



# An Algorithmic Study of Fully Dynamic Independent Sets for Map Labeling

SUJOY BHOORE, Indian Institute of Science Education and Research, Bhopal, India  
GUANGPING LI and MARTIN NÖLLENBURG, TU Wien, Austria

Map labeling is a classical problem in cartography and geographic information systems that asks to place labels for area, line, and point features, with the goal to select and place the maximum number of independent (i.e., overlap-free) labels. A practically interesting case is point labeling with axis-parallel rectangular labels of common size. In a fully dynamic setting, at each timestep, either a new label appears or an existing label disappears. Then, the challenge is to maintain a maximum cardinality subset of pairwise independent labels with sublinear update time. Motivated by this, we study the maximal independent set (MIS) and maximum independent set (MAX-IS) problems on fully dynamic (insertion/deletion model) sets of axis-parallel rectangles of two types: (i) uniform height and width and (ii) uniform height and arbitrary width; both settings can be modeled as rectangle intersection graphs.

We present the first deterministic algorithm for maintaining an MIS (and thus a 4-approximate MAX-IS) of a dynamic set of uniform rectangles with polylogarithmic update time. This breaks the natural barrier of  $\Omega(\Delta)$  update time (where  $\Delta$  is the maximum degree in the graph) for *vertex updates* presented by Asadi et al. (STOC 2018). We continue by investigating MAX-IS and provide a series of deterministic dynamic approximation schemes. For uniform rectangles, we first give an algorithm that maintains a 4-approximate MAX-IS with  $O(1)$  update time. In a subsequent algorithm, we establish the trade-off between approximation quality  $2(1 + \frac{1}{k})$  and update time  $O(k^2 \log n)$ , for  $k \in \mathbb{N}$ . We conclude with an algorithm that maintains a 2-approximate MAX-IS for dynamic sets of unit-height and arbitrary-width rectangles with  $O(\log^2 n + \omega \log n)$  update time, where  $\omega$  is the maximum size of an independent set of rectangles stabbed by any horizontal line. We implement our algorithms and report the results of an experimental comparison exploring the trade-off between solution quality and update time for synthetic and real-world map labeling instances. We made several major observations in our empirical study. First, the original approximations are well above their respective worst-case ratios. Second, in comparison with the static approaches, the dynamic approaches show a significant speedup in practice. Third, the approximation algorithms show their predicted relative behavior. The better the solution quality, the worse the update times. Fourth, a simple greedy augmentation to the approximate solutions of the algorithms boost the solution sizes significantly in practice.

CCS Concepts: • **Theory of computation** → **Computational geometry**; Data structures design and analysis; *Dynamic graph algorithms*;

Additional Key Words and Phrases: Independent sets, dynamic algorithms, rectangle intersection graphs, approximation algorithms, experimental evaluation

A preliminary version of this article appeared in the Proceedings of the 28th Annual European Symposium on Algorithms (ESA 2020) Track B. Source code and the benchmark data are available at <https://dyna-mis.github.io/dynaMIS/>.

Authors' addresses: S. Bhoore, Indian Institute of Science Education and Research, Bhopal, Bhopal Bypass Road, Bhauri, Madhya Pradesh, Bhopal, India, 462066; email: [sujoy.bhoore@gmail.com](mailto:sujoy.bhoore@gmail.com); G. Li and M. Nöllenburg, Algorithms and Complexity Group, TU Wien, Favoritenstraße 9-11, Vienna, Austria, 1040; emails: [guangping.noellenburg@ac.tuwien.ac.at](mailto:guangping.noellenburg@ac.tuwien.ac.at).

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](mailto:permissions@acm.org).

