

Scalable Kernelization for Maximum Independent Sets

DEMIAN HESPE, Karlsruhe Institute of Technology

CHRISTIAN SCHULZ, University of Vienna

DARREN STRASH, Hamilton College

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. Current algorithms are either slow but produce a small kernel or fast and give a large kernel. Yet it is of critical importance for these algorithms that kernelization is fast and returns a small 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 possible previously. 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.

CCS Concepts: • **Mathematics of computing** → **Graph algorithms**; *Combinatorial optimization*; • **Computing methodologies** → **Shared memory algorithms**;

Additional Key Words and Phrases: Maximum independent set, parallelization, shared-memory, kernelization, graphs

ACM Reference format:

Demian Hespe, Christian Schulz, and Darren Strash. 2019. Scalable Kernelization for Maximum Independent Sets. *J. Exp. Algorithmics* 24, 1, Article 1.16 (September 2019), 22 pages.

<https://doi.org/10.1145/3355502>

This article appeared in preliminary and shorter form in Proceedings of the 20th Workshop on Algorithm Engineering and Experiments [25].

Our research was supported by the Gottfried Wilhelm Leibniz Prize 2012, the Large-Scale Data Management and Analysis (LSDMA) project in the Helmholtz Association; received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP/2007-2013)/ERC grant agreement no. 340506; and was supported by DFG grant SCHU 2567/1-2.

Authors’ addresses: D. Hespe, Karlsruhe Institute of Technology, Kaiserstraße 12, 76131, Karlsruhe, Germany; email: hespe@kit.edu; C. Schulz, University of Vienna, Faculty of Computer Science, Universitätsring 1, 1010, Vienna, Austria; email: christian.schulz@univie.ac.at; D. Strash, Hamilton College, Department of Computer Science, 198 College Hill Road, Clinton, NY 13323, USA; email: dstrash@hamilton.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1084-6654/2019/09-ART1.16 \$15.00

<https://doi.org/10.1145/3355502>

1 INTRODUCTION

The maximum independent set (MIS) problem is a classic NP-hard problem [22] with applications spanning many fields, such as classification theory, information retrieval, computer vision [19], computer graphics [38], map labeling [23], and routing in road networks [31]. Given a graph $G = (V, E)$, our goal is to compute a maximum cardinality set of vertices $I \subseteq V$ such that no vertices in I are adjacent to one another. Such a set is called a *maximum independent set*. As a concrete application, independent sets are essential in labeling strategies for maps [23], where the objective is to maximize the number of visible nonoverlapping 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 an MIS 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 a 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 [42]. 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 [9, 14, 15, 37, 42, 45], where speeding up kernelization directly speeds up the algorithm. However, kernelization may also be applied repeatedly as part of an algorithm [2, 12, 32]. 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 [32]. 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]. Fast and simple reductions can even be used to solve many large-scale instances exactly [42] 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 MIS 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 MIS problem: a maximum clique in the complement graph \bar{G} is an MIS in G , and a minimum vertex cover C in G is the complement of an MIS $V \setminus C$ 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 [40, 41, 44], which use vertex reordering and pruning techniques based on approximate graph coloring [44] or MaxSAT [33], 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 MIS problems [1, 13, 21, 43], for efficient exact algorithms and heuristics alike.

2.1 Exact Algorithms

Butenko et al. [9] and Butenko and Trukhanov [11] were able to find exact MIS in graphs with thousands of vertices by first applying reductions. Further works have introduced reductions to more quickly solve the maximum clique problem [37, 45] and enumerate k -plexes [14]. Although 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 [43] 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 $O^*(1.1996^n)$ by Xiao and Nagamochi [48]. These algorithms apply reductions during recursion, only branching when the graph can no longer be reduced [20]—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, most instances cannot be solved by the MCS clique solver [44] within a 24-hour time limit [2]. However, as later shown by Strash [42], 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 as well. Lamm et al. [32] 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, and they found 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 implementing 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 $O(\Delta m)$, 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 [2], 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 article, we assume that G is *simple*: it has no multiedges or selfloops. 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 *MIS problem* is that of finding a maximum cardinality independent set, which is called a *maximum independent set*.

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 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 MIS problem. Note that Akiba and Iwata further use packing [2] and alternative [47] 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

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

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 [36]. 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. Note that there is a version of this reduction [28] that computes a solution whose half-integral part is minimal. However, preliminary experiments showed that in practice no additional vertices can be removed.

Unconfined [47]. Although 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, as there always exists an MIS that contains no unconfined vertices.

Diamond. Although not mentioned in their work, Akiba and Iwata [2] extend the unconfined reduction in their implementation [27]. 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 [47]. 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. [10]. We describe one reduction that was not included in the algorithm by Akiba and Iwata [2] but was shown by Butenko et al. [10] to be highly effective on medium-size graphs derived from error-correcting codes.

Isolated Vertex Removal [10]. If a vertex v forms a single clique C with all of 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, 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 $O(m)$. 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 work [12]. The degree-two path rule is a specialization of the vertex folding rule explained earlier 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 graphs 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 MIS 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 Algorithm 1.)

ALGORITHM 1: Algorithm Overview

```

G ← input graph
{V1, . . . , Vk} ← partition(G, k)
while G changed in last iteration do
  for all blocks Vi in parallel do
    G ← localReductions(G, Vi)
  end for
  G ← parallelLinearProgrammingReduction(G)
end while

```

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 (e.g., 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-grain parallelism would require locks, as 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. As reductions may still be applied, we call the resulting graph a quasi kernel instead of a kernel. 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 and only mark vertices as removed from the graph. We only change the adjacency list of vertices when performing vertex contractions.

