# A METAHEURISTIC ALGORITHM FOR LARGE MAXIMUM WEIGHT INDEPENDENT SET PROBLEMS 

YUANYUAN DONG, ANDREW V. GOLDBERG, ALEXANDER NOE, NIKOS PAROTSIDIS, MAURICIO G.C. RESENDE, AND QUICO SPAEN


#### Abstract

Motivated by a real-world vehicle routing application, we consider the maximum-weight independent set problem: Given a node-weighted graph, find a set of independent (mutually nonadjacent) nodes whose node-weight sum is maximum. Some of the graphs airsing in this application are large, having hundreds of thousands of nodes and hundreds of millions of edges.

To solve instances of this size, we develop a new local search algorithm, which is a metaheuristic in the greedy randomized adaptive search (GRASP) framework. This algorithm, which we call METAMIS, uses a wider range of simple local search operations than previously described in the literature. We introduce data structures that make these operations efficient. A new variant of path-relinking is introduced to escape local optima and so is a new alternating augmenting-path local search move that improves algorithm performance.

We compare an implementation of our algorithm with a state-of-the-art openly available code on public benchmark sets, including some large instances with hundreds of millions of vertices. Our algorithm is, in general, competitive and outperforms this openly available code on large vehicle routing instances. We hope that our results will lead to even better MWIS algorithms.


## 1. Introduction

Given an undirected graph $G=(V, E)$, where $V$ is the set of nodes and $E$ the set of edges, an independent set $S \subseteq V$ is a set of mutually non-adjacent nodes of graph $G$. If each node $v \in V$ is assigned a weight $w_{v}$, a maximum-weight independent set (MWIS) of nodes $S^{*} \subseteq V$ is an independent set whose sum of weights,

$$
W\left(S^{*}\right)=\sum_{v \in S^{*}} w_{v}
$$

is maximum. We denote $n=|V|$ and $m=|E|$.
MWIS is a classical optimization problem that has been extensively studied and has many applications (Butenko, 2003). It is one of Karp's original NP-complete problems (Garey and Johnson, 1979; Karp, 1972). The problem is also hard to approximate (Håstad, 1999).

One can state MWIS as an Integer Linear Program - ILP (see Section 1.1). We can solve small MWIS problems exactly using IP solvers, e.g., CPLEX, GUROBI, or XPRESS, or a partial enumeration algorithm, such as the one proposed by Carraghan and Pardalos (1990). However, these methods do not scale to large graphs. Over the years, heuristics have been the workhorse for solving large instances of the

[^0]maximum independent set problem approximately (Pelillo, 2009). In particular, the most successful heuristics have been the ones based on metaheuristic algorithms, such as GRASP (Feo et al., 1994), tabu search (Friden et al., 1989), and iterated local search (Andrade et al., 2012; Nogueira et al., 2018).

In this paper we introduce $M E T A M I S$, a new metaheuristic algorithm for the MWIS problem. METAMIS is based on the greedy randomized adaptive search procedure - GRASP (Resende and Ribeiro, 2016), with truncated path-relinking. Our motivation is a long-haul vehicle routing (VR) application that yields large MWIS problems, some with close to 900 thousand nodes. Compared to benchmark instances used in previously published work, the VR-MWIS instances are often larger and have a very different structure. We conduct experiments with METAMIS on MWIS instances arising in different applications, including on our VR-MWIS instances and on other publicly available ones.

The paper is organized as follows. In Section 2 we give a high-level description of the algorithm. Section 3 introduces the data structure and gives low-level implementation details. We present experimental results in Section 4 and make concluding remarks in Section 5.
1.1. ILP Formulation. Next we discuss several ILP formulations of MWIS. Let $x_{v}$ be a binary decision variable such that $x_{v}=1$ if node $v \in S \subseteq V$ and $x_{v}=0$ otherwise, where $S$ is an independent set of nodes. A simple integer programming (IP) formulation for selecting a maximum-weight independent set of nodes is

$$
\max \sum_{v \in V} w_{v} x_{v}
$$

subject to

$$
\begin{aligned}
& x_{u}+x_{v} \leq 1, \forall(u, v) \in E \\
& x_{v} \in\{0,1\}, \forall v \in V
\end{aligned}
$$

The objective is to maximize the sum of the weights of the nodes selected to be in the independent set $S$. The constraints guarantee that all selected nodes are mutually non-adjacent, i.e. for all edges in the graph, at most one node can be in the independent set. Stronger formulations add clique inequalities (Padberg, 1973) to the above formulation. For a 2-clique, or clique of size 2, denoted by $C_{2}$ we have $x_{u}+x_{v} \leq 1$, for all $(u, v) \in E$. A 3-clique $\left(C_{3}\right)$ inequality is $x_{u}+x_{v}+x_{t} \leq 1$, for all triangles, i.e. $u, v, t \in V$, such that $(u, v) \in E,(u, t) \in E$, and $(v, t) \in E$. In general, for any clique $Q$, we have a constraint

$$
\sum_{v \in Q} x_{v} \leq 1
$$

Let $C_{2}, C_{3}, \ldots, C_{k}$ be, respectively, the sets of 2 -clique, 3 -clique, $\ldots$, and $k$-clique inequalities. The clique IP formulation for maximum weight independent set is:

$$
\max \sum_{v \in V} w_{v} x_{v}
$$

subject to

$$
\begin{aligned}
& C_{2}, C_{3}, \ldots, C_{k} \\
& x_{v} \in\{0,1\}, \forall v \in V
\end{aligned}
$$

Other tight formulations for independent set can be found in Buchanan and Butenko (2014).

In the linear programming $(L P)$ relaxation of the problem, we allow fractional solutions: $x_{v} \in[0,1]$. The relaxed problem is much simpler and can be solved by an LP solver in reasonable time. We use an optimal solution to the LP relaxation to improve the performance of our local search algorithm.

The fact that for some problems an LP relaxation can help to find a good feasible solution is well-known. For example, for a class of IP problems, randomized rounding of the optimal relaxed solution yields a solution with a provably good approximation ratio (Raghavan and Tompson, 1987). However, for MWIS, a rounded solution may not be feasible. Our use of the relaxed solutions can be viewed as another way of extracting information about good integral solutions from the relaxed solution

Note that as long as all 2-clique inequalities are present, for any set of additional clique inequalities, the set of feasible integer solutions to the corresponding IP is the same as for the simple formulation. Additional clique constraints make the linear programming (LP) relaxation stronger.

Although our algorithm is a general-purpose heuristic, our motivation comes from vehicle routing (VR) (Dong et al., 2021a;b). A variant of our algorithm takes advantage of the application-specific structure for performance and solution quality. In this application, we have a good initial solution. One can use this solution to warm-start our MWIS algorithm. Our experiments show how much the warm-start improves solution quality.

## 2. High Level Description

The MWIS algorithm is an iterative local search algorithm based on the Greedy Randomized Adaptive Search Procedure (GRASP) metaheuristic, which is a general metaheuristic for combinatorial optimization (Feo and Resende, 1989; 1995; Resende and Ribeiro, 2016). The algorithm also uses path relinking to escape local optima (Laguna and Martí, 1999; Resende and Ribeiro, 2016).

Figure 1 gives a high-level view of the algorithm. In addition to the graph, the input to the algorithm includes a stopping criterion, e.g., a time limit, and an initial solution. In our application we have a good initial solution. When no such solution is available, one can find a solution using the randomized greedy algorithm described later in this section. The algorithm applies local search to improve the initial solution and enters the main loop. At termination of the local search procedure, we are at a local optimum.

The algorithm maintains a set of elite solutions $\mathcal{E S}$, which are the best solutions we have seen so far. We add a solution to $\mathcal{E S}$ immediately after a local search, so the elite solutions are always locally optimal. At each iteration of the loop, we first attempt to escape the local optimum corresponding to the elite solution. In the process, we can decrease the objective function. To escape a local optimum, we first find a randomized greedy solution $S_{G}$. Optionally, we apply local search to improve $S_{G}$. Then we apply path relinking to $S_{G}$ and a random elite solution from $\mathcal{E S}$ to find a new solution $S^{\prime}$. Then we apply local search to improve $S^{\prime}$, and update $S^{*}$ if we find a better solution.

In our experiments, we omit the optional call to local search immediately after path relinking because this variant of the algorithm seems to work better for the

```
Algorithm 1 Algorithm Overview
    procedure \(\operatorname{MWIS}\left(G=(V, E, w), \operatorname{maxTime}, S_{0}\right)\)
        \(S \leftarrow \operatorname{localSearch}\left(G, S_{0}\right)\)
        \(\mathcal{E S} \leftarrow\} \quad \triangleright\) Empty set of elite solutions
        \(\mathcal{E S} . \operatorname{add}(S)\)
        while \(t \leq\) maxTime do
            \(S_{G} \leftarrow\) findRandomizedGreedySolution(G)
            if LsBeforeRelinking then \(\triangleright\) Optional local search
                    \(S_{G} \leftarrow \operatorname{localSearch}\left(G, S_{G}\right)\)
            end if
            \(S_{e} \leftarrow \mathcal{E S}\).randomEliteSolution()
            \(S^{\prime} \leftarrow \operatorname{pathRelinking}\left(G, S_{G}, S_{e}\right)\)
            \(S^{\prime} \leftarrow \operatorname{localSearch}\left(G, S^{\prime}\right)\)
            \(\mathcal{E S}\).tryToAddAndEvict \(\left(S^{\prime}\right) \quad \triangleright\) Add solution to elite set, if full evict
    similar solution of lesser value (or don't insert if no worse elite solution exists)
        end while
        return \(\mathcal{E}\) S.bestSolution()
    end procedure
```

VR-MWIS instances. We also set the size of the elite set $\mathcal{E S}$ to 1 , so we only retain the best solution. This setting works best for the VR-MWIS instances. For other problem families, different parameter choices were found to work better (Kummer, 2020; Kummer et al., 2020).
2.1. Greedy Algorithm. We construct a greedy solution $S$ as follows. To initialize $S$, we add to it all zero-degree nodes and delete these nodes from the graph. Next, we define $\eta(v)=w(v) / \operatorname{degree}(v)$ and create a list of the nodes sorted by $\eta(v)$. Initially the list contains all nodes. At every step, we remove node $v$ with the largest $\eta(v)$ from the list, add $v$ to $S$, and remove neighbors of $v$ from the list. We terminate when the list is empty.