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1084-6654/2022/07-ART1.8 \$15.00

<https://doi.org/10.1145/3514240>

**ACM Reference format:**

Sujoy Bhore, Guangping Li, and Martin Nöllenburg, 2022. An Algorithmic Study of Fully Dynamic Independent Sets for Map Labeling. *J. Exp. Algor.* 27, 1, Article 1.8 (July 2022), 36 pages.  
<https://doi.org/10.1145/3514240>

---

**1 INTRODUCTION**

MAP LABELING is a classical problem in cartography and geographic information systems (GIS), that has received significant attention in the past few decades and is concerned with selecting and positioning labels on a map for area, line, and point features. The focus in the computational geometry community has been on labeling point features [3, 23, 48, 49]. The labels are typically modeled as the bounding boxes of short names, which correspond precisely to unit-height but arbitrary-width rectangles; alternatively, labels can be standardized icons or symbols, which correspond to rectangles of uniform size. In map labeling, a key task is in fact to select an independent (i.e., overlap-free) set of labels from a given set of candidate labels. Commonly, the optimization goal is related to maximizing the number of labels. Given a set  $\mathcal{R}$  of rectangular labels, MAP LABELING is essentially equivalent to the problem of finding a maximum independent set (MAX-IS) in the intersection graph induced by  $\mathcal{R}$ .

The independent set problem is a fundamental graph problem with a wide range of applications. Given a graph  $G = (V, E)$ , a set of vertices  $M \subset V$  is *independent* if no two vertices in  $M$  are adjacent in  $G$ . A maximal independent set (MIS) is an independent set that is not a proper subset of any other independent set. An MAX-IS is a maximum cardinality independent set. While MAX-IS is one of Karp's 21 classical NP-complete problems [35], computing an MIS can easily be done by a simple greedy algorithm in  $O(|E|)$  time. The MIS problem has been studied in the context of several other prominent problems, such as graph coloring [37], maximum matching [34], and vertex cover [42]. However, MAX-IS serves as a natural model for many real-life optimization problems that arise in the fields of cartography, scheduling, computer graphics, and information retrieval, among others (see [3, 43, 45, 47]).

Stronger results for independent set problems in geometric intersection graphs are known in comparison to general graphs. For instance, it is known that MAX-IS on general graphs cannot be approximated better than  $|V|^{1-\epsilon}$  in polynomial time for any  $\epsilon > 0$  unless  $NP = P$  [51]. In contrast, a randomized polynomial-time algorithm exists that computes for rectangle intersection graphs an  $O(\log \log n)$ -approximate solution to MAX-IS with high probability [12], as well as QPTASs [2, 17]. Very recently, the constant factor approximation schemes have been developed for the MAX-IS on rectangle intersection graphs (see [26, 41]). The MAX-IS problem is already NP-hard on unit square intersection graphs [24]; however, it admits a polynomial-time approximation scheme (PTAS) for unit square intersection graphs [22] and more generally for pseudo disks [13]. Moreover, for rectangles with either uniform size or at least uniform height and bounded aspect ratio, the size of an MIS is not arbitrarily worse than the size of a MAX-IS. For instance, any MIS of a set of uniform rectangles is a 4-approximate solution to the MAX-IS problem, since each rectangle can have at most four independent neighbors.

Past research has mostly considered static label sets in static maps [3, 23, 48, 49] and in dynamic maps allowing zooming [6] or rotations [30], but not fully dynamic label sets with insertions and deletions of labels. Recently, Klute et al. [36] proposed a framework for semi-automatic label placement, where domain experts can interactively insert and delete labels. In their setting, an initially computed large independent set of labels can be interactively modified by a cartographer, who can easily take context information and soft criteria such as interactions with the background map or surrounding labels into account. Standard map labeling algorithms typically do not handle such

aspects well [23, 44]. Based on these modifications (e.g., deletion, forced selection, translation, or resizing), the solution is updated by a dynamic algorithm while adhering to the new constraints. Another scenario for dynamic labels is maps, in which features and labels (dis-)appear over time, such as those based on a stream of geotagged, uniform-size photos posted on social media or, more generally, maps with labels of dynamic spatio-temporal point sets [25]. For instance, a geo-located event that happens at time  $t$  triggers the availability of a new label for a certain period of time, after which it vanishes again. Examples beyond social media are reports of earthquakes, forest fires, or disease incidences. Traditional geographic map labeling deals with small and relatively static label sets, whereas labeling of social network data, especially the ones used in anomaly detection and visual analytics, usually deals with vast and dynamic label sets (see [38, 46]). Furthermore, note that these applications often run on devices with limited computational resources, such as mobile devices. Therefore, it is desirable to design dynamic algorithms that can handle the changes in an efficient and robust manner. Motivated by this, we study the independent set problem for dynamic sets of axis-parallel rectangles of two types:

- Rectangles of uniform height and width
- Rectangles of uniform height and arbitrary width.

