

Parameterized Complexity of Database Resilience for Cyclic Queries

Henry Smith, Class of 2027

This past summer, I worked with Professor Freire on investigating the parameterized complexity of the Database Resilience problem for cyclic queries. This problem determines how resilient a database is by finding the minimum number of facts that can be removed to change the database's result when prompted by a certain query. We focused solely on cyclic queries, which are queries whose structure forms a cycle.

Through the lens of classical complexity theory, this problem is NP-hard, and therefore intractable-inefficient to solve. In the worst-case, the runtimes of its algorithms increase exponentially with the size of the input database. However, in parameterized complexity theory, if we can consider a separate parameter like the size of the solution, k , along with the size of the input, then we may be able to derive an algorithm with a more efficient runtime. The aim of this research was to derive a more efficient algorithm than the current ones for the parameterized version of the problem when the query is cyclic.

The current most efficient algorithm we could find for the parameterized version of this problem was a reduction from Database Resilience to the Hitting Set problem (Bodirsky et al). We attempted to find a different reduction, where the problem we reduce to would have a faster runtime than Hitting Set, hoping that the cyclic nature of the query could be leveraged to obtain such an algorithm. Unfortunately, we found that many of the problems we could reduce to that focused on the cyclic nature, like Feedback Vertex Set, were ultimately slower than Hitting Set. Thus, we considered improving the reduction to Hitting Set.

Our main contribution was finding that when the cyclic query is of length 3, the graph produced through our reduction to Hitting Set has its degree bounded by k , as any vertex with degree greater than k must be part of our solution. This allows us to create a problem kernel, which is a smaller but equivalent instance of the problem. Since our graph has bounded degree k , we can determine that Hitting Set on this graph admits a problem kernel of size $2k^2 + k$ with potential for better, possibly linear bounds (Cai). Our kernel is better than the best known general case, which is of size $5k^2 + k$, and the method behind this kernelization suggests that it has even more room for improvement as algorithms for bounded degree Hitting Set improve.

Faculty Mentor: Cibele Freire

Funded by the Bowdoin Computer Science Summer Fellowship

References:

Bodirsky, Manuel, et al. "The Complexity of Resilience Problems via Valued Constraint Satisfaction." *ArXiv.org*, 2023, arxiv.org/abs/2309.15654. Accessed 24 July 2025.

Cai, Xuan. "Linear Kernelizations for Restricted 3-Hitting Set Problems." *Information Processing Letters*, vol. 109, no. 13, 11 Mar. 2009, pp. 730–738, <https://doi.org/10.1016/j.ipl.2009.03.004>.