To better explore the solution space, we need diverse solutions, so we randomize greedy algorithm as follows. Instead of removing the node with the largest $\eta(v)$ value, we remove a random node from the set of $k$ nodes with the largest values. We set $k$ to be a fraction of $n$, e.g., $10 \%$.

The intuition behind the greedy algorithm is that a node with a large $\eta(v)$ value is more likely to be in an independent set of high weight. However, as we add nodes to $S$ and delete these nodes and their neighbors from the graph, residual node degrees change. This motivates the adaptive greedy algorithm.

The adaptive algorithm maintains node degrees in the graph with nodes in $S$ and their neighbors deleted. We maintain a priority queue of nodes in the graph. If we delete one or more neighbor of a node $v$, do the following. If degree of $v$ becomes zero, we add $v$ to $S$ and delete it from the graph. Otherwise, $\eta(v)$ increases, and we use increase-key operation to update the priority queue. Assuming that the increase-key operation takes constant time and other priority queue operations take logarithmic time (e.g., we use Fibonacci heaps (Fredman and Tarjan, 1986)), the complexity of the greedy algorithm is $O(m+n \log n)$.

In practice, for large graphs the adaptive algorithm is expensive. To mitigate the problem, we use it in the first local search (line 2 of Algorithm 1), which results
in a substantial improvement over using the deterministic algorithm. However we use the latter algorithm in the main loop and rely on path relinking to improve solution quality. We do not randomize the adaptive greedy algorithm as we only use it once, in the first local search.

```
Algorithm 2 Local Search Procedure
    procedure LocalSearch \((G=(V, E, w), S\), numIterations)
        \(i \leftarrow 1\)
        \(S^{*} \leftarrow S\)
        while \(i \leq\) numIterations do
            \(S_{i} \leftarrow\{ \} \quad \triangleright\) Empty solution
            while \(w\left(S_{i}\right)<w(S)\) do \(\quad \triangleright\) Repeat until no improvement is found
                    \(S_{i} \leftarrow S\)
                    \(S \leftarrow \operatorname{starOneMoves}(G, S)\)
                \(S \leftarrow \operatorname{AAPMoves}(G, S)\)
                \(S \leftarrow\) oneStarMoves \((G, S)\)
                if \(w\left(S_{i}\right)<w(S)\) then return \(\triangleright\) Solution improved
                end if
                \(S \leftarrow\) twoStarMoves \((G, S)\)
            end while
            if \(w(S)>w\left(S^{*}\right)\) then
                    \(S^{*} \leftarrow S\)
                \(i \leftarrow 1\)
            else
                    \(S \leftarrow \operatorname{perturb}(S)\)
            end if
        end while
    end procedure
```

2.2. Local Search. The local search procedure, outlined in Figure 2, repeatedly performs local moves with positive gain. We aim to find positive gain (improving) moves until we reach a local optimum, and then we perform a random perturbation of the solution. If we find an improving move, we apply it immediately. We use a subset of $(x, y)$ moves and alternating augmenting path moves (AAP-moves). While the $(x, y)$ moves have been studied previously, the AAP moves are new. We describe the moves at a high level in this section, and give a detailed description in Section 3.

An $(x, y)$ move removes $x$ nodes from the solution and adds $y$ nodes to it while maintaining solution independence. We use $*$ instead of $x$ or $y$ to denote an arbitrary positive integer. Note that the number of applicable moves increases significantly as $x$ and $y$ increase. Previous algorithms used $(x, y)$ moves for small values of $x$ and $y$. In particular, the algorithm of Nogueira et al. $(2018)$ uses $(*, 1)$ and $(1, *)$ moves. Our algorithm uses $(*, 1),(1, *)$, and $(2, *)$ moves. The number of $(2, *)$ moves is large. We use data structures and operation ordering that make improving moves more likely, which makes our algorithm more efficient.

If an $(x, y)$ move renders $S$ non-maximal, we add nodes without a neighbor in $S$ to the independent set in random order until $S$ becomes maximal. Note that through this update sequence, $S$ remains an independent set.

A $(*, 1)$ move is the simplest. It inserts a single node $u$ into the current solution $S$ and removes its neighbors from $S$. Procedure starOneMoves $(G, S)$ applies the $(*, 1)$ moves while no such improving move.

A $(1, *)$ move removes a node $v$ from $S$ and adds to $S$ an independent subset $I$ of the nodes whose only neighbor in $S$ before the removal is $v$. Usually one has multiple choices of independent sets to add. A good heuristic is to add a maximum weight set of the neighbors that maintains independence. This is done when the number of neighbors is small (at most seven in our experiments). We use a naive recursive algorithm: Pick a node $u$ in the neighborhood and recursively solve two subproblems. The first subproblem results by adding $u$ to $S$ and deleting its neighbors from the graph. We get the second subproblem by deleting $u$ without adding it to $S$. The better of the two corresponding solutions is returned. When the neighborhood is big, a variant of the greedy algorithm of Section 2.1 is used. In this case, however, we found that it is better to process nodes in descending order of their weight $w(v)$ and not of $\eta(v)=w(v) /$ degree $(v)$. Procedure oneStarMoves $(G, S)$ applies the $(1, *)$ moves while no such improving move.

A $(2, *)$ move removes two nodes, $u$ and $v$, from $S$ and adds to $S$ an independent subset $I$ of the nodes whose only neighbors in $S$ before the removal is $u$, or $v$, or both $u$ and $v$. Generally, this set is significantly larger than the corresponding set for the $(1, *)$ moves, and the recursive operation used for the $(1, *)$ moves is too expensive. One could use greedy addition, but in our experiments a random addition, that adds to $S$ a random node from $I$ that has no neighbors in $S$, was better. We also tried biased random addition that picks a node with probability proportional to its weight, but uniformly random selection was better. Procedure twoStarMoves $(G, S)$ applies the $(2, *)$ moves until it finds an improving move or there is no improving $(2, *)$ move. Note that unlike the corresponding procedures for other moves, twoStarMoves exits as soon as it finds an improving move.

Our idea for AAP moves comes from matching algorithms (Edmonds, 1965), although we use a somewhat different definition. Given an independent set $S$, we define an AAP $P$ as follows. Let $I=S \cap P$ and $O=P-S$ be nodes of $P$ that are in and out of $S$, respectively.
(1) if $v \in I$, then the neighbors of $v$ on $P$ are in $O$,
(2) if $v \in O$, then the neighbors of $v$ on $P$ are in $I$,
(3) if we flip the path, i.e., set $S=S-I+O, S$ remains an independent set.

An AAP move finds an alternating augmenting path, flips it, and looks at the change in $w(S)$. If the change is positive, we accept the AAP move; otherwise we reject the move. For efficiency, we apply a limited number of AAP moves. We postpone the details of finding the AAP moves until Section 3. Procedure AAPMoves $(G, S)$ applies the AAP moves while no such improving move.

During an execution of the algorithm, most local search moves do not improve solution quality and thus do not change the solution. Note that complexity of evaluating $(2, *)$ moves is significantly higher than those for the other moves. Our local search repeatedly applies starOneMoves, AAPMoves, and oneStarMoves procedures while these procedures find improving moves. If we find an improving move, an immediate application of these procedures may find additional improvements due to neighborhood changes, so we iterate. Only when these procedures fail to find improving moves we call twoStarMoves. If twoStarMoves fails to improve the solution, we perform a random perturbation.

The perturbation adds a small set of random nodes to $S$ and removes their neighbors. After perturbing, we resume local search. The local search algorithm terminates if there has been no improvement to the best solution after a predefined number of iterations.
2.3. Using the Relaxed LP Solution. In our VR application, we use clique information and get a relaxed LP solution to the relaxed problem (1). The solution assigns a value $x_{v} \in[0,1]$ to each node $v$. We use these values to bias random node selection in the perturbation step of the local search. When performing a random perturbation in Algorithm 2, we add a node $v$ to the solution with probability proportional to $x_{v}+\epsilon$. Here $\epsilon$ is a positive value (set to $\epsilon=0.005$ ) that ensures that each node can be picked, even if $x_{v}=0$. This guides the local search by biasing route selection toward nodes with higher fractional relaxed solution value. The idea is that a node with a high value in the relaxed solution is more likely to be part of a good solution for the MWIS. Using prefix sums we can pick a random node in time $\mathcal{O}(\log |V|)$ : We draw a random floating-point number $z \in\left[0, \sum_{v \in V} x_{v}\right)$ and use binary search on the prefix sum array to pick a node such that the sum up to but excluding the node is less than $z$, and the sum up to and including the node is greater or equal to $z$.
2.4. Adaptive Path Relinking. Path relinking is a technique for escaping local optima. We discuss this technique in the context of MWIS, assuming that the local search moves are symmetric: the reverse of a move is a valid move. Define an undirected graph associated with the search space MWIS, where the nodes correspond to feasible solutions and the edges correspond to local search moves that transform the solution corresponding to the tail of the edge to the solution corresponding to the head. A path in this graph corresponds to a sequence of the moves that transform the solution at one end of the path into a solution at the other end. Note that the moves need not improve the objective function value. The underlying assumption of path relinking is that if the end-points of a path correspond to high quality solutions, then the path will contain previously undiscovered high-quality solutions.

For our local search, given two solutions $S$ and $T$, we can transform $S$ into $T$ as follows. Initialize $S^{\prime}=S$. At every step, we do either a $(*, 1)$ move or a $(1, *)$ move. In the former case, pick a node $v \in T-S^{\prime}$, add $v$ to $S^{\prime}$, and remove neighbors of $v$ from $S^{\prime}$. In the latter case, pick a node $v \in S^{\prime}, v \notin T$ and remove $v$ from $S^{\prime}$. Let $N(v)$ denote the set of neighbors of $v$. Then we iterate over nodes $u$ in $N(v) \cap T$. If $N(u) \cap S^{\prime}=\emptyset$, we add $u$ to $S^{\prime}$.

