

Improved Exact and Heuristic Algorithms for Maximum Weight Clique

Roman Erhardt* Kathrin Hanauer† Nils Kriege†‡ Christian Schulz*
 Darren Strash§

Abstract

We propose improved exact and heuristic algorithms for solving the maximum weight clique problem, a well-known problem in graph theory with many applications. Our algorithms interleave successful techniques from related work with novel data reduction rules that use local graph structure to identify and remove vertices and edges while retaining the optimal solution. We evaluate our algorithms on a range of synthetic and real-world graphs, and find that they outperform the current state of the art on most inputs. Our data reductions always produce smaller reduced graphs than existing data reductions alone. As a result, our exact algorithm, **MWCRedu**, finds solutions orders of magnitude faster on naturally weighted, medium-sized map labeling graphs and random hyperbolic graphs. Our heuristic algorithm, **MWCPeel**, outperforms its competitors on these instances, but is slightly less effective on extremely dense or large instances.

1 Introduction

Finding cliques in graphs is a classic problem in graph theory with many applications. In social networks, group behavior can be predicted with the help of cliques [47]. In biochemistry, cliques can be used to study the interaction between molecules, which can inform drug discovery [33]. Vertex-weighted graphs, and the analogous *maximum weight clique problem* (MWC), can be used in an even wider variety of applications including video object co-segmentation [51], coding theory [52], combinatorial auctions [49], and genomics [4].

Solving the maximum (unweighted) clique problem has been the subject of extensive research [9, 31, 41, 42, 48, 53], with the most effective solvers combining branch-and-bound with MaxSAT reasoning for pruning [30, 38]. However, state-of-the-art algorithms still struggle to find solutions for certain instances in a reasonable time limit. Indeed, there are still unsolved instances, and recently closed instances have required over

a year of computation [50]. Recent work has focused on solving weighted variants of **NP**-hard graph problems [7, 29, 45], which are more difficult in practice.

One powerful technique for tackling **NP**-hard graph problems is to use *data reduction rules*, which remove or contract local graph structures, to reduce the input instance to an equivalent, smaller instance. Originally developed as a tool for parameterized algorithms [13], data reduction rules have been effective in practice for computing an (unweighted) maximum independent set [11, 28, 39] / minimum vertex cover [2], maximum clique [10, 43], and maximum k -plex [12, 25], as well as solving graph coloring [32, 43] and clique cover problems [19, 40], among others [1]. However, recent work has only scratched the surface for *weighted* problems. Lamm et al. [29], Gellner et al. [17], and Gu et al. [20] recently introduce an extensive collection of effective data reductions for maximum weight independent set problem (MWIS), and Wang et al. [45] perform data reduction for weighted graph coloring.

However, to our knowledge, the only data reduction rules for MWC remove vertices simply based on the weight of a neighborhood or the largest weight of a neighbor [7]. Thus, there is untapped potential for reducing input instances further, making them more amenable to exact solving. One strategy is to apply MWIS reductions to the *complement* of the input; however, MWIS reductions are most effective on large, sparse instances and the complements of the graphs considered here are dense and unlikely to fit in memory.

Our Results. We develop a suite of novel exact and heuristic data reduction rules for MWC, with the goal of reducing the number of vertices and edges in the input graph while maintaining solution quality. To the best of our knowledge our data reduction rules are the first to exploit local graph structures for the MWC problem. We also present data reduction rules that are solely aimed at removing *edges* in a graph, which to the best of our knowledge has not been done before for similar problems. After reducing the graph, we apply either heuristic or exact algorithms on the remaining instance to obtain a solution to the original input. We extend the recent reduce-and-peel framework

*Heidelberg University

†University of Vienna, Faculty of Computer Science, Vienna Austria

‡University of Vienna, Research Network Data Science, Vienna, Austria

§Department of Computer Science, Hamilton College

introduced for the MIS and MWIS problems, engineering methods for how and when to apply the reductions and switch to the exact solver. Our experiments show that our algorithms outperform the state of the art.

2 Preliminaries

2.1 Basic Concepts. We consider a simple, weighted, undirected graph $G = (V, E, w)$ with $n = |V|$ and $m = |E|$, where $V = \{1, \dots, n\}$ is the set of vertices, $E \subseteq \{\{u, v\} \mid u, v \in V\}$ is the set of dyadic edges, and $w: V \rightarrow \mathbb{R}_{>0}$ is a function that assigns a positive real-valued weight to each vertex. We extend w to sets, such that for $V' \subseteq V$, $w(V') = \sum_{v \in V'} w(v)$. The *maximum weight* of V' is denoted by $w^*(V') = \max_{v \in V'} w(v)$. Two vertices u and v are *adjacent* (also *neighbors*) if $\{u, v\} \in E$. The (*open*) *neighborhood* $N(v)$ of a vertex $v \in V$ is defined as $N(v) = \{u \in V \mid \{u, v\} \in E\}$, and its *closed neighborhood* is $N[v] = N(v) \cup \{v\}$. Both definitions extend straightforwardly to the neighborhood $N(V')$ of a set of vertices $V' \subseteq V$, i.e., $N(V') = \cup_{v \in V'} N(v) \setminus V'$ and $N[V'] = N(V') \cup V'$. The *degree* of a vertex $\deg(v)$ is the number of its neighbors $\deg(v) = |N(v)|$, and $\Delta := \Delta(G)$ denotes the maximum degree $\max_{v \in V} \deg(v)$. The *complement* of G is defined as $\overline{G} = (V, \overline{E})$, where $\overline{E} = \{\{u, v\} \mid u, v \in V \wedge u \neq v \wedge \{u, v\} \notin E\}$ is the set of edges not present in G . The *density* $\rho := \rho(G)$ of G is the ratio of the number of edges present to those that could exist, $\rho(G) = \frac{2m}{n(n-1)}$. The subgraph *induced* by the subset $V' \subseteq V$ is denoted by $G[V'] = (V', E')$, where $E' = \{\{v_i, v_j\} \in E \mid v_i, v_j \in V'\}$. A set $V' \subseteq V$ is called *independent* if for all pairs of vertices $u, v \in V'$, $\{u, v\} \notin E$.

A *clique* is a set $Q \subseteq V$ where all vertices are pairwise adjacent. A clique in the complement graph \overline{G} corresponds to an *independent set* in the original graph G and vice-versa. The *maximum weight clique problem* (MWC) consists in finding a clique of maximum weight. If $w \equiv 1$, we obtain the *maximum cardinality clique problem* (MCC) (more succinctly referred to as the maximum clique problem). The *maximum independent set problem* (MIS) is that of finding an independent set of maximum cardinality, whereas the *maximum weight independent set problem* (MWIS) asks for an independent set of maximum total weight. The complement of an independent set is a *vertex cover*, i.e. a subset $C \subseteq V$ such that every edge $e \in E$ is incident to at least one vertex in C . The *minimum vertex cover problem*, which asks for a vertex cover with minimum cardinality, is thus complementary to the maximum independent set problem. The maximum clique problem is also dual to the maximum independent set problem and the minimum vertex cover problem via the complement graph \overline{G} . By

extension, the weighted versions of independent set and clique are also dual to each other.

The *vertex coloring problem* asks to assign a color label $c \in \mathbb{Z}$ to each vertex such that no two adjacent vertices have the same label and the number of different colors is minimal. All vertices in a clique must receive different colors. Thus, if a graph has a vertex coloring with k colors, any clique can have cardinality at most k . All these problems are **NP-hard**.

2.2 Related Work. This paper is a summary and extension the master thesis [14]. A lot of research has been done for both the MCC and the MWC problem. As our focus in this work is on the weighted version, we only mention results for MWC and largely omit solvers and results for the cardinality version unless they were extended to the weighted case. A detailed review on approaches for MCC can be found in Wu and Hao [48] as well as in Abu-Khzam et al. [1] in the context of data reductions.

2.2.1 Exact Solvers. Most exact solvers for the MCC use a B&B framework [9], which maintains a current clique C and a candidate set $P = N(C)$ of vertices for extending C . Fast solvers prune the search space by quickly computing a tight upper bound on the clique size that can be found by including vertices from P into C . One successful technique to do so is to compute a greedy heuristic vertex coloring on $G[P]$ and use the number of colors as an upper bound. This approach was subsequently extended to MWC by Kumlander [27] as follows: Given a valid vertex coloring of $G[P]$ that uses k colors and partitions V into color classes $\mathcal{D} = D_1 \sqcup D_2 \sqcup \dots \sqcup D_k$, an upper bound can be computed as $ub(\mathcal{D}) = \sum_{j=1}^k w^*(D_j)$, assuming each color class contributes a vertex of maximum weight.

Fang et al. [15] were the first to implement the idea of MaxSAT reasoning introduced by the MCC solver MaxCLQ [31] for MWC. Jiang et al. [24] also rely on MaxSAT reasoning and contributed an efficient preprocessing step that computes an initial clique \hat{C} as well as a vertex branching ordering. It furthermore computes a simple upper bound on the maximum weight clique that each vertex v can be part of as $w(N[v])$ and removes v if $w(N[v]) \leq w(\hat{C})$. TSM-MWC [23] refines the approach further with a two-stage MaxSAT reasoning approach that applies less expensive MaxSAT techniques to reduce the number of branching vertices before exhaustively looking for disjoint conflicting soft clauses. TSM-MWC currently achieves the best results for a wide spectrum of graph instances, most notably large sparse real-world graph instances, and is the current state-of-the-art exact solver for maximum weight clique.

