

Scalable Kernelization for Maximum Independent Sets*

Demian Hesse†

Christian Schulz‡

Darren Strash§

Abstract

The most efficient algorithms for finding maximum independent sets in both theory and practice use reduction rules to obtain a much smaller problem instance called a *kernel*. The kernel can then be solved quickly using exact or heuristic algorithms—or by repeatedly kernelizing recursively in the branch-and-reduce paradigm. It is of critical importance for these algorithms that kernelization is fast and returns a small kernel. Current algorithms are either slow but produce a small kernel, or fast and give a large kernel. We attempt to accomplish both of these goals simultaneously, by giving an efficient parallel kernelization algorithm based on graph partitioning and parallel bipartite maximum matching.

We combine our parallelization techniques with two techniques to accelerate kernelization further: dependency checking that prunes reductions that cannot be applied, and reduction tracking that allows us to stop kernelization when reductions become less fruitful. Our algorithm produces kernels that are orders of magnitude smaller than the fastest kernelization methods, while having a similar execution time. Furthermore, our algorithm is able to compute kernels with size comparable to the smallest known kernels, but up to two orders of magnitude faster than previously possible. Finally, we show that our kernelization algorithm can be used to accelerate existing state-of-the-art heuristic algorithms, allowing us to find larger independent sets faster on large real-world networks and synthetic instances.

1 Introduction

The maximum independent set problem is a classic NP-hard problem [21] with applications spanning many fields, such as classification theory, information retrieval, computer vision [18], computer graphics [38], map labeling [22] and routing in road networks [30]. Given a graph $G = (V, E)$, our goal is to compute a

maximum cardinality set of vertices $\mathcal{I} \subseteq V$ such that no vertices in \mathcal{I} are adjacent to one another. Such a set is called a *maximum independent set* (MIS). As a concrete application, independent sets are essential in labeling strategies for maps [22], where the objective is to maximize the number of visible non-overlapping labels on a map. This problem can be solved by constructing the *label conflict graph*, in which any two conflicting/overlapping labels are connected by an edge, and then computing a maximum independent set in this graph.

One of the most powerful techniques for solving the MIS problem in practice is *kernelization*—reducing the input to its most difficult part, the *kernel*. A kernel (for the MIS problem) of a graph G is a graph $r(G)$ of smaller or equal size, obtained by applying a specified polynomial time algorithm to G that reduces its size while preserving the information required to find an MIS in G . The algorithm is often composed of a set of algorithms (so called *reduction* rules), which are applied exhaustively. After finding an MIS in $r(G)$ we “undo” the kernelization to find an MIS of G . Fixed-parameter tractable algorithms for the MIS problem are exponential in the size of the kernel, and therefore the MIS problem is considered “hard” for a particular instance when its kernel size is large [40]. Thus, it is often desirable to apply *many* different reduction rules to reduce the input size as much as possible when solving the problem exactly.

In practice, kernelization is used as a preprocessing step to other algorithms [10, 14, 15, 35, 40, 43], where speeding up kernelization directly speeds up the algorithm. However, kernelization may also be applied repeatedly as part of an algorithm [2, 12, 31]. In either case, the smallest kernels (or seemingly equivalently, the most varied reductions) give the best chance at finding solutions. For instance, the reductions used by Akiba and Iwata [2] are the *only* ones known to compute an exact MIS on certain large-scale graphs, and these reductions are further successful in computing exact solutions in an evolutionary approach [31]. However it is not always beneficial to compute the smallest kernel possible. Fast and simple reductions can compute kernels that are “small enough” for local search to quickly find high-quality, and even exact, solutions much faster than the reductions used to find the smallest kernels [12, 15].

*Our research was supported by the Gottfried Wilhelm Leibniz Prize 2012, and the Large-Scale Data Management and Analysis (LSDMA) project in the Helmholtz Association.

†Institute for Theoretical Informatics, Karlsruhe Institute of Technology, Karlsruhe, Germany, hespe@kit.edu

‡Institute for Theoretical Informatics, Karlsruhe Institute of Technology, Karlsruhe, Germany and University of Vienna, Vienna, Austria, christian.schulz@{kit.edu, univie.ac.at}

§Department of Computer Science, Colgate University, Hamilton, NY, USA, dstrash@cs.colgate.edu

Fast and simple reductions can even be used to solve many large-scale instances exactly [40] just as quickly as the algorithm by Akiba and Iwata [2].

Thus, for kernelization, there is a trade-off between kernel size and kernelization time. The smallest kernels are necessary to solve the most number of instances to optimality, but the fastest reductions have just enough power to solve most instances quickly. However, when run on the largest instances, the large kernels given by simple rules may make it prohibitive to solve these instances exactly, or even near-optimally with heuristic methods. Thus, to be effective for a majority of applications, kernelization routines should compute a kernel that is as *small* as possible as *quickly* as possible.

1.1 Our Results. To this end, we develop an efficient shared-memory parallel kernelization algorithm based on graph partitioning and parallel bipartite maximum matching. We combine our parallelization with *dependency checking*—a strategy for pruning inapplicable reductions—as well as *reduction tracking* that allows us to stop kernelization when reductions become less fruitful. These pruning techniques achieve large additional speedups over the kernelization of Akiba and Iwata [2], which computes similarly sized kernels. Our experimental evaluation shows that on average our algorithm finds kernels that are seven times smaller than the algorithms of Chang et al. [12], while having similar a running time. At the same time our algorithms are 41 times faster on average than other algorithms that are able to find kernels of similar size. In further experiments we apply our kernelization algorithm to state-of-the-art heuristic maximum independent set algorithms and find that our kernels can be used to find larger independent sets faster in large real-world networks and synthetic instances.

2 Related Work

The *maximum clique* and *minimum vertex cover* problems are equivalent to the maximum independent set problem: a maximum clique in the complement graph \bar{G} is a maximum independent set in G , and a minimum vertex cover C in G is the complement of a maximum independent set $V \setminus C$ in G . Thus, an algorithm that solves one of these problems can be used to solve the others. Many branch-and-bound algorithms have been developed for the maximum clique problem [36, 37, 42], which use vertex reordering and pruning techniques based on approximate graph coloring [42] or MaxSAT [32], and can be further sped up by applying local search to obtain an initial solution of high quality [7].

A common theme among algorithms for these (and other) NP-hard problems is that of *kernelization*—

reducing the input to a smaller instance that, when solved optimally, optimally solves the original instance. Rules that are used to reduce the graph while retaining the ability to compute an optimal solution are called *reductions*. Reductions and kernelization have long been used in algorithms for the minimum vertex cover and maximum independent set problems [1, 13, 20, 41], for efficient exact algorithms and heuristics alike.

2.1 Exact Algorithms. Butenko et al. [10] and Butenko and Trukhanov [11] were able to find exact maximum independent sets in graphs with thousands of vertices by first applying reductions. Further works have introduced reductions to more quickly solve the maximum clique problem [35, 43] and enumerate k -plexes [14]. Though these works apply reduction techniques as a preprocessing step, many works apply reductions as a natural step of the algorithm. Reductions were originally used by Tarjan and Trojanowski [41] to reduce the running time of the brute force $O(n^2 2^n)$ algorithm to time $O(2^{n/3})$, and reductions are further used to give the fastest known polynomial space algorithm with running time of $O^*(1.9996^n)$ by Xiao and Nagamochi [46]. These algorithms apply reductions during recursion, only branching when the graph can no longer be reduced [19]—known as the *branch-and-reduce* method.

Akiba and Iwata [2] were the first to show the effectiveness of the branch-and-reduce method for solving the minimum vertex cover problem in practice for large sparse real-world graphs. Using a large collection of reductions, they solve graphs with millions of vertices within seconds. In contrast, the vast majority of instances can not be solved by the MCS clique solver [42] within a 24-hour time limit [2]. However, as later shown by Strash [40], many of these same instances can be solved just as quickly by first kernelizing with two simple standard reductions (namely, isolated vertex removal and vertex folding reductions) and then running MCS.