For large graphs, finding good solutions is expensive. Instead of combining two good solutions, we apply path relinking to combine the randomized greedy solution $S_{G}$ with the current best solution $S^{*}$, which is locally optimal. While $S^{*}$ is a good solution, $S_{G}$ may not be good, and the solutions on the path far from $S^{*}$ are usually not good either. We modify path relinking so that it examines only a prefix of the path close to $S^{*}$. The prefix is small enough so that the solution quality remains good, yet big enough so that the subsequent local search will not end up with a locally optimal solution equivalent to $S^{*}$. This an adaptive variant of the truncated greedy path relinking described in Resende et al. (2010).

The first modification is to choose the node $x$ to add to $S$ or to remove from $S$ greedily. We pick a node that maximizes the weight of the solution we get after
the move. A second modification is to do a truncated path relinking: we stop the process after a certain number of steps, which we adjust adaptively. We start with a small limit on the number of steps and increase the limit if the algorithm gets stuck in a local optimum of weight $w\left(S^{*}\right)$.

More precisely, we stop if the weight factor $f=w(S) / w\left(S^{*}\right)$ is smaller than a predefined limit, or we performed more than $c_{n}$ negative gain changes, or made more than $c_{p}<c_{n}$ positive gain moves. Positive gain moves help us escape local optima, as they will not be undone by the $(1, *)$ moves of the subsequent local search. We set the initial parameter values as follows: $f=f_{0}=0.9998, c_{n}=c_{n 0}=1.0$ and $c_{p}=c_{p 0}=0.1$. If $w(S)=w\left(S^{*}\right)$ after local search on $S$, we multiply $f^{\prime}=f \cdot 0.9998$, $c_{n}^{\prime}=c_{n} \cdot 1.5$ and $c_{p}=c_{p} \cdot 1.5$. If $w(S) \neq w\left(S^{*}\right)$, we reset the limits to their initial values. If the local search algorithm finds a solution $S$ with $w(S)=w\left(S^{*}\right)$, we check whether the solutions are equivalent and terminate local search if they are. Two solutions are equivalent if they are the same or if we can transform one into another by zero-gain $(1, *)$ and $(*, 1)$.

## 3. Data Structures and Optimizations

As we are interested in solving problems on graphs with hundreds of millions of edges, the choice of data structures is important for the efficiency of the algorithm. When making trade-offs between performance on sparse and dense graphs we favor the former because our motivating application yields relatively sparse graphs.

Several of our data structures use sets of objects. We use a representation of sets based on hashing. This representation allows constant time addition, deletion, and membership query, and linear time iteration over all set elements. We also assume that if we add an element to the set that already contains the element, the set does not change. Similarly, if we delete an element not in the set, the set does not change.
3.1. Input Graph. The input graph is static: it does not change throughout the execution. We assign to the nodes of the graph integer IDs from $[0, \ldots, n-1]$ and place them in an array, with node $i$ in position $i$. Each node has an array of edges incident to it. This places the edges incident to a node in contiguous memory locations, assuring that a common operation of scanning an edge list has a good memory locality. We sort edges by IDs of the head node. This allows us to do neighborhood queries (e.g., "Is $v$ in $N(u)$ ?") in time logarithmic in the degree of $u$ using binary search.

Note that using sets to represent neighborhoods would give constant neighborhood queries and linear time edge list scan. However, the constant factors, both in terms of running time and memory consumption, associated with hashing are large. In addition, we lose the locality in edge list scans. For graphs arising from our motivating application, the array-based implementation is significantly faster than the one based on sets.
3.2. Interstate Graph. The interstate graph makes the local search operations more efficient. To describe this graph, we need a few definitions.

For a node $u \in S,(u, v) \in E$, we say that $v$ is a 1-tight neighbor of $u$ if $N(v) \cap S=$ $\{u\}$ (Andrade et al., 2012). Note that if we remove $u$ from $S$, we can add to $S$ any 1-tight neighbor of $u$.

Two nodes $u, v \in S$ are mates if for at least one node $w \notin S, w$ has exactly two neighbors in $S: N(w) \cap S=\{u, v\}$. We call the node $w$ a 2-tight neighbor of $u$ and $v$. We say that $w$ is a 2-tight neighbor of $u$ if $u$ has a mate $v$ such that $w$ is a 2-tight neighbor of $u$ and $v$. If we delete $u$ and $v$ from $S$, we can replace them by an independent set of the union of three sets: the set of the 1-tight neighbors of $u$, the set of 1-tight neighbors of $v$, and the set of the shared 2-tight neighbors of $u$ and $v$.

Our main data structure is the interstate graph $G_{I S}=\left(V, E_{I S}, w\right)$. For $G_{I S}$, the nodes and node weights are the same as in the input graph $G$. The edge set $E_{I S}$ is changed dynamically depending on the nodes in the current independent set $S$. $E_{I S}$ has three types of edges:
(1) $e=(u, v) \in E$, where $u \in S$ and $v$ is a 1-tight neighbor of $u$;
(2) $e=(u, w) \in E$, where $u \in S$ and $w$ is a 2-tight neighbor of $u$;
(3) $e=(u, v)$, where $u, v \in S$ are mates.

We represent the three edge types separately.
(1) For every $u \in S$, we represent its 1-tight neighbors as sets. For $v \notin S$ that is a 1-tight neighbor of $u$ we add the 1 -tight edge $(v, u)$.
(2) For every pair of mates $u$ and $v$, we maintain a set of 2-tight neighbors of $u$ and $v$. For every 2-tight neighbor $w \notin S$, we add the pair of 2 -tight edges $(w, u)$ and $(w, v)$.
(3) For every node $v$ in $S$, we maintain a set $M_{v}$ of its mates. Every mate $w \in M_{v}$ corresponds to a mate edge $(v, w)$.
3.3. Efficient Implementation of $(x, y)$ Moves. In this section we show how to efficiently implement $(x, y)$ moves using the interstate graph and two additional optimizations, one for the $(1, *)$ moves and another for $(*, y)$ moves. We discuss maintenance of the interstate graph in Section 3.2.

To implement $(*, 1)$ operations efficiently, we use an idea from Nogueira et al. (2018). For every $u \notin S$, we maintain a value

$$
\Delta(u)=w(u)-\sum_{v \in S \cap N(u)} w(v)
$$

to speed up the $(*, 1)$ moves. Such a move is an improving move when $\Delta(u)>$ 0 . We keep a set $S^{+}$of the nodes $u$ with $\Delta(u)>0$. Note that for an efficient implementation of $(*, 1)$ moves, we need to update the vector $\Delta(\cdot)$ and the set $S^{+}$. We do this as follows. Every time we add a node $u$ to $S$, we remove $u$ from $S^{+}$. Then for each $v \in N(u), v \notin S$, we set $\Delta(v)=\Delta(v)-w(u)$. Every time we remove $u$ from $S$, we scan the edge list of $u$ and compute $\Delta(u)$. If $\Delta(u)>0$, we add $u$ to $S^{+}$. Also during the scan, for every neighbor $v$ of $u$ such that $v \notin S$, we increase $\Delta(v)$ by $w(u)$, and if $\Delta(v)$ becomes positive, we add $v$ to $S^{+}$. We have an improving $(*, 1)$ move if and only if $S^{+}$is non-empty. In this case, we can pick a node $u$ from $S^{+}$and apply the $(*, 1)$ move to it.

Since for every $u \in S$ we maintain a set of its 1-tight neighbors as a hash set, we can efficiently run the recursive or the greedy algorithm described in Section 2 on this set. Similarly, since for every $u \in S$ we maintain the set of its mates, we can iterate over all mates of $u$. Furthermore, for a pair of mates $u$ and $v$, we have the set of the common 2-tight neighbors, and we can apply the randomized algorithm to this set.

Next we describe an optimization that prunes $(1, *)$ and $(2, *)$ moves that are unlikely to improve the solution. For the $(1, *)$ move that removes $u$, we evaluate the move only if the 1-tight neighborhood of $u$ changed since the last time we evaluated the move but failed to improve the solution. We say that the neighborhood changed if we add $u$ to $S$ and $u$ has a non-trivial 1-tight neighborhood. Since our implementation of the $(1, *)$ move is deterministic and depends only on the 1-tight neighborhood, we know that the move will fail. We maintain the set $S_{1}$ of nodes $u \in S$ whose 1-tight neighborhood changed but is not empty. We pick nodes for $(1, *)$ moves from $S_{1}$. While initializing $G_{I N}$, we initialize $S_{1}$ to include all nodes with non-trivial 1-tight neighborhoods. When we update $G_{I N}$, we also update $S_{1}$ (see Section 3.5).

For the $(2, *)$ move, we maintain a set $S_{2}$ of mate pairs $\{u, v\}$ which are eligible for the move. We delete a pair from $S_{2}$ and evaluate the move that removes this pair from $S$. We add a pair $\{u, v\}$ to $S_{2}$ when they become 2-tight mates, or when $\{u, v\}$ are 2 -tight mates and their 2 -tight neighborhood changes, or when they are 2-tight mates and the 1-tight neighborhood of either $u$ or $v$ changes. Our implementation of the $(2, *)$ move depends only on the 2-tight neighborhood of the mates. However, the implementation is randomized. Although it is possible that one evaluation of the move succeeds and another fails when the 2-tight neighborhood stays the same, we assume this is unlikely and prune the move. We maintain the set $S_{2}$ of mates whose 2-tight neighborhood changed. We pick mates for $(2, *)$ moves from $S_{2}$. While initializing $G_{I N}$, we initialize $S_{2}$ to all pairs of mates. When we update $G_{I N}$, we update $S_{2}$ as well.
3.4. AAP Moves. For efficiency, we only look for alternating augmenting paths (AAPs) in the interstate graph. The only edges on any AAP are either edges from members of $S$ to their 1-tight and 2-tight neighbors (as edges between 2-tight mates would not yield an alternating path). To limit the number of AAP move evaluations, we start a search for an AAP from a 1-tight neighbor of $v \in S_{1}$ ( $S_{1}$ was introduced in Section 3.3). This way we guarantee that the move will not decrease the cardinality of $S$, making the move more likely to succeed. The alternating path initially contains $v$ and its single neighbor $u \in S$. We grow the path as follows. Let $u \in S$ be the last node on the current AAP, and let $U$ be the set of nodes on the AAP that are in $S$ and $\bar{U}$ be the set of nodes on the AAP that are not in $S$. We pick a mate $w$ and a 2-tight neighbor $x$ of $u$ such that

- $x$ is not a neighbor of any node of $\bar{U}$ in the input graph (so that the extended path will be an AAP),
- neither $x$ nor $w$ are already in AAP,
- the gain of flipping the extended path is maximized.