2.2.2 Heuristic Solvers. The general scheme of a local search algorithm for MCC is as follows: A clique C is constructed by starting with a single vertex and repeatedly adding vertices that are adjacent to all vertices in C using some evaluation function. Again, candidate vertices are those vertices that could potentially be added to C . Once no more add operations can be performed, some vertices can be removed in an attempt to construct a larger clique.

Gendrau et al. [18] proposed two algorithms for MCC based on this strategy: One is a deterministic scheme which adds the vertex with the highest degree first and when no further vertex can be added, the vertex that results in the largest set of candidate vertices is removed. The second algorithm randomly selects which vertex to add to the current solution. Pullan [35] proposed to include a swap operator in the main search procedure. This operator looks for a vertex that is connected to all but one vertex of the current candidate clique C . Furthermore, the algorithm perturbs the current candidate clique by adding a random vertex and removing all non-adjacent vertices from the clique.

This algorithm has been extended to MWC by Pullan [36] by adding a vertex which is randomly chosen only among the vertices of highest weight. Wang et al. [46] added a prohibition rule based on configuration checking. Cai [5] further improved this algorithm by using a better strategy to decide which vertex from the candidate set to add next. This strategy works by randomly sampling k different candidate vertices and choosing the best vertex with respect to some benefit estimation function. Cai and Lin [7] combined the algorithm with data reduction rules in their solver **FastWCLq**. The reductions they use compute upper bounds for each vertex and remove a vertex if one of the computed upper bounds is less than the weight of the current best clique. Every time an improved solution is found by local search, the reductions are reapplied, which in turn improves the chance of local search finding the optimal solution.

SCCWALK41 [44] adopts the previously seen configuration checking strategies as well as data reductions. The authors furthermore introduce a technique called walk perturbation, which adds a random vertex to the solution when the search stagnates and removes all vertices from the candidate set that become invalid by this perturbation. Cai et al. [8] improved **FastWCLq** further to also apply a reduction-and-hill-climbing method based on vertex coloring.

SCCWALK41 and **FastWCLq** are the current state-of-the-art for heuristic MWC solvers, with the former being especially dominant in small dense networks, such as graphs from the DIMACS and BHOSLIB challenge [44],

and the latter showing the best results in large sparse real-world networks [8].

3 Data Reductions

So far, only few reductions are known that can be used for the MWC. However, especially for large instances, applying exact data reductions is a very important technique to decrease the problem size. In general, reductions allow the classification of vertices as either (1) part of a solution, (2) non-solution vertices, or (3) deferred, i.e. the decision for this vertex depends on additional information about neighboring vertices that will be obtained later. We denote by \mathcal{K} the resulting *reduced graph*, where no reduction rule applies anymore. In the following, we review existing and introduce a large set of new reductions for the MWC.

3.1 Neighborhood Weight Reduction. A simple but effective reduction often seen in literature [7, 8, 23, 24, 44] is based on the upper bound $w(N[v])$ for any clique containing $v \in V$.

REDUCTION RULE 1. ([7]) *Let \hat{C} be the highest-weight clique found so far and let $v \in V$ s.t. $w(N[v]) \leq w(\hat{C})$. Then v can be removed from the graph without reducing the maximum solution weight.*

The rule can be applied on a vertex $v \in V$ in $\mathcal{O}(1)$ time, given that the neighborhood weight is stored and maintained throughout the reductions.

3.2 Largest-Weight Neighbor Reduction. Cai et al. [7] tighten the neighborhood weight reduction rule by either including or excluding the highest weight vertex u^* in the neighborhood.

REDUCTION RULE 2. ([7]) *Let \hat{C} be the highest-weight clique found so far, let $v \in V \setminus \hat{C}$, and let $u^* = \arg \max_{u \in N(v)} w(u)$. If $\max\{w(N[v]) - w(u^*), w(N[v] \cap N[u^*])\} \leq w(\hat{C})$, then v can be removed from the graph without reducing the maximum solution weight.*

For applying the rule on a vertex $v \in V$, first its highest weight neighbor u^* is identified in $\mathcal{O}(\deg(v))$ and then the intersection of their neighborhoods is computed in $\mathcal{O}(\min\{\deg(v), \deg(u^*)\})$, resulting in overall $\mathcal{O}(\deg(v))$ time. Computing the intersection of neighborhoods is a crucial operation for the application of this reduction rule as well as several others described in the following. The running time for computing $N(u) \cap N(v)$ depends on the graph representation. Assuming constant time for checking whether two vertices are adjacent, we can iterate over the smaller set and identify those that are also adjacent to the other vertex in $\mathcal{O}(\min\{\deg(u), \deg(v)\})$ time. For the application to

large sparse graphs we use an adjacency list and realize the operation using indicators by iterating over the neighbors of both vertices in $\mathcal{O}(\deg(u) + \deg(v))$ time.

3.3 Twin Reduction. We now introduce our first new data reduction rule, based on twins. Consider two adjacent vertices u and v that share the same closed neighborhood. Such vertices are called *twins*. If either one of them is in the solution, then the other one must also be in it. Figure 1 gives an illustration.

REDUCTION RULE 3. *Let $u, v \in V$, $u \neq v$, and $N[u] = N[v]$. Then u and v can be contracted to a new vertex $\{u, v\}$ with weight $w(\{u, v\}) = w(u) + w(v)$ and $N(\{u, v\}) = N(u) \cap N(v)$ without reducing the maximum solution weight.*

Proof. Suppose there is an optimal solution C^* that, w.l.o.g., contains u , but not v . Then it is always possible to add v to the solution, as it is connected to all neighbors of u , resulting in a solution of larger weight. Hence, each optimal solution contains either both u and v or neither. \square

To check the precondition for two vertices $u, v \in V$ where $\deg(u) = \deg(v)$, the intersection of their neighborhoods can be obtained in time $\mathcal{O}(\deg(v))$ using a marking scheme.

3.4 Domination Reduction. Vertex $u \in V$ is said to dominate $v \in V$ when $N(v) \subseteq N(u)$. Furthermore, if $w(v) \leq w(u)$, then a maximal clique containing u would have a weight greater or equal to one that contains v . This observation leads to the following reduction rule:

REDUCTION RULE 4. *Let $u, v \in V$, $\{u, v\} \notin E$, $N(v) \subseteq N(u)$, and $w(v) \leq w(u)$. Then, v can be removed from the graph without reducing the maximum solution weight.*

Proof. Suppose there is an optimal solution C^* that, w.l.o.g., contains v , but not u . As u is adjacent to all neighbors of v , it is always possible to substitute v with u in the solution, resulting in a solution with at least the same weight since $w(v) \leq w(u)$. As $\{u, v\} \notin E$, no clique can contain both u and v . Hence, there is at least one optimal solution that does not contain v . \square

Given v , we find vertices u with $N(u) \supseteq N(v)$ as follows: We choose $x \in N(v)$ arbitrarily and iterate over all $u' \in N(x)$. If $\{u', v\} \notin E$, $\deg(u') \geq \deg(v)$, and $w(u') \geq w(v)$, we test whether $N(v) \subseteq N(u')$ in $\mathcal{O}(\deg(v))$ time. The approach identifies all vertices u' for a given vertex v satisfying the conditions of the reduction rule of Lemma 4 in $\mathcal{O}(\deg(v) \cdot \Delta)$ time.

We now introduce our first reduction that is designed to remove *edges* from the graph. A similar reduction is applicable if u and v are adjacent. However, simply removing v is not possible, as v may be part of a clique containing u . Therefore, we add the weight of u to v and then remove the edge $\{u, v\}$, thus preserving the best solution achievable by v and u being in the same clique while reducing the graph at the same time.

REDUCTION RULE 5. *Let $u, v \in V$, $\{v, u\} \in E$, and $N(v) \subseteq N[u]$. Then, increasing $w(v)$ to $w'(v) = w(v) + w(u)$ and removing the edge $\{u, v\}$ from the graph does not reduce the maximum solution weight.*

Proof. Let C^* be an optimal solution in the *original* graph. Assume that C^* contains v , but not u . Then u can be added to C^* leading to a higher weight, contradicting the assumption that C^* is optimal. Hence, if C^* contains v , it also contains u . There are two cases left to consider:

Case 1: If C^* contains both u and v , then $w(C^*) \leq w(u) + w(v) + w(N(v) \setminus \{u\}) = w'(v) + w(N(v) \setminus \{u\})$, so there exists an equivalent solution only containing v in the reduced graph.

Case 2: If C^* contains u but not v , then $w(C^*) \leq w(u) + w(N(u))$, and the same solution exists in the reduced graph. \square

The reduction can be implemented analogously to the twin reduction (Reduction Rule 3).

3.5 Edge Bounding Reduction. This rule is a natural extension to Reduction Rule 2, using the computed bounds not only to decide whether a vertex can be removed, but also the edge that connects it with its highest-weight neighbor. Given a vertex $v \in V$ and its highest-weight neighbor $u^* \in N(v)$, let $ub_{inc}(v, u^*)$ denote the *including upper bound* $w(v) + w(u^*) + w(N(v) \cap N(u^*))$ and let ub_{exc} be the *excluding upper bound* $w(N[v]) - w(u^*)$. Reduction Rule 2 states that v can be removed if both $ub_{inc}(v, u^*) \leq w(\hat{C})$ and $ub_{exc}(v, u^*) \leq w(\hat{C})$, where \hat{C} is the currently best solution. The extension provided by the edge bounding reduction is based on the observation that if $ub_{exc}(v, u^*) > w(\hat{C})$, but $ub_{inc}(v, u^*) \leq w(\hat{C})$, it is possible to remove the edge $\{v, u^*\}$. We extend this rule to apply to all neighbors of v :

REDUCTION RULE 6. *Let $v \in V$, $u \in N(v)$, and let \hat{C} be the best clique found so far. If $ub_{inc}(v, u) < w(\hat{C})$, the edge $\{v, u\}$ can be removed from the graph without reducing the maximum solution weight.*

Proof. The value $ub_{inc}(v, u)$ is an upper bound on the weight of any clique containing both v and u . If a clique




Figure 1: Twin reduction (left) and simplicial vertex removal reduction (right) for MWC.