2.2 Heuristic Algorithms. Kernelization and reductions play an important role in heuristic algorithms too. Lamm et al. [31] showed that including reductions in a branch-and-reduce inspired evolutionary algorithm enables finding exact solutions much faster than provably exact algorithms. Dahlum et al. [15] further showed how to effectively combine reductions with local search. Dahlum et al. find that standard kernelization techniques are too slow to be effective for local search and show that applying simple reductions in an online fashion improves the speed of local search. Chang et al. [12] improved on this result, by implement-

ing reduction rules to reduce the lead time for kernelization for local search. They introduce two kernelization techniques: a reduction rule to collapse maximal degree-two paths in a single shot, resulting in a fast linear-time kernelization algorithm (**LinearTime**), and a near linear-time algorithm (**NearLinear**) that uses triangle counting to detect when the domination reduction can be applied. **NearLinear** has running time $\mathcal{O}(\Delta m)$ where Δ is the maximum degree of the graph. They further introduce “reducing–peeling” to find a large initial solution for local search. This technique can be viewed as computing one path through the search space of a branch-and-reduce algorithm: they repeatedly exclude high-degree vertices and kernelize the graph until it is empty, then take the independent set found as an initial solution for local search. Their **NearLinear** algorithm is able to find kernels small enough and fast enough to be effectively used with local search¹; however, their kernels are much larger than those of Akiba and Iwata, who use many more advanced reduction rules. Hence, their technique may not be effective for solving large instances exactly.

3 Preliminaries

Basic Concepts. Let $G = (V, E)$ be an undirected graph on $n = |V|$ nodes and $m = |E|$ edges. We assume that $V = \{0, \dots, n - 1\}$, and to eliminate ambiguity, we at times denote by $V[G]$ and $E[G]$ the sets V and E , respectively, for a particular graph G . Throughout this paper, we assume that G is *simple*: it has no multi-edges or self loops. The set $N(v) = \{u \mid \{v, u\} \in E\}$ denotes the *open* neighborhood (also simply called the *neighborhood*) of v . We further define the open neighborhood of a set of nodes $U \subseteq V$ to be $N(U) = \cup_{v \in U} N(v)$. We similarly define the *closed* neighborhood as $N[v] = N(v) \cup \{v\}$ and $N[U] = N(U) \cup U$. We sometimes use N_G to denote the neighborhood in a particular graph G . A graph $H = (V_H, E_H)$ is said to be a *subgraph* of $G = (V, E)$ if $V_H \subseteq V$ and $E_H \subseteq E$. We call H an *induced* subgraph when $E_H = \{\{u, v\} \in E \mid u, v \in V_H\}$. For a set of nodes $U \subseteq V$, $G[U]$ denotes the subgraph induced by U . A set $\mathcal{I} \subseteq V$ of vertices, is said to be an *independent set* if all nodes in \mathcal{I} are pairwise nonadjacent; that is, $E[G[\mathcal{I}]] = \emptyset$. The *maximum independent set problem* is that of finding a maximum cardinality independent set which is called a *maximum independent set* (MIS).

The *graph partitioning problem* is to partition V into k blocks $V_1 \cup \dots \cup V_k = V$ with $V_i \cap V_j = \emptyset$, $\forall i \neq j$ while optimizing a given cost function—typically

¹Although their implementation of **NearLinear** has $\mathcal{O}(\sqrt{nm})$ time, as it includes the linear programming reduction.

the number of edges with end vertices in different blocks. Additionally, a balance constraint is applied, which demands that the blocks have approximately equal size with respect to the number of vertices or, alternatively, the sum of weights associated with the vertices. *Boundary vertices* are adjacent to vertices in other blocks and *cut edges* cross block boundaries.

3.1 Reductions. We now briefly describe the reduction rules that we consider. Each reduction allows us to choose vertices that are in *some* MIS by following simple rules. If an MIS is found on the kernel graph \mathcal{K} , then each reduction may be undone, producing an MIS in the original graph.

3.1.1 Reductions of Akiba and Iwata [2]. Akiba and Iwata [2] use a full suite of advanced reduction rules, which can efficiently solve the minimum vertex cover problem for a variety of instances. Here, we briefly describe the reductions we use, but for the maximum independent set problem. Note that Akiba and Iwata further use packing [2], and alternative [45] reductions. For brevity, we do not describe them here.

Vertex folding [13]: For a vertex v with degree two whose neighbors u and w are not adjacent, either v is in some MIS, or both u and w are in some MIS. Therefore, we can contract u , v , and w to a single vertex v' and decide which vertices are in the MIS later. If v' is in the computed MIS, then u and w are added to the independent set, otherwise v is added. Thus, a vertex fold contributes a vertex to an independent set.

Linear programming relaxation [34]: A well-known linear programming relaxation for the MIS problem with a half-integral solution (i.e., using only values 0, 1/2, and 1) can be solved using bipartite matching: maximize $\sum_{v \in V} x_v$ such that $\forall (u, v) \in E$, $x_u + x_v \leq 1$ and $\forall v \in V$, $x_v \geq 0$. Vertices with value 1 must be in the MIS and can thus be removed from G along with their neighbors. We use a version [27] that computes a solution whose half-integral part is minimal using maximum bipartite matching.

Unconfined [45]: Though there are several definitions of *unconfined* vertex in the literature, we use the simple one from Akiba and Iwata [2]. A vertex v is *unconfined* when determined by the following simple algorithm. First, initialize $S = \{v\}$. Then find a $u \in N(S)$ such that $|N(u) \cap S| = 1$ and $|N(u) \setminus N[S]|$ is minimized. If there is no such vertex, then v is confined. If $N(u) \setminus N[S] = \emptyset$, then v is unconfined. If $N(u) \setminus N[S]$ is a single vertex w , then add w to S and repeat the algorithm. Otherwise, v is confined. Unconfined vertices can be removed from the graph, since there always exists a MIS that contains no unconfined vertices.

Diamond: Although not mentioned in their paper, Akiba and Iwata [2] extend the unconfined reduction in their implementation [26]. Let S be the set constructed in the unconfined reduction for a vertex v that is not unconfined. If there are nonadjacent vertices u_1, u_2 in $N(S)$ such that $N(u_1) \setminus N(S) = N(u_2) \setminus N(S) = \{v_1, v_2\}$, then we can remove v from the graph because there always exists a MIS that does not contain v . Note that this implies that $\{v_1, v_2\} \subseteq S$.

Twin [45]: Let u and v be vertices of degree three with $N(u) = N(v)$. If $G[N(u)]$ has edges, then add u and v to \mathcal{I} and remove $u, v, N(u), N(v)$ from G . Otherwise, some vertices in $N(u)$ may belong to some MIS \mathcal{I} . We still remove $u, v, N(u)$ and $N(v)$ from G , and add a new gadget vertex w to G with edges to u 's two-neighborhood (vertices at a distance 2 from u). If w is in the computed MIS, then none of u 's two-neighbors are in \mathcal{I} , and therefore $N(u) \subseteq \mathcal{I}$. Otherwise, if w is not in the computed MIS, then some of u 's two-neighbors are in \mathcal{I} , and therefore u and v are added to \mathcal{I} .

3.1.2 The Reduction of Butenko et al. [9]. We describe one reduction that was not included in the algorithm by Akiba and Iwata [2], but was shown by Butenko et al. [9] to be highly effective on medium-sized graphs derived from error-correcting codes.

Isolated Vertex Removal [9]: If a vertex v forms a single clique C with all its neighbors, then v is called *isolated* (*simplicial* is also used in the literature) and is always contained in some MIS. To see this, at most one vertex from C may be in any MIS. Either it is v or, if a neighbor of v is in an MIS, then we select v instead. Note that this reduction rule is completely contained in the unconfined reduction rule as every neighbor $u \in N(v)$ of a simplicial vertex v is unconfined, leaving only v without any neighbors in the graph. As it can be implemented more efficiently than the unconfined reduction, we apply the reduction by isolated vertex removal before removing unconfined vertices.

3.1.3 The Linear Time Algorithm by Chang et al. [12]. Chang et al. [12] present a kernelization algorithm `LinearTime` that runs in time $\mathcal{O}(m)$. It removes vertices of degree zero and one and uses a reduction rule using maximal paths of degree two. They split the rule into five cases depending on the length of the maximal degree two path and the endpoints of the path. The full description can be found in their paper [12]. The degree two path rule is a specialization of the vertex folding rule explained above, and does not cover the case of a vertex with two neighbors of degree higher than two. However, in contrast to the vertex folding rule, it has linear

time complexity. This algorithm often removes a large fraction of a graph's vertices in very little time; however, it still leaves the possibility to apply more powerful, but time-consuming reduction rules. We therefore run `LinearTime` as a preprocessing step of our algorithm.

4 Parallel Kernelization

As current machines usually have more than one processor and kernelization can run for hours on large instances, parallelization is a promising way to make larger graphs feasible for maximum independent set algorithms. In this section, we describe how we parallelize kernelization: we partition the graph into blocks so that “local” reductions can be run on blocks in parallel, and perform parallel maximum bipartite matching for the “global” reduction by linear programming. Our algorithm first applies the reductions parallelized by partitioning exhaustively. We then apply the reduction by linear programming. These steps are repeated until no more vertices can be removed from the graph. (See pseudocode in Appendix A.)

