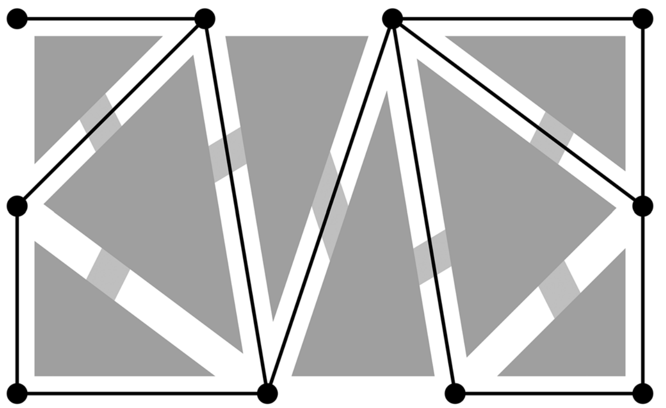

Figure 1. An outerplanar graph (black) and a tree of triangles forming a tree-decomposition (gray).

Treewidth has many equivalent definitions. It is the maximum clique size (minus one) in a chordal supergraph minimizing this size. It is the maximum number k of cops that a robber can escape in a certain pursuit-evasion game, the maximum k such that some function (called a *haven*) maps positions of $\leq k$ cops to an unoccupied subgraph into which the robber can escape, and the maximum k for which in some family of connected subgraphs, all touching each other (called a *bramble*), $\leq k$ cops always leave an unoccupied subgraph into which the robber can escape [ST93]. Tree-decompositions and chordal supergraphs are in some sense dual to havens and brambles: a structure of the former type provides an upper bound on

David Eppstein is a distinguished professor of computer science at the University of California, Irvine. His email address is eppstein@uci.edu.

His work is supported in part by grant CCF-2212129 of the National Science Foundation.

DOI: https://doi.org/10.1090/noti3043

treewidth, whereas a haven or bramble provides a lower bound. Treewidth is the top element in a certain lattice of well-behaved functions from graphs to numbers [Hal76]. Many other graph parameters are both upper bounded and lower bounded by a function of treewidth [HW17].

Structure

Tree-decompositions, with bag size replaced by other measures, play a central role in the work of Neil Robertson and Paul Seymour on families of graphs closed under *minors*, the results of edge contractions and vertex or edge deletions. In 23 papers published from 1983 to 2012, Robertson and Seymour found that the graphs in any minor-closed class have tree-decompositions in which each bag induces a *torso*, a graph of bounded genus with certain constrained embellishments, and in which adjacent bags share few vertices. This structural result was the key to their proof of Wagner's conjecture, that every minor-closed family can be characterized by finitely many *forbidden minors*, in the same way that the planar graphs are exactly the graphs that do not have K_5 or $K_{3,3}$ as minors [RS04].

Edge contraction and deletion cannot increase the treewidth of a graph: treewidth is *minor-monotone*. Thus, the graphs of treewidth $\leq w$ have finitely many forbidden minors. As Robertson and Seymour also proved, every minor-closed family of graphs that has a planar forbidden minor has bounded treewidth. The planar graphs themselves have unbounded treewidth—an $n \times n$ grid graph has treewidth n—and every graph of sufficiently high treewidth contains an $n \times n$ grid minor [RS84, DH08]. It was in the context of this work on planar forbidden minors that Robertson and Seymour gave the name "treewidth" (or, as they wrote it, "tree-width") to the concept we now study [RS84], although it appeared earlier under other names [BB72, Hal76].

More recent work has shown that the graphs in any minor-closed family, and in some families that go beyond graph minors, have a different structure: each graph in the family is a subgraph of a *strong product of graphs* $H \boxtimes P$ where H has bounded treewidth and P is a path. Here, the strong product has as vertex set the Cartesian product of the vertices of the two given graphs. It has an edge between two vertices, when, in each factor graph, the corresponding two vertices are adjacent or identical. This product structure, in turn, has been used to solve many old open problems in graph theory, on queue layouts of graphs, on nonrepetitive and centered colorings of graphs, on implicit label schemes for graphs, and on universal graphs for families of graphs [DHJ+21].

Computation

The use of treewidth in graph algorithms can be illustrated by graph coloring, a problem that is NP-complete for arbitrary graphs but polynomial-time solvable for graphs of bounded treewidth. Any graph of treewidth w can be colored with $w+1$ colors (not necessarily optimal) by choosing a root bag in a tree-decomposition, processing all bags in order by distance from the root, and assigning each uncolored vertex in each bag a color distinct from its neighbors in the same bag. Any other neighbors of such a vertex can only be in bags that come later in the processing order, so the resulting coloring is proper.