Figure 2: Domination reduction for MWC by applying Reduction Rule 4 (left) and Reduction Rule 5 (right).

\hat{C} with weight $w(\hat{C}) > ub_{inc}(v, u)$ is known, then there is at least one optimal solution C^* that does not contain both v and u . The edge $\{v, u\}$ is thus irrelevant in the search for a solution of higher weight. \square

Given an edge $\{v, u\}$, the time complexity is $\mathcal{O}(\min\{\deg(v), \deg(u)\})$, as with Reduction Rule 2.

3.6 Simplicial Vertex Removal Reduction. A vertex v is called *simplicial* if its closed neighborhood forms a clique C_v , i.e. $\forall x_1, x_2 \in N[v], \{x_1, x_2\} \in E$. Simplicial vertices may be removed before applying a maximum weight clique solver as well: Once a simplex v has been identified, the largest clique it can be part of is C_v with $w(C_v) = w(N[v])$. If this weight is larger than the currently known highest-weight clique, the lower bound is updated.

REDUCTION RULE 7. *Let $v \in V$ be a simplicial vertex and let \hat{C} be the best clique found so far. Only if $w(N[v]) > w(\hat{C})$, set $\hat{C} = N[v]$. In any case, removing v from the graph then does not reduce the maximum solution weight.*

Proof. If $w(N[v]) \leq w(\hat{C})$, v cannot be part of a strictly better solution. Otherwise, if $w(N[v]) > w(\hat{C})$, the same holds after the currently best solution has updated to $\hat{C} = N[v]$. \square

Testing the adjacency of each pair of vertices in $N[v]$ takes $\mathcal{O}(\deg(v)^2)$ in the worst case.

Algorithm 1 Reduce graph G via data reductions and improve clique \hat{C} , using vertex degree limits if *lim* is set.

```

REDUCE( $G = (V, E, w), \hat{C}, \text{lim}$ )
  if lim then set vertex degree limit to  $0.1\Delta$ 
  initialize  $D_i$  for each reduction rule  $r_i$ 
  repeat
    for all reductions  $r_i$  do
      if  $r_i$  is not paused and  $D_i \neq \emptyset$  then
        apply  $r_i$  on all vertices in  $D_i$   $\triangleright$  Section 3.7
        update  $D_i$ 
        if reduction rate not achieved then
          pause  $r_i$ 
          update  $\hat{C}$  via local search  $\triangleright$  Section 4.1
        else if  $r_i$  paused  $\wedge$   $G$  reduced enough then
          unpause  $r_i$   $\triangleright$  Section 3.7
    if lim and all reductions paused then
      increase limit on vertex degrees, update  $D_i$ 's
  until all reductions paused and degree unlimited
  return  $G, \hat{C}$ 

```

Observe that in contrast to the other reductions, the simplicial Vertex reduction may directly improve the currently best solution \hat{C} .

3.7 Applying the Reductions. For applying the exact reduction rules proposed in this section, an adapted version of the strategy from Hespe et al. [22] that entails both dependency checking and reduction tracking is used. Specifically, the set of reductions $\{r_i\}$

is iterated, where each rule r_i is tried on its set of viable vertices D_i , which is initially set to $D_i = V$. After preliminary experiments, we settled on the following order of reductions: neighborhood weight, twin, simplicial vertex, edge bounding (which includes largest-weight neighbor), domination case 1, and domination case 2. Every time a rule r_i fails to *reduce* a vertex, i.e. to remove it from the graph, this vertex is removed from the set of viable candidates D_i . Otherwise, the set of each rule r_j is updated to $D_j = D_j \cup N(v)$ and the applicable vertices or edges are removed from the graph. This minimizes redundant computations without affecting the final size of the reduced graph [22].

Reduction tracking aims at tracking the effectiveness of reductions. Slightly different from the original strategy, reduction tracking is implemented by pausing a reduction once it fails to achieve a reduction rate of at least 1% of the current number of vertices or edges per second, until other reductions reduced the graph by that amount. Reduction tracking is checked both in between the application of different reduction rules as well as periodically during the iteration over candidate vertices, in order to prevent single reductions to delay the solver and allow either more efficient reductions or the exact solver to take over. Another addition to the strategy by Hespe et al. is to set a dynamic limitation on the degree of vertices that are tried in the reductions. The limit is set to 10% of the highest degree initially and is increased by 10% whenever the reductions have been exhaustively applied in the previous level. This guarantees that reductions applicable on low degree vertices, which are typically more efficient, are applied first. The loop terminates once the degree is no longer limited and all reductions are paused, at which point we run either an exact or heuristic solver on the reduced graph.

4 MWCRedu: A New Exact Algorithm

Our exact algorithm MWCRedu works in two stages: First, the set of exact reduction rules from Section 3 is used to reduce the graph. Second, the reduced graph is passed to an exact B&B solver to compute the final solution.

4.1 Computing a Lower Bound. Reduction Rules 1, 2 and 6 depend on the currently best solution \hat{C} to be applicable. For computing bounds, fast heuristics are generally preferred, since spending more time on improving the initial solution typically gives diminishing returns. A well-suited heuristic for computing an initial lower bound is the one employed in Jiang et al. [24]: Repeatedly remove the vertex with the smallest vertex degree from the graph until all remaining vertices are pairwise adjacent and form the initial clique \hat{C} , which yields an initial lower bound of $w(\hat{C})$.

Afterwards, \hat{C} is continuously improved by the simplicial vertex reduction (Reduction Rule 7) and the local search algorithm from FastWCLq [8], the latter being applied on the reduced graph in between checking each reduction rule. Subsequently, \hat{C} provides the lower bound in the Reduction Rules 1, 2 and 6, and it also serves as the initial solution for the solver that is applied on the reduced graph. Algorithm 1 gives an outline.

4.2 Branch and Bound. The reduced graph is solved using the branch and bound paradigm. As the procedure has exponential time complexity, it is important to choose a good ordering and to reduce the set of branching vertices by computing tight upper bounds. We use the same ordering as Jiang et al. [24], i.e. the ordering of the vertices is given as $v_1 < v_2 < \dots < v_n$, where v_1 has the smallest vertex degree, v_2 has the smallest vertex degree after v_1 is removed, etc. Such an ordering is called a *degeneracy ordering* of the graph.

To compute tight upper bounds and reduce the set of branching vertices, we apply efficient MIS- and MaxSAT-based approaches from [24, 23] throughout the search. Recall from Section 2.2 that for any vertex coloring that partitions V into color classes $\mathcal{D} = D_1 \sqcup D_2 \sqcup \dots \sqcup D_k$, each color class forms an independent set and $ub(\mathcal{D}) = \sum_{j=1}^k w^*(D_j)$ is an upper bound on the maximum clique weight. The set of branching vertices is then further reduced via the two-stage MaxSAT reasoning approach from TSM-MWC [23].

In the first stage, which the authors refer to as binary MaxSAT reasoning, the set of branching vertices is reduced by inserting as many vertices as possible into the independent sets s.t. $\sum_{j=1}^{k'} w^*(D_j) \leq w(\hat{C})$. As these vertices cannot form a clique with a weight larger than $w(\hat{C})$ by themselves, they can be removed from the set of branching vertices. If a vertex $v_i \in V$ has neighbors in all existing independent sets but $ub + w(v_i) \leq w(\hat{C})$ holds, it is inserted as a new independent set. Otherwise we try to split its weight among independent sets that do not contain any of its neighbors by adding v_i with weight $w^*(S_j)$ into independent S_j and updating the weight to $w(v_i) = w(v_i) - w^*(S_j)$ for $j = 1, 2, \dots, k'$, until its remaining weight is given as $\delta = w(v_i) - \sum_{j=1}^{k'} w^*(S_j)$. If $\delta > 0$ and $ub + \delta \leq w(\hat{C})$, v_i is inserted as a new independent set with weight δ , otherwise the weight splitting procedure is undone and v_i is kept in the set of branching vertices.

In the second stage, called ordered MaxSAT reasoning, the set of branching vertices is reduced further by detecting disjoint conflicting subsets of independent sets. Firstly, the weight of a branching vertex v_i is again split among the independent sets $\{S_1, S_2, \dots, S_{k'}\}$ that do not contain any of its neighbors, resulting in the re-

maining weight $w(v_i) = \delta > 0$, since the vertex was not removed from the set of branching vertices in the first stage. After that, the algorithm tries to find a set of independent sets $\{U_1, U_2, \dots, U_r\}$ that each contain exactly one neighbor u of v_i . It then looks for an independent set D_q s.t. $D_q \cap N(v_i) \cap N(u) = \emptyset$ for any U_j , proving that the sets $\{\{v_i\}, U_j, D_q\}$ are conflicting. In this case, ub can be further improved to $ub + \delta - \beta$, where $\beta = \min(\delta, w^*(U_j), w^*(D_q))$ [23].