4.1 Blockwise Reductions. Many reductions have an element of locality. In particular, we call a reduction *local* if it is applied one vertex at a time, if determining that the reduction can be applied is based on local graph structure (for example, by its neighborhood or by neighbors of neighbors), and the reduction itself modifies only local graph structure. A challenge in parallelizing local reductions is in how to apply them simultaneously. Fine-grained parallelism would require locks, since attempting to simultaneously remove or contract (near-)neighboring vertices in the graph results in a race condition: these (near-)neighbors may *both* be mistakenly added to the independent set or the graph may be modified incorrectly. However, with locks, local reductions become more expensive—reductions must wait if they overlap other reductions in progress.

To avoid locks altogether, we partition the graph into vertex-disjoint blocks and perform local reductions on each block in parallel (i.e., *blockwise*). Note that the only way for two blockwise reductions to simultaneously (mistakenly) reduce neighbors is if they are incident to a cut edge. We therefore avoid race conditions by restricting reductions to only read and write to vertices and neighborhoods within a single block. By using a high-quality partitioning that minimizes the number of cut edges, we expect the number of vertices excluded from these local reductions to be small. To avoid race conditions when removing boundary vertices from the graph, we leave the adjacency lists of neighboring vertices unchanged. We only change the adjacency list of vertices when performing vertex contractions.




Figure 1: Vertices v and v' are simplicial but lie in different blocks, they are not removed from the graph. Vertex u is simplicial and is not a boundary vertex, so $N[u]$ is removed.

We now explain how to apply each local reduction in our parallel framework. Let V_i be the block in which we are applying the reduction. Let a reduction on a vertex $v \in V_i$ only depend on (and modify) vertices $R(v) \subseteq V_i$. Then no other vertex $u \in V_j$ for all $i \neq j$ is traversed or modified as the result of this reduction. Thus, we apply this reduction to v correctly: changes to it and/or adjacency lists of vertices in $R(v)$ do not affect vertices of other blocks.

Further, let B_k denote the set of vertices of distance at most k from some boundary vertex in our partitioning. Note that B_0 is the set of boundary vertices and $B_k = N[B_{k-1}]$.

Vertex Folding: Let $v \in V_i$ be a vertex with neighborhood $\{u, w\} \subset V_i$. Contracting v, u, w into v' will cause a race condition whenever $u, w \in V_i \cap B_0$, as their neighbors in some other block V_j must have their adjacency lists updated to include v' . We handle vertex folding with two cases. First, for $u, w \in V_i \setminus B_0$ (or equivalently, $v \in V_i \setminus B_1$), we apply the reduction normally. Then $N(\{u, w\}) \subseteq V_i$ and there is no race condition. Secondly, without loss of generality, if $u \in V_i \cap B_0$ and $w \in V_i \setminus B_0$ then we still apply vertex folding, using u as the new vertex v' . Neighborhoods of vertices in $N(u) \setminus \{v\} \not\subseteq V_i$ remain unchanged.

Isolated Vertex Removal: Let $v \in V_i \setminus B_0$ be an isolated vertex. Then we add v to \mathcal{I} and remove $N[v] \subseteq V_i$ from the graph as usual. (See Figure 1.)

Twin: Let $u, v \in V_i$ such that $N(u) = N(v) \subseteq V_i$, and note that u, v will not be boundary vertices as otherwise $N(u) \not\subseteq V_i$. We have two cases:

$G[N(u)]$ has edges: Since u, v are not boundary vertices, we add u, v to \mathcal{I} and remove $\{u, v\} \cup N(u) \subseteq V_i$ from the graph.

$G[N(u)]$ has no edges: We only apply this reduction when $N[N[u]] \subseteq V_i$: we remove $\{u, v\} \cup N(u)$, and create a new vertex $w \in V_i$ with neighborhood $N(w) = N[N[u]]$; otherwise we would modify the

adjacency list of a vertex in a different block.

Unconfined: Unlike other blockwise reductions, every vertex $v \in V_i$ is eligible for the unconfined reduction, including boundary vertices. If a vertex is unconfined, we mark it as excluded from the independent set and remove it from the graph (by setting a flag if v is a boundary vertex). However, the algorithm for finding unconfined vertices must be adapted—it does not simply rely on a (two-)neighborhood, but depends on an expanding set of vertices S , which should be drawn from V_i in order to avoid a race condition. In particular, a vertex $u \in N(S)$ can only be used if $u \in V_i$ and $S \subseteq V_i$ must hold. This way, we ensure all vertices that we classify as unconfined are truly unconfined and can be removed from the graph. We might, however, falsely classify some vertices as confined.

Diamond: As with the unconfined reduction, we can safely remove even boundary vertices from the graph by using the diamond reduction. However, since vertices in $V \setminus V_i$ cannot be inserted into S during the blockwise unconfined reduction, there might be u_1, u_2 such that $N(u_1) \setminus N(S) = N(u_2) \setminus N(S) = \{v_1, v_2\}$ and $\{v_1, v_2\} \not\subseteq S$ because they are located in different blocks, so we have to check that $v_1, v_2 \in S$. If not, they might be removed by another reduction which can lead to race conditions.

4.2 Parallel Linear Programming. Unlike the local reductions, the reduction by linear programming is not applied to single vertices and their (near-)neighbors. It instead relies on a *global view* of the graph to find a set of vertices that can be removed at once. Therefore our parallelization strategy for local reductions cannot be applied to the linear programming reduction. The computationally expensive part of this reduction is finding a maximum bipartite matching of the bi-double graph: $B(G) = (L_V \cup R_V, E')$, where $L_V = \{l_v \mid v \in V\}$, $R_V = \{r_v \mid v \in V\}$, and $E' = \{\{l_u, r_v\} \mid \{u, v\} \in E\}$.

Azad et al. [4] give a parallel augmenting path based algorithm for maximum bipartite matching. Their algorithm requires a maximal matching as input, which we first compute using the maximal matching algorithm by Karp and Sipser [29], which was parallelized by Azad et al. [5]. For better performance when repeatedly applying the reduction, we reuse the parts of the previous matching which are still part of the graph. If the graph changed only slightly since the last application, this is still close to a maximum matching, which results in less work for the maximum matching algorithm. This technique is also used by Akiba and Iwata [2]. To obtain the half-integral result of the linear program, we use the set of vertices reachable by alternating paths starting from matched vertices in L_V . To find these, we start a depth

first search from each vertex $v \in L_V$ in parallel and mark all reached vertices. We then obtain the result by iterating in parallel over all vertices in the original graph and checking whether their respective vertices in L_V and L_R are marked.

5 Pruning Reductions

5.1 Dependency Checking. To compute a kernel, Akiba and Iwata [2] apply their reductions r_1, \dots, r_j by iterating over all reductions and trying to apply the current reduction r_i to all vertices. If r_i reduces at least one vertex, they restart with reduction r_1 . When reduction r_j is executed, but does not reduce any vertex, all reductions have been applied exhaustively, and a kernel is found. Trying to apply every reduction to all vertices can be expensive in later stages of the algorithm where few reductions succeed. The algorithm may repeatedly attempt to apply the same reduction to a vertex even though the graph has not changed sufficiently to allow the reduction to succeed. For example, let G' be a graph obtained by applying reductions to a graph G . If vertex v is not isolated in G and $N_{G'}[v] = N_G[v]$, then v is still not isolated in G' and can be pruned from further attempts.

We define a scheme for checking dependencies between reductions, which allows us to avoid applying isolated vertex removal, vertex folding, and twin reductions when they will provably not succeed. After unsuccessfully trying to apply one of these reductions to a vertex v , one only has to consider v again for reduction after its neighborhood has changed. We therefore keep a set D of *viable* candidate vertices: vertices whose neighborhood has changed and vertices that have never been considered for reductions. Initially we set $D = V$. Then for each $v \in D$, we remove v from D and try to apply our reductions to v . If v is removed from the graph (or a new vertex w is inserted), we set $D = D \cup N(v)$ (or $D = D \cup N[w]$). We repeat until D is empty. Figure 2 shows an example for isolated vertex removal.




Figure 2: After removing isolated vertex v and $N(v)$, u is isolated. Orange vertices are in D .