We consider fully dynamic algorithms for maintaining independent sets under insertions and deletions of rectangles—that is, vertex insertions and deletions in the corresponding dynamic rectangle intersection graph.

*Dynamic* graphs are subject to discrete changes over time (i.e., insertions or deletions of vertices or edges) [21]. A dynamic graph algorithm solves a computational problem, such as the independent set problem, on a dynamic graph by updating efficiently the previous solution as the graph changes over time rather than recomputing it from scratch. A dynamic graph algorithm is called *fully dynamic* if it allows both insertions and deletions, and *partially dynamic* if only insertions or only deletions are allowed. Although general dynamic independent set algorithms can obviously also be applied to rectangle intersection graphs, our goal is to exploit their geometric properties to obtain more efficient algorithms.

**Related work.** There has been a lot of work on dynamic graph algorithms in the past decade, and dynamic algorithms still receive considerable attention in theoretical computer science. We point out some of these works, such as those on spanners [8], vertex cover [9], set cover [1], graph coloring [10], and maximal matching [27]. In particular, the MIS problem on dynamic graphs with edge updates has attracted significant attention in the past two years [4, 4, 7, 15, 18]. For vertex insertion/deletion, an MIS can be maintained dynamically in  $O(\Delta)$  update time by using the recent algorithm of Assadi et al. [4], where  $\Delta$  is the maximum degree of the intersection graph.

Recently, Henzinger et al. [32] studied the MAX-IS problem for intervals, hypercubes, and hyperrectangles in  $d$  dimensions, with special assumptions. They assumed that the objects are axis-parallel and contained in the space  $[0, N]^d$ ; the value of  $N$  is given in advance, and each edge of an input object has length at least 1 and at most  $N$ . They have presented dynamic approximation schemes with the update time  $\text{polylog}(n, N)$ , where  $n$  is the instance size. We note that in general,  $N$  might be exponential in  $n$  or even unbounded, thus those bounds are not sublinear in  $n$  in the general case. Subsequently, Bhore et al. [11] designed a dynamic approximation scheme for dynamic intervals that maintains a  $(1 + \epsilon)$ -approximate MAX-IS in  $O_\epsilon(\log n)$  update time, where  $\epsilon > 0$  is any positive constant and the notation  $O_\epsilon$  hides terms depending only on  $\epsilon$ . Gavruskin et al. [28] studied the MAX-IS problem for dynamic proper intervals (intervals cannot contain one another) and showed how to maintain a MAX-IS with polylogarithmic update time.

There is a long history of the empirical study of map-labeling problems. This chain of research started with the work of Christensen et al. [16]. They proposed two methods: one based on a

discrete form of gradient descent and the other on simulated annealing. An alternative approach was presented by Wagner and Wolff [49] for the labeling problem, who used the sample data in the experimental evaluation that consists of three different classes of random problems and a selection of problems arising in the production of groundwater quality maps by the authorities of the City of Munich. Nascimento and Eades [20] proposed a practically motivated framework, called *user hints*, and proposed an interactive map-labeling system based on this along with its evaluation. This type of user-interactive approach was empirically studied by Klute et al. [36]. Moreover, other aspects of dynamic map labeling, such as rotation and zooming, have been studied over the years (see [5, 6, 29]). De Berg and Gerrits [19] developed and experimentally evaluated a heuristic for labeling moving points on static maps.