Finally, if after considering all $U_j \in \{U_1, U_2, \dots, U_r\}$ ub is still higher than the lower bound, ub is reduced by identifying conflicting subsets via unit propagation as first implemented for maximum weight clique [15]. Unit propagation works from the idea that clauses with more literals are more likely to be satisfied and are thus considered *weaker* clauses. A unit clause is thus the strongest clause since it only has one possibility of evaluating to **true**. The algorithm repeatedly satisfies such a clause, removing all occurrences of the contained literal from the other clauses. If an empty clause remains, the set of clauses is identified as conflicting. Each time a set of conflicting clauses $\{S_0, S_1, \dots, S_r\}$ is identified, the upper bound can be reduced by $\delta = \min\{w^*(S_1), \dots, w^*(S_r)\}$. To tighten the bound further, each S_j ($0 \leq j \leq r$) is split into S'_j and S''_j so that $w^*(S'_j) = \delta$ and $w^*(S''_j) = w^*(S_j) - \delta$. S'_j then represents the conflicting subset found so far, whereas further conflicts can be deduced from S''_j [15].

The procedure is run at every branch of the solver in order to reduce the amount of work to be done. The algorithm terminates when all branches are either explored or pruned or when the time limit is reached, in which case the best solution found is reported.

5 MWCPeel: A New Heuristic Algorithm

For our new heuristic algorithm **MWCPeel**, we investigate vertex peeling techniques, which remove vertices from the graph that are assigned the lowest scores by some heuristic rule. This rule must therefore capture the likelihood of a vertex belonging to the solution as well as possible. Using the vertex degree is an obvious choice for MCC, since a vertex with a high degree is more likely to form a large clique. Furthermore, a vertex v cannot be part of a clique larger than $\deg(v)$. For the measure to remain an upper bound in the context of MWC, the weight of the neighborhood of each vertex is taken into account. The resulting simple and intuitive scoring measure $w(N[v])$ is used in our peeling step.

Overall, our heuristic solver works similarly to the exact approach **MWCRedu** described in Section 4, but implements the peeling reduction on top of the previously introduced exact reductions: We first run exact reductions exhaustively. On the reduced graph,

Algorithm 2 Heuristic Solver MWCPeel

```

MWCPEEL( $G = (V, E, w)$ )
  compute initial clique  $\hat{C}$  ▷ Section 4.1
  repeat ▷ Algorithm 1
     $G, \hat{C} \leftarrow \text{REDUCE}(G, \hat{C}, \text{isFirstIteration})$ 
     $\mathcal{N} \leftarrow \# \text{vertices to peel off}$  ▷ Section 5.1
    remove  $\mathcal{N}$  vertices  $v$  with lowest score  $w(N[v])$ 
  until stopping criteria met ▷ Section 5.2
  return TSM-MWC( $G, \hat{C}$ )

```

we apply our peeling strategy that removes vertices that are unlikely to be part of a large clique. We repeat the process until the remaining graph is small or the scores of the peeling reductions are not sufficiently large, and then apply the exact algorithm on the remaining graph. Algorithm 2 gives an overview.

5.1 Peeling Strategy. Chang et al. [11] introduced a reduce-and-peel heuristic technique to repeatedly remove the minimum degree vertex from a graph, adding it to a growing independent set. For MWC, a straightforward approach is to remove the vertices with the lowest score and exclude them from the solution. More precisely, we remove a fixed percentage of the currently remaining vertices in each peeling step. The number of vertices to be peeled off in one step, \mathcal{N} , is dynamically determined as follows:

$$\mathcal{N} = \begin{cases} 0.1n & \text{if } n > 50,000, \\ \max\{0.01n, \frac{0.01}{50,000}n\} & \text{otherwise,} \end{cases}$$

where n always refers to the current number of vertices and the threshold of 50,000 has proven itself suitable in preliminary experiments. Without the differentiation between larger and smaller graphs, the exact reductions would often be reapplied on many vertices, which would significantly slow down the solver. Furthermore, as the vertex degrees often follow a power-law distribution in real-world graph instances [21], the size of the optimal solution makes up a smaller portion of the graph for large graphs. After each peeling step, the viable candidate sets are updated and exact reductions are applied again.

5.2 Stopping Criteria. Another important decision is when to stop applying the peeling reduction; stopping too early could result in a much higher running time for the solver applied on the reduced graph, whereas stopping late might negatively impact the solution quality. Since the optimal amount of vertices to reduce is highly dependent on the graph structure, a static stopping criterion is unlikely to be a good strategy. For this reason, we employ a dynamic strategy that

works by comparing the current computed score with previously computed scores.

The first stopping criterion is the deterioration of the maximum score value below a certain threshold relative to the total maximum score value. This indicates that the peeling reduction begins to reduce the maximum solution.

A second stopping criterion takes effect if the difference between the minimum and maximum score shrinks below a certain threshold. This shows that the scoring model can no longer clearly distinguish high quality vertices from low quality vertices.

We set both thresholds to 90 % to achieve a good balance between speed-up and solution quality. As a fail-safe, a backup of the current graph state is created before applying the heuristic reduction, which can be reloaded in the case the graph is reduced to zero. After the reduction procedure, the branch-and-bound solver is applied on the reduced graph to obtain the final result.

6 Experimental Evaluation

We implemented our new solvers **MWCRedu** and **MWCPeel** and evaluate them against the state-of-the-art solvers in their class on an extensive and diverse set of instances. More precisely, we compare our exact solver **MWCRedu** with the currently best exact solver **TSM-MWC** on each dataset, and our heuristic solver **MWCPeel** with the currently best heuristic solvers **FastWCLq** and **SCWalk41**.

Methodology. The experiments were run on an Intel Xeon Silver 4216 CPU @2.10GHz with 16 cores under Linux with 95 GB of RAM. All solvers are implemented in C/C++ and compiled using GNU g++ with full optimization (-O3). Each solver was executed on up to 16 graph instances in parallel. As the solvers were run exclusively on the machine, there is no relevant difference to solving the graph instances sequentially. We always report the solution quality $w(\hat{C})$ and the time to find that solution t_{sol} . For exact solvers, we additionally give the time needed to prove optimality of the solution t_{prv} . Solvers that use random number generation are run five times with different seeds and we report their average solutions to better capture their general performance. If an exact algorithm did not finish within a time limit of 3,600 seconds, it is halted and the best solution found so far is output. Heuristic algorithms are stopped after 1,000 seconds.

Instances. We evaluate our algorithms on a broad selection of graphs, covering different sizes, densities, weightings and areas of application. Some of the graphs are originally unweighted and thus were assigned weights artificially. For each unweighted graph, weights




Figure 3: Original graph sizes and reduced graph sizes for old and old + new reductions.

are drawn uniformly from the range $[1, 200]$.¹

We compiled four sets of instances, with 58 instances altogether: **OSM** contains 12 naturally-weighted map labeling instances from Cai et al. [6], generated from OpenStreetMap data using the technique of Barth et al. [3]. The 10 instances in **REP** are real-world graphs from the network data repository [37], and the 23 instances in **DIMACS** were taken from the second DIMACS implementation challenge [26]. Moreover, we use 13 random hyperbolic graphs (**RHG**). These are randomly generated graphs such that the vertex degrees follow a power-law distribution [34] and were generated by the KaGen framework [16]. We varied the power-law exponent between 1.75 and 2.25 and chose the average degree between 100 and 500. For **REP**, **DIMACS**, and **RHG**, we assigned artificial weights as described above. See Table 8 in Appendix D for detailed per-instance statistics.

6.1 Impact of New Data Reduction Rules. We first investigate the impact of the reduction rules on the instances and compare the effect of adding our “new” rules to the “old” ones that are described in current literature. Table 3 shows reduced graph sizes on all instances, and Table 1 shows reduced graph sizes for a subset of instances.

On the **DIMACS** instances, the new data reduction rules do not help to compute smaller reduced graphs (hence they are excluded from the table). This is expected as these instances are dense and data reduction rules tend to work well on sparse instances. On the other instances, reduced graphs are significantly smaller when the new data reduction rules are employed additionally.

The largest reduction in the **REP** instance set is observed on **web-wikipedia.link_it**, where the new reduction rules result in an empty reduced graph, i.e., the instance is fully solved by the reductions only. The

¹Other weight distributions such as power-law and exponential gave similar results and were excluded due to space constraints.

Graph	Reduced Graph Size			
	old+new reductions		old reductions only	
	absolute	% of n_0	absolute	% of n_0
REP				
bio-human-gene1	3,915	17.57	4,485	20.13
sc-TSOPF-RS-b2383	16,123	42.29	37,737	98.99
soc-orkut	1,264,963	42.21	1,521,404	50.76
web-wikipedia-link.it	0	0.00	1,214	0.04
web-wikipedia-growth	83,724	4.48	637,483	34.08
RHG				
rhg_250k_100_1.75	7	0.00	1,061	0.42
rhg_500k_500_2.25	0	0.00	1,761	0.35
rhg_750k_250_2.25	15	0.00	1,062	0.14
rhg_750k_500_1.75	4,445	0.59	7,341	0.98
rhg_750k_500_2.25	12	0.00	2,651	0.35
OSM				
district-of-columbia-AM2	0	0.00	759	5.58
greenland-AM3	0	0.00	1,768	35.46
idaho-AM3	0	0.00	2,293	56.42
massachusetts-AM3	0	0.00	802	21.66
virginia-AM3	0	0.00	907	14.66

Table 1: Selected instances and reduced graph sizes (number of nodes) when both old and new data reductions rules are applied vs. reduced graph sizes obtained when only running reductions from the current literature. Smaller is better. n_0 refers to the initial number of nodes.

biggest improvement occurred on `sc-TSOPF-RS-b2383`, where the old rules were barely effective and reduced the number of nodes by only roughly 1%. In combination with the new rules, however, the computed reduced graph contains only 42.29% of the nodes of the original instance. On all instances, using the new rules in addition to the old ones always resulted in smaller reduced graphs than when just using the old ones. On average, the old rules alone reduced the graph size by about 67%, which improved to over 80% when combined with our new rules.