If we succeed in finding such a $\{w, x\}$ pair, we add $w$ and $x$ to the path. Then we redefine $u$ to be $x$ and continue growing the path. To introduce additional randomness, we increase the gain for every $\{w, x\}$ pair by a random real number $\epsilon \in[-\delta, \delta]$ and maximize the perturbed gains. We use $\delta=50$ in our experiments. We terminate the search if the length of the path exceeds a threshold or the gain of flipping the path falls below a (negative) threshold. We then perform the highest positive gain move that flips a prefix of the final path. If no positive gain move is encountered, we do nothing (the move fails).
3.5. Maintaining the Interstate Graph. The vast majority of the local search moves we evaluate do not improve the solution and $G_{I N}$ does not change. We need to update the graph only when a move succeeds, which happens rarely. Our data structures speed up move evaluations and support move pruning. The added overhead is in data structure initialization and updates. The update complexity is non-trivial, but for sparse graphs the complexity is much smaller than the time we save due to the improved move efficiency and pruning.

Let $\rho(u)=|N(u) \cap S|$ denote the number of the neighbors of $u$ in $S$. Note that for nodes $u \in S, \rho(u)=0$. We maintain $\rho(u)$ for all nodes $u \in V$.

Given an initial solution $S$, we build $G_{I N}, S_{1}$, and $S_{2}$ as follows. We process all nodes $u \notin S$. For each $u$, we scan its edge list in $G$ and initialize $\rho(u)$. If $\rho(u)=1$, we let $N(u) \cap S=\{v\}$, add the 1-tight edge $(u, v)$ to the edge list of $u$ in $G_{I N}$, and add $u$ to the set of 1-tight neighbors of $v$. If $\rho(u)=2$, we let $N(u) \cap S=\{v, w\}$, add $v$ to the set of mates of $w$ and add $w$ to the set of mates of $v$. We also add the pair of 2-tight edges $(u, v)$ and $(u, w)$ to $G_{I N}$. Finally, we add $u$ to the set of 2-tight neighbors of the mates $\{v, w\}$. We initialize $S_{1}$ to the set of all nodes $u \in S$ with non-empty set of 1-tight neighbors. We initialize $S_{2}$ to the set of all mate pairs $\{u, v\}$. The initialization takes linear time.

Our algorithm updates $S$ by removing a set of nodes $S^{-}$and adding a set of nodes $S^{+}$. We break the update into a sequence of single-node updates: first we remove nodes of $S^{-}$one by one, then we add nodes of $S^{+}$one by one. We update $G_{I N}$ after each individual update of $S$.

After removing a node $u$ from $S$, we empty its set of 1-tight neighbors and remove $u$ from $S_{1}$. For each mate $v$ of $u$, we set the corresponding set of 2 -tight neighbors to empty and remove $u$ from the set of mates of $v$. We also remove the pair $\{u, v\}$ from $S_{2}$. Afterwards, we empty the set of mates of $u$. We then visit its neighbors $v \in V \backslash S$. For each neighbor $v$, we decrement $\rho(v)$. We need to update $G_{I N}$ if $\rho(v)$ becomes zero, one, or two.

Cases for zero and two are simpler. If the value is zero, we set the 1-tight neighbor of $v$ to null. If the value is two, let $N(v) \cap S=\{a, b\}$. We can find $a$ and $b$ by scanning the edge list of $v$ in $G$. We add $a$ to the set of mates of $b$ and vice versa. We also add $v$ to the set of 2 -tight neighbors of $\{a, b\}$. Finally, we add the 2-tight pair of edges $(v, a)$ and $(v, b)$ to $G_{I N}$.

If the value is one, we have to update both the old 2-tight neighborhood and the new 1-tight neighborhood. For the latter, we set the 1-tight neighbor of $v$ to the unique neighbor $w \in S$, and add $v$ to the 1-tight neighbor set of $v$. For the former update, note that $v$ was a 2-tight neighbor for mates $\{v, w\}$ for some $w \in S$ before the removal of $v$. We remove $v$ from the set of 2-tight neighbors of $w$ and delete the 2-tight edge pair $(v, u)$ and $(v, w)$ from $G_{I N}$.

Now consider the addition of a node $u$ to $S$ that maintains the independence of $S$. We scan the edge list of $u$ and for all neighbors $v$ (guaranteed not to be in $S$ ) and increment $\rho(v)$. We need to update $G_{I N}$ if $\rho(v)$ becomes one, two, or three.

Cases for one and three are simpler. If the value is one, we add the 1-tight edge $(v, u)$ to $G_{I N}$, add $v$ to the set of 1-tight neighbors of $u$, and add $u$ to $S_{1}$. If the value is three, $v$ has a pair of 2-tight edges $(v, a)$ and $(v, b)$, where $a$ and $b$ are mates. We delete $(v, a)$ and $(v, b)$ from $G_{I N}$. Then we remove $v$ from the set of 2 -tight neighbors of $a$ and $b$. If the set becomes empty, $a$ and $b$ are no longer mates,
so we remove $a$ from the set of mates of $b$, remove $b$ from the list of mates of $a$, and remove $\{a, b\}$ from $S_{2}$.

If the value is two, we have to update both the old 1-tight neighborhood and the new 2-tight neighborhood. For the former, let $(v, w)$ be the 1-tight edge. We remove the edge and remove $v$ from the set of 1-tight neighbors of $w$. If the set becomes empty, we remove $w$ from $S_{1}$. In the latter case, $N(v) \cap S=\{v, w\}$ for some $w \in S$. We add $u$ to the set of mates of $w$ and vice versa. We also add $v$ to the set of 2-tight neighbors of $v$ and $w$. Finally, we add $\{a, b\}$ to $S_{2}$.

Note that since when we add or remove $u$ to or from $S$, we may need to scan edge lists of multiple neighbors of $u$, updating $G_{I N}$ when $G$ is dense may be expensive.

## 4. Experimental results

4.1. Algorithms and Computational Environment. We implemented our algorithm, which we call METAMIS, in Java because it is used in a production system at Amazon and Java is a requirement. For the same reason, we use doubles for node weights. Furthermore, due to licensing restrictions, we use only standard Java libraries. We compiled our code using Java 8. Our choice of programming language probably affected performance. A C++ implementation, for example, is expected to be faster (Alnaser et al., 2012).

Although one can tune our algorithm for specific problem families, we use fixed parameter settings in all experiments.

We compare our implementation to the $I L S V N D$ algorithm of Nogueira et al. (2018). The publicly available code of Nogueira et al. (2018) is implemented in C++ and represents weights using integers. We made one modification to ILSVND: added the ability to warm start from an initial solution. Given a solution in the input, we initialize the current solution of ILSVND to the input solution. We compiled ILSVND using full optimization (-03) on an AWS r3.4xlarge instance (AWS, 2021).

For a given instance, algorithm time-quality plots give a lot of information about relative performance of the algorithms. For example, one algorithm may dominate another, or one can converge to a better solution but take longer to converge, etc. The algorithms we compare are stochastic and algorithm performance depends on the pseudo-random seed we use. Furthermore, the algorithms we compare do not know if and when they reach an optimal solution. Usually there is a chance that a solution may improve. However, the algorithms converge in a sense that it may reach a point of diminishing returns when a substantial improvement becomes unlikely. To compare the two algorithms, we put a time limit $T$ on their executions. For different problem families, the limit may be different. We run each instance with five different pseudo-random seeds and report the best solution value the algorithm finds. In many cases the algorithms converge. However, for harder problems this may take too long, and the algorithms do not converge within the time limit.

For representative instances, we give the time-quality plots, but we have too many instances to give all the plots. Therefore, we report solution quality at times $T / 10$ and $T / 2$. In addition, we report the time $t^{*}$ defined as follows. For a given problem instance, consider the set of final solution values over all algorithms and seed values. Let $s$ be the smallest one of these values. For a given algorithm, consider the run producing the best final solution value. For this algorithm, we define $t^{*}$ to be the earliest time this run reaches the value of $s$ or higher, Intuitively,

Table 1. VR instances

| Graph | $\|V\|$ | $\|E\|$ | Initial Sol. | LP bound |
| :---: | :---: | :---: | :---: | :---: |
| MT-D-01 | 979 | 3841 | 228874404 | 238166485 |
| MT-D-200 | 10880 | 547529 | 286750411 | 287228467 |
| MT-D-FN | 10880 | 645026 | 290723959 | 290881566 |
| MT-W-01 | 1006 | 3140 | 299132358 | 312121568 |
| MT-W-200 | 12320 | 554288 | 383620215 | 384099118 |
| MT-W-FN | 12320 | 593328 | 390596383 | 390869891 |
| MW-D-01 | 3988 | 19522 | 465730126 | 477563775 |
| MW-D-20 | 10790 | 718152 | 522485254 | 531510712 |
| MW-D-40 | 33563 | 2169909 | 533938531 | 543396252 |
| MW-D-FN | 47504 | 4577834 | 542182073 | 549872520 |
| MW-W-01 | 3079 | 48386 | 1268370807 | 1270311626 |
| MW-W-05 | 10790 | 789733 | 1328552109 | 1334413294 |
| MW-W-10 | 18023 | 2257068 | 1342415152 | 1360791627 |
| MW-W-FN | 22316 | 3495108 | 1350675180 | 1373020454 |
| MR-D-01 | 14058 | 60738 | 1664446852 | 1695332636 |
| MR-D-03 | 21499 | 168504 | 1739544141 | 1763685757 |
| MR-D-05 | 27621 | 295700 | 1775123794 | 1796703313 |
| MR-D-FN | 30467 | 367408 | 1794070793 | 1809854459 |
| MR-W-FN | 15639 | 267908 | 5386472651 | 5386842781 |
| CW-T-C-1 | 266403 | 162263516 | 1298968 | 1353493 |
| CW-T-C-2 | 194413 | 125379039 | 933792 | 957291 |
| CW-T-D-4 | 83091 | 43680759 | 457715 | 463672 |
| CW-T-D-6 | 83758 | 44702150 | 457605 | 463946 |
| CW-S-L-1 | 411950 | 316124758 | 1622723 | 1677563 |
| CW-S-L-2 | 443404 | 350841894 | 1692255 | 1759158 |
| CW-S-L-4 | 430379 | 340297828 | 1709043 | 1778589 |
| CW-S-L-6 | 267698 | 191469063 | 1159946 | 1192899 |
| CW-S-L-7 | 127871 | 89873520 | 589723 | 599271 |
| CR-T-C-1 | 602472 | 216862225 | 4605156 | 4801515 |
| CR-T-C-2 | 652497 | 240045639 | 4844852 | 5032895 |
| CR-T-D-4 | 651861 | 245316530 | 4789561 | 4977981 |
| CR-T-D-6 | 381380 | 128658070 | 2953177 | 3056284 |
| CR-T-D-7 | 163809 | 49945719 | 1451562 | 1469259 |
| CR-S-L-1 | 863368 | 368431905 | 5548904 | 5768579 |
| CR-S-L-2 | 880974 | 380666488 | 5617351 | 5867579 |
| CR-S-L-4 | 881910 | 383405545 | 5629351 | 5869439 |
| CR-S-L-6 | 578244 | 245739404 | 3841538 | 3990563 |
| CR-S-L-7 | 270067 | 108503583 | 1969254 | 2041822 |

we are comparing best runs of the algorithms being evaluated. The choice of $s$ assures that for each algorithm, $t^{*}$ is well-defined.