Using this technique we reduce the amount of work for kernelization, especially in the later stages of the algorithm, where only few reductions are left to ap-

ply. Dependency checking can also help finding a kernel faster after finding a quasi kernel using our parallel algorithm: as most parts of the graph are already fully reduced, we expect dependency checking to quickly prune these parts and focus further kernelization on the boundaries when running the sequential version of our algorithm on the quasi kernel. Note that this strategy does not support unconfined and diamond reductions, as they depend on a set S that can grow arbitrarily large, and include vertices with large distances from the starting vertex. Thus a vertex can become unconfined due to a change in the graph outside of its neighborhood. Neither does it support the linear programming reduction, which operates on the entire graph instead of a single vertex. However, when performing these reductions we continue to add vertices whose neighborhoods have changed to D , saving effort when next attempting isolated vertex removal, vertex folding, and twin reductions.

We briefly mention that targeted forms of dependency checking have been used before. Previous works, including Akiba and Iwata [2] and Chang et al. [12] perform so-called “iterated” reductions, which allow for the repeated application of successful reductions. These include, for example, iteratively removing degree-one and -zero vertices (including any newly introduced degree-one and -zero vertices) until none remain, and applying the domination reduction when triangle counts change [12]. Unlike these previous works, our focus is on eliminating reductions that cannot be applied, as is not targeted at any particular reduction, but a collection of reductions. Strash [40] implements similar dependency checking for isolated vertex and vertex folding reductions, though it is not mentioned in his paper. We are the first to introduce such a strategy that can be used with *any* collection of reductions.

5.2 Reduction Tracking: Counteracting Diminishing Returns. It is not *always* ideal to apply reductions exhaustively—for example if only few reductions will succeed and they are costly. We note that, during later stages of our algorithm, local reductions may lead to very few graph changes, while the linear programming reduction often significantly reduces the graph size. Therefore, it may be better to stop local reductions early before applying the linear programming reduction, as any remaining local reductions can still be applied afterwards. Furthermore, in our parallel algorithm, applying the local reductions exhaustively can take significantly longer for some blocks than for others. That is, the total graph size is not significantly reduced once the first threads finish their blocks.

We therefore implement reduction tracking to de-

tect and stop local reductions when they are not quickly reducing the graph. Once the first thread finishes applying local reductions, we assign it to sample the current graph size at fixed time intervals. We then stop local reductions on all threads when the change in graph size becomes small relative to the rate of vertex removals and switch to the linear programming reduction. We continue local reductions afterwards. For the sequential case, we start sampling the current size immediately when starting the local reductions. In our implementation, sampling is performed by an additional thread; however, it does not introduce significant overhead and can be done in the same thread. Although the resulting graph is kernel-like, it may be possible to reduce it further. We call such a graph a *quasi kernel*.

6 Experimental Evaluation

Methodology. We implemented our algorithm using C++ and compiled all code using gcc 5.4.0 with full optimizations turned on (-O3 flag). For shared memory parallelization we use OpenMP 4.0. Our implementation includes the parallel application of reduction rules, the dependency checking scheme and the reduction tracking technique. For graph partitioning we use ParHIP [33], the parallel version of the KaHIP graph partitioner [39]. We compare against several existing sequential kernelization techniques. For fast reduction strategies, we compare against the kernelization routines `LinearTime` and `NearLinear` recently introduced by Chang et al. [12]. We use the authors' original implementation, written in C. For extensive reduction strategies, we use the full reduction suite of Akiba and Iwata's `VCSolver` [2]². We modified their code to stop execution after kernelization and output the kernel size. For all instances, we perform three independent runs of each algorithm. Their code was compiled and run sequentially with Java 1.8.0_102. All results are averages over three runs on a machine with 512 GB RAM and two Intel Xeon E5-2683 v4 processors with 16 cores running at 2.1 GHz each, running Ubuntu 14.04.5 LTS.

For further experiments studying the impact of the partition quality on parallelization as well as the impact of individual reduction rules, see Hespe's Master's thesis [24].

Algorithm Configuration. We run our algorithm with all reduction rules explained in this paper but restrict the isolated vertex removal reduction to cliques of size 3 or less. We use the *ultrafast* configuration of the parallel partitioner and default values for all other parameters. When running our algorithm in parallel on p threads, we partition the graph into p

blocks. We stop applying local reduction rules when the reduction in graph size per time during the last time interval is less than 5% of the average size reduction per time since starting to apply local reductions (i.e., since the last application of the linear programming reduction). An experimental evaluation of this technique can be found in Appendix C. As the `LinearTime` algorithm by Chang et al. [12] has very low running times and reduces the initial graph size, we run it as a preprocessing step using the original implementation. We then partition the resulting kernel and process it with our parallel kernelization algorithm. Throughout this section, we will refer to sequential runs of our algorithm as `FastKer` and to parallel runs (32 threads, unless otherwise stated) as `ParFastKer`. All repetitions of `ParFastKer` use the same partitioning of the input.

Instances. We perform experiments on large web [8] and road networks [6, 17], random (hyper)-geometric graphs [25, 44] and Delaunay triangulations [6, 33]. Basic instance properties can be found in Appendix B. These instances are all large (> 10M vertices) and kernelization takes a considerable amount of time on them. As our methods introduce some overhead compared to other kernelization algorithms, we focus our attention on speeding up kernelization for these hard instances.

6.1 Comparison with State-of-the-Art. We now compare our implementation to the implementations of `VCSolver` by Akiba and Iwata [2] and the `LinearTime` and `NearLinear` algorithm by Chang et al. [12]. Table 1 and Figure 3 give an overview. Figure 3 normalizes running time and kernel size on each instance by the result of `VCSolver`.

First note that `LinearTime`'s running time is almost negligible compared to that of `VCSolver`, almost never surpassing 1% of `VCSolver`'s time. `LinearTime` also decreases the graph size significantly for most graphs (except for the Delaunay triangulations, where `LinearTime` is not able to reduce the graph size at all), however, the `LinearTime` kernel is still orders of magnitude larger than `VCSolver`'s kernel. Due to fast running time and graph size reduction, we use `LinearTime` as a preprocessing step to our algorithm.

The `NearLinear` algorithm by Chang et al. [12] uses fewer reduction rules than our algorithm, so it finds larger kernels, often orders of magnitude larger than the kernels by `VCSolver` and our algorithms. The largest relative difference to the smallest kernel size of `NearLinear` is the 1329923-vertex kernel for `europa.osm`. This is 159 times larger than the smallest kernel and 94 times larger than the quasi kernel found by `ParFastKer`. For the Delaunay triangulations and the random geometric

²https://github.com/wata-orz/vertex_cover

Graph		LinearTime		NearLinear		VCSolver		FastKer		ParFastKer		
name	n	$ \mathcal{K} $	time	$ \mathcal{K} $	time	$ \mathcal{K} $	time	$ \mathcal{K} $	time	$ \mathcal{K} $	time	su
it-2004	41M	27.1M	3.3	11.3M	1544.6	1.6M	6749.0	1.7M	439.3	1.7M	151.1	44.7
webbase-2001	118M	51.7M	13.0	17.3M	121.1	0.7M	4207.8	0.8M	303.5	0.9M	54.9	76.7
asia.osm	12M	626.7K	0.8	594.4K	1.4	15.2K	204.7	34.9K	1.4	34.9K	1.2	172.3
europa.osm	51M	1500.0K	4.1	1329.9K	6.1	8.4K	302.4	14.1K	5.4	14.2K	5.0	60.0
rgg26	67M	67.1M	1.0	51.3M	172.6	49.6M	9887.7	49.8M	6329.1	49.8M	148.9	66.4
rhg	100M	*	*	*	*	0	124.0	0	138.1	16	65.1	1.9
del26	67M	67.1M	0.7	62.5M	53.3	49.9M	20728.7	51.7M	506.3	51.7M	176.1	117.7

Table 1: Running times and kernel sizes for all algorithms on selected instances. $|\mathcal{K}|$ is the (quasi) kernel size for each algorithm, and “su” is the speedup of ParFastKer over VCSolver. A * indicates that an instance cannot be processed due to a 32-bit implementation. All times are in seconds.

graph, the relative kernel size difference is comparatively low. This is because the kernel for these graphs is still very large compared to the input size, but we find quasi kernels much closer to the size found by VCSolver than NearLinear. LinearTime actually cannot remove any vertices from the Delaunay instances and only very few from the random geometric instance. In the geometric mean, LinearTime’s kernel is a factor 12 larger than ParFastKer’s quasi kernel and NearLinear’s kernel is a factor 7 larger. Due to NearLinear’s fast worst-case running time, it runs faster than FastKer on 8 out of 12 instances and on 2 instances even faster than ParFastKer. As LinearTime is a preprocessing step of our algorithm, it is of course always faster.

