

# 1 Introduction

Real-world data can often be modeled as a graph, e.g., social networks, biological networks, or infrastructure networks. Covering problems on such networks are often NP-hard. We consider infrastructure networks coming from public transit systems. Solving covering problems in such graphs can potentially help to let these systems run properly, which is important for a reliable, fast, and safe transport of passengers. Weihe [1] considered the problem STATION COVER that defines connections (e.g., trains or buses) as paths in a graph of stations. His algorithm selects the minimum number of vertices, such that every path contains at least one selected vertex. Despite the fact that covering problems are typically NP-hard, his algorithm is surprisingly fast. Thus, the algorithm uses the fact that the NP-hardness is usually based on an unrealistic variant of the problem, while real-world instances typically have certain structural properties that make them easier.

Weihe's algorithm works by first applying a simple set of reduction rules and running a brute-force algorithm on the remaining instance. On real-world instances, applying the reduction rules typically leads to a surprisingly small core. This raises the question, why these reduction rules are so effective. One approach to explain this, is to consider parameterized algorithms that exploit structural properties of the input. Such an algorithm can run in time  $f(k)n^{O(1)}$ , where  $k$  is a parameter that is assumed to be small for realistic instances. Understanding why the reduction rules are so effective is important to close the gap between theory and practice and can help to get more efficient algorithms, also for other covering problems.

We show that the reduction rules reduce the graph at least to its 2-core. Moreover, for every graph, there exists an instance that is completely solved by the reduction rules. On the other hand, there is an instance for every graph where the reduction rules only lead to the 2-core. Beyond that, we show that the problem remains NP-hard even for graphs of treewidth 3. We observe that real-world instances are heterogeneous and have high clustering. Thus, we run experiments on generated instances with varying heterogeneity and clustering. We observe that both properties help the reduction rules to be effective but the clustering appears to be the deciding factor.