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

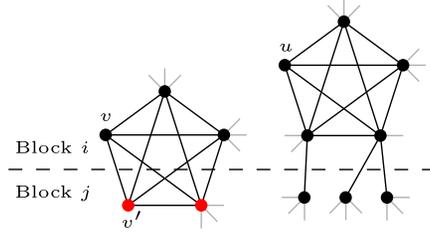


Fig. 1. Vertices v and v' are simplicial but lie in different blocks; they are not removed from the graph. Vertex u is simplicial and is not a boundary vertex, so $N[u]$ is removed.

(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' . In these cases, we do not apply the vertex folding reduction. 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. Second, without loss of generality, if $u \in V_i \cap B_0$ and $w \in V_i \setminus B_0$ 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 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 that 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 that 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 [30], which was parallelized by Azad et al. [5]. For better performance when repeatedly applying the reduction, we reuse the parts of the previous matching that are still part of the graph. If the graph has 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 R_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.

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 apply. 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

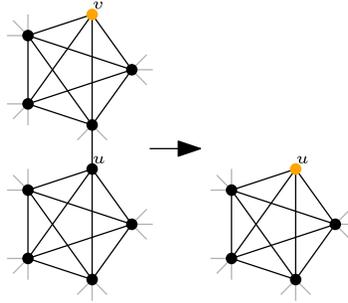


Fig. 2. After removing isolated vertex v and $N(v)$, u is isolated. Orange vertices are in D .

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 [42] implements similar dependency checking for isolated vertex and vertex folding reductions, although it is not mentioned in his work. 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, whereas 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 afterward. Furthermore, in our parallel algorithm, applying the local reductions exhaustively can take significantly longer for some blocks than for others. In other words, the total graph size is not significantly reduced once the first threads finish their blocks.

We therefore implement reduction tracking to detect 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 afterward. 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.

6 EXPERIMENTAL EVALUATION

Methodology. We implement our algorithm using C++ and compile 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. Our source code is available on GitHub,² and a sequential version of our algorithm has been integrated into the KaMIS software for finding high-quality independent sets.³ For graph partitioning, we use ParHIP [35], 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 VCSolver, the full reduction suite of Akiba and Iwata [2].⁴ We modify 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 512GB RAM and two Intel Xeon E5-2683 v4 processors with 16 cores running at 2.1GHz each.

Data Structure Details. We represent our graph using adjacency lists. For every vertex, we store an array of its neighbors ids. When a vertex is removed from the graph, it is not removed from the adjacency lists of its neighbors but instead is marked as removed in an array accessible by all threads. Note that we cannot store all edges of the graph consecutively as the vertex folding and twin reductions can increase the number of neighbors of a vertex. To efficiently check the degree of a vertex or its number of incident cut edges, we store these values using atomic integers.

For each block, we additionally store the following data structures that are only used by the thread that is handling the respective block. To efficiently iterate over vertices that have not been removed from the graph, we keep a consecutive array of vertex ids from the respective block. For constant time removal of vertices from this array, we store additional pointers from the vertex id to the position in the array. This data structure is also used in our dependency checking technique to store the vertices that have to be considered for reduction.

Algorithm Configuration. We run our algorithm with all reduction rules explained in this article 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 Section 6.3. 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 [26, 46] and Delaunay triangulations [6, 35]. Basic instance properties can be

²<https://github.com/Hespian/ParFastKer>.

³<http://algo2.iti.kit.edu/kamis/>.

⁴https://github.com/wata-orz/vertex_cover.

Table 1. Basic Properties of the Graphs Used in Our Evaluation

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	[46]
del24	Delaunay	16.8M	50.3M	[6]
del26	Delaunay	67.1M	201.3M	[35]

found in Table 1. 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 2 and Figure 3 give an overview. Figure 3 normalizes running time and kernel size on each instance by the result of VCSolver. See Appendix A for more detailed results.

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 1,329,923-vertex kernel for europe.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 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.

Table 2. Running Times and Kernel Sizes ($|\mathcal{K}|$) for All Algorithms

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	60.1	0.3M	11.8	28.4
arabic-2005	23M	15.6M	2.6	6.7M	246.1	0.6M	1,033.2	0.6M	148.0	0.6M	25.7	40.2
gsh-2015-tpd	31M	2.0M	11.6	1.2M	97.4	0.4M	372.3	0.4M	66.4	0.5M	32.0	11.7
uk-2005	39M	28.2M	2.5	5.9M	60.5	0.8M	541.4	0.9M	131.9	0.9M	53.3	10.1
it-2004	41M	27.1M	3.3	11.3M	1,544.6	1.6M	6,749.0	1.7M	499.7	1.7M	151.8	44.4
sk-2005	51M	*	*	*	*	3.2M	10,010.5	3.3M	2,349.8	3.5M	178.3	56.1
uk-2007-05	106M	*	*	*	*	3.5M	18,829.4	3.6M	2,073.4	3.7M	372.4	50.6
webbase-2001	118M	51.7M	13.0	17.3M	121.1	0.7M	4,207.8	0.8M	290.8	0.9M	54.9	76.6
asia.osm	12M	626.7K	0.8	594.4K	1.4	15.2K	204.7	34.9K	1.6	34.9K	1.2	169.8
road_usa	24M	2.5M	2.5	2.4M	4.1	0.2M	310.0	0.2M	8.0	0.2M	4.1	76.0
europa.osm	51M	1,500.0K	4.1	1,329.9K	6.1	8.4K	302.4	14.1K	5.8	14.2K	4.9	61.3
rgg26	67M	67.1M	1.0	51.3M	172.6	49.6M	9,887.7	49.8M	13,572.6	49.8M	150.3	65.8
rhg	100M	*	*	*	*	0	124.0	0	164.5	16	64.6	1.9
del24	17M	16.8M	0.2	15.6M	12.7	12.4M	4,789.5	12.9M	142.0	12.9M	51.5	93.1
del26	67M	67.1M	0.7	62.5M	53.3	49.9M	20,728.7	51.7M	718.9	51.7M	179.0	115.8

The column “SU” is the speedup of ParFastKer over VCSolver. Instances marked with an asterisk (*) cannot be processed by the NearLinear and LinearTime implementations due to the 32-bit implementation. All times are in seconds. Quasi kernel sizes that differ from VCSolver’s kernel size by at most 0.5% of the graph size are emphasized in **bold**.