**Results and organization.** We study MIS and MAX-IS problems for dynamic sets of  $O(n)$  axis-parallel rectangles of two types: (i) congruent rectangles of uniform height and width and (ii) rectangles of uniform height and arbitrary width.

In this work we design and implement algorithms for dynamic MIS and MAX-IS that demonstrate the trade-off between update time and approximation factor, both from a theoretical perspective and in an experimental evaluation. In contrast to the recent dynamic MIS algorithms, which are randomized [4, 4, 7, 15], our algorithms are deterministic.

In Section 3, we present an algorithm that maintains an MIS of a dynamic set of unit squares in  $O(\log n \log \log n)$  update time or, alternatively, with sublogarithmic amortized update time, improving the best-known update time  $\Omega(\Delta)$  by Assadi et al. [4], where  $\Delta$  is the maximum degree of the intersection graph. A major but generally unavoidable bottleneck of that algorithm is that the entire graph is stored explicitly, and thus insertions/deletions of vertices take  $\Omega(\Delta)$  time. We use structural geometric properties of the unit squares along with a dynamic orthogonal range searching data structure to bypass the explicit intersection graph and overcome this bottleneck.

In Section 4, we study the MAX-IS problem. For dynamic unit squares, we give an algorithm that maintains a 4-approximate MAX-IS with  $O(1)$  update time. We generalize this algorithm and improve the approximation factor to  $2(1 + \frac{1}{k})$ , which increases the update time to  $O(k^2 \log n)$ . We conclude with an algorithm that maintains a 2-approximate MAX-IS for a dynamic set of unit-height and arbitrary-width rectangles (in fact, for a dynamic interval graph, which is of independent interest) with  $O(\log^2 n + \omega \log n)$  update time, where  $\omega$  is the maximum size of an independent set of rectangles stabbed by any horizontal line.

Finally, Section 5 provides an experimental evaluation of the proposed MAX-IS approximation algorithms on synthetic and real-world map labeling datasets. The experiments explore the trade-off between solution size and update time, as well as the speedup of the dynamic algorithms over their static counterparts. See the supplemental material<sup>1</sup> for source code and benchmark data.

## 2 MODEL AND NOTATION

For every  $N \in \mathbb{N}$ ,  $[N]$  denotes the set  $\{1, 2, \dots, N\}$ . Let  $R$  be a dynamic set of axis-parallel, unit-height rectangles in the plane, which is dynamically updated by a sequence of  $N \in \mathbb{N}$  insertions and deletions. Let  $R_i$  denote the set of rectangles at step  $i \in [N]$ , and let  $n = \max\{|R_i| \mid i \in [N]\}$  be the maximum number of rectangles over all steps. The rectangle intersection graph defined by  $R_i$  at timestep  $i$  is denoted as  $G_i = (R_i, E_i)$ , where two rectangles  $r, r' \in R_i$  are connected by an edge  $\{r, r'\} \in E_i$  if and only if  $r \cap r' \neq \emptyset$ . We use  $M_i$  to denote an MIS in  $G_i$ , and  $OPT_i$  to denote a MAX-IS in  $G_i$ . For a graph  $G = (V, E)$  and a vertex  $v \in V$ , let  $N(v)$  denote the set of neighbors of  $v$  in  $G$ . This notation also extends to any subset  $U \subseteq V$  by defining  $N(U) = \bigcup_{v \in U} N(v)$ . We use  $\deg(v)$

<sup>1</sup>Source code and the benchmark data are available at <https://dyna-mis.github.io/dynaMIS/>.

to denote the degree of a vertex  $v \in V$ . For any vertex  $v \in V$ , let  $N^r(v)$  be the  $r$ -neighborhood of  $v$ —that is, the set of vertices that are within distance at most  $r$  from  $v$  (excluding  $v$ ).

