


PACE Solver Description: Hydra Prime

Yosuke Mizutani ✉ 

University of Utah, Salt Lake City, UT, USA

David Dursteler ✉ 

University of Utah, Salt Lake City, UT, USA

Blair D. Sullivan ✉ 

University of Utah, Salt Lake City, UT, USA

Abstract

This note describes our submission to the 2023 PACE Challenge on the computation of twin-width. Our solver **Hydra Prime** combines modular decomposition with a collection of upper- and lower-bound algorithms, which are alternatingly applied on the prime graphs resulting from the modular decomposition. We introduce two novel approaches which contributed to the solver's winning performance in the Exact Track: *timeline encoding* and *hydra decomposition*. Timeline encoding is a new data structure for computing the width of a given contraction sequence, enabling faster local search; the hydra decomposition is an iterative refinement strategy featuring a small vertex separator.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis

Keywords and phrases Twin-width, PACE 2023

Digital Object Identifier 10.4230/LIPIcs.IPEC.2023.36

Supplementary Material *Software*: <https://github.com/TheoryInPractice/hydraprime>
archived at `swh:1:dir:eb6788de444b4d5277f0a400dea4a1affa0e6df7`

Funding *Blair D. Sullivan*: This work was supported in part by the Gordon & Betty Moore Foundation under award GBMF4560.

1 Introduction

The goal of the 2023 Parameterized Algorithms & Computational Experiments (PACE) Challenge (<https://pacechallenge.org/2023/>) was to compute twin-width [2], a structural graph parameter which measures how close a given graph is to a *cograph* – a graph which can be reduced to a single vertex by repeatedly merging (*contracting*) pairs of *twins* – vertices with identical open neighborhoods. More generally, twin-width measures the minimum number of “mistakes” made in such a process when the pairs being contracted are no longer twins. If u and v are being merged, we say uy becomes a *red edge* if y is a neighbor of u but not v (and analogously for edges vy). The width of a contraction sequence is then the maximum number of red edges incident to any vertex (*red degree*) at any time during the process, and the twin-width of a graph is the minimum width of all valid contraction sequences. While graphs with bounded twin-width admit many FPT algorithms, computing the parameter is NP-hard, and prior to the PACE challenge its exact computation had remained impractical even on relatively small graphs.

Most twin-width solvers naturally begin by removing twins, as all groups of twins can be collapsed without incurring any red edges, making it a safe operation. In **Hydra Prime**, we employ a stronger notion of this via *modular decompositions* [4], which decompose a graph into a hierarchy of maximal modules. A key property of these decompositions is that the twin-width of the original graph is exactly the maximum of the twin-width of the twin-free, prime quotient graphs (Theorem 3.1 from [5]). We thus begin by running a



© Yosuke Mizutani, David Dursteler, and Blair D. Sullivan;
licensed under Creative Commons License CC-BY 4.0

18th International Symposium on Parameterized and Exact Computation (IPEC 2023).

Editors: Neeldhara Misra and Magnus Wahlström; Article No. 36; pp. 36:1–36:5

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

re-implemented linear-time modular decomposition solver based on [6], then process each prime graph separately, maintaining a global lower bound. If a prime graph is a tree, we run `PrimeTreeSolver`, otherwise we run a series of lower- and upper-bound algorithms (listed at the end of this section) alternatively until the bounds match, from the quickest algorithms to the slowest. Those algorithms marked with (*) use a SAT solver as a subroutine; the implementation submitted to PACE uses the Kissat solver [1].

Algorithm List.

- Exact algorithms
 - `PrimeTreeSolver`: Linear-time exact solver for trees without twins.
 - `BranchSolver`: Brute-force solver equipped with caching mechanism and reduction rules.
 - `DirectSolver` (*): SAT-based solver implementing the relative encoding presented in [5].
- Lower-bound algorithms
 - `LBGreedy`: Greedily removes a vertex u from the graph G such that $|\Delta(u, v)|$ is minimized for some v . Reports the maximum value of $\min_{u, v \in V(G), u \neq v} |\Delta_G(u, v)|$.
 - `LBCore` (*): SAT-based algorithm to find $\max_{S \subseteq V(G)} \min_{u, v \in S, u \neq v} |\Delta_{G[S]}(u, v)|$.
 - `LBSample`: Sampling-based algorithm. Finds a connected induced subgraph G' of G by random walk and computes the exact or lower-bound twin-width of G' .
 - `LBSeparate` (*): Similar to `LBSample`, but uses the hydra decomposition to find an induced subgraph to check for the lower-bound.
- Upper-bound algorithms
 - `UBGreedy`: Iteratively contract a vertex pair minimizing the weak red potential.
 - `UBLocalSearch`: Using the timeline encoding, we make small changes to the elimination ordering and the contraction tree to see if there is a better solution.
 - `UBSeparate` (*): Iterative refinement algorithm using the hydra decomposition.