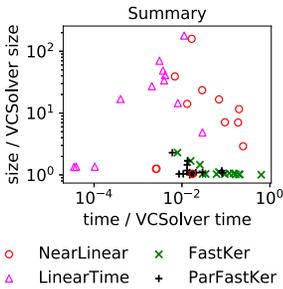


Fig. 3. Comparison of kernelization algorithms against VCSolver.⁵

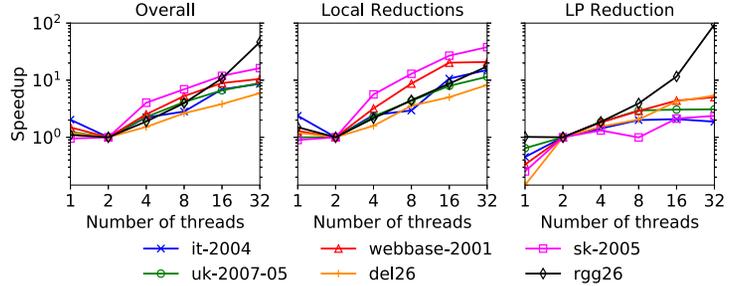


Fig. 4. Scaling experiments on the six hardest instances of the benchmark set for the overall algorithm (left), blockwise reductions (center), and the reduction by linear programming (right). Speedups are relative to two threads.

As VCSolver implements a larger set of reduction rules, adding the desk and funnel reductions by Xiao and Nagamochi [47], as well as the *packing* reduction rule of Akiba and Iwata [2], 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 among all graphs we tested is 2.9% on del24 (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 and rgg26. ParFastKer, however, is faster than VCSolver on these instance. On 11 of 15 instances, FastKer is faster by a factor of more than 5

⁵As empty kernels lead to a division by zero for this plot, graphs with an empty kernel are not shown here.

than VCSolver and on 5 instances even by a factor of more than 28. The largest speedup of FastKer over VCSolver is 129 on `asia.osm`, and the geometric mean of the speedups is 10. 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 more than 10 and on 9 instances more than 50.

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. [2], 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 1.7 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 2 threads is between 6 and 16.3, with 16 being perfect speedup. The speedup relative to the single-threaded case is between 3.3 and 13.1 (42.1 for `rgg26`).

Figure 4 shows that local reductions parallelized by partitioning are faster single threaded than on 2 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 2 threads, the speedup for 32 threads compared to 2 threads for these reductions is between 8 and 37 (between 4.8 and 32 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, as it does not scale as well as the rest of the reductions. In many cases, about half of the reduction time is spend on this reduction rule alone.

6.3 Reduction Tracking: Counteracting Diminishing Returns

Table 3 shows the effect of our reduction tracking technique described in Section 5.2 on ParFastKer. It shows the algorithmic speedup over ParFastKer achieved by enabling reduction tracking. We also show the relative quasi kernel size increase caused by using reduction tracking.

Our experiments show that stopping long-lasting reductions early can lead to significant speedups on some graphs with close to no penalty on the quasi kernel size. The quasi kernel size found with reduction tracking enabled is less than 0.1% larger than without it on all but two of our test instances. And even for these two instances, the difference is only minor (at most 0.6%). 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.

Table 3. Speedup and Relative Change in Kernel Size Change Achieved by Using Reduction Tracking

Graph	Speedup	Δ Size
uk-2002	1.1	+0.0%
arabic-2005	2.2	-0.0%
gsh-2015-tpd	1.0	-0.0%
uk-2005	1.0	-0.0%
it-2004	1.8	-0.0%
sk-2005	135.0	-0.0%
uk-2007-05	2.4	+0.1%
webbase-2001	1.3	+0.0%
asia.osm	1.0	-0.0%
road_usa	1.0	+0.0%
europa.osm	1.0	+0.6%
rgg26	1.2	+0.0%
rhg	1.0	+0.0%
del24	1.0	+0.0%
del26	1.0	+0.0%

6.4 Impact of Partitioning

In this section, we assess the impact of the partitioning quality on ParFastKer. We compare ParHIP’s fastest configuration (ultrafast), which we used for our other experiments, with a higher-quality configuration (fast), as well as a size-constrained label propagation algorithm (SC-LPA) [34]. SC-LPA is a simple graph partitioning algorithm with fast running time; however, it usually gives low-quality output. We set the imbalance for all three algorithms to 3%.

Table 4 shows that with ParHIP’s fast configuration, the total time for kernelization increases by a factor of 1.52 in the geometric mean compared to the ultrafast configuration. The kernel size, however, remains largely unchanged. Focusing on the difference in the number of edges cut, we clearly see that ParHIP’s fast configuration gives only minimally better partitions than the ultrafast configuration: on all instances except for *sk-2005*, the difference is under 0.03%. Although this might be important for certain applications, it does not seem to be worth the longer running time for our algorithm, as the quasi kernel size changes only slightly, at most 1% on most instances. Note that *rhg* has an empty kernel and ParFastKer using ParHIP’s ultrafast configuration finds a quasi kernel of size 16, so the size difference reported in Table 4 is negligible.

However, using the size-constrained label propagation algorithm, the kernel size increases drastically due to the much larger amount of cut edges. We see that SC-LPA produces up to 12% larger cuts than ParHIP’s ultrafast configuration, resulting in quasi kernels larger by a factor of 2.4 in the geometric mean as more reductions are skipped because they lie on boundaries between blocks. It is also important to note that our implementation of SC-LPA is sequential—it is possible that total kernelization time would be faster with a parallel size-constrained label propagation than with ParHIP’s ultrafast configuration. However, these experiments show that this simple partitioning algorithm does not produce partitions of high enough quality to be used by ParFastKer—even with faster running times—as quasi kernel sizes become too large.