We study the independent set problem for dynamic sets of axis-parallel rectangles of two types: (i) unit rectangles and (ii) rectangles of unit height and arbitrary width. In this work, we may assume that the unit rectangles are unit squares. If the rectangles of  $R$  are of uniform height and width, we can use an affine transformation to map  $R$  to a set of unit squares  $S$  and map  $R_i$  to unit square set  $S_i$  for  $i \in [N]$ . We further define the set  $C_i$  be the corresponding centers of squares of  $S_i$ .

### 3 ALGORITHMS FOR DYNAMIC MIS

In this section, we study the MIS problem for dynamic uniform rectangles. As stated before, we can assume w.l.o.g. that the rectangles are unit squares. We design an algorithm that maintains an MIS for a dynamic set of  $O(n)$  unit squares in polylogarithmic update time. Assadi et al. [4] presented an algorithm for maintaining an MIS on general dynamic graphs with  $O(\Delta)$  update time, where  $\Delta$  is the maximum degree in the graph. In the worst case, however, that algorithm takes  $O(n)$  update time. In fact, it seems unavoidable for an algorithm that explicitly maintains the (intersection) graph to perform an MIS update in less than  $\Omega(\deg(v))$  time for an insertion/deletion of a vertex  $v$ . In contrast, our proposed algorithm in this section does not explicitly maintain the intersection graph  $G_i = (S_i, E_i)$  (for any  $i \in [N]$ ), but rather only the set of squares  $S_i$  in a suitable dynamic geometric data structure. For the ease of explanation, however, we do use graph terms at times.

Let  $i \in [N]$  be any time point in the sequence of updates. For each square  $s_v \in S_i$ , let  $s_v^a$  be a square of side length  $a$  concentric with  $s_v$ . Further, let  $M_i$  denote the MIS that we compute for  $G_i = (S_i, E_i)$ , and let  $C(M_i) \subseteq C_i$  be their corresponding square centers. We maintain two fully dynamic orthogonal range searching data structures, which maintain a set of points dynamically and support efficient deletions and insertions of points, throughout: (i) a dynamic range tree  $T(C_i)$  for the entire point set  $C_i$  and (ii) a dynamic range tree  $T(C(M_i))$  for the point set  $C(M_i)$  corresponding to the centers of  $M_i$ . They can be implemented with dynamic fractional cascading [39], which yields  $O(\log n \log \log n)$  update time and  $O(k + \log n \log \log n)$  query time for reporting  $k$  points.

We compute the initial MIS  $M_1$  for  $G_1 = (S_1, E_1)$  by using a simple linear-time greedy algorithm. First, we initialize the range tree  $T(C_1)$ . Then, we iterate through the set  $S_1$  as long as it is not empty, select a square  $s_v$  for  $M_1$  and insert its center into  $T(C(M_1))$ , find its neighbors  $N(s_v)$  by a range query in  $T(C_1)$  with the concentric square  $s_v^2$ , and delete  $N(s_v)$  from  $S_1$ . It is clear that once this process terminates,  $M_1$  is an MIS.

When we move in the next step from  $G_i = (S_i, E_i)$  to  $G_{i+1} = (S_{i+1}, E_{i+1})$ , either a square is inserted into  $S_i$  or deleted from  $S_i$ . Let  $s_x$  be the square that is inserted or deleted.

**LEMMA 1.** *Given an arbitrary set of pairwise overlap-free unit squares  $S$  and an arbitrary square  $r$  of side length 2,  $r$  contains at most four centers of unit squares of  $S$ .*

**PROOF.** We split  $r$  equally into four parts where each quarter corresponds to a unit square. Consider two squares  $s_1, s_2 \in S$ . If their centers lie in the same quarter, then they must overlap. Hence, by a simple packing argument the claim holds.  $\square$

**INSERTION:** When we insert a square  $s_x$  into  $S_i$  to obtain  $S_{i+1}$ , we do the following operations. First, we obtain  $T(C_{i+1})$  by inserting the center of  $s_x$  into  $T(C_i)$ . Next, we have to detect whether  $s_x$  can be included in  $M_{i+1}$ . If there exists a square  $s_u$  from  $M_i$  intersecting  $s_x$ , we should not include  $s_x$ ; otherwise, we will add it to the MIS. To check this, we search with the range  $s_x^2$  in  $T(C(M_i))$ . By Lemma 1, we know that no more than four points (the centers of four independent squares) of  $C(M_i)$  can be in the range  $s_x^2$ . If the query returns such a point, then  $s_x$  would intersect with