The new rules also work very well on the RHG instances and consistently produced smaller reduced graphs than when just using the old ones. Generally, the reductions are very efficient on these instances. If using only the old rules, the resulting reduced graphs are reduced to between 0.14% and 1.64% of the original graph sizes. Combined with the new rules, the range is between 0% and 0.59%. Two RHG instances were reduced to zero nodes when using the new rules in addition to the old ones. On average, the reduced graphs obtained by old and new rules together were only 0.05% of the original graph sizes, whereas the average

for the old rules alone was more than ten times larger.

The new and old rules together computed empty reduced graphs on all OSM instances, which never happened when using only the rules from the literature. On average, the old rules reduced the number of vertices down to 25.4%, where the range is relatively large and between 5.58% on `district-of-columbia-AM2` and 56.42% on `idaho-AM3`.

In summary, our new reduction rules distinctly and consistently produce smaller reduced graphs on all REP, RHG, and OSM instances and even compute empty reduced graphs on 15 instances, which the old ones alone never accomplished on any instance of our collection. Figure 3 summarizes this visually.

6.2 Exact Algorithms. We discuss the aggregated results for each of the four instance sets (see Table 2).

Our algorithm `MWCRedu` is more than an order of magnitude faster in the geometric mean than `TSM-MWC` on the OSM instances (Table 4), both with respect to time to find the solution t_{sol} and to prove optimality t_{prv} . It is also consistently faster than `TSM-MWC` on each of the twelve instances in the set. As both are exact

Instance Set	t_{sol}		t_{prv}		$w(\hat{C})$	
	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu
Exact Results						
DIMACS	1,106.99	946.34	1,714.98	1,650.88	6,460	6,587
REP	117.15	134.60	190.45	259.14	14,092	14,321
RHG	95.93	14.55	128.16	17.73	106,210	106,781
OSM	27.62	1.55	31.01	2.43	537,149	542,993
	136.15	41.14	189.82	65.55	47,738	48,359
Instance Set	t_{sol}			$w(\hat{C})$		
	FastWCLq	SCCWalk41	MWCPeel	FastWCLq	SCCWalk41	MWCPeel
Heuristic Results						
DIMACS	193.11	4.46	91.31	6,792	6,968	6,547
REP	190.36	345.88	52.12	14,190	9,382	14,056
RHG	40.36	513.37	10.69	106,781	63,647	106,731
OSM	4.45	64.29	1.45	542,993	528,956	542,993
	50.68	84.49	16.49	48,620	38,516	48,056

Table 2: Overview of results for exact (top) and heuristic (bottom) algorithms as geometric mean per graph set.

algorithms, the solution weights are identical except for two cases, where TSM-MWC failed to find the optimal solution within the time limit and stopped prematurely with a worse result. Thus, MWCRedu dominates here.

On DIMACS (Table 4), no major difference in performance between the two solvers is observable. MWCRedu was able to finish on nine of the 23 instances within the time limit, whereas TSM-MWC finished on only eight instances. The running times generally lie very close together, and the solution weights are identical except for seven cases. The reason for the similar behavior is that none of the new exact reductions employed by MWCRedu is able to remove vertices or edges for any instance in this set. Thus, the solver quickly proceeds to apply the B&B solver, which uses the same techniques as TSM-MWC. The overhead from applying the reduction rules is only notable for the easier instances. On average over those instances, where both finished regularly, MWCRedu performs slightly better, which is likely due to better initial solutions obtained from running local search during the reduction phase.

On the REP instances (Table 5), the results are mixed. MWCRedu and TSM-MWC both outperform the respective other algorithm for some instances. On three instances, TSM-MWC failed to prove optimality of a solution and terminated with a suboptimal result twice. TSM-MWC is very efficient for large instances with more than 1 000 000 vertices, whereas MWCRedu outperforms TSM-MWC on the smaller, more dense biology graphs.

On RHG (Table 5), MWCRedu outperforms its competitor TSM-MWC clearly. While TSM-MWC runs into a time-out twice and terminates with a suboptimal solution, MWCRedu always finishes regularly and is the faster algorithm except on one instance. Its dominance in running time is pronounced and up to two orders of magnitude. The reason for MWCRedu’s good performance is likely the structure of the instances, which allows it to remove most vertices quickly using very efficient reductions.

In summary, MWCRedu is clearly the better algorithm on the OSM and RHG instances and on par with TSM-MWC on the DIMACS graphs. On instances that are small and dense, such as in the REP set, TSM-MWC may be the faster algorithm, whereas MWCRedu can play out its strengths on very large ones. Notably, MWCRedu finished within the time limit on the same instances as TSM-MWC plus some more, making it the *more reliable candidate*.

6.3 Heuristic Algorithms. We now compare our heuristic solver MWCPeel against the state-of-the-art solvers FastWCLq and SCCWalk41 and discuss the differences on each of the four instance sets. Aggregated results are presented in Table 2.

As shown in Table 6, MWCPeel performs best for 11 out of 12 OSM instances. Both MWCPeel and FastWCLq find the optimal solution to all instances.

For the DIMACS graphs (Table 6), SCCWalk41 clearly dominates its competitors. Between FastWCLq and MWCPeel, FastWCLq mostly computes slightly higher

weight solutions, though it takes longer to compute them. Looking at the instances where TSM-MWC fails to find the optimal solution, both **FastWCLq** and **MWCPeel** achieve higher weight solutions in a much smaller amount of time for most of them.

As shown in Table 7, performance on **REP** graphs is very competitive among the heuristic solvers. While all algorithms compute the best solution an approximately equal amount of times, the solution quality of **SCCWalk41** is the lowest on average. Taking speed into account, **MWCPeel** shows a good performance in comparison. On average, **MWCPeel** is a factor 3.7 faster than the second fastest algorithm **FastWCLq** which computing 0.9% better solutions on average than **MWCPeel**. It should be noted, however, that our exact solver **MWCPeel** computes even higher weight solutions than **FastWCLq**, while also being faster on average.

The results for **RHG** are presented in Table 7. Here, **MWCPeel** outperforms the other solvers in 31 out of 39 instances. While **FastWCLq** sometimes finds a slightly higher weight solution than **MWCPeel**, it has a higher running time on average (a factor 3.8). **SCCWalk41** is clearly outperformed both in speed and solution quality.

7 Conclusion

We presented an exact algorithm called **MWCRedu** and a heuristic algorithm called **MWCPeel** for solving the maximum weight clique problem. Our algorithms interleave successful techniques from related work with novel data reduction rules that use local graph structures to identify and remove vertices and edges while maintaining the optimal solution. In experiments on a large range of graphs, we find that they outperform the current state-of-the-art solvers on most inputs. In particular, **MWCRedu** is faster by orders of magnitude on naturally weighted, medium-sized street network graphs and random hyperbolic graphs. **MWCPeel** outperforms its competitors on these instances, but is slightly less effective on extremely dense or large instances. In future work, we want to consider parallelization of our approaches. Given the good results of our algorithms, we plan to release them as open source.

Acknowledgments. We acknowledge support by DFG grant SCHU 2567/3-1. N.K. was supported by the Vienna Science and Technology Fund (WWTF) through project VRG19-009.

References

- [1] Faisal N. Abu-Khzam, Sebastian Lamm, Matthias Mnich, Alexander Noe, Christian Schulz, and Darren Strash. Recent advances in practical data reduction. In Hannah Bast, Claudius Korzen, Ulrich Meyer, and Manuel Penschuck, editors, *Algorithms for Big Data: DFG Priority Program 1736*, pages 97–133. Springer Nature Switzerland, Cham, 2022.
- [2] T. Akiba and Y. Iwata. Branch-and-reduce exponential/FPT algorithms in practice: A case study of vertex cover. *Theor. Comput. Sci.*, 609, Part 1:211–225, 2016.
- [3] Lukas Barth, Benjamin Niedermann, Martin Nöllenburg, and Darren Strash. Temporal map labeling: a new unified framework with experiments. In Siva Ravada, Mohammed Eunus Ali, Shawn D. Newsam, Matthias Renz, and Goce Trajcevski, editors, *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, GIS 2016, Burlingame, California, USA, October 31 - November 3, 2016*, pages 23:1–23:10. ACM, 2016.
- [4] Sergiy Butenko and Wilbert E Wilhelm. Clique-detection models in computational biochemistry and genomics. *European Journal of Operational Research*, 173(1):1–17, 2006.
- [5] Shaowei Cai. Balance between complexity and quality: Local search for minimum vertex cover in massive graphs. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [6] Shaowei Cai, Wenying Hou, Jinkun Lin, and Yuanjie Li. Improving local search for minimum weight vertex cover by dynamic strategies. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 1412–1418. International Joint Conferences on Artificial Intelligence Organization, 7 2018.
- [7] Shaowei Cai and Jinkun Lin. Fast solving maximum weight clique problem in massive graphs. In *IJCAI*, pages 568–574, 2016.
- [8] Shaowei Cai, Jinkun Lin, Yiyuan Wang, and Darren Strash. A semi-exact algorithm for quickly computing a maximum weight clique in large sparse graphs. *Journal of Artificial Intelligence Research*, 72:39–67, 2021.
- [9] Randy Carraghan and Panos M Pardalos. An exact algorithm for the maximum clique problem. *Operations Research Letters*, 9(6):375–382, 1990.
- [10] Lijun Chang. Efficient maximum clique computation and enumeration over large sparse graphs. *VLDB J.*, 29(5):999–1022, 2020.
- [11] Lijun Chang, Wei Li, and Wenjie Zhang. Computing a near-maximum independent set in linear time by reducing-peeling. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 1181–1196, 2017.
- [12] Alessio Conte, Donatella Firmani, Maurizio Patrignani, and Riccardo Torlone. A meta-algorithm for finding large k -plexes. *Knowl. Inf. Syst.*, 63(7):1745–1769, 2021.
- [13] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [14] R. Erhardt. Engineering Algorithms for the Weighted