For graph algorithms, C++ is usually faster than Java by a factor from three to six. We expect this to hold for our algorithm as well, especially since we make heavy use of standard Java hash set library, which incurs significant overhead compared

TABLE 2. Problem kernels of computer, road, and social (CRS) networks

| Graph | $\|V\|$ | $\|E\|$ |
| :--- | ---: | ---: |
| web-Google | 1172 | 3469 |
| web-Stanford | 3167 | 8805 |
| as-Skitter | 9078 | 45485 |
| web-NotreDame | 25999 | 260681 |
| soc-LiveJournal1 | 41917 | 240765 |
| web-BerkStan | 71315 | 162587 |
| roadNet-PA | 264199 | 414923 |
| roadNet-TX | 315459 | 492722 |
| roadNet-CA | 505103 | 787106 |
| soc-Pokec | 906926 | 10188576 |

to $\mathrm{C}++$. Although we do not adjust the runtimes we report, one has to keep this in mind that if re-implemented in $\mathrm{C}++$, our algorithm would be faster.

We run our experiments on an AWS r3.4xlarge instance with 122GiB RAM and 16 virtual CPUs on Intel Xeon Ivy Bridge processors.
4.2. Benchmark Families. We use three problem families as benchmarks. The first one, VR Instances (Dong et al., 2021a), comes from our vehicle routing application. In this application, the MWIS problem comes up in several contexts, and we have several instances for each of these contexts. Table 1 lists the instances with their sizes. The number of nodes in these instances ranges from 979 to 883,238 ; the number of edges ranges from 3,140 to $389,304,424$. The instances are moderately sparse, but the density tends to grow with the problem size. The average degree is below 4 on some small instances and over 400 on some large ones.

Table 1 also gives values for the initial solutions we use and upper bounds on optimal solution values. We obtain the upper bounds by solving the corresponding LP relaxation problems to optimality. The initial solution are good: their values are close to the upper bound. Note that an optimal solution may not achieve the upper bound.

For VR instances, we have additional information: relaxed LP solutions and initial solutions. We use this information in practice as it yields better results. In our experiments, we give results both for runs with and runs without initial solutions. We also run our algorithm with initial solutions but without the relaxed solutions to see how much a good initial solution matters, and to have an apples to apples comparison with ILSVND, which does not use this information.

The second problem family contains the CRS Instances. We derive these instances from the computer, social, and road network instances studied by Lamm et al. (2019). The underlying graphs are natural, but the MWIS instances do not correspond to any real application. Lamm et al. (2019) describe an exact algorithm that uses local transformations producing equivalent problems on smaller graphs. For the Kernel instances, preprocessing based on these local transformations reduces the graph size significantly. The resulting graphs are the kernel graphs. Although the study of Lamm et al. (2019) focuses on exact algorithms, the authors mention that one can combine the preprocessing with a heuristic to get a fast heuristic algorithm. This is what we do. Our instances are kernel graphs for the graphs

Table 3. Map labeling instance sizes

| Name | $\|V\|$ | $\|E\|$ | degree |
| :--- | ---: | ---: | ---: |
| florida_AM3 | 1069 | 62088 | 62.3 |
| alabama_AM3 | 1614 | 117426 | 72.8 |
| rhodeisland_AM2 | 1103 | 81688 | 74.1 |
| dc_AM2 | 6360 | 592457 | 93.2 |
| virginia_AM3 | 3867 | 485330 | 125.5 |
| northcarolina_AM3 | 1178 | 189362 | 160.7 |
| massachusetts_AM3 | 2008 | 373537 | 186.0 |
| kansas_AM3 | 1605 | 408108 | 254.3 |
| washington_AM3 | 8030 | 2120696 | 264.1 |
| vermont_AM3 | 2630 | 811482 | 308.5 |
| dc_AM3 | 33367 | 17459296 | 523.3 |
| oregon_AM3 | 3670 | 1958180 | 533.6 |
| greenland_AM3 | 3942 | 2348539 | 595.8 |
| idaho_AM3 | 3208 | 2864466 | 892.9 |
| rhodeisland_AM3 | 13031 | 11855557 | 909.8 |
| hawaii_AM3 | 24436 | 40724109 | 1666.6 |
| kentucky_AM3 | 16871 | 54160431 | 3210.3 |

studied in Lamm et al. (2019). We use only the kernels with at least 1000 nodes and disregard smaller instances. Table 2 lists the CRS kernel graphs we use. The smallest graph has 1,172 nodes and 3,469 edges, and the largest has 906,926 nodes and $10,188,576$ edges. The graphs are sparse, with the average degree between 3 and 23.

We remark that we tried applying the preprocessing step to VR instances, but the reduction in problem size was very small, below $2 \%$. Therefore we do not use the preprocessing for VR instances.

One difference for our CRS instances is the way we assign weights to nodes. Lamm et al. use unweighted problem instances and assign uniform random weights. The exact weights they use are not publicly available. This makes it hard to reproduce the exact instances, especially when one uses a different language platform with a different built-in random generator. We assign a node the weight equal its ID modulo a constant (we use 200), which is reproducible.

We treat the kernel graph as an input for the algorithms we compare. When reporting runtimes, we do not include the time to reduce a graph to the kernel graph, and report solution values for the kernel graphs.

The third problem family we use is the Map Labeling Instances. These are also kernel instances, but from a real map matching application, so we treat them as a separate problem family. We obtained the instances from Sebastian Lamm (Lamm, 2020). The motivating problem is to place a subset of labels (names of cities, points of interest, etc.) on a map at a given magnification level so that the labels do not intersect. The assumption is that locations and geometry of the labels are fixed. In the graph that models the problem, the labels correspond to nodes, and we connect two nodes if their labels intersect. An independent set of the labels has no overlaps. Node weights correspond to label importance. See Strijk et al. (2000) for a discussion of unweighted version of the problem. Higher weight subsets of

TABLE 4. Results on VR instances with no additional information.

|  | METAMIS |  |  |  | ILSVND |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Name | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ |
| MT-D-01 | 238166485 | 238166485 | 238166485 | 0.948 | 238166485 | 238166485 | 238166485 | 1.290 |
| MT-D-200 | 286976422 | 287048909 | 287048909 | 188.1 | 286838210 | 286838210 | 286943799 | 2276 |
| MT-D-FN | 290866943 | 290866943 | 290866943 | 104.4 | 290393532 | 290666380 | 290666380 | 561.6 |
| MT-W-01 | 312121568 | 312121568 | 312121568 | 0.278 | 312121568 | 312121568 | 312121568 | 0.080 |
| MT-W-200 | 383818136 | 383961099 | 383961323 | 1433 | 383865836 | 383896403 | 383896403 | 1036 |
| MT-W-FN | 390688944 | 390830057 | 390854593 | 568.1 | 390715890 | 390798842 | 390798842 | 709.2 |
| MW-D-01 | 476099262 | 476164209 | 476334711 | 267.9 | 475653439 | 475906790 | 475906790 | 1173 |
| MW-D-20 | 524255389 | 525036493 | 525124575 | 85.40 | 520854115 | 522415092 | 523138978 | 2685 |
| MW-D-40 | 533934442 | 535707479 | 536520199 | 81.36 | 530227261 | 532272896 | 532400878 | 1830 |
| MW-D-FN | 539754400 | 541372345 | 541918916 | 98.34 | 532663872 | 537238784 | 537674129 | 2466 |
| MW-W-01 | 1270305952 | 1270305952 | 1270305952 | 0.500 | 1246949460 | 1246949460 | 1246949460 | 23.66 |
| MW-W-05 | 1328958047 | 1328958047 | 1328958047 | 19.96 | 1327687399 | 1328707787 | 1328707787 | 984.8 |
| MW-W-10 | 1340878388 | 1342899725 | 1342899725 | 1204 | 1331002512 | 1341482310 | 1342067985 | 1876 |
| MW-W-FN | 1349369736 | 1350818543 | 1350818543 | 527.7 | 1334835589 | 1348128240 | 1350159705 | 3584 |
| MR-D-01 | 1689074331 | 1689520690 | 1689781114 | 15.52 | 1683529331 | 1686091786 | 1687842856 | 2906 |
| MR-D-03 | 1753188475 | 1753968167 | 1754110286 | 20.34 | 1743429914 | 1747269072 | 1749972580 | 3257 |
| MR-D-05 | 1784519403 | 1785664042 | 1786342921 | 19.56 | 1770832093 | 1774407092 | 1777876780 | 3595 |
| MR-D-FN | 1795912642 | 1797284091 | 1797573192 | 22.65 | 1779897201 | 1785545729 | 1788331878 | 3388 |
| MR-W-FN | 5357026363 | 5358386615 | 5358386615 | 1442 | 5352347338 | 5370471580 | 5371649721 | 461.6 |
| CW-T-C-1 | 1310223 | 1315122 | 1317775 | 94.52 | 1290974 | 1299279 | 1302478 | 3585 |
| CW-T-C-2 | 924664 | 929626 | 931802 | 189.7 | 914736 | 921021 | 922858 | 3599 |
| CW-T-C-4 | 454769 | 456565 | 457185 | 324.4 | 452035 | 453741 | 454544 | 2365 |
| CW-T-D-6 | 455823 | 457382 | 457790 | 70.48 | 452366 | 454254 | 454254 | 1582 |
| CW-S-L-1 | 1623280 | 1630417 | 1634950 | 261.9 | 1603051 | 1615247 | 1620756 | 3597 |
| CW-S-L-2 | 1695131 | 1704424 | 1708820 | 225.3 | 1670836 | 1685870 | 1690536 | 3596 |
| CW-S-L-4 | 1712553 | 1722542 | 1725591 | 173.7 | 1689318 | 1701309 | 1706264 | 3599 |
| CW-S-L-6 | 1150229 | 1156916 | 1158925 | 138.4 | 1136356 | 1142720 | 1145694 | 3086 |
| CW-S-L-7 | 582925 | 585929 | 587288 | 125.2 | 577087 | 581583 | 581583 | 1278 |
| CR-T-C-1 | 4617204 | 4644635 | 4654419 | 58.16 | 4508901 | 4558780 | 4576695 | 3598 |
| CR-T-C-2 | 4834040 | 4863054 | 4874346 | 62.29 | 4715023 | 4772847 | 4789909 | 3600 |
| CR-T-D-4 | 4778868 | 4808490 | 4817281 | 56.91 | 4663588 | 4716258 | 4734674 | 3598 |
| CR-T-D-6 | 2945721 | 2964007 | 2970011 | 94.09 | 2896260 | 2921540 | 2929671 | 3574 |
| CR-T-D-7 | 1431915 | 1438896 | 1440281 | 148.4 | 1411061 | 1423279 | 1426400 | 3581 |
| CR-S-L-1 | 5547038 | 5575602 | 5588489 | 72.42 | 5400658 | 5464532 | 5487254 | 3595 |
| CR-S-L-2 | 5652928 | 5680688 | 5691892 | 57.91 | 5491814 | 5561766 | 5586973 | 3580 |
| CR-S-L-4 | 5634886 | 5671369 | 5681336 | 65.09 | 5477340 | 5550943 | 5572856 | 3573 |
| CR-S-L-6 | 3833391 | 3851432 | 3859513 | 92.45 | 3751019 | 3793995 | 3808314 | 3599 |
| CR-S-L-7 | 1977161 | 1986354 | 1989879 | 90.90 | 1940573 | 1957872 | 1963579 | 3584 |