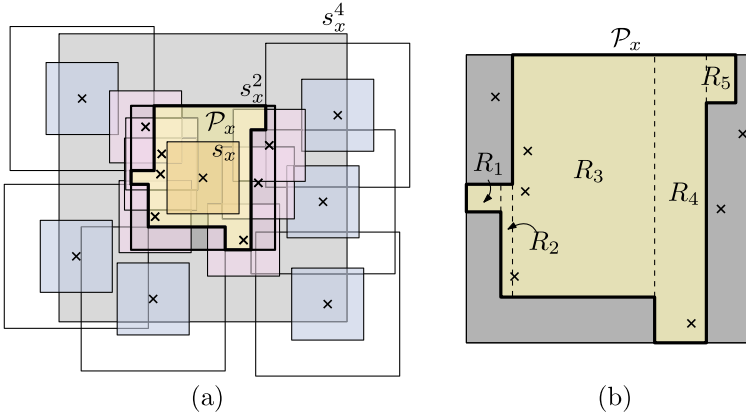


Fig. 1. Example for the deletion of a square  $s_x$ . (a) Square  $s_x$ , its neighborhood with centers in  $s_x^2$ , its 2-neighborhood with centers in  $s_x^4$ , and the polygon  $\mathcal{P}_x$ . (b) Vertical slab partition of  $\mathcal{P}_x$ .

another square in  $M_i$  and we set  $M_{i+1} = M_i$ . Otherwise, we add  $s_x$  to the current solution  $M_i$  and to the tree  $T(C(M_i))$  to obtain  $M_{i+1}$  and  $T(C(M_{i+1}))$ .

**DELETION:** When we delete a square  $s_x$  from  $\mathcal{S}_i$ , it is possible that  $s_x \in M_i$ . In this case, we may have to add squares from  $N(s_x)$  into  $M_{i+1}$  to keep it maximal. Since any square can have at most four independent neighbors, we can add in this step up to four squares to  $M_{i+1}$ .

First, we check if  $s_x \in M_i$ . If not, then we simply delete  $s_x$  from  $T(C_i)$  to get  $T(C_{i+1})$  and set  $M_{i+1} = M_i$ . Otherwise, we delete again  $s_x$  from  $T(C_i)$  and also from  $T(C(M_i))$ . To detect which neighbors of  $s_x$  can be added to  $M_i$ , we use suitable queries in the data structures  $T(C(M_i))$  and  $T(C_i)$ . Figure 1(a) illustrates the next observations. The centers of all neighbors in  $N(s_x)$  must be contained in the square  $s_x^2$ . But some of these neighbors may intersect other squares in  $M_i$ . In fact, these squares would by definition belong to the 2-neighborhood (i.e., be in the set  $Q_x = N^2(s_x) \cap M_i$ ). We can obtain  $Q_x$  by querying  $T(C(M_i))$  with the range  $s_x^4$ . Since  $s_x \in M_i$ , we know that no center point of squares in  $M_i$  lie in  $s_x^2$ . Hence, the center points of the squares in  $Q_x$  lie in the annulus  $s_x^4 - s_x^2$ . A simple packing argument (similar to the proof of Lemma 1) implies that  $|Q_x| \leq 12$ , and therefore querying  $T(C(M_i))$  will return at most 12 points.

Next, we define the rectilinear polygon  $\mathcal{P}_x = s_x^2 - \bigcup_{s_y \in Q_x} s_y^2$ , which contains all possible center points of squares that are neighbors of  $s_x$  but do not intersect any square  $s_y \in M_i \setminus \{s_x\}$ .