With this bound on the chromatic number in hand, an optimal coloring can be found by processing the bags in the opposite order, bottom-up from leaves to the root. At each bag, for each c from 2 to $w+1$, store a set of *valid c-colorings*, assignments of c colors to the vertices of the bag that are compatible with a proper coloring of the induced subgraph of all vertices in descendants of the bag. To find the valid c-colorings, check for each color assignment whether it is a coloring of the induced subgraph of the bag and is consistent with an already-computed valid c-coloring of each child bag. The whole graph is c-colorable when there exists a valid c-coloring at the root bag, as determined after processing all bags in this way.

Because this algorithm considers all color assignments within a bag, it takes super-exponential time (in the bag size) per bag. However, the number of bags in a tree-decomposition can be assumed to be linear in the number of vertices, so for any constant bound on width, this algorithm takes time linear in the graph size. This type of time bound, polynomial of fixed degree whenever some parameter such as treewidth is bounded, is called *fixed-parameter tractable*. Constructing optimal or approximately-optimal tree-decompositions is also fixed-parameter tractable [KL23], so we need not assume that the tree-decomposition is given as input. Fixed-parameter tractable algorithms are the main topic of study in *parametrized complexity theory*.

The same sort of bottom-up processing of a tree-decomposition applies to many other problems. A powerful algorithmic meta-theorem known as *Courcelle's theorem* provides a fixed-parameter tractable algorithm for testing whether a given graph models any quantified Boolean formula whose variables represent vertices, edges, sets of vertices, or sets of edges [Cou90]. Here, the parameters are the treewidth and the length of the formula. For instance, c-coloring has a formula involving the existence of c sets of vertices, one for each color class.

In a different direction, treewidth-based algorithms can often be extended beyond graphs of bounded treewidth. For instance, in planar graphs, one can find a bounded number of bounded-treewidth subgraphs that cover the whole graph with little overlap, or that cover most of the graph disjointly, and then use these systems of subgraphs for fast and accurate optimization algorithms for many hard graph problems [Bak94]. Another approach uses Robertson and Seymour's grid-minor theorem: if a graph has a large grid minor one can use this somehow (perhaps either solving the problem directly or removing an irrelevant vertex from the middle of the grid), and if it does not, then one can apply algorithms that assume bounded treewidth. This combination of grids and treewidth in algorithm design has been formalized in the theory of *bidimensionality* [DH08]. Both Courcelle's theorem and bidimensionality won their authors the Nerode Prize, an annual award for outstanding research in parametrized complexity, in 2022 and 2015, respectively.

An area of active research is the extension of treewidth to other related width parameters that can be bounded on dense graphs, apply to directed graphs, or characterize the graphs on which efficient testing of logically defined properties is possible. As well as longer-studied variants such as bandwidth, branchwidth, carving width, clique-width, cutwidth, and pathwidth, newer alternatives include *twin-width*, which characterizes the ordered graphs (graphs equipped with a total order on vertices) for which first-order model checking is efficient [BKTW22], and *flip-width*, defined through a pursuit-evasion game like the cops-and-robbers game for treewidth, and conjectured to characterize the unordered graphs on which first-order model checking is efficient [Tor23].

Applications

Graphs are a frequent abstraction in many other problem domains, and the strong relation between treewidth and algorithmic complexity is reflected in these domains. For instance:

- In sparse numerical linear algebra, the method of *generalized nested dissection* allows for the fast solution of systems of linear equations defined on graphs of low or sublinear treewidth [AY10].
- The treewidth of a given belief network controls the complexity of *junction tree algorithms* for machine learning and Bayesian inference, and of algorithms for inferring the belief network when it is not given [CG07]. Similar methods can also be applied to stochastic optimal control problems [KGO12].
- In game theory, finding equilibria in certain coalition forming games is intractible in general, but becomes efficient when a graph underlying the game has bounded treewidth [Pet16].
- Extensions of Courcelle's theorem to triangulated manifolds with low-treewidth dual graphs have been used to compute invariants in low-dimensional topology involving taut angle structures, discrete Morse theory, and Turaev–Viro invariants [BD17].
- In network science, low treewidth allows many types of queries on network databases to be answered more efficiently than otherwise. Although large real-world networks modeling roads and infrastructure, social interactions, biological subunits, and interrelated pieces of knowledge tend to have high treewidth, these networks can often be decomposed into smaller low-treewidth subnetworks to which these fast query techniques apply [MSJ19].

The structural aspects of treewidth can also be reflected in other areas of mathematics that are less directly computational. As one example, in algebraic geometry and tropical geometry, the treewidth of certain graphs can be used as a lower bound for the gonality of these graphs, and hence for the gonality of Riemann surfaces and tropical curves associated to these graphs [vDdBG20].