- Maximum Clique Problem. Master's Thesis, Heidelberg University, 2022.
- [15] Zhiwen Fang, Chu-Min Li, and Ke Xu. An exact algorithm based on maxsat reasoning for the maximum weight clique problem. *Journal of Artificial Intelligence Research*, 55:799–833, 2016.
 - [16] Daniel Funke, Sebastian Lamm, Ulrich Meyer, Manuel Penschuck, Peter Sanders, Christian Schulz, Darren Strash, and Moritz von Looz. Communication-free massively distributed graph generation. *Journal of Parallel and Distributed Computing*, 131:200–217, 2019.
 - [17] Alexander Gellner, Sebastian Lamm, Christian Schulz, Darren Strash, and Bogdán Zaválnij. Boosting data reduction for the maximum weight independent set problem using increasing transformations. In *2021 Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 128–142. SIAM, 2021.
 - [18] Michel Gendreau, Patrick Soriano, and Louis Salvail. Solving the maximum clique problem using a tabu search approach. *Annals of operations research*, 41(4):385–403, 1993.
 - [19] Jens Gramm, Jiong Guo, Falk Hüffner, and Rolf Niedermeier. Data reduction and exact algorithms for clique cover. *ACM J. Exp. Algorithmics*, 13, feb 2009.
 - [20] Jiewei Gu, Weiguo Zheng, Yuzheng Cai, and Peng Peng. Towards computing a near-maximum weighted independent set on massive graphs. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD '21*, pages 467–477, New York, NY, USA, 2021. Association for Computing Machinery.
 - [21] Luca Gugelmann, Konstantinos Panagiotou, and Ueli Peter. Random hyperbolic graphs: degree sequence and clustering. In *International Colloquium on Automata, Languages, and Programming*, pages 573–585. Springer, 2012.
 - [22] Demian Hesse, Christian Schulz, and Darren Strash. Scalable kernelization for maximum independent sets. *Journal of Experimental Algorithmics (JEA)*, 24:1–22, 2019.
 - [23] Hua Jiang, Chu-Min Li, Yanli Liu, and Filip Manyà. A two-stage maxsat reasoning approach for the maximum weight clique problem. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
 - [24] Hua Jiang, Chu-Min Li, and Filip Manyà. An exact algorithm for the maximum weight clique problem in large graphs. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31, 2017.
 - [25] Hua Jiang, Dongming Zhu, Zhichao Xie, Shaowen Yao, and Zhang-Hua Fu. A new upper bound based on vertex partitioning for the maximum k -plex problem. In Zhi-Hua Zhou, editor, *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 1689–1696. International Joint Conferences on Artificial Intelligence Organization, 8 2021. Main Track.
 - [26] David S Johnson and Michael A Trick. *Cliques, coloring, and satisfiability: second DIMACS implementation challenge, October 11-13, 1993*, volume 26. American Mathematical Soc., 1996.
 - [27] Deniss Kumlander. A new exact algorithm for the maximum-weight clique problem based on a heuristic vertex-coloring and a backtrack search. In *Proc. 5th Int'l Conf. on Modelling, Computation and Optimization in Information Systems and Management Sciences*, pages 202–208, 2004.
 - [28] Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Renato F. Werneck. Finding near-optimal independent sets at scale. *J. Heuristics*, 23(4):207–229, 2017.
 - [29] Sebastian Lamm, Christian Schulz, Darren Strash, Robert Williger, and Huashuo Zhang. Exactly solving the maximum weight independent set problem on large real-world graphs. In *2019 Proceedings of the Twenty-First Workshop on Algorithm Engineering and Experiments (ALENEX)*, pages 144–158. SIAM, 2019.
 - [30] Chu-Min Li, Hua Jiang, and Filip Manyà. On minimization of the number of branches in branch-and-bound algorithms for the maximum clique problem. *Computers & Operations Research*, 84:1–15, 2017.
 - [31] Chu-Min Li and Zhe Quan. An efficient branch-and-bound algorithm based on maxsat for the maximum clique problem. In *Twenty-fourth AAAI conference on artificial intelligence*, 2010.
 - [32] Jinkun Lin, Shaowei Cai, Chuan Luo, and Kaile Su. A reduction based method for coloring very large graphs. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 517–523, 2017.
 - [33] Ingo Muegge and Matthias Rarey. Small molecule docking and scoring. *Reviews in computational chemistry*, 17:1–60, 2001.
 - [34] Manuel Penschuck, Ulrik Brandes, Michael Hamann, Sebastian Lamm, Ulrich Meyer, Ilya Safro, Peter Sanders, and Christian Schulz. Recent advances in scalable network generation. *arXiv preprint arXiv:2003.00736*, 2020.
 - [35] Wayne Pullan. Phased local search for the maximum clique problem. *Journal of Combinatorial Optimization*, 12(3):303–323, 2006.
 - [36] Wayne Pullan. Approximating the maximum vertex/edge weighted clique using local search. *Journal of Heuristics*, 14(2):117–134, 2008.
 - [37] Ryan A. Rossi and Nesreen K. Ahmed. The network data repository with interactive graph analytics and visualization. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
 - [38] Pablo San Segundo, Fabio Furini, David Álvarez, and Panos M. Pardalos. CliSAT: A new exact algorithm for hard maximum clique problems. *European Journal of Operational Research*, 2022.
 - [39] Darren Strash. On the power of simple reductions for the maximum independent set problem. In Thang N. Dinh and My T. Thai, editors, *Computing and Com-*

- binatorics - 22nd International Conference, COCOON 2016, Ho Chi Minh City, Vietnam, August 2-4, 2016, Proceedings*, volume 9797 of *Lecture Notes in Computer Science*, pages 345–356. Springer, 2016.
- [40] Darren Strash and Louise Thompson. *Effective Data Reduction for the Vertex Clique Cover Problem*, pages 41–53. SIAM, 2022.
 - [41] Etsuji Tomita and Tomokazu Seki. An efficient branch-and-bound algorithm for finding a maximum clique. In *International conference on discrete mathematics and theoretical computer science*, pages 278–289. Springer, 2003.
 - [42] Etsuji Tomita, Yoichi Sutani, Takanori Higashi, Shinya Takahashi, and Mitsuo Wakatsuki. A simple and faster branch-and-bound algorithm for finding a maximum clique. In *International Workshop on Algorithms and Computation*, pages 191–203. Springer, 2010.
 - [43] Anurag Verma, Austin Buchanan, and Sergiy Butenko. Solving the maximum clique and vertex coloring problems on very large sparse networks. *INFORMS Journal on Computing*, 27(1):164–177, 2015.
 - [44] Yiyuan Wang, Shaowei Cai, Jiejiang Chen, and Minghao Yin. Scwalk: An efficient local search algorithm and its improvements for maximum weight clique problem. *Artificial Intelligence*, 280:103230, 2020.
 - [45] Yiyuan Wang, Shaowei Cai, Shiwei Pan, Ximing Li, and Monghao Yin. Reduction and local search for weighted graph coloring problem. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03):2433–02441, Apr. 2020.
 - [46] Yiyuan Wang, Shaowei Cai, and Minghao Yin. Two efficient local search algorithms for maximum weight clique problem. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
 - [47] Stanley Wasserman, Katherine Faust, et al. Social network analysis: Methods and applications. 1994.
 - [48] Qinghua Wu and Jin-Kao Hao. A review on algorithms for maximum clique problems. *European Journal of Operational Research*, 242(3):693–709, 2015.
 - [49] Qinghua Wu and Jin-Kao Hao. Solving the winner determination problem via a weighted maximum clique heuristic. *Expert Systems with Applications*, 42(1):355–365, 2015.
 - [50] Jingen Xiang, Cong Guo, and Ashraf Aboulnga. Scalable maximum clique computation using mapreduce. In *2013 IEEE 29th International Conference on Data Engineering (ICDE)*, pages 74–85, 2013.
 - [51] Dong Zhang, Omar Javed, and Mubarak Shah. Video object co-segmentation by regulated maximum weight cliques. In *European Conference on Computer Vision*, pages 551–566. Springer, 2014.
 - [52] Hootan Zhian, Masoud Sabaei, Nastoo Taheri Javan, and Omid Tavallaie. Increasing coding opportunities using maximum-weight clique. In *2013 5th Computer Science and Electronic Engineering Conference (CEECE)*, pages 168–173. IEEE, 2013.
 - [53] Patric R.J. Östergård. A fast algorithm for the maximum clique problem. *Discrete Applied Mathematics*, 120(1):197–207, 2002. Special Issue devoted to the 6th Twente Workshop on Graphs and Combinatorial Optimization.