In this paper we focus on two additional contributions to solving twin-width which are used in the `LocalSearch` and `Separate` algorithms implemented in `Hydra Prime`: “timeline encoding” and “hydra decomposition”. *Timeline encoding* is a novel data structure which enables faster computation of twin-width by storing red “sources” and “intervals” indicating the cause and window of each red edge. In the `Separate` upper- and lower-bound algorithms, we introduce *hydra decomposition*, an iterative refinement strategy using small vertex separators. After defining necessary notation, we briefly describe these in Sections 2 and 3, respectively. Additional details are in the appendix available on the code repository.

Notation. We follow standard graph-theoretic notation (e.g. found in [3]), the original definition of twin-width [2], and terminology introduced by Schidler and Szeider [5]. Refer to [4] and [6] for the definitions of a *module*, *modular decomposition*, a *prime* graph, etc. We write $u \leftarrow v$ when vertex v is contracted into vertex u . Given a trigraph G , the *weak red potential* of $u, v \in V(G), u \neq v$ is the red degree of u after contraction $u \leftarrow v$. We further define the *unshared neighbors* of vertices u and v , denoted by $\Delta(u, v)$ as $N(u) \Delta N(v) \setminus \{u, v\}$, where Δ denotes the symmetric difference of two sets. We write $[n]$ for $\{1, \dots, n\}$.

2 Timeline Encoding

In this work we developed the *timeline encoding*, a data structure to compute the width of a given contraction sequence. An instance of the timeline encoding stores the following data:

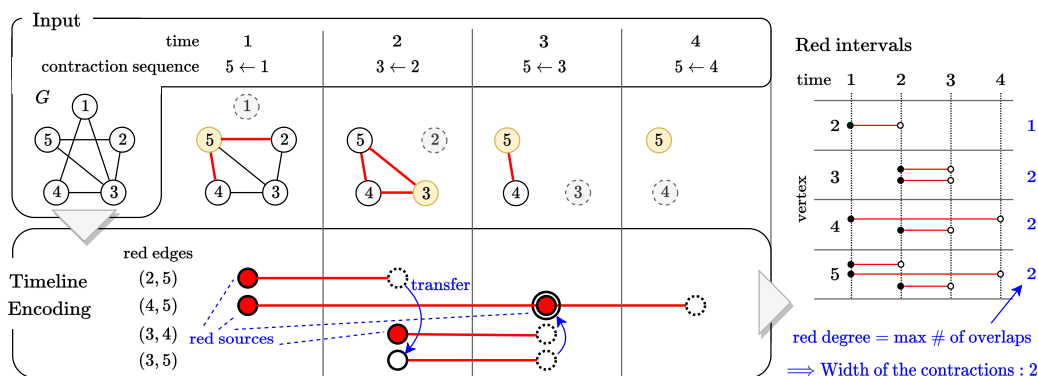


Figure 1 An illustration of the timeline encoding given a graph and its contraction sequence. Vertex labels show the elimination ordering. For each time i with contraction $j \leftarrow i$ ($i < j$), we create red sources $\{k, j\}$ for every $k \in \Delta_{>}(j, i)$, which determines red intervals $[i, \min\{k, j\}]$ that will then disappear or transfer at time $\min\{k, j\}$. The red degree corresponds to the number of overlaps of red intervals aggregated by vertices, and its maximum value is the width of the contraction sequence.

- G : input graph with n vertices.
- $\phi : V(G) \rightarrow [n]$: bijection that encodes an elimination ordering (vertex v is eliminated at time $\phi(v)$ if $\phi(v) < n$).
- $p : [n - 1] \rightarrow [n]$: encoding of a contraction tree. For $i < j$, $p(i) = j$ if vertex $\phi^{-1}(i)$ is merged into vertex $\phi^{-1}(j)$ (i.e. j is the *parent* of i in the contraction tree).


For internal data structures, we introduce a few terms. First, define $\Delta_{>}(j, i) := \{\phi(w) \mid w \in \Delta(\phi^{-1}(i), \phi^{-1}(j)), \phi(w) > i\}$. Then, the *red sources* at time t are a set of red edges introduced at time t , defined as $\{\{p(t), k\} \mid k \in \Delta_{>}(p(t), t)\}$. Red sources determine the *red intervals* – non-overlapping, continuous intervals where an edge is red, defined as follows: for $i < j$, red source (i, j) at time t creates an interval $[t, i)$ (red edge ij disappears at time i). If $p(i) \neq j$, then we recurse this process as if red source $\{p(i), j\}$ was created (red edge ij transfers to $\{p(i), j\}$), as illustrated in Figure 1.