the labels convey more useful information. Table 3 lists the map labeling graphs we use, ordered by their average degrees. The number of nodes in these instances ranges from 1,068 to 33,367 ; the number of edges ranges from 62,088 to $40,724,109$. Average node degrees vary from 62 to 3,210 .

Note that VR graphs, graphs from the map labeling application have natural cliques: For a fixed label $\ell$, nodes corresponding to $\ell$ and the labels intersection $\ell$ form a clique. We had access to the map labeling graph instances but not to the underlying raw data, so we do not use clique information for this problem family. In addition, we use kernels of the map labeling graphs. It is unclear if the kernelforming transformations can be adopted to generate cliques in the kernel graph from those in the original one.
4.3. Results for VR Family. For VR instances, we run experiments on the standard version of the problem and on the augmented version, where we have a good initial solution and clique information, as we do in our application. We set the time limit $\mathrm{T}=3600$ seconds.

Table 4 gives results for the standard version. For each instance in the table, column $w_{10 \%}$ shows the best solution value found at time point $T / 10$, column $w_{50 \%}$ shows the best solution value found at time point $T / 2$, and column $w$ shows the best solution value found when the process is finished at time $T$. METAMIS finds better
solutions than ILSVND except for three instances. For two instances, MT-D-01 and MT-W-01, solution quality is the same. On MW-W-01, the ILSVND solution is better, but only by $0.8 \%$. All three exceptions happen on smaller instances and both algorithms converge quickly. There is no improvement after time $T / 10$.

An interesting observation is that on MT-D-01, MT-W-01 and MT-W-FN, solution values match the corresponding upper bounds given in Table 1, so the solutions are optimal. Since the upper bound need not be tight, it is possible that we solve other instances to optimality, but do not have a proof.

On larger instances, METAMIS has better final values as well as better values at times $T / 10$ and $T / 2$. On the problem with the highest number of nodes, CR-S-$\mathrm{L}-3$, the difference in the final values is $2.1 \%$. Note that on large instances, neither algorithms converged in time $T$.

Tables 5, 6, and 7 shows results for the VR instances for METAMIS+LP, METAMIS, and ILSVND, respectively. On the instances where ILSVND does not find any improvement, $t^{*}$ is undefined. While both algorithms allow a warm start from a given solution, the METAMIS+LP version of our algorithm uses clique information to compute the relaxed LP solution, and uses it to guide local search., We evaluate both versions of METAMIS to see how much the LP relaxation helps. As in the case of no initial solution, the algorithms converge on most of the small instances and do not converge on larger instances.


Figure 1. Time-quality plot for the standard CR-S-L-4 instance with $95 \%$ confidence intervals, initial solution and LP bound. Note that the plots for METAMIS+Init and METAMIS+Init+LP are very close

Recall that with no initial solution, we found optimal solutions for MT-D-01 and MT-W-01. With the initial solution, METAMIS+LP finds an optimal solution for two more instances, MT-W-FN and MR-W-FN. METAMIS finds an optimal

Table 5. METAMIS+LP results on VR instances.

| Name | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ |
| :---: | :---: | :---: | :---: | :---: |
| MT-D-01 | 238166485 | 238166485 | 238166485 | 0.109 |
| MT-D-200 | 287038328 | 287048081 | 287048081 | 69.51 |
| MT-D-FN | 290771450 | 290771450 | 290771450 | - |
| MT-W-01 | 312121568 | 312121568 | 312121568 | 0.122 |
| MT-W-200 | 383971124 | 383985408 | 383985408 | 893.0 |
| MT-W-FN | 390828160 | 390869891 | 390869891 | 139.9 |
| MW-D-01 | 475886356 | 475987082 | 475987082 | 270.2 |
| MW-D-20 | 525052532 | 525402318 | 525486034 | 8.694 |
| MW-D-40 | 535705687 | 536210247 | 536735155 | 0.434 |
| MW-D-FN | 543098071 | 543622238 | 543857187 | - |
| MW-W-01 | 1269314742 | 1269344846 | 1269344846 | 672.0 |
| MW-W-05 | 1328958047 | 1328958047 | 1328958047 | 0.431 |
| MW-W-10 | 1342915691 | 1342915691 | 1342915691 | 0.511 |
| MW-W-FN | 1350814699 | 1350814699 | 1350818543 | - |
| MR-D-01 | 1688024106 | 1688777944 | 1689278470 | 7.245 |
| MR-D-03 | 1756186736 | 1756989875 | 1757227519 | 5.123 |
| MR-D-05 | 1787220357 | 1787666207 | 1787849777 | 19.91 |
| MR-D-FN | 1798215807 | 1798926794 | 1799452160 | 17.40 |
| MR-W-FN | 5386842781 | 5386842781 | 5386842781 | 0.503 |
| CW-T-C-1 | 1334884 | 1336953 | 1338064 | 30.69 |
| CW-T-C-2 | 944404 | 945748 | 945886 | 25.86 |
| CW-T-D-4 | 460643 | 461027 | 461056 | 2.000 |
| CW-T-D-6 | 460982 | 461223 | 461312 | 2.717 |
| CW-S-L-1 | 1656404 | 1660475 | 1660815 | 46.27 |
| CW-S-L-2 | 1731077 | 1735964 | 1738128 | 85.11 |
| CW-S-L-4 | 1748029 | 1752354 | 1753803 | 91.08 |
| CW-S-L-6 | 1174005 | 1175931 | 1177156 | 27.48 |
| CW-S-L-7 | 593045 | 593744 | 593891 | 4.825 |
| CR-T-C-1 | 4730533 | 4739684 | 4743040 | 17.92 |
| CR-T-C-2 | 4954613 | 4966121 | 4968952 | 25.80 |
| CR-T-D-4 | 4898377 | 4908285 | 4911646 | 19.69 |
| CR-T-D-6 | 3017902 | 3022448 | 3024523 | 28.67 |
| CR-T-D-7 | 1458949 | 1459958 | 1460240 | 16.88 |
| CR-S-L-1 | 5672398 | 5686939 | 5692891 | 36.79 |
| CR-S-L-2 | 5763866 | 5780859 | 5784034 | 24.90 |
| CR-S-L-4 | 5756016 | 5771410 | 5777081 | 24.08 |
| CR-S-L-6 | 3926517 | 3933476 | 3936137 | 22.24 |
| CR-S-L-7 | 2014584 | 2018371 | 2019428 | 34.09 |

solution for the latter instance, but not for the former. ILSVND does not find any new optimal solutions.