As VCSolver implements a larger set of reduction rules, adding the desk and funnel reductions by Xiao and Nagamochi [45] as well as Akiba and Iwata’s own *packing* reduction rule, it achieves smaller kernel sizes. In the geometric mean, ParFastKer’s quasi kernel’s are 20% larger than VCSolver’s (excluding rhg, which has an empty kernel). However, comparing the kernel sizes to the size of the input network, these differences in size are negligible. The largest obtained difference relative to the size of the input network from the graphs in Table 1 is 2.7% on del26. Among all graphs we tested, including the results in Appendix D, this rises to 2.9% on del124 (0.6% on sk-2005 when only considering the real-world instances). In addition, VCSolver only applies a scheme similar to our dependency checking for the removal of degree zero and one vertices, so our algorithm runs faster on all instances except rhg where FastKer takes 138 seconds for kernelization and VCSolver finishes after 124 seconds. ParFastKer, however, is faster than VCSolver on this instance, taking 65 seconds. On 11 out of 15 instances, FastKer is faster by a factor of over 6 than VCSolver and on 4 instances even by a factor of

over 40. The largest speedup of FastKer over VCSolver is 142 on asia.osm and the geometric mean of the speedups is 12. As FastKer is the sequential version of our algorithm, this is a purely algorithmic speedup. Using parallelization, ParFastKer achieves speedups of 41 over VCSolver in the geometric mean, combining the algorithmic speedup with parallel speedup. On all instances, except for rhg, the speedup is over 10 and on 9 instances over 50. For exact running times and kernel sizes for all benchmark instances, see the appendix.

6.2 Scalability. Figure 4 shows the parallel speedup of our algorithm on the six hardest instances of our benchmark set (i.e., those with the longest sequential running time). The left plot shows the total speedup relative to two threads for all parts of our algorithm combined: LinearTime preprocessing, partitioning and parallel reductions with dependency checking and inexact reduction pruning. The center and right plots show the speedups for the reductions parallelized by partitioning and the reduction by linear programming, respectively. The preprocessing step of our algorithm, the LinearTime algorithm by Chang et al., is sequential and thus limits the possible scalability of our parallelization, however running times are very short.

We observe that, due to the overhead caused by having to find a partition of the graph, the single threaded execution is on average 1.9 times faster than the parallelization using 2 threads. However, our algorithm scales well so that parallelization brings better performance for higher numbers of threads. Compared to the two-threaded case, our highest speedup is 46.5 for rgg26 on 32 threads. The main reason for this is that reductions on this graph are so slow that, for low thread counts, our inexact reduction pruning technique stops local reductions early, switching to a very long lasting reduction by linear programming. For the other graphs, the speedup on 32 threads compared to

³As empty kernels lead to a division by zero for this plot, graphs with an empty kernel are not shown here.

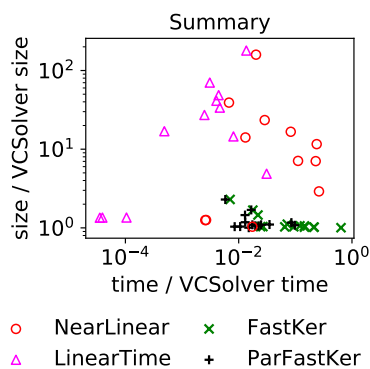


Figure 3: Comparison of kernelization algorithms against VCSolver.³

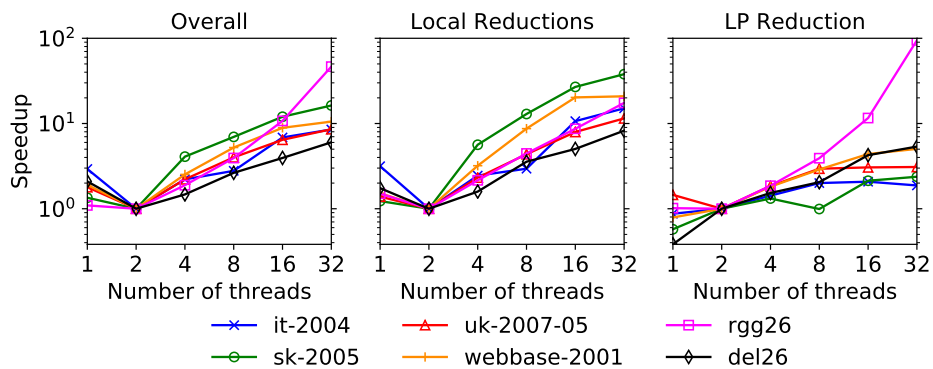


Figure 4: Scaling experiments on the six hardest instances of the benchmark set for (left) the overall algorithm, (center) blockwise reductions and (right) the reduction by linear programming. Speedups are relative to two threads.

2 threads is between 6 and 16.3 with 16 being perfect speedup. The speedup relative to the single threaded case is between 2.9 and 12 (42.5 for *rgg26*).

Figure 4 shows that local reductions parallelized by partitioning are faster single threaded than on two threads. This is caused by our inexact reduction pruning technique which starts after the first thread finishes reductions. When the number of threads is low, reductions might already have become too slow when the first thread finishes, causing longer times of slow size reduction. For higher thread counts, there is always a thread that finishes while other threads are still applying reductions fast and thus less time is wasted by slow reductions. After the drop at two threads, the speedup for 32 threads compared to 2 threads for these reductions is between 8 and 37 (between 5 and 31 compared to 1 thread). For some graphs, the reduction by linear programming, which we parallelized using the parallel maximum bipartite matching algorithm by Azad et al. [4], is a bottleneck of our algorithm.

6.3 Local Search on the Quasi Kernel. We now demonstrate the impact that quickly finding a small quasi kernel has on algorithms for finding large independent sets with local search. Currently, the algorithms with the best trade-off between speed and solution quality are *LinearTime* and *NearLinear* by Chang et al. [12]. In our previous experiments, we compared only against the *LinearTime* and *NearLinear* kernelization. However, we now run the full algorithm of Chang et al., which first kernelizes the graph, then invokes “reducing-peeling” to compute an initial solution for local search, and then runs local search [3] on the kernel. We compare their original algorithms against variants that first kernelize the graph with *ParFastKer* and then run on the quasi kernel. We use a time limit of 30 minutes, including

kernelization and finding an initial solution. Figure 5 shows the size of the independent set found over time for the largest web graph, road network and generated graph from our benchmark set (excluding graphs that cannot be processed due to the original 32-bit implementation of *LinearTime* and *NearLinear*). We ran local search three times using a different input seed for each run; however, we use the same input graph (either the original graph or a quasi kernel found by *ParFastKer*) for each run. The plots show at any given time the geometric mean of the current best solution of all runs. Results for the other graphs from our benchmark set can be found in Table 6 in Appendix D.

For web graphs we see that using *ParFastKer*’s quasi kernel, the independent set found is much larger (80 009 858 for *webbase-2001*) than the one found by *LinearTime* (18 286 vertices less) and *NearLinear* only converges to approximately the size found using the quasi kernel after several hundred seconds. On road networks, we observe interesting behavior: local search seemingly converges for all algorithms, but to *different* independent set sizes: the smaller the initial kernel size, the larger independent set size. On *europa.osm*, the final solution size is 25 633 238 for *LinearTime*, 84 more for *NearLinear* and 188 more for both versions that use *ParFastKer*’s quasi kernel. Also, using *ParFastKer*’s quasi kernel, the algorithm converges much faster. In particular, after an initial improvement over the starting solution, that takes about 0.1 seconds plus the time for kernelization and finding an initial solution (which is about 5 seconds), very few changes occur with *ParFastKer*’s quasi kernel. *LinearTime* and *NearLinear*, on the other hand, make an increase of several hundred vertices for the first 30 to 40 seconds. On the Delaunay triangulation graphs, the smaller quasi kernel enables local search to find larger independent sets.




Figure 5: Solution size over time of different algorithms.

7 Conclusion

We presented an efficient parallel kernelization algorithm based on graph partitioning and parallel bipartite maximum matching, vertex pruning as well as reduction tracking. On the one hand, our algorithm produces kernels that are orders of magnitude smaller than the fastest kernelization methods, while having a similar execution time. On the other hand, our algorithm is able to compute kernels with size comparable to the smallest known kernels, but up to two orders of magnitude faster than previously possible. Experiments with local search algorithms show that we find larger independent sets faster. In future work, we want to apply our parallel kernelization techniques in more MIS algorithms, such as exact branch-and-reduce [2], and transfer our techniques to other problems that use kernelization [16, 28, 23].

