PACE Solver Description: Zygosity

Emmanuel Arrighi
University of Trier, Germany

Pål Grønås Drange
University of Bergen, Norway

Kenneth Langedal
University of Bergen, Norway

Farhad Vadiee
University of Bergen, Norway

Martin Vatshelle
University of Bergen, Norway

Petra Wolf
University of Bergen, Norway

Abstract

The graph parameter twin-width was recently introduced by Bonnet et al. Twin-width is a parameter that measures a graph's similarity to a cograph, which is a graph that can be reduced to a single vertex by repeatedly contracting twins. This brief description introduces Zygosity, a heuristic for computing a low-width contraction sequence that achieved second place in the 2023 edition of Parameterized Algorithms and Computational Experiments Challenge (PACE). Zygosity starts by repeatedly contracting twins. Then, any attached trees are contracted down to a single pendant vertex. The remaining graph is then contracted using a randomized greedy algorithm.

2012 ACM Subject Classification Mathematics of computing → Graph algorithms

Keywords and phrases Twin-width, randomized greedy algorithm

Digital Object Identifier 10.4230/LIPIcs.IPEC.2023.39

Supplementary Material *Software (Source Code)*: `https://doi.org/10.5281/zenodo.8370343`

1 Problem description

Let $G = (V, E)$ be an undirected graph with vertex set V and edge set E, where the edge set is partitioned into black edges $B(E)$ and red edges $R(E)$. The neighborhood of a vertex u is denoted by $N(u)$, formally $N(u) = \{v \in V \mid uv \in E\}$. A neighbor v of a vertex u can be black or red, depending on whether uv is in $B(E)$ or $R(E)$. Formally, for a vertex $u \in V$, we let $N^r(u) = \{v \in V \mid uv \in R(E)\}$ be the red neighborhood of u and $N^b(u) = \{v \in V \mid uv \in B(E)\}$ be the black neighborhood of u. The degree of a vertex u is the size of the neighborhood $|N(u)|$ and the black and red degree refers to the size of the black and red neighborhood ($|N^b(u)|$ and $|N^r(u)|$), respectively.

A contraction between two vertices u and v is performed by removing u and v from G, and adding a new vertex w as follow:

- $N^r(w) = N^r(u) \cup N^r(v) \cup (N^b(u) \Delta N^b(v)) \setminus \{u, v\}$ where $N^b(u) \Delta N^b(v)$ is the symmetric difference.
- $N^b(w) = N^b(u) \cap N^b(v)$.

A contraction sequence starts from the original graph and ends with a single vertex. This original graph starts with no red edges. The *width* of a contraction sequence is the highest red degree of any vertex along the way. And finally, the twin-width parameter introduced by Bonnet et al. [1] is the lowest width among all possible contraction sequences for G.

2 Solver description

Zygosity is a randomized greedy heuristic for computing low-width contraction sequences. Before the randomized part, it performs two preprocessing steps. First, all twins are contracted. Secondly, as described by Bonnet et al. [1], trees are contracted by either contracting two leaves with the same parent or a leaf with its parent. If the whole graph is a tree, this gives a sequence with a width of at most two. Otherwise, the tree will contract down to a single pendant vertex connected by a red edge.

The remaining sequence is constructed by considering multiple random contractions at each step and selecting the best one. The exact number of contractions to consider is updated continuously during the construction. The number is updated every second based on the remaining time and the average time it takes to check one contraction. Zygosity employs two measures to evaluate the quality of a contraction.

- **Red degree:** This red degree refers to the highest red degree of a vertex that increased its red degree due to this contraction. The newly merged vertex is considered increased if its red degree is greater than both of the merged vertices. Intuitively, the red degree is the most important thing to keep low when constructing a low-width contraction sequence.
- **Intersection size:** This is the size of the intersection between the neighborhoods of the two merged vertices. This number also says how many edges will be removed due to this contraction. If an edge connects the two merged vertices, this counts one to the intersection size. The intuition for this measure is that fewer edges make for fewer problems later.

Zygosity prioritizes the smallest red degree and uses intersection size as a tiebreaker. The only exception to this rule is when both contractions being compared have a red degree lower than the current width of the sequence generated so far. In such cases, only the intersection size is taken into account.

Instead of randomly picking a pair of vertices, Zygosity randomly picks the first vertex and then does a random walk from there. The random walk performs either one or two steps, depending on a biased coin toss. The odds on this coin are decided based on the density of the input graph, favoring distance one walks on sparser graphs and distance two on dense ones.

3 Implementation details

For this randomized procedure, it is crucial to quickly compute the intersection and red degree of a potential contraction. Therefore, spending more time elsewhere can be beneficial if it speeds up the contraction checking. For this reason, we don't use any associative data structure but instead store the neighborhood of a vertex continuously in a single array. The black and red edges are kept separate and sorted separately. With this setup, computing the intersection and red degree can be done by scanning through the two neighborhoods in a cache-friendly manner. It is possible to make a single pass using four pointers. However, we noticed that doing a separate pass over the red neighborhoods first was faster. This pass copies the difference to a separate array. Then, a second pass over the black parts and the temporarily stored difference from the red ones. This captures every relation a vertex can have with the contraction in question, and the intersection and max degree are registered accordingly.

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.