Next we discuss the effect of a good initial solution, comparing results for METAMIS and ILSVND from Tables 4 and 5-7. Comparing initial solution values from Table 1 with solutions obtained by solving the problems from scratch, we see

Table 6. METAMIS results on VR instances.

| Name | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ |
| :---: | :---: | :---: | :---: | :---: |
| MT-D-01 | 238166485 | 238166485 | 238166485 | 0.373 |
| MT-D-200 | 287010847 | 287018324 | 287036715 | 122.6 |
| MT-D-FN | 290752054 | 290771450 | 290771450 | - |
| MT-W-01 | 312121568 | 312121568 | 312121568 | 0.320 |
| MT-W-200 | 383804298 | 383986483 | 383986483 | 1343 |
| MT-W-FN | 390787880 | 390848998 | 390856179 | 710.1 |
| MW-D-01 | 475549969 | 475814986 | 475955989 | 2278 |
| MW-D-20 | 524574519 | 525068939 | 525192291 | 7.699 |
| MW-D-40 | 535436892 | 535711417 | 536092070 | 0.474 |
| MW-D-FN | 542740347 | 543253226 | 543374394 | - |
| MW-W-01 | 1269344846 | 1269344846 | 1269344846 | 0.603 |
| MW-W-05 | 1328958047 | 1328958047 | 1328958047 | 0.447 |
| MW-W-10 | 1342915691 | 1342915691 | 1342915691 | 1.255 |
| MW-W-FN | 1350771010 | 1350818542 | 1350818543 | - |
| MR-D-01 | 1687486503 | 1687807619 | 1688118984 | 16.84 |
| MR-D-03 | 1755768835 | 1756154528 | 1756337669 | 12.31 |
| MR-D-05 | 1786084687 | 1786734327 | 1786755817 | 73.22 |
| MR-D-FN | 1798075911 | 1798571155 | 1798661823 | 38.60 |
| MR-W-FN | 5386842781 | 5386842781 | 5386842781 | 0.855 |
| CW-T-C-1 | 1333129 | 1335297 | 1336563 | 22.44 |
| CW-T-C-2 | 943366 | 944785 | 945565 | 27.72 |
| CW-T-D-4 | 460554 | 460852 | 461025 | 1.828 |
| CW-T-D-6 | 460815 | 461057 | 461174 | 2.706 |
| CW-S-L-1 | 1656404 | 1660475 | 1660815 | 90.11 |
| CW-S-L-2 | 1730208 | 1734736 | 1736245 | 109.3 |
| CW-S-L-4 | 1746941 | 1751474 | 1751988 | 84.05 |
| CW-S-L-6 | 1174169 | 1175886 | 1176233 | 33.79 |
| CW-S-L-7 | 593077 | 593744 | 593947 | 6.622 |
| CR-T-C-1 | 4725855 | 4735644 | 4738289 | 18.10 |
| CR-T-C-2 | 4950818 | 4962045 | 4964446 | 19.83 |
| CR-T-D-4 | 4896504 | 4906792 | 4909999 | 17.86 |
| CR-T-D-6 | 3016890 | 3022046 | 3023349 | 40.35 |
| CR-T-D-7 | 1458571 | 1459653 | 1459948 | 8.115 |
| CR-S-L-1 | 5668764 | 5685884 | 5690515 | 21.79 |
| CR-S-L-2 | 5759512 | 5775002 | 5780449 | 22.53 |
| CR-S-L-4 | 5755282 | 5771391 | 5775704 | 24.18 |
| CR-S-L-6 | 3923574 | 3932059 | 3935089 | 19.00 |
| CR-S-L-7 | 2013466 | 2017034 | 2017836 | 40.24 |

that in many cases, the initial solution is better than the solution computed from scratch. In fact, for ILSVND, most solutions are worse than the corresponding initial solution. This confirms that our initial solutions are good.

TABLE 7. ILSVND results on VR instances.

| Name | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ |
| :---: | :---: | :---: | :---: | :---: |
| MT-D-01 | 238166485 | 238166485 | 238166485 | 1.473 |
| MT-D-200 | 286949274 | 286973561 | 286973561 | 363.1 |
| MT-D-FN | 290723959 | 290723959 | 290723959 |  |
| MT-W-01 | 312121568 | 312121568 | 312121568 | 0.063 |
| MT-W-200 | 383808376 | 383979962 | 383979962 | 1721 |
| MT-W-FN | 390805960 | 390805960 | 390805960 | 196.2 |
| MW-D-01 | 475523699 | 475732519 | 475825497 | 2134 |
| MW-D-20 | 523248884 | 523248884 | 523248884 | 26.98 |
| MW-D-40 | 534040009 | 534040009 | 534040009 | 7.797 |
| MW-D-FN | 542182073 | 542182073 | 542182073 | - |
| MW-W-01 | 1269344846 | 1269344846 | 1269344846 | 1.247 |
| MW-W-05 | 1328955871 | 1328955871 | 1328955871 | 4.266 |
| MW-W-10 | 1342847887 | 1342847887 | 1342847887 | 19.39 |
| MW-W-FN | 1350675180 | 1350675180 | 1350675180 | - |
| MR-D-01 | 1684211854 | 1686046636 | 1686452467 | 2763 |
| MR-D-03 | 1751006933 | 1752345436 | 1752769459 | 3305 |
| MR-D-05 | 1782046226 | 1782560957 | 1783836981 | 3525 |
| MR-D-FN | 1794949819 | 1794949819 | 1796037791 | 3564 |
| MR-W-FN | 5386838669 | 5386838669 | 5386838669 | 10.01 |
| CW-T-C-1 | 1322410 | 1326551 | 1327556 | 3501 |
| CW-T-C-2 | 939568 | 940356 | 940356 | 701.3 |
| CW-T-D-4 | 458360 | 458360 | 458360 | 48.65 |
| CW-T-D-6 | 459096 | 459096 | 459096 | 80.73 |
| CW-S-L-1 | 1644241 | 1649006 | 1651483 | 3585 |
| CW-S-L-2 | 1714923 | 1722672 | 1724930 | 3452 |
| CW-S-L-4 | 1733007 | 1739992 | 1742459 | 3553 |
| CW-S-L-6 | 1167611 | 1169914 | 1170096 | 1886 |
| CW-S-L-7 | 591398 | 591398 | 591398 | 161.2 |
| CR-T-C-1 | 4665849 | 4687422 | 4696568 | 3591 |
| CR-T-C-2 | 4891697 | 4912140 | 4920058 | 3585 |
| CR-T-D-4 | 4836312 | 4859311 | 4867272 | 3597 |
| CR-T-D-6 | 2990174 | 2999852 | 3004067 | 3593 |
| CR-T-D-7 | 1455226 | 1456048 | 1456048 | 752.0 |
| CR-S-L-1 | 5590089 | 5617916 | 5630437 | 3596 |
| CR-S-L-2 | 5670522 | 5701371 | 5715430 | 3589 |
| CR-S-L-4 | 5676163 | 5701271 | 5715256 | 3598 |
| CR-S-L-6 | 3883092 | 3898898 | 3905831 | 3597 |
| CR-S-L-7 | 1998320 | 2006129 | 2007794 | 3488 |

With the warm start, both variants of our algorithm, METAMIS and METAMIS+LP, dominate ILSVND, producing same or (in most cases) better quality solutions. ILSVND is also slower on all instances except one.

To evaluate the benefit of using LP relaxation, we compare METAMIS+LP with METAMIS. On most instances, METAMIS+LP dominates METAMIS. The latter

Table 8. CRS family results

| Name | METAMIS |  |  |  | ILSVND |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ |
| web-Google | 50143 | 50143 | 50143 | 0.398 | 50143 | 50143 | 50143 | 0.105 |
| web-Stanford | 140528 | 140528 | 140528 | 1.263 | 140524 | 140528 | 140528 | 627.5 |
| as-Skitter | 403178 | 403220 | 403243 | 1443 | 403100 | 403159 | 403254 | 1254 |
| web-NotreDame | 882498 | 883034 | 883073 | 98.10 | 882347 | 882347 | 882347 | 131.0 |
| soc-LiveJournal1 | 1639058 | 1639150 | 1639150 | 58.56 | 1638224 | 1638537 | 1638619 | 1467 |
| web-BerkStan | 3202724 | 3203853 | 3203896 | 50.64 | 3192742 | 3196624 | 3196865 | 1478 |
| roadNet-PA | 12327884 | 12381128 | 12386969 | 909.7 | 12165765 | 12330319 | 12342736 | 7143 |
| roadNet-TX | 14706445 | 14798448 | 14809215 | 1323 | 14466017 | 14726128 | 14758276 | 7193 |
| roadNet-CA | 23447525 | 23688219 | 23731691 | 1601 | 22837850 | 23400764 | 23576536 | 7200 |
| soc-Pokec | 40093941 | 40154755 | 40167652 | 119.5 | 39193450 | 39677478 | 39882658 | 7200 |

never finds a better solution. For about $1 / 3$ of the instances, solution quality is the same, and for the remaining $2 / 3$, METAMIS+LP performs better. The same holds for intermediate times $T / 10$ and $T / 2$ except for one instance at $T / 2$ where METAMIS solution value is slightly better.

Plots help visualize relative algorithm performance. We plot the data for 2hour runs of all algorithms on one of the largest instances in Figure 1. Neither algorithm converges. Both with and without an initial solution, METAMIS dominates ILSVND. With the initial solution, the algorithms converge to a better value and dominate the corresponding algorithms without an initial solution. Comparing METAMIS+LP with METAMIS, we see that the relaxed solution helps. The former algorithm dominates the latter.
4.4. Results for CRS Problems. Table 8 gives results for the CRS family. On this family we set the time limit to 1500 seconds for all problems except for the four largest ones, as the algorithms did not converge on them in 1500 seconds. For the four largest problems, we set the limit to 7200 seconds.

On web-Google, both algorithms find solutions with the same value in under a second and fail to improve it afterwards (the solution may be optimal). On web-Stanford and as-Skitter, ILSVND performs worse on average and has higher variance. The best solution of ILSVND is the same for web-Stanford and slightly better for as-Skitter (see Figure 2). On three other smaller instances, METAMIS finds a better solution and dominates ILSVND. For example, see the plot for webBerkSt in Figure 3.

On the four larger instances METAMIS dominates ILSVND and converges to better values. See, for example, Figure 4 for a plot of two hour-long runs on RoadNet-TX.
4.5. Results for Map Labeling Problems. Table 9 gives our results for the Map Labeling instances. On these instances, we ran the two algorithms for 1,500 seconds. The quality of the solutions the algorithms find is the same except for four instances. METAMIS is better on rhodeisland_AM3 and hawaii_AM3, and ILSVND - on dc_AM3 and kentucky_AM3. Although both algorithms converge to a similar value, ILSVND usually converges faster.

Figure 5 illustrates convergence to an essentially the same value using greenland_AM3 as an example. Figure 6 gives the plot for rhodeisland_AM3, where the picture is similar but METAMIS converges to a slightly better value. Figure 7 gives the plot for kentucky_AM3, where ILSVND converges to a better value.