References

- [1] F. N. Abu-Khzam, M. R. Fellows, M. A. Langston, and W. H. Suters. Crown structures for vertex cover kernelization. *Theor. Comput. Syst.*, 41(3):411–430, 2007.
- [2] T. Akiba and Y. Iwata. Branch-and-reduce exponential/FPT algorithms in practice: A case study of vertex cover. *Theor. Comput. Sci.*, 609(1):211–225, 2016.
- [3] D. V. Andrade, M. G. Resende, and R. F. Werneck. Fast local search for the maximum independent set problem. *J. Heur.*, 18(4):525–547, 2012.
- [4] A. Azad, A. Buluç, and A. Pothen. Computing maximum cardinality matchings in parallel on bipartite graphs via tree-grafting. *IEEE Trans. Parallel Distrib. Syst.*, 28(1):44–59, 2017.
- [5] A. Azad, M. Halappanavar, S. Rajamanickam, E. G. Boman, A. Khan, and A. Pothen. Multithreaded algorithms for maximum matching in bipartite graphs. In *Proc. 26th International Parallel & Distributed Processing Symposium (IPDPS’12)*, pages 860–872. IEEE, 2012.
- [6] D. A. Bader, H. Meyerhenke, P. Sanders, C. Schulz, A. Kappes, and D. Wagner. Benchmarking for graph clustering and partitioning. In *Encyclopedia of Social Network Analysis and Mining*, pages 73–82. Springer, 2014.
- [7] M. Batsyn, B. Goldengorin, E. Maslov, and P. Pardalos. Improvements to MCS algorithm for the maximum clique problem. *J. Comb. Optim.*, 27(2):397–416, 2014.
- [8] P. Boldi, B. Codenotti, M. Santini, and S. Vigna. UbiCrawler: A scalable fully distributed web crawler. *Software: Practice & Experience*, 34(8):711–726, 2004.
- [9] S. Butenko, P. Pardalos, I. Sergienko, V. Shylo, and P. Stetsyuk. Finding maximum independent sets in graphs arising from coding theory. In *Proc. 2002 ACM Symposium on Applied Computing (SAC’02)*, pages 542–546. ACM, 2002.
- [10] S. Butenko, P. Pardalos, I. Sergienko, V. Shylo, and P. Stetsyuk. Estimating the size of correcting codes using extremal graph problems. In C. Pearce and E. Hunt, editors, *Optimization: Structure and Applications*, pages 227–243. Springer, 2009.
- [11] S. Butenko and S. Trukhanov. Using critical sets to solve the maximum independent set problem. *Oper. Res. Lett.*, 35(4):519–524, 2007.
- [12] L. Chang, W. Li, and W. Zhang. Computing a near-maximum independent set in linear time by reducing-peeling. In *Proc. 2017 ACM International Conference on Management of Data (SIGMOD 2017)*, pages 1181–1196. ACM, 2017.
- [13] J. Chen, I. A. Kanj, and W. Jia. Vertex cover: Further observations and further improvements. *Journal of Algorithms*, 41(2):280–301, 2001.
- [14] A. Conte, D. Firmani, C. Mordente, M. Patrignani, and R. Torlone. Fast enumeration of large k -plexes. In *Proc. 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2017)*, pages 115–124. ACM, 2017.
- [15] J. Dahlum, S. Lamm, P. Sanders, C. Schulz, D. Strash, and R. F. Werneck. Accelerating local search for the maximum independent set problem. In A. V. Goldberg and A. S. Kulikov, editors, *Experimental Algorithms (SEA’16)*, volume 9685 of *LNCS*, pages 118–133. 2016.
- [16] F. Dehne, M. Fellows, M. A. Langston, F. Rosamond, and K. Stevens. An $O(2^{O(k)}n^3)$ FPT algorithm for the undirected feedback vertex set problem. In L. Wang,

- editor, *Computing and Combinatorics (COCOON'15)*, volume 3595 of *LNCS*, pages 859–869. 2005.
- [17] C. Demetrescu, A. V. Goldberg, and D. S. Johnson. *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, volume 74. AMS, 2009.
- [18] T. A. Feo, M. G. C. Resende, and S. H. Smith. A greedy randomized adaptive search procedure for maximum independent set. *Oper. Res.*, 42(5):860–878, 1994.
- [19] F. Fomin and D. Kratsch. *Exact Exponential Algorithms*. Springer, 2010.
- [20] J. Gajarský, P. Hliněný, J. Obdržálek, S. Ordyniak, F. Reidl, P. Rossmanith, F. Sánchez Villaamil, and S. Sikdar. Kernelization using structural parameters on sparse graph classes. In H. L. Bodlaender and G. F. Italiano, editors, *Algorithms (ESA '13)*, volume 8125 of *LNCS*, pages 529–540. Springer, 2013.
- [21] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [22] A. Gamsa, M. Nöllenburg, and I. Rutter. Evaluation of labeling strategies for rotating maps. In *Experimental Algorithms (SEA'14)*, volume 8504 of *LNCS*, pages 235–246. Springer, 2014.
- [23] J. Guo and R. Niedermeier. Invitation to data reduction and problem kernelization. *ACM SIGACT News*, 38(1):31–45, 2007.
- [24] D. Hespe. Scalable kernelization for the maximum independent set problem. Master's thesis, Karlsruhe Institute of Technology, 2017.
- [25] M. Holtgrewe, P. Sanders, and C. Schulz. Engineering a Scalable High Quality Graph Partitioner. *Proc. 24th International Parallel and Distributed Processing Symposium*, pages 1–12, 2010.
- [26] Y. Iwata. Personal communication, August 24, 2016.
- [27] Y. Iwata, K. Oka, and Y. Yoshida. Linear-time fpt algorithms via network flow. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms, SODA '14*, pages 1749–1761. SIAM, 2014.
- [28] B. M. P. Jansen and S. Kratsch. Data reduction for graph coloring problems. *Information and Computation*, 231:70–88, 2013.
- [29] R. M. Karp and M. Sipser. Maximum matching in sparse random graphs. In *Proc. 22nd Annual Symposium on Foundations of Computer Science (SFCS'81)*, pages 364–375. IEEE, 1981.
- [30] T. Kieritz, D. Luxen, P. Sanders, and C. Vetter. Distributed time-dependent contraction hierarchies. In P. Festa, editor, *Experimental Algorithms (SEA'10)*, volume 6049 of *LNCS*, pages 83–93. Springer, 2010.
- [31] S. Lamm, P. Sanders, C. Schulz, D. Strash, and R. F. Werneck. Finding near-optimal independent sets at scale. *J. Heur.*, 23(4):207–229, Aug 2017.
- [32] C. Li, Z. Fang, and K. Xu. Combining MaxSAT Reasoning and Incremental Upper Bound for the Maximum Clique Problem. In *Proc. 25th International Conference on Tools with Artificial Intelligence (ICTAI)*, pages 939–946, Nov 2013.
- [33] H. Meyerhenke, P. Sanders, and C. Schulz. Parallel graph partitioning for complex networks. *IEEE Transactions on Parallel and Distributed Systems*, 28(9):2625–2638.
- [34] G. Nemhauser and L. E. Trotter Jr. Vertex packings: Structural properties and algorithms. *Mathematical Programming*, 8(1):232–248, 1975.
- [35] P. San Segundo, A. Lopez, and P. M. Pardalos. A new exact maximum clique algorithm for large and massive sparse graphs. *Computers & Operations Research*, 66:81–94, 2016.
- [36] P. San Segundo, F. Matia, D. Rodriguez-Losada, and M. Hernando. An improved bit parallel exact maximum clique algorithm. *Optim. Lett.*, 7(3):467–479, 2013.
- [37] P. San Segundo, D. Rodríguez-Losada, and A. Jiménez. An exact bit-parallel algorithm for the maximum clique problem. *Comput. Oper. Res.*, 38(2):571–581, 2011.
- [38] P. V. Sander, D. Nehab, E. Chlamtac, and H. Hoppe. Efficient traversal of mesh edges using adjacency primitives. *ACM Trans. Graph.*, 27(5):144:1–144:9, Dec. 2008.
- [39] P. Sanders and C. Schulz. Think locally, act globally: Highly balanced graph partitioning. In V. Bonifaci, C. Demetrescu, and A. Marchetti-Spaccamela, editors, *Experimental Algorithms (SEA'13)*, volume 7933 of *LNCS*. Springer, 2013.
- [40] D. Strash. On the power of simple reductions for the maximum independent set problem. In T. N. Dinh and M. T. Thai, editors, *Computing and Combinatorics (COCOON'16)*, volume 9797 of *LNCS*, pages 345–356. Springer, 2016.
- [41] R. E. Tarjan and A. E. Trojanowski. Finding a maximum independent set. *SIAM J. Comput.*, 6(3):537–546, 1977.
- [42] E. Tomita, Y. Sutani, T. Higashi, S. Takahashi, and M. Wakatsuki. A simple and faster branch-and-bound algorithm for finding a maximum clique. In M. S. Rahman and S. Fujita, editors, *Algorithms and Computation (WALCOM'10)*, volume 5942 of *LNCS*, pages 191–203. Springer, 2010.
- [43] A. Verma, A. Buchanan, and S. Butenko. Solving the maximum clique and vertex coloring problems on very large sparse networks. *INFORMS Journal on Computing*, 27(1):164–177, 2015.
- [44] M. von Looz, M. S. Özdayi, S. Laue, and H. Meyerhenke. Generating massive complex networks with hyperbolic geometry faster in practice. In *High Performance Extreme Computing Conference (HPEC), 2016 IEEE*, pages 1–6. IEEE, 2016.
- [45] M. Xiao and H. Nagamochi. Confining sets and avoiding bottleneck cases: A simple maximum independent set algorithm in degree-3 graphs. *Theor. Comput. Sci.*, 469:92–104, 2013.
- [46] M. Xiao and H. Nagamochi. Exact algorithms for maximum independent set. *Information and Computation*, 255(1):126–146, 2017.