A Reduced Graph Sizes

Graph	Reduced Graph Size			
	old+new reductions		old reductions only	
	absolute	% of n_0	absolute	% of n_0
REP				
aff-digg	74,864	8.58	131,624	15.08
bio-human-gene1	3,915	17.57	4,485	20.13
bio-human-gene2	3,353	23.38	3,787	26.41
bio-mouse-gene	4,840	10.73	13,888	30.79
sc-TSOPF-RS-b2383	16,123	42.29	37,737	98.99
soc-flickr-und	9,465	0.55	46,728	2.72
soc-orkut	1,264,963	42.21	1,521,404	50.76
soc-orkut-dir	1,360,796	44.29	1,736,368	56.51
web-wikipedia-link_it	0	0.00	1,214	0.04
web-wikipedia-growth	83,724	4.48	637,483	34.08
arithmetic mean		19.41		33.55
RHG				
rhg_250k_100_1.75	7	0.00	1,061	0.42
rhg_250k_100_2.25	62	0.02	782	0.31
rhg_250k_250_1.75	11	0.00	2,170	0.87
rhg_250k_250_2.25	14	0.01	983	0.39
rhg_250k_500_1.75	32	0.01	4,103	1.64
rhg_250k_500_2.25	39	0.02	1,250	0.50
rhg_500k_250_1.75	2	0.00	3,782	0.76
rhg_500k_250_2.25	122	0.02	1,138	0.23
rhg_500k_500_2.25	0	0.00	1,761	0.35
rhg_750k_250_1.75	0	0.00	4,365	0.58
rhg_750k_250_2.25	15	0.00	1,062	0.14
rhg_750k_500_1.75	4,445	0.59	7,341	0.98
rhg_750k_500_2.25	12	0.00	2,651	0.35
arithmetic mean		0.05		0.58
OSM				
district-of-columbia-AM2	0	0.00	759	5.58
district-of-columbia-AM3	0	0.00	1,513	3.27
greenland-AM3	0	0.00	1,768	35.46
hawaii-AM3	0	0.00	11,130	39.74
idaho-AM3	0	0.00	2,293	56.42
kentucky-AM3	0	0.00	9,447	49.47
massachusetts-AM3	0	0.00	802	21.66
oregon-AM3	0	0.00	1,662	29.74
rhode-AM3	0	0.00	2,913	19.26
vermont-AM3	0	0.00	498	14.49
virginia-AM3	0	0.00	907	14.66
washington-AM3	0	0.00	1,503	15.00
arithmetic mean		0.00		25.40

Table 3: Reduced graph sizes (number of nodes) with when both old and new data reductions rules are applied vs. reduced graph sizes obtained when only running reductions from the current literature. n_0 refers to the initial number of nodes. Smaller is better. On the DIMACS instances reduced graph sizes did not change.

B Detailed Results for Exact Algorithms

Graph	t_{sol}		t_{prv}		$w(\hat{C})$	
	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu
OSM Exact Results						
district-of-col-AM2	0.80	0.15	0.88	0.21	235,777	235,777
district-of-col-AM3	1,937.83	5.64	1,938.06	6.52	545,969	545,969
greenland-AM3	10.31	0.82	11.91	3.30	604,575	604,575
hawaii-AM3	3,598.87	30.88	-	54.68	1,110,978	1,229,741
idaho-AM3	218.66	4.47	220.42	5.55	1,101,721	1,101,721
kentucky-AM3	3,580.19	114.09	-	144.51	1,808,419	1,860,308
massachusetts-AM3	0.81	0.03	1.01	0.13	115,636	115,636
oregon-AM3	8.93	1.70	11.07	2.47	557,634	557,634
rhode-island-AM3	81.30	11.03	93.69	18.40	1,162,925	1,162,925
vermont-AM3	4.57	0.49	4.90	0.56	604,213	604,213
virginia-AM3	0.13	0.06	0.21	0.06	207,457	207,457
washington-AM3	12.22	1.01	12.75	1.06	356,314	356,314
Geo. Mean	27.62	1.55	31.01	2.43	537,149	542,993
Graph	t_{sol}		t_{prv}		$w(\hat{C})$	
	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu
DIMACS Exact Results						
brock800.1	1,122.28	1,191.32	2,667.43	2,820.98	3,006	3,006
brock800.2	2,601.19	2,752.57	3,111.08	3,285.93	3,074	3,074
brock800.3	1,318.78	1,400.10	2,735.77	2,896.13	2,984	2,984
brock800.4	2,019.67	2,137.40	-	-	3,059	3,059
C1000.9	2,889.65	1,429.16	-	-	7,338	7,459
C2000.5	1,867.49	1,908.40	-	-	2,395	2,395
C2000.9	2,248.93	697.11	-	-	7,898	8,284
C4000.5	2,804.98	2,488.93	-	-	2,460	2,437
C500.9	2,716.53	2,611.48	-	-	6,789	6,789
gen400_p0.9_55	2,359.72	2,378.30	-	-	6,654	6,654
gen400_p0.9_65	2,768.34	2,772.78	-	-	6,535	6,535
gen400_p0.9_75	1,420.53	1,388.95	2,522.75	2,511.52	7,492	7,492
hamming10-4	3,110.88	3,018.87	-	-	5,205	5,125
johnson32-2-4	1,983.31	2,070.46	-	-	2,935	2,935
keller5	2,801.07	2,821.09	-	-	3,807	3,827
keller6	2,569.15	2,301.13	-	-	5,617	6,175
MANN_a27	2.67	2.09	2.81	2.14	17,866	17,866
MANN_a45	121.72	59.03	126.40	63.31	49,459	49,459
MANN_a81	3,599.91	1,922.82	-	3,116.00	118,898	161,903
p_hat1000-3	2,177.09	2,068.98	-	-	8,261	8,261
p_hat1500-2	952.79	902.21	2,491.91	2,371.21	7,556	7,556
p_hat1500-3	1,212.92	1,111.24	-	-	10,796	10,796
sanr400_0.7	11.86	12.69	26.63	28.43	2,926	2,926
Geo. Mean	1,106.99	946.34	1,714.98	1,650.88	6,460	6,587

Table 4: OSM and DIMACS exact results for each graph and the geometric mean.

Graph	t_{sol}		t_{prv}		$w(\hat{C})$	
	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu
REP Exact Results						
aff-digg	16.55	45.40	244.04	273.67	3,829	3,829
bio-human-gene1	3,598.98	2,010.82	-	3,327.65	136,325	136,692
bio-human-gene2	2,298.36	474.49	-	1,380.66	131,904	131,904
bio-mouse-gene	1,789.70	198.47	-	240.97	50,785	59,476
sc-TSOPF-RS-b2383	9.70	27.59	9.70	365.22	913	913
soc-flickr-und	35.10	133.73	74.52	162.42	10,847	10,847
soc-orkut	56.42	135.79	67.48	144.89	5,832	5,832
soc-orkut-dir	46.26	116.30	60.22	128.33	5,261	5,261
web-wiki_link_it	170.90	45.99	171.08	46.15	87,175	87,175
web-wiki-growth	13.09	84.77	17.30	88.57	3,334	3,334
Geo. Mean	117.15	134.60	190.45	259.14	14,092	14,321
Graph	t_{sol}		t_{prv}		$w(\hat{C})$	
	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu	TSM-MWC	MWCRedu
RHG Exact Results						
rhg_25k_100_1.75	94.71	2.26	94.91	2.72	99,839	99,839
rhg_250k_100_2.25	2.53	2.02	2.75	2.10	37,947	37,947
rhg_250k_250_1.75	101.48	4.74	107.14	4.76	112,769	112,769
rhg_250k_250_2.25	21.87	4.53	22.60	4.66	71,001	71,001
rhg_250k_500_1.75	1,079.84	35.47	1,092.32	38.43	137,234	137,234
rhg_250k_500_2.25	51.02	9.05	52.12	9.25	102,364	102,364
rhg_500k_250_1.75	3,093.33	37.09	-	37.70	130,973	131,559
rhg_500k_250_2.25	88.77	11.36	89.79	11.71	88,512	88,512
rhg_500k_500_2.25	43.49	15.11	47.06	21.09	122,781	122,781
rhg_750k_250_1.75	3,599.25	33.80	-	55.44	150,676	160,845
rhg_750k_250_2.25	132.25	18.15	132.43	18.33	96,362	96,362
rhg_750k_500_1.75	9.74	311.73	226.52	1,173.96	207,197	207,197
rhg_750k_500_2.25	35.91	34.13	49.39	34.19	119,936	119,936
Geo. Mean	95.93	14.55	128.16	17.73	106,210	106,781

Table 5: REP and RHG exact results for each graph and the geometric mean.