6.5 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

Table 4. Comparison of ParFastKer’s Running Time and Quasi Kernel Size with the Ultrafast Configuration of ParHIP to the Fast Configuration and an SC-LPA

Graph	ParHIP (Fast)			SC-LPA		
	Δ Cut	Time	$ \mathcal{K} $	Δ Cut	Time ⁶	$ \mathcal{K} $
uk-2002	-0.00%	1.56	1.00	+3.59%	4.89	1.80
arabic-2005	-0.01%	1.75	1.00	+4.77%	4.26	1.73
gsh-2015-tpd	-0.00%	1.71	1.00	+6.47%	1.16	1.13
uk-2005	-0.02%	1.87	1.00	+4.68%	3.28	1.35
it-2004	-0.01%	1.36	1.00	+5.08%	1.94	1.51
sk-2005	-0.49%	1.68	0.97	*	*	*
uk-2007	-0.03%	1.68	0.99	*	*	*
webbase-2001	-0.00%	1.43	1.00	+2.84%	4.21	2.18
asia.osm	-0.00%	1.18	0.99	+9.41%	1.22	8.69
road_usa	0.00%	1.30	1.00	+9.70%	1.54	5.94
europa.osm	0.00%	1.14	0.99	+8.71%	1.21	33.02
rgg26	-0.01%	1.20	1.00	+10.79%	2.84	1.02
rhg	-0.00%	1.94	0.62	*	*	*
del24	-0.01%	1.70	1.00	+12.27%	1.11	1.16
del26	-0.01%	1.64	1.00	+12.30%	1.27	1.16

Times and kernel sizes are divided by the respective value for the ultrafast configuration. The column “ Δ cut” gives the difference in the number edges cut by the partition divided by the total number of edges in the graph (in comparison to ParHIP’s ultrafast configuration).

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. [2], 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.

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

⁶The implementation we use for SC-LPA is a 32-bit implementation. Graphs that cannot be processed by it are marked with an asterisk (*).

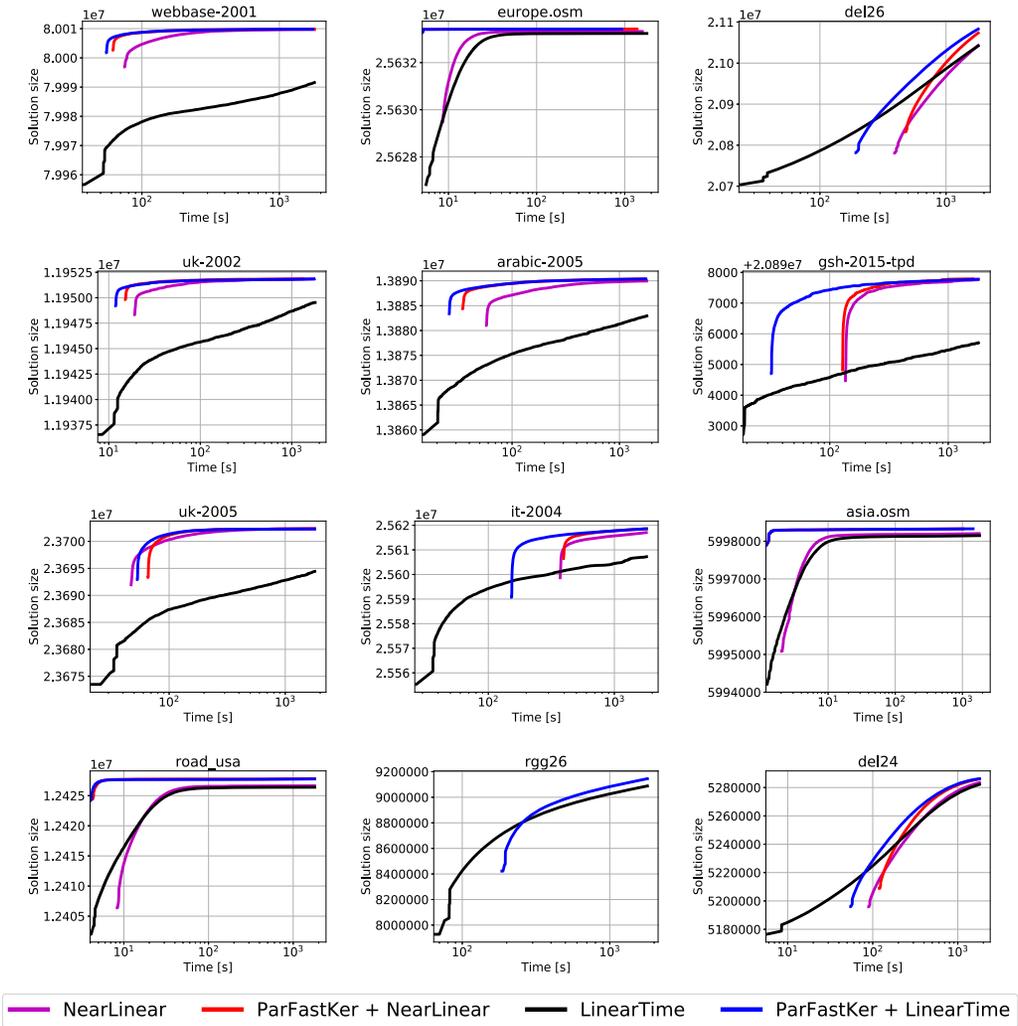


Fig. 5. Solution size over time of LinearTime and NearLinear in the original version and with ParFastKer as preprocessing step (marked with "ParFastKer +"). On rgg26, NearLinear did not find an initial solution within the time limit. Note that by LinearTime and NearLinear, we refer to the full local search algorithms, not just their kernelization parts.

kernel. In addition, using ParFastKer’s quasi kernel, the algorithm converges much faster. In particular, after an initial improvement over the starting solution, which 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, however, 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.