Now we aggregate red intervals by vertices. We maintain a *multiset* of intervals for each vertex such that a red interval of an edge accounts to its both endpoints. The maximum number of the overlaps of such intervals gives the maximum red degree at a vertex over time. Finally, we obtain the width of the contraction sequence by taking the maximum of the red degrees over all vertices.

A key observation is that we can dynamically compute the number of overlaps of a multiset of intervals efficiently with a balanced binary tree (e.g. modification in time $\mathcal{O}(\log n)$, getting the maximum number of overlaps in $\mathcal{O}(1)$, etc.). For local search, we implemented methods for modifying a contraction tree and also updating a bijection ϕ .

3 Hydra Decomposition

We also implemented an iterative refinement strategy which we term *hydra decomposition*, based on finding a small vertex separator. A *hydra* is a structured trigraph which consists of a (possibly empty) set of *heads* and a (possibly empty) vertex set *tail*. Each head is a set of vertices containing one *top vertex* and a nonempty set of *boundary vertices*. The neighbors of the top vertex must be a subset of the boundary vertices. All red edges in the trigraph must be incident to one of the top vertices. Heads must be vertex-disjoint, but the tail may contain boundary vertices (but not a top vertex). A *compact hydra* is a hydra consisting of its tail and



■ **Figure 2** Structure of the hydra and two examples of performing a round of hydra decomposition.

one extra vertex, with no heads. A head of a hydra H can additionally be viewed as a compact hydra C , where the boundary vertices of H are the tail of C . Now that we have defined the parts of a hydra, we will now show the operations performed in hydra decomposition:

1. *Separate*: partitions the vertices of a hydra into three parts S, A, B such that S separates A from B . The part S should not contain any vertices from the heads, and any tail vertices cannot be in A . Figure 2 shows two ways of choosing a separator S of a hydra.
2. *Contract*: takes a hydra and contracts all vertices but its tail. The output is a contraction sequence and the resulting compact hydra.
3. *Join*: combines a compact hydra C and another hydra H such that $V(C) \cap V(H)$ is the tail of C . The output is the union of C and H , where the heads and tail of H remain and C becomes an additional head.

We now present a description of `UBSeparate`. Given a graph H and a target width d for a contraction sequence, `UBSeparate` runs *contract* on the original graph without any heads or tails. The *contract* operation works as follows: If the input H is small enough, or a vertex separator of size at most d is not found, we directly search for a contraction sequence of width at most d for all but tail vertices, which can be done by modified `UBGreedy` and other exact algorithms. Otherwise, we perform *separate* to obtain a partition S, A, B . We recursively call *contract* with $H[A \cup S]$ with S being the tail. Then, we have a contraction sequence s_1 and a compact hydra C . Next, we *join* C with $H[B \cup S]$ and obtain a hydra H' . Notice that the tail of C must be S . We again call *contract* with H' and get a contraction sequence s_2 , resulting in a compact hydra C' with the original tail of H . Finally, C' is returned along with the concatenation of s_1 and s_2 as the result of the original *contract* operation.

A key observation is that since red edges reside only in heads and the size of separators are bounded by d , the red degree of a hydra is also upper-bounded by d , which helps construct a d -contraction sequence part by part. For $d = 1$ we use a linear-time algorithm to find a vertex separator, or a cut vertex (articulation point); for $d \geq 2$, we instead call a SAT solver.

References

- 1 Armin Biere and Mathias Fleury. Gimsatul, IsaSAT and Kissat entering the SAT Competition 2022. In *Proc. of SAT Competition 2022 – Solver and Benchmark Descriptions*, volume B-2022-1 of *Department of Computer Science Series of Publications B*, pages 10–11, 2022.
- 2 Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: tractable FO model checking. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 601–612, 2020.
- 3 Reinhard Diestel. *Graph Theory*. Springer Publishing Company, Inc., 5th edition, 2017.
- 4 Michel Habib and Christophe Paul. A survey of the algorithmic aspects of modular decomposition. *Computer Science Review*, 4(1):41–59, 2010.

- 5 André Schidler and Stefan Szeider. A SAT Approach to Twin-Width. In *2022 Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 67–77, 2022.
- 6 Marc Tedder, Derek Corneil, Michel Habib, and Christophe Paul. Simple, linear-time modular decomposition, 2008. arXiv:0710.3901.