Figure 2. Time-quality plot for as-Skitter with $95 \%$ confidence intervals


Figure 3. Time-quality plot for web-BerkSt with $95 \%$ confidence intervals

## 5. Concluding remarks

We developed METAMIS for a real-world VR application for which even a small improvement in solution quality yields substantial cost reduction. We published benchmark VR instances in Dong et al. (2021a;b). These instances are structurally different from other MWIS benchmarks and include large instances. Our study is


Figure 4. Time-quality plot for roadNet-TX with $95 \%$ confidence intervals
Table 9. Map Labeling family results.

| Name | METAMIS |  |  |  | ILSVND |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ | $w_{10 \%}$ | $w_{50 \%}$ | $w$ | $t^{*}[s]$ |
| florida_AM3 | 46132 | 46132 | 46132 | 3.080 | 46132 | 46132 | 46132 | 0.024 |
| alabama_AM3 | 45449 | 45449 | 45449 | 5.412 | 45449 | 45449 | 45449 | 0.104 |
| rhodeisland_AM2 | 43722 | 43722 | 43722 | 0.269 | 43722 | 43722 | 43722 | 0.048 |
| dc_AM2 | 100302 | 100302 | 100302 | 58.38 | 100302 | 100302 | 100302 | 107.3 |
| virginia_AM3 | 97873 | 97873 | 97873 | 9.080 | 97873 | 97873 | 97873 | 2.750 |
| northcarolina_AM3 | 13062 | 13062 | 13062 | 0.376 | 13062 | 13062 | 13062 | 0.040 |
| massachusetts_AM3 | 17224 | 17224 | 17224 | 2.223 | 17224 | 17224 | 17224 | 0.142 |
| kansas_AM3 | 5694 | 5694 | 5694 | 3.383 | 5694 | 5694 | 5694 | 0.209 |
| washington_AM3 | 118196 | 118196 | 118196 | 62.35 | 118196 | 118196 | 118196 | 119.8 |
| vermont_AM3 | 28349 | 28349 | 28349 | 9.332 | 28349 | 28349 | 28349 | 0.297 |
| dc_AM3 | 141785 | 142736 | 142910 | 1347 | 142917 | 143014 | 143014 | 111.8 |
| oregon_AM3 | 34471 | 34471 | 34471 | 11.56 | 34471 | 34471 | 34471 | 49.42 |
| greenland_AM3 | 11960 | 11960 | 11960 | 28.02 | 11959 | 11960 | 11960 | 178.4 |
| idaho_AM3 | 9224 | 9224 | 9224 | 21.23 | 9224 | 9224 | 9224 | 127.4 |
| rhodeisland_AM3 | 80897 | 81013 | 81013 | 449.7 | 80980 | 80980 | 80980 | 101.7 |
| hawaii_AM3 | 58394 | 58808 | 58819 | 1207 | 58779 | 58783 | 58814 | 757.0 |
| kentucky_AM3 | 30736 | 30777 | 30789 | 1387 | 31084 | 31101 | 31103 | 11.90 |

the first to include the new benchmark, and we show that METAMIS works well on the VR instances. METAMIS is also competitive on CRS and Map Matching instances.

METAMIS uses a more sophisticated set of local search moves and introduces data structures that facilitate efficient implementation of these moves. We also introduce a new variation of path relinking tailored to large problems. In addition, we show how to use a good relaxed solution to guide local search. These techniques add to the metaheuristic design toolset. We hope that our ideas will lead to even


Figure 5. Time-quality plot for greenland_AM3 with 95\% confidence intervals


Figure 6. Time-quality plot for rhodeisland_AM3 with $95 \%$ confidence intervals
more efficient MWIS algorithms. The ideas may also prove useful in methaheuristic algorithms for other problems.

We do not include DIMACS problem instances (Johnson and Trick, 1996) in our experiments. These instances have a special structure that is an artifact of how the instances are generated. The instances originate from clique instances,


Figure 7. Time-quality plot for kentucky_AM3 with $95 \%$ confidence intervals
both synthetic and real-life. One takes a compliment of a graph to obtain the corresponding unweighted maximum independent set instances. To get MWIS instances, one adds weights, which are either random or ID modulo a constant. The original clique graphs are sparse, and the resulting MWIS graphs are very dense. Even for the clique instances that correspond to real-life applications, the resulting MWIS instances do not. These problems are relatively small and easy. Many have been solved exactly (Nogueira et al., 2018).

Preliminary experiments show that our algorithm performs worse than ILSVND on DIMACS instances. We tuned our code for large VR instances and used the same parameter values in all experiments. A different choice of parameter values may improve performance on specific problem families, such as the DIMACS family (Kummer, 2020; Kummer et al., 2020). Also, for dense graphs, a different (e.g., adjacency matrix) representation of the input graph may be more efficient.

For the Map Labeling problem family, the fact that on most instances METAMIS and ILSVND converge to the same solution value suggests that in many cases the solutions may be optimal. One way to verify this conjecture is to prove optimality by using local search solutions to warm start an exact solver. Even if the solver cannot solve such a problem from scratch, it may be able to prove optimality of a given solution.

## References

A. Alnaser, O. AlHeyasat, A. Abu-Ein, H. Hatamleh, and A. Sharadqeh. Time comparing between Java and C++ software. J. of Software Engineering and Applications, 5:630-633, 2012.
D.V. Andrade, M.G.C. Resende, and R.F. Werneck. Fast local search for the maximum independent set problem. J. of Heuristics, 18:525-547, 2012.

AWS. Aws r3.4xlarge instance, 2021. URL https://tiny.amazon.com/ebd69zqv. Last visited January 21, 2021.
A. Buchanan and S. Butenko. Tight extended formulations for independent set. Technical report, Optimization Online, September 2014. URL http: //www.optimization-online.org/DB_FILE/2014/09/4540.pdf.
S. Butenko. Maximum independent set and related problems with applications. PhD thesis, U. of Florida, Gainesville, Florida, 2003.
R. Carraghan and P.M. Pardalos. An exact algorithm for the maximum clique problem. Operations Research Letters, 9:375-382, 1990.
Y. Dong, A.V. Goldberg, A. Noe, N. Parotsidis, M.G.C. Resende, and Q. Spaen. New instances for maximum weight independent set from a vehicle routing application. Technical Report arXiv:2105.12623, arXiv.org, 2021a.
Y. Dong, A.V. Goldberg, A. Noe, N. Parotsidis, M.G.C. Resende, and Q. Spaen. New instances for maximum weight independent set from a vehicle routing application. Operations Research Forum, 2(48), 2021b. doi: https://doi.org/10.1007/ s43069-021-00084-x.
J. Edmonds. Paths, trees, and flowers. Can. J. Math., 17:449-467, 1965.
T.A. Feo and M.G.C. Resende. A probabilistic heuristic for a computationally difficult set covering problem. Operations Research Letters, 8:67-71, 1989.
T.A. Feo and M.G.C. Resende. Greedy randomized adaptive search procedures. Journal of Global Optimization, 6:109-133, 1995.
T.A. Feo, M.G.C. Resende, and S.H. Smith. A greedy randomized adaptive search procedure for maximum independent set. Operations Research, 42:860-878, 1994.
M.L. Fredman and R.E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. Journal of the ACM, 34(3):596-615, 1986.
C. Friden, A. Hertz, and D. de Werra. STABULUS: A technique for finding stable sets in large graphs with tabu search. Computing, 42:35-44, 1989.
M.R. Garey and D.S. Johnson. Computers and Intractability: A guide to the theory of NP-completeness. W.H. Freeman and Company, San Francisco, 1979.
J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. Acta Mathematica, 182: 105-142, 1999.
D.S. Johnson and M.A. Trick. Cliques, Coloring, and Satisfiability: Second DIMACS Implementation Challenge. AMS, Boston, MA, USA, 1996.
R.M. Karp. Reducibility among combinatorial problems. In R.E. Miller and J.W. Thatcher, editors, Complexity of Computer Computations, pages 85-103. Plenum, New York, 1972.
A.F. Kummer. Personal communication, 2020.
A.F. Kummer, M.G.C. Resende, and M. Souto. Automatic algorithm configuration and selection of MetaMIS for maximum independent set. Technical report, Amazon MMPROS, Seattle, 2020.
M. Laguna and R. Martí. GRASP and path relinking for 2-layer straight line crossing minimization. INFORMS Journal on Computing, 11:44-52, 1999.
S. Lamm. Personal communication, 2020.
S. Lamm, C. Schulz, D. Strash, R. Williger, and Huashuo Zhang. Exactly solving the maximum weight independent set problem on large real-world graphs. In Proceedings of the Twenty-First Workshop on Algorithm Engineering and Experiments, ALENEX 2019, pages 144-158. SIAM, 2019. doi: 10.1137/1. 9781611975499.12. URL https://doi.org/10.1137/1.9781611975499.12.
B. Nogueira, R.G.S. Pinheiro, and A. Subramanian. A hybrid iterated local search heuristic for the maximum weight independent set problem. Optimization Letters, 12(3):567-583, 2018.
M.W. Padberg. On the facial structure of set packing polyhedra. Mathematical Programming, 5:199-215, 1973.
M. Pelillo. Heuristics for maximum clique and independent set. In C.A. Floudas and P.M. Pardalos, editors, Encyclopedia of Optimization, pages 1508-1520. Springer US, Boston, MA, 2009.
P. Raghavan and C.D. Tompson. Randomized rounding: A technique for provably good algorithms and algorithmic proofs. Combinatorica, 7:365-374, 1987.
M.G.C. Resende and C.C. Ribeiro. Optimization by GRASP: Greedy Randomized Adaptive Search Procedures. Springer, New York, 2016.
M.G.C. Resende, R. Martí, M. Gallego, and A. Duarte. GRASP and path relinking for the max-min diversity problem. Computers \& Operations Research, 37:498508, 2010.
T.W. Strijk, A.M. Verweij, and K.I. Aardal. Algorithms for maximum independent set applied to map labelling, 2000.
(Yuanyuan Dong) Dallas, TX USA. The work was done while the author was at AmaZON.COM

Email address: njdyy03@gmail.com
(Andrew V. Goldberg) Amazon.com, East Palo Alto, CA USA.
Email address: avg@alum.mit.edu
(Alexander Noe) Amazon.com, Bellevue, WA USA.
Email address: alexander.noe@univie.ac.at
(Nikos Parotsidis) Google, Zurich, Switzerland. The work was done while the author was at Amazon.com

Email address: nickparo1@gmail.com
(Mauricio G.C. Resende) Amazon.com and Industrial \& Systems Engineering, University of Washington, Seattle, WA USA.

Email address: mgcr@berkeley.edu
(Quico Spaen) Amazon.com, East Palo Alto, CA USA.
Email address: quico@spaen.nl


[^0]:    Date: March 28, 2022.
    Key words and phrases. GRASP, local search, maximum-weight independent set, path relinking, heuristic, metaheuristic.