References

[AY10] Noga Alon and Raphael Yuster, *Solving linear systems through nested dissection*, 2010 IEEE 51st Annual Symposium on Foundations of Computer Science—FOCS 2010, IEEE Computer Soc., Los Alamitos, CA, 2010, pp. 225–234. MR3024796

[Bak94] Brenda S. Baker, *Approximation algorithms for NP-complete problems on planar graphs*, J. Assoc. Comput. Mach. **41** (1994), no. 1, 153–180, `https://doi.org/10.1145/174644.174650`. MR1369197

[BB72] Umberto Bertelè and Francesco Brioschi, *Nonserial dynamic programming*, Mathematics in Science and Engineering, Vol. 91, Academic Press, New York-London, 1972. MR416607

[BKTW22] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant, *Twin-width I: Tractable FO model checking*, J. ACM **69** (2022), no. 1, Art. 3, 46, DOI 10.1145/3486655. MR4402362

[BD17] Benjamin A. Burton and Rodney G. Downey, *Courcelle's theorem for triangulations*, J. Combin. Theory Ser. A **146** (2017), 264–294, `https://doi.org/10.1016/j.jcta.2016.10.001`. MR3574232

[CG07] Anton Chechetka and Carlos Guestrin, *Efficient principled learning of thin junction trees*, Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems, Vancouver, British Columbia, Canada, December 3–6, 2007, 2007, pp. 273–280, `https://proceedings.neurips.cc/paper/2007/hash/0768281a05da9f27df178b5c39a51263-Abstract.html`.

[Cou90] Bruno Courcelle, *The monadic second-order logic of graphs. I. Recognizable sets of finite graphs*, Inform. and Comput. **85** (1990), no. 1, 12–75, DOI 10.1016/0890-5401(90)90043-H. MR1042649

[DH08] Erik D. Demaine and Mohammadtaghi Hajiaghayi, *Linearity of grid minors in treewidth with applications through bidimensionality*, Combinatorica **28** (2008), no. 1, 19–36, DOI 10.1007/s00493-008-2140-4. MR2399006

[DHJ+21] Zdeněk Dvořák, Tony Huynh, Gwenaël Joret, Chun-Hung Liu, and David R. Wood, *Notes on graph product structure theory*, 2019–20 MATRIX annals, MATRIX Book Ser., vol. 4, Springer, Cham, 2021, pp. 513–533, DOI 10.1007/978-3-030-62497-2_32. MR4294791

[Hal76] Rudolf Halin, *S-functions for graphs*, J. Geom. **8** (1976), no. 1-2, 171–186, https://doi.org/10.1007/BF01917434. MR444522

[HW17] Daniel J. Harvey and David R. Wood, *Parameters tied to treewidth*, J. Graph Theory **84** (2017), no. 4, 364–385, DOI 10.1002/jgt.22030. MR3623383

[KGO12] Hilbert J. Kappen, Vicenç Gómez, and Manfred Opper, *Optimal control as a graphical model inference problem*, Mach. Learn. **87** (2012), no. 2, 159–182, https://doi.org/10.1007/s10994-012-5278-7. MR2914010

[KL23] Tuukka Korhonen and Daniel Lokshtanov, *An improved parameterized algorithm for treewidth*, STOC'23—Proceedings of the 55th Annual ACM Symposium on Theory of Computing, ACM, New York, 2023, pp. 528–541, DOI 10.1145/3564246.3585245. MR4617406

[MSJ19] Silviu Maniu, Pierre Senellart, and Suraj Jog, *An experimental study of the treewidth of real-world graph data*, 22nd International Conference on Database Theory, LIPIcs. Leibniz Int. Proc. Inform., vol. 127, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2019, pp. Art. No. 12, 18. MR3936362

[Pet16] Dominik Peters, *Graphical hedonic games of bounded treewidth*, Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12–17, 2016, Phoenix, Arizona, USA, 2016, pp. 586–593, https://doi.org/10.1609/aaai.v30i1.10046.

[RS84] Neil Robertson and P. D. Seymour, *Graph minors. III. Planar tree-width*, J. Combin. Theory Ser. B **36** (1984), no. 1, 49–64, DOI 10.1016/0095-8956(84)90013-3. MR742386

[RS04] Neil Robertson and P. D. Seymour, *Graph minors. XX. Wagner's conjecture*, J. Combin. Theory Ser. B **92** (2004), no. 2, 325–357, https://doi.org/10.1016/j.jctb.2004.08.001. MR2099147

[ST93] P. D. Seymour and Robin Thomas, *Graph searching and a min-max theorem for tree-width*, J. Combin. Theory Ser. B **58** (1993), no. 1, 22–33, DOI 10.1006/jctb.1993.1027. MR1214888

[Tor23] Szymon Toruńczyk, *Flip-width: cops and robber on dense graphs*, 2023 IEEE 64th Annual Symposium on Foundations of Computer Science—FOCS 2023, IEEE Computer Soc., Los Alamitos, CA, 2023, pp. 663–700, DOI 10.1109/FOCS57990.2023.00045. MR4720287

[vDdBG20] Josse van Dobben de Bruyn and Dion Gijswijt, *Treewidth is a lower bound on graph gonality*, Algebr. Comb. **3** (2020), no. 4, 941–953, DOI 10.5802/alco.124. MR4145985

David Eppstein

Credits

Figure 1 is courtesy of David Eppstein.

Photo of David Eppstein is courtesy of Laurel Hungerford.