C Detailed Results for Heuristic Algorithms

Graph	t_{sol}			$w(\hat{C})$		
	FastWCLq	SCCWalk4l	MWCPeel	FastWCLq	SCCWalk4l	MWCPeel
OSM Heuristic Results						
district-of-columbia-AM2	0.32	4.91	0.16	235,777	234,219	235,777
district-of-columbia-AM3	16.87	208.46	5.26	545,969	545,969	545,969
greenland-AM3	2.96	39.33	0.85	604,575	604,575	604,575
hawaii-AM3	86.47	727.42	29.82	1,229,741	1,224,690	1,229,741
idaho-AM3	15.76	162.51	4.29	1,101,721	1,098,044	1,101,721
kentucky-AM3	374.57	997.09	102.17	1,860,308	1,437,770	1,860,308
massachusetts-AM3	0.25	50.59	0.03	115,636	113,381	115,636
oregon-AM3	6.06	239.62	1.72	557,634	546,314	557,634
rhode-island-AM3	34.16	252.86	9.34	1,162,925	1,162,920	1,162,925
vermont-AM3	0.37	2.32	0.39	604,213	602,793	604,213
virginia-AM3	0.20	6.38	0.05	207,457	207,457	207,457
washington-AM3	1.88	23.11	0.96	356,314	356,314	356,314
Geo. Mean	4.45	64.29	1.45	542,993	528,956	542,993
Graph	t_{sol}			$w(\hat{C})$		
	FastWCLq	SCCWalk4l	MWCPeel	FastWCLq	SCCWalk4l	MWCPeel
DIMACS Heuristic Results						
brock800_1	150.98	0.33	45.67	3,000	3,006	2,886
brock800_2	163.22	73.82	59.84	3,024	3,074	2,935
brock800_3	127.04	0.36	37.69	2,984	2,984	2,912
brock800_4	230.35	98.56	48.01	3,007	3,059	2,887
C1000.9	409.16	1.45	262.85	8,693	9,058	7,779
C2000.5	404.03	8.25	536.78	2,426	2,467	2,390
C2000.9	351.96	105.48	123.18	9,822	10,874	8,603
C4000.5	375.71	82.45	638.30	2,580	2,787	2,472
C500.9	251.22	0.28	474.51	7,277	7,313	6,964
gen400_p0.9.55	169.17	0.09	50.16	6,781	6,781	6,614
gen400_p0.9.65	312.92	0.28	59.97	6,869	6,881	6,654
gen400_p0.9.75	104.98	84.15	10.78	7,547	7,551	7,261
hamming10-4	287.81	1.42	527.12	5,727	5,917	5,279
johnson32-2-4	459.61	0.06	537.90	3,004	3,042	3,020
keller5	390.66	2.14	326.39	3,811	3,851	3,545
keller6	347.22	216.10	400.16	6,727	8,412	6,103
MANN_a27	16.84	168.23	1.63	17,866	17,864	17,710
MANN_a45	79.35	43.89	93.78	49,459	49,459	49,312
MANN_a81	180.29	376.13	7.17	161,903	161,895	161,648
p_hat1000-3	375.61	0.51	404.89	8,248	8,295	8,223
p_hat1500-2	423.40	1.22	475.25	7,519	7,556	7,546
p_hat1500-3	356.96	4.08	206.37	10,725	10,926	10,801
sanr400_0.7	6.94	0.06	0.91	2,926	2,926	2,874
	193.11	4.46	91.31	6,792	6,968	6,547

Table 6: OSM and DIMACS heuristic results for each graph and the geometric mean.

Graph	t_{sol}			$w(\hat{C})$		
	FastWCLq	SCCWalk41	MWCPeel	FastWCLq	SCCWalk41	MWCPeel
REP Heuristic Results						
aff-digg	240.01	30.73	47.68	3,514	3,829	3,829
bio-human-gene1	719.13	640.92	493.36	136,581	136,647	136,713
bio-human-gene2	457.56	534.75	89.59	131,763	131,862	131,904
bio-mouse-gene	593.17	412.13	13.43	59,439	59,473	59,146
sc-TSOPF-RS-b2383	33.02	243.56	1.43	913	900	870
soc-flickr-und	601.63	252.98	44.40	10,806	8,968	10,847
soc-orkut	135.88	526.16	183.51	5,832	4,552	5,582
soc-orkut-dir	157.93	521.46	185.02	5,261	4,080	5,116
web-wikipedia_link_it	71.06	972.97	36.59	87,175	2,903	87,175
web-wikipedia-growth	44.03	343.25	66.06	3,334	2,960	3,136
Geo. Mean	190.36	345.88	52.12	14,190	9,382	14,056
Graph	t_{sol}			$w(\hat{C})$		
	FastWCLq	SCCWalk41	MWCPeel	FastWCLq	SCCWalk41	MWCPeel
RHG Heuristic Results						
rhg_250k_100_1.75	10.56	134.23	2.67	99,839	99,839	99,839
rhg_250k_100_2.25	2.76	59.67	2.09	37,947	37,947	37,947
rhg_250k_250_1.75	42.01	501.91	4.24	112,769	112,074	112,756
rhg_250k_250_2.25	9.82	204.35	4.58	71,001	71,001	71,001
rhg_250k_500_1.75	125.03	869.09	18.39	137,234	86,773	136,884
rhg_250k_500_2.25	36.98	624.17	10.15	102,364	100,087	102,364
rhg_500k_250_1.75	90.36	854.58	22.69	131,559	67,352	131,100
rhg_500k_250_2.25	22.04	647.54	11.40	88,512	85,244	88,512
rhg_500k_500_2.25	71.83	995.47	15.33	122,781	46,201	122,781
rhg_750k_250_1.75	122.79	998.01	22.72	160,845	36,006	160,845
rhg_750k_250_2.25	28.23	687.98	17.00	96,362	95,558	96,362
rhg_750k_500_1.75	375.30	1,017.89	21.41	207,197	27,079	207,197
rhg_750k_500_2.25	72.82	1,002.84	36.00	119,936	43,470	119,936
Geo. Mean	40.36	513.37	10.69	106,781	63,647	106,731

Table 7: REP and RHG heuristic results for each graph and the geometric mean.

D Detailed Properties of Instances

Dataset	Instance	$ V $	$ E $	ρ
DIMACS	brock800.1	800	207,505	6.49×10^{-1}
	brock800.2	800	208,166	6.51×10^{-1}
	brock800.3	800	207,333	6.49×10^{-1}
	brock800.4	800	207,643	6.50×10^{-1}
	C1000.9	1,000	450,079	9.01×10^{-1}
	C2000.5	2,000	999,836	5.00×10^{-1}
	C2000.9	2,000	1,799,532	9.00×10^{-1}
	C4000.5	4,000	4,000,268	5.00×10^{-1}
	C500.9	500	112,332	9.00×10^{-1}
	gen400_p0.9.55	400	71,820	9×10^{-1}
	gen400_p0.9.65	400	71,820	9×10^{-1}
	gen400_p0.9.75	400	71,820	9×10^{-1}
	hamming10-4	1,024	434,176	8.29×10^{-1}
	johnson32-2-4	496	107,880	8.79×10^{-1}
	keller5	776	225,990	7.52×10^{-1}
	keller6	3,361	4,619,898	8.18×10^{-1}
	MANN_a27	378	70,551	9.90×10^{-1}
	MANN_a45	1,035	533,115	9.96×10^{-1}
	MANN_a81	3,321	5,506,380	9.99×10^{-1}
	p_hat1000-3	1,000	371,746	7.44×10^{-1}
OSM	p_hat1500-2	1,500	568,960	5.06×10^{-1}
	p_hat1500-3	1,500	847,244	7.54×10^{-1}
	sanr400.0.7	400	55,869	7.00×10^{-1}
	district-of-columbia-AM2	13,597	1,609,795	8.05×10^{-2}
	district-of-columbia-AM3	46,221	27,729,137	1.93×10^{-1}
	greenland-AM3	4,986	3,652,361	1.74×10^{-2}
	hawaii-AM3	28,006	49,444,921	1.87×10^{-1}
	idaho-AM3	4,064	3,924,080	4.75×10^{-1}
	kentucky-AM3	19,095	59,533,630	2.94×10^{-1}
	massachusetts-AM3	3,703	551,491	1.10×10^{-1}
	oregon-AM3	5,588	2,912,701	2.60×10^{-2}
	rhode-island-AM3	15,124	12,622,219	1.26×10^{-1}
REP	vermont-AM3	3,436	1,136,164	3.27×10^{-1}
	virginia-AM3	6,185	665,903	3.48×10^{-2}
	washington-AM3	10,022	2,346,213	4.67×10^{-2}
	aff-digg	872,622	22,501,699	5.91×10^{-5}
	bio-human-gene1	22,283	12,323,680	4.96×10^{-2}
	bio-human-gene2	14,340	9,027,024	8.78×10^{-2}
	bio-mouse-gene	45,101	14,461,095	1.42×10^{-2}
	sc-TSOPF-RS-b2383	38,121	16,115,324	2.22×10^{-2}
	soc-flickr-und	1,715,256	15,555,040	1.06×10^{-5}
RHG	soc-orkut	3,072,442	117,185,082	2.48×10^{-5}
	soc-orkut-dir	2,997,167	106,349,208	2.37×10^{-5}
	web-wikipedia.link.it	2,936,414	86,754,663	2.01×10^{-5}
	web-wikipedia-growth	1,870,710	36,532,530	2.09×10^{-5}
	rhg_250k_100.1.75	250,000	7,755,473	2.48×10^{-4}
	rhg_250k_100.2.25	250,000	10,546,938	3.38×10^{-4}
	rhg_250k_250.1.75	250,000	17,828,988	5.71×10^{-4}
	rhg_250k_250.2.25	250,000	24,036,880	7.69×10^{-4}
	rhg_250k_500.1.75	250,000	35,161,098	1.13×10^{-3}
	rhg_250k_500.2.25	250,000	47,230,197	1.51×10^{-3}
	rhg_500k_250.1.75	500,000	35,493,799	2.84×10^{-4}
	rhg_500k_250.2.25	500,000	49,954,694	4.00×10^{-4}
	rhg_500k_500.2.25	500,000	92,901,492	7.43×10^{-4}
	rhg_750k_250.1.75	750,000	53,201,080	1.89×10^{-4}
	rhg_750k_250.2.25	750,000	73,667,026	2.62×10^{-4}
	rhg_750k_500.1.75	750,000	102,363,505	3.64×10^{-4}
	rhg_750k_500.2.25	750,000	139,633,569	4.96×10^{-4}

Table 8: Instance Properties