A High Level Algorithm Overview

Algorithm 1 Algorithm Overview

```

 $G \leftarrow$  input graph
 $\{V_1, \dots, V_k\} \leftarrow \text{partition}(G, k)$ 
while  $G$  changed in last iteration do
  for all blocks  $V_i$  in parallel do
     $G \leftarrow \text{localReductions}(G, V_i)$ 
  end for
   $G \leftarrow \text{parallelLinearProgrammingReduction}(G)$ 
end while

```

B Benchmark Properties

name	type	# vertices	# edges	from
uk-2002	web	18.5M	261.8M	[6]
arabic-2005	web	22.7M	553.9M	[8]
gsh-2015-tpd	web	30.8M	489.7M	[8]
uk-2005	web	39.5M	783.0M	[8]
it-2004	web	41.3M	1 027.5M	[8]
sk-2005	web	50.6M	1 810.1M	[8]
uk-2007-05	web	105.9M	3 301.9M	[8]
webbase-2001	web	118.1M	854.8M	[8]
asia.osm	road	12.0M	12.7M	[6]
road_usa	road	23.9M	28.9M	[17]
europa.osm	road	50.9M	54.1M	[6]
rgg26	rgg	67.1M	574.6M	[6]
rhg	rhg	100.0M	1 999.5M	[44]
del24	delaunay	16.8M	50.3M	[6]
del26	delaunay	67.1M	201.3M	[33]

Table 2: Basic properties of the graphs used in our evaluation.

C Reduction Tracking: Counteracting Diminishing Returns

graph	speedup
uk-2002	1.1
arabic-2005	2.2
it-2004	1.8
sk-2005	137.7
uk-2007-05	2.4
webbase-2001	1.3
rgg26	1.2

Table 3: Overview of our reduction tracking technique for graphs with speedup ≥ 1.0 .

Table 3 shows the effect of our reduction tracking technique described in Section 5.2 on ParFastKer. It

shows the algorithmic speedup achieved by enabling reduction tracking. Only graphs with speedup ≥ 1.0 are shown. All other speedups are 1.0.

Our experiments show that stopping long lasting reductions early can lead to significant speedups with close to no penalty on the quasi kernel size. The quasi kernel size found with reduction tracking enabled is never more than 0.1% larger than without it. In fact, the quasi kernel is sometimes even slightly smaller. The reason for this is that different orders of reduction application can lead to different kernel sizes.

D Detailed Results

Here, we provide detailed results of our experiments. In addition to the time to reach a quasi kernel, we also provide the time it takes to reach a full kernel. We do this by first applying our algorithm as described throughout the paper to find a quasi kernel. We then apply the remaining reductions by running sequentially and disabling the inexact reduction pruning technique described in Section 5.2.

In the comparisons to the other algorithms, we also provide columns for a “same size comparison”. This is found by logging the current time and size throughout the algorithms. When comparing two algorithms with different kernel sizes, the time column of the same size comparison then reports the first time stamp at which the algorithm with the smaller kernel size logged a size smaller than (or equal to) the final size of the algorithm with the larger kernel.

The implementation by Chang et al. uses 32-bit integers as edge identifiers, so they cannot process graphs with $2^{32} \approx 4.29\text{B}$ or more edges. Respective entries in the tables are marked with a star (*). As our algorithm uses their LinearTime implementation as a preprocessing step, for graphs with too many edges, we use the original graph as input to our algorithm instead of the kernel found by LinearTime.

Graph		LinearTime		NearLinear		VCSolver		FastKer		ParFastKer		
name	n	$ \mathcal{K} $	time	$ \mathcal{K} $	time	$ \mathcal{K} $	time	$ \mathcal{K} $	time	$ \mathcal{K} $	time	su
uk-2002	19M	11.7M	1.5	4.0M	28.0	0.2M	336.9	0.3M	48.9	0.3M	11.8	28.6
arabic-2005	23M	15.6M	2.6	6.7M	246.1	0.6M	1033.2	0.6M	124.7	0.6M	25.7	40.3
gsh-2015-tpd	31M	2.0M	11.6	1.2M	97.4	0.4M	372.3	0.4M	54.0	0.5M	31.7	11.7
uk-2005	39M	28.2M	2.5	5.9M	60.5	0.8M	541.4	0.9M	116.2	0.9M	52.8	10.2
it-2004	41M	27.1M	3.3	11.3M	1 544.6	1.6M	6 749.0	1.7M	439.3	1.7M	151.1	44.7
sk-2005	51M	*	*	*	*	3.2M	10 010.5	3.3M	2 113.2	3.5M	174.8	57.3
uk-2007-05	106M	*	*	*	*	3.5M	18 829.4	3.6M	1 768.9	3.7M	369.5	51.0
webbase-2001	118M	51.7M	13.0	17.3M	121.1	0.7M	4 207.8	0.8M	303.5	0.9M	54.9	76.7
asia.osm	12M	626.7K	0.8	594.4K	1.4	15.2K	204.7	34.9K	1.4	34.9K	1.2	172.3
road_usa	24M	2.5M	2.5	2.4M	4.1	0.2M	310.0	0.2M	6.8	0.2M	4.0	77.3
europa.osm	51M	1 500.0K	4.1	1 329.9K	6.1	8.4K	302.4	14.1K	5.4	14.2K	5.0	60.0
rgg26	67M	67.1M	1.0	51.3M	172.6	49.6M	9 887.7	49.8M	6 329.1	49.8M	148.9	66.4
rhg	100M	*	*	*	*	0	124.0	0	138.1	16	65.1	1.9
del24	17M	16.8M	0.2	15.6M	12.7	12.4M	4 789.5	12.9M	128.0	12.9M	50.9	94.2
del26	67M	67.1M	0.7	62.5M	53.3	49.9M	20 728.7	51.7M	506.3	51.7M	176.1	117.7

Table 4: Running times and kernel sizes ($|\mathcal{K}|$) for all algorithms. The column “su” is the speedup of ParFastKer over VCSolver. Instances marked with a star (*) cannot be processed by the NearLinear and LinearTime implementations due to the 32-bit implementation. All times are in seconds.

graph	quasi kernel				kernel		
	$ \mathcal{K} $	LinearTime [s]	all reductions [s]	total [s]	$ \mathcal{K} $	time [s]	total [s]
uk-2002	255 498	1.5	47.4	48.9	255 498	1.5	50.3
arabic-2005	610 643	2.6	122.2	124.7	610 643	2.9	127.7
gsh-2015-tpd	425 736	11.6	42.4	54.0	425 736	1.8	55.8
uk-2005	854 511	2.5	113.7	116.2	854 511	3.2	119.4
it-2004	1 650 982	3.3	436.0	439.3	1 645 596	629.8	1 069.1
sk-2005	3 264 231	*	2 113.2	2 113.2	3 256 661	691.7	2 804.9
uk-2007-05	3 631 632	*	1 768.9	1 768.9	3 628 009	2 566.3	4 335.2
webbase-2001	821 094	13.0	290.5	303.5	820 948	24.3	327.8
asia.osm	34 930	0.8	0.6	1.4	34 930	0.1	1.5
road_usa	247 395	2.5	4.3	6.8	247 395	0.3	7.1
europa.osm	14 066	4.1	1.3	5.4	14 066	0.1	5.5
rgg26	49 843 889	1.0	6 328.0	6 329.1	49 838 878	919.4	7 248.5
rhg	0	*	138.1	138.1	0	7.8	145.9
del24	12 884 246	0.2	127.8	128.0	12 877 158	159.2	287.3
del26	51 674 080	0.7	505.6	506.3	51 624 274	626.1	1 132.5