6.6 Improving High-Quality Heuristic Algorithms

We now show the impact that our fast kernelization makes on a heuristic algorithm that is tailored toward very high quality solutions. In particular, we consider ReduMIS by Lamm et al. [32].

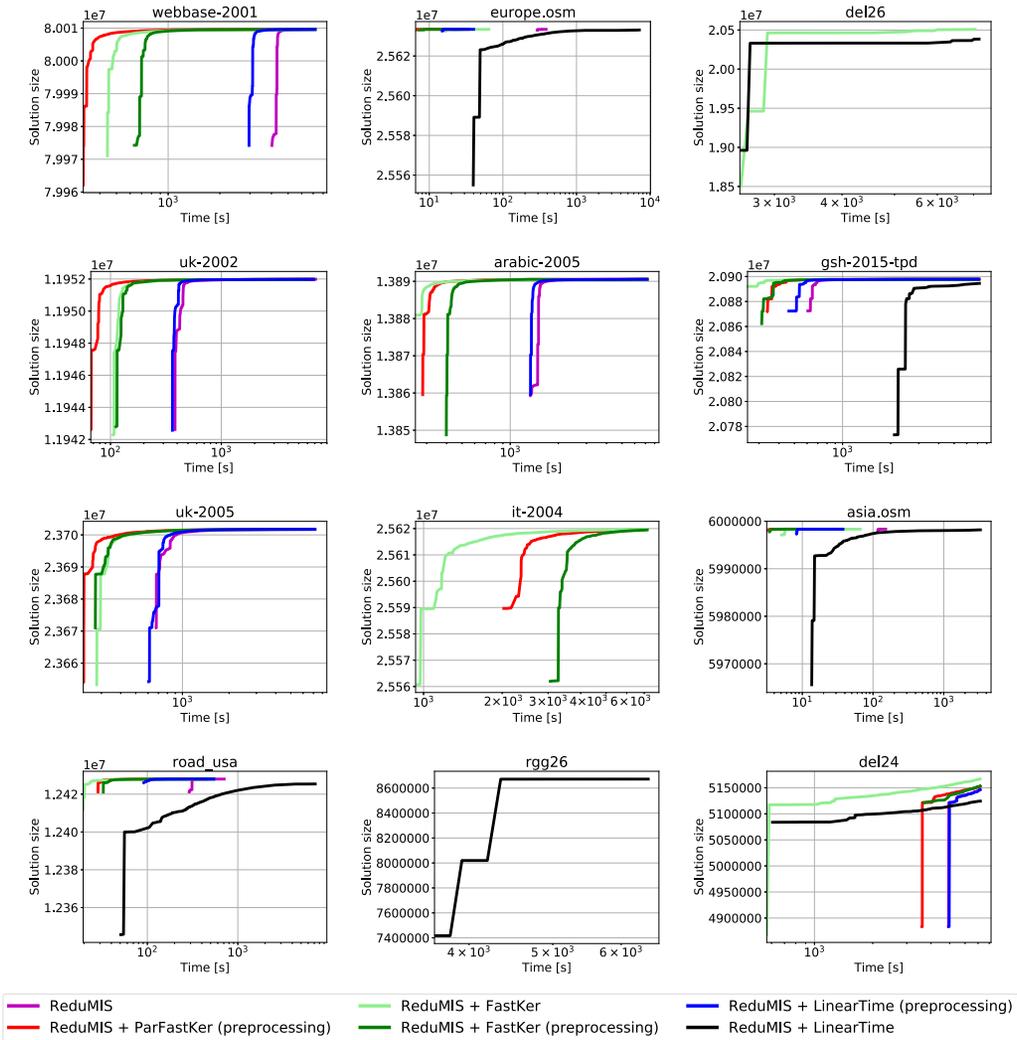


Fig. 6. Experimental results for different variations of ReduMIS. Variants marked with (preprocessing) use the indicated kernelization algorithm as preprocessing and then run the original ReduMIS algorithm on the resulting kernel. The other variants replace the kernelization algorithm used inside of ReduMIS. All runs were done with a time limit of 2 hours.

ReduMIS first finds the kernel of the input graph using VCSolver’s kernelization algorithm and then uses an evolutionary algorithm to find a large independent set of the kernel. It then fixes the 10% lowest degree vertices from the independent set into the solution, removing them and their neighborhood. After removal, the graph has changed and kernelization can be run again. This is repeated until a time limit is met or the graph has been fully reduced. We show results for two different experiments: replacing the kernelization algorithm used by ReduMIS internally by integrating a different kernelization algorithm, and first reducing the input graph using different kernelization algorithms as a preprocessing step and then using the resulting kernel as input to the original version of ReduMIS that uses VCSolver for kernelization. The time limit for all experiments in this section was set to 2 hours. Figure 6 shows the results of these experiments.

We see that for many instances, the versions with FastKer and ParFastKer outperform the other versions. The algorithm starts finding solutions earlier and thus has more time to improve its solution until the time limit is met. On some graphs, the version that integrates FastKer into ReduMIS performs significantly better than the versions that use our algorithms as a preprocessing step: on 6 of 12 instances ReduMIS + FastKer performs best among all other variants and on 5 of 12 instances ReduMIS + ParFastKer (preprocessing) outperforms the other variants—only slightly in some cases, however. The advantage of the integrated version is especially noticeable on `del26`, where the nonintegrated versions fail to find a solution within the time limit, and `it-2004` and `del24`, where the nonintegrated versions start finding solutions much later than the integrated version and thus have less time to improve their initial solution. Possible explanations for this are that VCSolver (which is used by ReduMIS in all versions that just use a different kernelization algorithm as preprocessing step) is slow in applying the remaining reductions that FastKer and ParFastKer did not apply, or that there are large graph size reductions in later stages of the algorithm that can be sped up in the integrated version. As on most graphs, the plots for the integrated FastKer version and the preprocessing versions behave quite similarly after finding a first solution; we assume the former to be the case.