**OBSERVATION 2.** *The polygon  $\mathcal{P}_x$  has at most 28 corners.*

**PROOF.** We know that  $Q_x$  contains at most 12 squares  $s_y$ , for each of which we subtract  $s_y^2$  from  $s_x^2$ . Since all squares have the same side length, at most two new corners can be created in  $\mathcal{P}_x$  when subtracting a square  $s_y^2$ . Initially,  $\mathcal{P}_x$  had four corners, which yields the claimed bound of at most 28 corners.  $\square$

Next, we want to query  $T(C_i)$  with the range  $\mathcal{P}_x$ , which we do by vertically partitioning  $\mathcal{P}_x$  into rectangular slabs  $R_1, \dots, R_c$  for some  $c \leq 28$  (see Figure 1(b)). For each slab  $R_j$ , where  $1 \leq j \leq c$ , we perform a range query in  $T(C_i)$ . If a center  $p$  is returned, we can add the corresponding square  $s_p$  into  $M_{i+1}$ , and  $p$  into  $T(C(M_i))$  to obtain  $T(C(M_{i+1}))$ . Moreover, we have to update  $\mathcal{P}_x \leftarrow \mathcal{P}_x - s_p$ , refine the slab partition, and continue querying  $T(C_i)$  with the slabs of  $\mathcal{P}_x$ . Observe that after cutting the square  $s_p$  from the rectilinear polygon  $\mathcal{P}_x$ , the number of sides of the remaining region

of  $\mathcal{P}_x$  can be increased by at most 4. We know that the deleted square  $s_x$  can have at most four independent neighbors. So after adding at most four new squares to  $M_{i+1}$ , we know that there is no center point in  $C_i$  in the range  $\mathcal{P}_x$  and can stop searching.

LEMMA 3. *The set  $M_i$  is an MIS set of  $G_i = (S_i, E_i)$  for each step  $i \in [N]$ .*

PROOF. The correctness proof is inductive. By construction, the initial set  $M_1$  is an MIS for  $G_1$ . Let us consider some step  $i > 1$  and assume by induction that  $M_{i-1}$  is an MIS for  $G_{i-1}$ . If a new square  $s_x$  is inserted in step  $i$ , we add it to  $M_i$  if it does not intersect any other square in  $M_{i-1}$ ; otherwise, we keep  $M_{i-1}$ . In either case,  $M_i$  is an MIS of  $G_i$ . If a square  $s_x$  is deleted in step  $i$  and  $s_x \notin M_{i-1}$ , then  $M_i = M_{i-1}$  is an MIS of  $G_i$ . Finally, let  $s_x \in M_{i-1}$ . Assume for contradiction that  $M_i$  is not an MIS (i.e., some square  $s_q$  could be added to  $M_i$ ). Since  $M_{i-1}$  was an MIS,  $s_q \in N(s_x)$  and thus its center must lie in the region  $\mathcal{P}_x$ . But then we would have found  $s_q$  in our range queries with the slabs of  $\mathcal{P}_x$ . Hence,  $M_i$  is indeed an MIS of  $G_i$ .  $\square$

THEOREM 4. *We can maintain an MIS of a dynamic set of unit squares, deterministically, in  $O(\log n \log \log n)$  update time and  $O(n)$  space.*

PROOF. The correctness follows from Lemma 3. It remains to show the running time for the fully dynamic updates. At each step  $i$ , we perform either an INSERTION or a DELETION operation. Let us first discuss the update time for the insertion of a square. As described earlier, an insertion performs one or two insertions of the center of the square into the range trees and one range query in  $T(C(M_{i-1}))$ , which will return at most four points. Since we use the data structure of Mehlhorn and Näher [39], the update time for inserting a square is  $(\log n \log \log n)$ , which corresponds to the time required for inserting a new point into their range searching data structure and one range query. The deletion of a square triggers either just a single deletion from the range tree  $T(C_{i-1})$  or, if it was contained in the MIS  $M_{i-1}$ , two deletions, up to four insertions, and a sequence of range queries: one query in  $T(C(M_{i-1}))$ , which can return at most 12 points and a constant number of queries in  $T(C_{i-1})$  with the constant-complexity slab partition of  $\mathcal{P}_x$ . Note that although the number of points in  $\mathcal{P}_x$  can be large, for our purpose it is sufficient to return a single point in each query range if it is not empty. Therefore, the update time for a deletion is again  $O(\log n \log \log n)$  with dynamic fractional cascading [39].

