UAIC_Twin_Width: A Heuristic Twin-Width Algorithm

Andrei Arhire ✉
Alexandru Ioan Cuza University of Iași, Romania

Matei Chiriac ✉
Alexandru Ioan Cuza University of Iași, Romania

Radu Timofte ✉
ETH Zürich, Switzerland
University of Würzburg, Germany
Alexandru Ioan Cuza University of Iași, Romania

Abstract

Twin-width is a graph parameter that provides an intuitive measure of the similarity between a graph and a co-graph [1]. To understand this concept, let's consider a tri-graph initially consisting of only black edges. We repeatedly apply an operation until only one vertex remains in the graph: two vertices are merged into a new vertex, and the edges are updated based on specific rules. The neighborhood of the new vertex is formed by combining the neighborhoods of the contracted vertices. Within this neighborhood union, an edge is considered black if it appears as a black edge in both original neighborhoods; otherwise, it becomes red. The twin-width is defined as the highest red degree encountered at any vertex during the contraction steps that minimize this degree.

In this short paper, we provide an overview of the main techniques utilized in the heuristic twin-width algorithm, which we have submitted as the UAIC_Twin_Width solver for the heuristic track of the 2023 PACE challenge.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis

Keywords and phrases twin-width

Supplementary Material The source code is available on Zenodo (10.5281/zenodo.8010510) and GitHub (https://github.com/AndreiiArhire/PACE2023).

1 Preliminaries

Consider a trigraph $G = (V, B, R)$ consisting of a set of vertices V, black edges B, and red edges R. Here, B and R are subsets of $\binom{V}{2}$, and they do not have any common elements ($B \cap R = \emptyset$). Without any specific indication, a graph $H = (V, E)$ is treated as the trigraph $H' = (V, E, \emptyset)$.

In this context, we define the red degree of a vertex as the number of incident red edges connected to that vertex. When two vertices are compressed, a new vertex is formed by considering the edges associated with the original vertices, and subsequently, the original vertices are removed from the graph.

Formally, the union of vertices x and y, with B_x, B_y, R_x, and R_y representing the sets of incident red and black edges, gives rise to a vertex z with the following definitions:

- A vertex $u \in R_z$ if $u \in B_x \cup B_y \cup R_x \cup R_y \setminus B_x \cap B_y$.
- A vertex $u \in B_z$ if $u \in B_x \cap B_y$.

As per the requirements of the challenge format, the solver assigns the same label to z as that of x. By following this process of reducing the graph to a single vertex through these operations, a sequence of $|V| - 1$ pairs of vertices is generated. The width of such a sequence is defined as the maximum red degree encountered at any vertex during the contraction process.

The twin-width represents the minimum width achievable through a sequence of contractions that reduces the graph.

2 Solver summary

The solver uses a greedy algorithm that systematically adds promising pairs of vertices to the solution. When dealing with small-sized graphs, the algorithm selects the pair that produces the smallest maximum degree by red edges. However, due to the time limitations and the presence of large graphs, the algorithm is simplified. Instead of considering all possible options in each step ($|V| - 1$ steps), the algorithm now follows a modified approach. It performs a logarithmic number of steps, specifically $log_{\frac{k}{(k-1)}}(|V|)$, where the top evaluated $(\frac{|V|}{k})$ vertex pairs, satisfying certain conditions, are contracted out of a pool of $k \cdot |V|$ pairs. Here, k takes a minimum value of 2.

To identify the potential pairs for contraction, a breadth-first search algorithm is applied, starting from the vertex with the highest number of neighbors. The algorithm always prioritizes visiting the neighbor with the highest degree first. Each pair consists of a vertex and another vertex placed at most k positions to the right in the order of visitation. Initially, k is set to a maximum value of 10 and is increased if there is sufficient remaining time.

For the largest test cases, a faster but less accurate method is applied. In this method, the vertex with the lowest degree is always contracted with one of its neighbors or with a vertex that shares a neighbor in common.

Before running the algorithm, vertices that, when contracted, do not create a red edge are identified. This identification process utilizes polynomial hashes on adjacency lists and Fermat's Little Theorem.

References

1 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width i: tractable fo model checking. *ACM Journal of the ACM (JACM)*, 69(1):1–46, 2021.