The version with integrated LinearTime kernelization cannot reduce the graph enough for ReduMIS to find any solution on the graphs with high-degree vertices (road networks and the geometric graph instances shown here usually do not have high-degree vertices).

7 CONCLUSION

We presented an efficient parallel kernelization algorithm based on graph partitioning and parallel bipartite maximum matching, vertex pruning, and 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 parallelize the LinearTime algorithm by Chang et al. [12] so that our algorithm is fully parallel; apply our parallel kernelization techniques in more MIS algorithms, such as exact branch-and-reduce [2]; explore techniques for further parallelizing the LP reduction; and transfer our techniques to other problems that use kernelization [16, 18, 24, 29].

APPENDIX

A DETAILED RESULTS

Here we provide detailed results of our experiments. In addition to the time to reach a quasi kernel, Tables 5 and 6 provide the time it takes to reach a full kernel. We do this by first applying our algorithm as described throughout the article 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 VCSolver, Table 7 provides 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 timestamp 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.

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)	All Reductions (s)	Total (s)	$ \mathcal{K} $	Time (s)	Total (s)
uk-2002	255,497	1.5	58.6	60.1	255,497	1.6	61.7
arabic-2005	610,715	2.6	145.5	148.0	610,697	4.4	152.5
gsh-2015-tpd	425,645	11.6	54.8	66.4	425,645	1.8	68.2
uk-2005	854,511	2.5	129.3	131.9	854,511	3.7	135.6
it-2004	1,651,630	3.3	496.4	499.7	1,645,591	741.7	1,241.5
sk-2005	3,265,615	*	2,349.8	2,349.8	3,256,645	793.0	3,142.8
uk-2007-05	3,631,546	*	2,073.4	2,073.4	3,627,912	2,830.7	4,904.1
webbase-2001	821,492	13.0	277.8	290.8	821,092	59.2	350.0
asia.osm	34,930	0.8	0.8	1.6	34,930	0.1	1.6
road_usa	247,395	2.5	5.4	8.0	247,395	0.4	8.3
europa.osm	14,066	4.1	1.7	5.8	14,066	0.1	5.9
rgg26	49,843,887	1.0	13,571.6	13,572.6	49,838,878	1,024.9	14,597.6
rhg	0	*	164.5	164.5	0	11.3	175.8
del24	12,884,514	0.2	141.8	142.0	12,877,164	173.4	315.4
del26	51,701,698	0.7	718.2	718.9	51,624,241	708.3	1,427.2

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	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

The implementation by Chang et al. [12] 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 an asterisk (*). 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.

Table 7. Comparison between VCSolver and ParFastKer

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

"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.

REFERENCES

- [1] Faisal N. Abu-Khzam, Michael R. Fellows, Michael A. Langston, and W. Henry Suters. 2007. Crown structures for vertex cover kernelization. *Theor. Comput. Syst.* 41, 3 (2007), 411–430. DOI: <https://doi.org/10.1007/s00224-007-1328-0>
- [2] Takuya Akiba and Yoichi Iwata. 2016. Branch-and-reduce exponential/FPT algorithms in practice: A case study of vertex cover. *Theor. Comput. Sci.* 609, 1 (2016), 211–225. DOI: <https://doi.org/10.1016/j.tcs.2015.09.023>
- [3] Diogo V. Andrade, Mauricio G. C. Resende, and Renato F. Werneck. 2012. Fast local search for the maximum independent set problem. *J. Heur.* 18, 4 (2012), 525–547. DOI: <https://doi.org/10.1007/s10732-012-9196-4>
- [4] Ariful Azad, Aydin Buluç, and Alex Pothen. 2017. Computing maximum cardinality matchings in parallel on bipartite graphs via tree-grafting. *IEEE Trans. Parallel Distrib. Syst.* 28, 1 (2017), 44–59. DOI: <https://doi.org/10.1109/TPDS.2016.2546258>
- [5] Ariful Azad, Mahantesh Halappanavar, Sivasankaran Rajamanickam, Erik G. Boman, Arif Khan, and Alex Pothen. 2012. Multithreaded algorithms for maximum matching in bipartite graphs. In *Proceedings of the 26th International Parallel and Distributed Processing Symposium (IPDPS'12)*. IEEE, Los Alamitos, CA, 860–872. DOI: <https://doi.org/10.1109/IPDPS.2012.82>
- [6] David A. Bader, Henning Meyerhenke, Peter Sanders, Christian Schulz, Andrea Kappes, and Dorothea Wagner. 2014. Benchmarking for graph clustering and partitioning. In *Encyclopedia of Social Network Analysis and Mining*. Springer, 73–82. DOI: https://doi.org/10.1007/978-1-4614-6170-8_23
- [7] Mikhail Batsyn, Boris Goldengorin, Evgeny Maslov, and Panos M. Pardalos. 2014. Improvements to MCS algorithm for the maximum clique problem. *J. Comb. Optim.* 27, 2 (2014), 397–416. DOI: <https://doi.org/10.1007/s10878-012-9592-6>
- [8] Paolo Boldi, Bruno Codenotti, Massimo Santini, and Sebastiano Vigna. 2004. UbiCrawler: A scalable fully distributed web crawler. *Software Pract. Exper.* 34, 8 (2004), 711–726. DOI: <https://doi.org/10.1002/spe.587>
- [9] Sergiy Butenko, Panos Pardalos, Ivan Sergienko, Vladimir Shylo, and Petro Stetsyuk. 2009. Estimating the size of correcting codes using extremal graph problems. In *Optimization: Structure and Applications*, C. Pearce and E. Hunt (Eds.). Springer, 227–243. https://doi.org/10.1007/978-0-387-98096-6_12
- [10] Sergiy Butenko, Panos M. Pardalos, Ivan Sergienko, Vladimir Shylo, and Petro Stetsyuk. 2002. Finding maximum independent sets in graphs arising from coding theory. In *Proceedings of the 2002 ACM Symposium on Applied Computing (SAC'02)*. ACM, New York, NY, 542–546. DOI: <https://doi.org/10.1145/508791.508897>
- [11] Sergiy Butenko and Svyatoslav Trukhanov. 2007. Using critical sets to solve the maximum independent set problem. *Oper. Res. Lett.* 35, 4 (2007), 519–524. DOI: <https://doi.org/10.1016/j.orl.2006.07.004>
- [12] Lijun Chang, Wei Li, and Wenjie Zhang. 2017. Computing a near-maximum independent set in linear time by reducing-peeling. In *Proceedings of the 2017 ACM International Conference on Management of Data (SIGMOD'17)*. ACM, New York, NY, 1181–1196. DOI: <https://doi.org/10.1145/3035918.3035939>