Table 5: Kernel sizes and kernelization times for FastKer to reach a quasi kernel and to reach a full kernel by running our algorithm without stopping the reduction application on the quasi kernel.

graph	quasi kernel						kernel		
	$ \mathcal{K} $	LinearTime [s]	part. [s]	all reductions [s]	total [s]	$ \mathcal{K} $	time [s]	total [s]	
uk-2002	266 328	1.5	4.9	5.4	11.8	255 594	4.8	16.6	
arabic-2005	628 850	2.6	8.3	14.8	25.7	610 288	21.0	46.6	
gsh-2015-tpd	486 328	11.6	13.5	6.6	31.7	425 751	19.7	51.5	
uk-2005	901 896	2.5	35.5	14.8	52.8	854 383	18.1	70.9	
it-2004	1 697 934	3.3	30.0	117.8	151.1	1 645 643	148.2	299.3	
sk-2005	3 504 786	*	70.0	104.8	174.8	3 256 591	211.7	386.6	
uk-2007-05	3 735 056	*	71.0	298.5	369.5	3 629 214	510.0	879.5	
webbase-2001	869 443	13.0	18.3	23.6	54.9	821 131	14.8	69.7	
asia.osm	34 851	0.8	0.3	0.1	1.2	34 823	0.1	1.3	
road_usa	246 939	2.5	1.2	0.4	4.0	246 939	0.3	4.3	
europa.osm	14 152	4.1	0.7	0.2	5.0	14 096	0.1	5.2	
rgg26	49 847 428	1.0	44.2	103.7	148.9	49 838 810	780.7	929.6	
rhg	16	*	44.4	20.7	65.1	0	28.0	93.0	
del24	12 901 142	0.2	31.7	19.0	50.9	12 877 029	135.0	185.8	
del26	51 668 286	0.7	86.9	88.5	176.1	51 624 361	611.0	787.0	

Table 6: Kernel sizes and kernelization times for ParFastKer to reach a quasi kernel and to reach a full kernel by running our algorithm sequentially and without stopping the reduction application on the quasi kernel.

graph	VCSolver		ParFastKer			Same size comparison	
	$ \mathcal{K} $	time [s]	$ \mathcal{K} $	time [s]	speedup	time [s]	speedup
uk-2002	241 517	336.9	266 328	11.8	28.6	199.5	16.9
arabic-2005	574 878	1 033.2	628 850	25.7	40.3	509.5	19.9
gsh-2015-tpd	417 031	372.3	486 328	31.7	11.7	139.5	4.4
uk-2005	835 480	541.4	901 896	52.8	10.2	137.1	2.6
it-2004	1 602 560	6 749.0	1 697 934	151.1	44.7	4 108.9	27.2
sk-2005	3 200 806	10 010.5	3 504 786	174.8	57.3	2 822.2	16.1
uk-2007-05	3 514 783	18 829.4	3 735 056	369.5	51.0	11 828.5	32.0
webbase-2001	736 842	4 207.8	869 443	54.9	76.7	2 626.6	47.9
asia.osm	15 201	204.7	34 851	1.2	172.3	159.5	134.3
road_usa	169 808	310.0	246 939	4.0	77.3	91.6	22.8
europa.osm	8 366	302.4	14 152	5.0	60.0	214.5	42.6
rgg26	49 590 973	9 887.7	49 847 428	148.9	66.4	1 637.8	11.0
rhg	0	124.0	16	65.1	1.9	113.4	1.7
del24	12 417 301	4 789.5	12 901 142	50.9	94.2	1 002.4	19.7
del26	49 864 448	20 728.7	51 668 286	176.1	117.7	4 713.6	26.8

Table 7: Comparison between VCSolver and ParFastKer. "Same size comparison" compares the time that the algorithm with the smaller kernel size takes to reach the final size of the algorithm with the larger kernel.

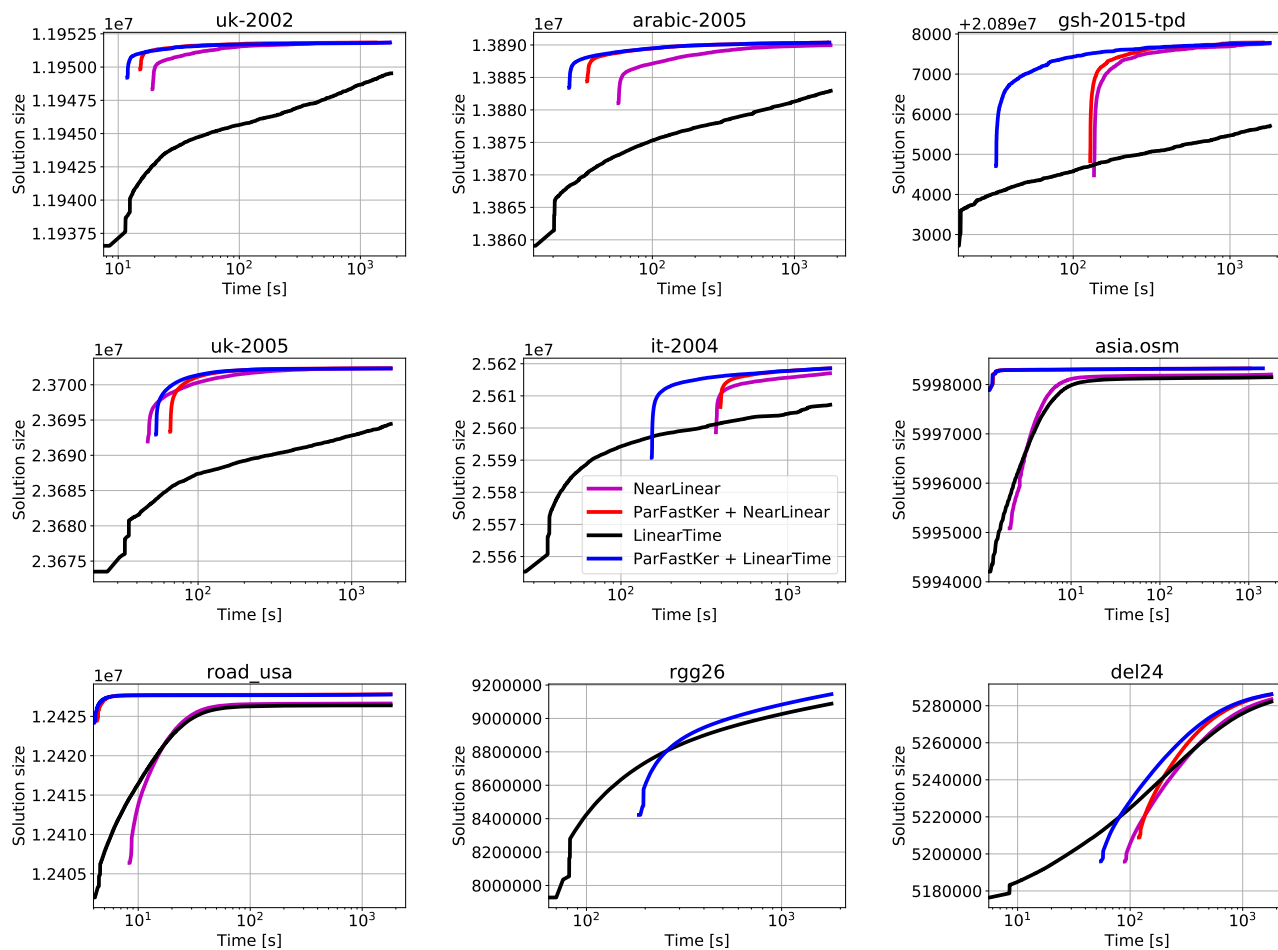


Figure 6: Solution size over time of LinearTime and NearLinear in the original version and with ParFastKer as preprocessing step. On rgg26, NearLinear did not find an initial solution within the time limit.