- [13] Jianer Chen, Iyad A. Kanj, and Weijia Jia. 2001. Vertex cover: Further observations and further improvements. *J. Algorithms* 41, 2 (2001), 280–301. DOI : <https://doi.org/10.1006/jagm.2001.1186>
- [14] Alessio Conte, Donatella Firmani, Caterina Mordente, Maurizio Patrignani, and Riccardo Torlone. 2017. Fast enumeration of large k -plexes. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'17)*. ACM, New York, NY, 115–124. DOI : <https://doi.org/10.1145/3097983.3098031>
- [15] Jakob Dahlum, Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Renato F. Werneck. 2016. Accelerating local search for the maximum independent set problem. In *Experimental Algorithms*. Lecture Notes in Computer Science, Vol. 9685. Springer, 118–133. DOI : https://doi.org/10.1007/978-3-319-38851-9_9
- [16] Frank Dehne, Michael Fellows, Michael Langston, Frances Rosamond, and Kim Stevens. 2007. An $O(2^O(k)n^3)$ FPT algorithm for the undirected feedback vertex set problem. *Theory Comput. Syst.* 41, 3 (2007), 479–492. DOI : <https://doi.org/10.1007/s00224-007-1345-z>
- [17] Camil Demetrescu, Andrew V. Goldberg, and David S. Johnson. 2009. *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*. Vol. 74. AMS.
- [18] Michael Etscheid and Matthias Mnich. 2018. Linear kernels and linear-time algorithms for finding large cuts. *Algorithmica* 80, 9 (2018), 2574–2615. DOI : <https://doi.org/10.1007/s00453-017-0388-z>
- [19] Thomas A. Feo, Mauricio G. C. Resende, and Stuart H. Smith. 1994. A greedy randomized adaptive search procedure for maximum independent set. *Oper. Res.* 42, 5 (1994), 860–878. DOI : <https://doi.org/10.1287/opre.42.5.860>
- [20] Fedor V. Fomin and Dieter Kratsch. 2010. *Exact Exponential Algorithms*. Springer. DOI : <https://doi.org/10.1007/978-3-642-16533-7>
- [21] Jakub Gajarský, Petr Hliněný, Jan Obdržálek, Sebastian Ordyniak, Felix Reidl, Peter Rossmanith, Fernando Sánchez Villaaamil, and Somnath Sikdar. 2013. Kernelization using structural parameters on sparse graph classes. In *Algorithms*. Lecture Notes in Computer Science, Vol. 8125. Springer, 529–540. DOI : https://doi.org/10.1007/978-3-642-40450-4_45
- [22] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman.
- [23] Andreas Gemsa, Martin Nöllenburg, and Ignaz Rutter. 2014. Evaluation of labeling strategies for rotating maps. In *Experimental Algorithms*. Lecture Notes in Computer Science, Vol. 8504. Springer, 235–246. DOI : https://doi.org/10.1007/978-3-319-07959-2_20
- [24] Jiong Guo and Rolf Niedermeier. 2007. Invitation to data reduction and problem kernelization. *ACM SIGACT News* 38, 1 (2007), 31–45. DOI : <https://doi.org/10.1145/1233481.1233493>
- [25] Demian Hesse, Christian Schulz, and Darren Strash. 2018. Scalable kernelization for maximum independent sets. In *Proceedings of the 20th Workshop on Algorithm Engineering and Experiments (ALENEX'18)*. 223–237. DOI : <https://doi.org/10.1137/1.9781611975055.19>
- [26] Manuel Holtgrewe, Peter Sanders, and Christian Schulz. 2010. Engineering a scalable high quality graph partitioner. In *Proceedings of the 24th International Parallel and Distributed Processing Symposium*. 1–12. DOI : <https://doi.org/10.1109/IPDPS.2010.5470485>
- [27] Yoichi Iwata. 2016. Personal communication. August 24.
- [28] Yoichi Iwata, Keigo Oka, and Yuichi Yoshida. 2014. Linear-time FPT algorithms via network flow. In *Proceedings of the 25th ACM-SIAM Symposium on Discrete Algorithms (SODA'14)*. 1749–1761. <http://dl.acm.org/citation.cfm?id=2634074.2634201>
- [29] Bart M. P. Jansen and Stefan Kratsch. 2013. Data reduction for graph coloring problems. *Inform. Comput.* 231 (2013), 70–88. DOI : <https://doi.org/10.1016/j.ic.2013.08.005>
- [30] Richard M. Karp and Michael Sipser. 1981. Maximum matching in sparse random graphs. In *Proceedings of the 22nd Annual Symposium on Foundations of Computer Science (SFCS'81)*. IEEE, Los Alamitos, CA, 364–375. DOI : <https://doi.org/10.1109/SFCS.1981.21>
- [31] Tim Kieritz, Dennis Luxen, Peter Sanders, and Christian Vetter. 2010. Distributed time-dependent contraction hierarchies. In *Experimental Algorithms*. Lecture Notes in Computer Science, Vol. 6049. Springer, 83–93. DOI : https://doi.org/10.1007/978-3-642-13193-6_8
- [32] Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Renato F. Werneck. 2017. Finding near-optimal independent sets at scale. *J. Heuristics* 23, 4 (2017), 207–229. DOI : <https://doi.org/10.1007/s10732-017-9337-x>
- [33] Chu-Min Li, Zhiwen Fang, and Ke Xu. 2013. Combining MaxSAT reasoning and incremental upper bound for the maximum clique problem. In *Proceedings of the 25th International Conference on Tools with Artificial Intelligence (ICTAI'13)*. 939–946. DOI : <https://doi.org/10.1109/ICTAI.2013.143>
- [34] Henning Meyerhenke, Peter Sanders, and Christian Schulz. 2016. Partitioning (hierarchically clustered) complex networks via size-constrained graph clustering. *J. Heuristics* 22, 5 (2016), 759–782. DOI : <https://doi.org/10.1007/s10732-016-9315-8>
- [35] Henning Meyerhenke, Peter Sanders, and Christian Schulz. 2017. Parallel graph partitioning for complex networks. *IEEE Trans. Parallel Distrib. Syst.* 28, 9 (2017), 2625–2638. DOI : <https://doi.org/10.1109/TPDS.2017.2671868>

- [36] George L. Nemhauser and Leslie E. Trotter Jr. 1975. Vertex packings: Structural properties and algorithms. *Math. Program.* 8, 1 (1975), 232–248. DOI : <https://doi.org/10.1007/BF01580444>
- [37] Pablo San Segundo, Alvaro Lopez, and Panos M. Pardalos. 2016. A new exact maximum clique algorithm for large and massive sparse graphs. *Comput. Oper. Res.* 66 (2016), 81–94. DOI : <https://doi.org/10.1016/j.cor.2015.07.013>
- [38] Pedro V. Sander, Diego Nehab, Eden Chlamtac, and Hugues Hoppe. 2008. Efficient traversal of mesh edges using adjacency primitives. *ACM Trans. Graph.* 27, 5, Article 144 (2008), 9 pages. DOI : <https://doi.org/10.1145/1409060.1409097>
- [39] Peter Sanders and Christian Schulz. 2013. Think locally, act globally: Highly balanced graph partitioning. In *Experimental Algorithms*. Lecture Notes in Computer Science, Vol. 7933. Springer. DOI : https://doi.org/10.1007/978-3-642-38527-8_16
- [40] Pablo San Segundo, Fernando Matia, Diego Rodríguez-Losada, and Miguel Hernando. 2013. An improved bit parallel exact maximum clique algorithm. *Optim. Lett.* 7, 3 (2013), 467–479. DOI : <https://doi.org/10.1007/s11590-011-0431-y>
- [41] Pablo San Segundo, Diego Rodríguez-Losada, and Agustín Jiménez. 2011. An exact bit-parallel algorithm for the maximum clique problem. *Comput. Oper. Res.* 38, 2 (2011), 571–581. DOI : <https://doi.org/10.1016/j.cor.2010.07.019>
- [42] Darren Strash. 2016. On the power of simple reductions for the maximum independent set problem. In *Computing and Combinatorics*. Lecture Notes in Computer Science, Vol. 9797. 345–356. DOI : https://doi.org/10.1007/978-3-319-42634-1_28
- [43] Robert E. Tarjan and Anthony E. Trojanowski. 1977. Finding a maximum independent set. *SIAM J. Comput.* 6, 3 (1977), 537–546. DOI : <https://doi.org/10.1137/0206038>
- [44] Etsuji Tomita, Yoichi Sutani, Takanori Higashi, Shinya Takahashi, and Mitsuo Wakatsuki. 2010. A simple and faster branch-and-bound algorithm for finding a maximum clique. In *Algorithms and Computation*. Lecture Notes in Computer Science, Vol. 5942. Springer, 191–203. DOI : https://doi.org/10.1007/978-3-642-11440-3_18
- [45] Anurag Verma, Austin Buchanan, and Sergiy Butenko. 2015. Solving the maximum clique and vertex coloring problems on very large sparse networks. *INFORMS J. Comput.* 27, 1 (2015), 164–177. DOI : <https://doi.org/10.1287/ijoc.2014.0618>
- [46] Moritz von Looz, Mustafa S. Özdayı, Sören Laue, and Henning Meyerhenke. 2016. Generating massive complex networks with hyperbolic geometry faster in practice. In *Proceedings of the 2016 IEEE High Performance Extreme Computing Conference (HPEC'16)*. IEEE, Los Alamitos, CA, 1–6. DOI : <https://doi.org/10.1109/HPEC.2016.7761644>
- [47] Mingyu Xiao and Hiroshi Nagamochi. 2013. Confining sets and avoiding bottleneck cases: A simple maximum independent set algorithm in degree-3 graphs. *Theor. Comput. Sci.* 469 (2013), 92–104. DOI : <https://doi.org/10.1016/j.tcs.2012.09.022>
- [48] Mingyu Xiao and Hiroshi Nagamochi. 2017. Exact algorithms for maximum independent set. *Inform. Comput.* 255, 1 (2017), 126–146. <https://doi.org/10.1016/j.ic.2017.06.001>

Received July 2018; revised March 2019; accepted August 2019