In this approach, we maintain two dynamic range trees for the center points of rectangles. We use the dynamic range tree structure by Mehlhorn and Näher [39], whose space requirement is linear in the number of elements stored. Thus, the space requirement of this approach is  $O(n)$ .  $\square$

For unit square intersection graphs, recall that any square in an MIS can have at most four mutually independent neighbors. Therefore, maintaining a dynamic MIS immediately implies maintaining a dynamic 4-approximate MAX-IS.

Note that the update time of the dynamic data structure for orthogonal range queries dominates the update time of this algorithm. Its update time can be improved by using a state-of-the-art dynamic range query structure. The best-known dynamic data structure for orthogonal range reporting requires  $O(\log^{2/3+\epsilon} n)$  amortized update time, where  $\epsilon$  denotes an arbitrarily small positive constant, and  $O(k + \frac{\log n}{\log \log n})$  amortized query time, where  $k$  is the number of reported points [14]. From this, we conclude the following corollary.

COROLLARY 5. *We can maintain a 4-approximate MAX-IS of a dynamic set of unit squares, in amortized  $O(\frac{\log n}{\log \log n})$  update time.*

#### 4 APPROXIMATION ALGORITHMS FOR DYNAMIC MAX-IS

In this section, we study the MAX-IS problem for dynamic unit squares as well as for unit-height and arbitrary-width rectangles. In a series of dynamic schemes proposed in this section, we establish the trade-off between the update time and the solution size (i.e., the approximation factors). First, we design a 4-approximation algorithm with  $O(1)$  update time for MAX-IS of dynamic unit squares (Section 4.1). We generalize this to an algorithm that maintains a  $2(1 + \frac{1}{k})$ -approximate MAX-IS with  $O(k^2 \log n)$  update time, for any integer  $k > 1$  (Section 4.2). Finally, we conclude with an algorithm that deterministically maintains a 2-approximate MAX-IS with  $O(\log^2 n + \omega \log n)$  update time, where  $\omega$  is the maximum size of an independent set of the unit-height rectangles stabbed by any horizontal line (Section 4.3).

Let  $\mathcal{B}$  be a bounding square of the dynamic set of  $1 \times 1$ -unit squares  $\bigcup_{i \in [N]} \mathcal{S}_i$  of side length  $\sigma \times \sigma$ . Let  $H = \{h_1, \dots, h_\sigma\}$  and  $L = \{l_1, \dots, l_\sigma\}$  be a set of top-to-bottom and left-to-right ordered equidistant horizontal and vertical lines partitioning  $\mathcal{B}$  into a square grid of side-length-1 cells (Figure 2). Let  $E_H = \{h_i \in H \mid i \equiv 0 \pmod{2}\}$  and  $O_H = \{h_i \in H \mid i \equiv 1 \pmod{2}\}$  be the set of even and odd horizontal lines, respectively.

##### 4.1 4-Approximation Algorithm with Constant Update Time

We design a 4-approximation algorithm for the MAX-IS problem on dynamic unit square intersection graphs with constant update time. Our algorithm is based on a grid partitioning approach. Consider the square grid on  $\mathcal{B}$  induced by the sets  $H$  and  $L$  of horizontal and vertical lines. We denote the grid points as  $g_{p,q}$  for  $p, q \in [\sigma]$ , where  $g_{p,q}$  is the intersection point of lines  $h_p$  and  $l_q$ . We assign each unit square in any set  $\mathcal{S}_i$  to a grid point (denoted by its associated grid point) in the following deterministic way. Due to the unit grid construction, each unit square intersects at least one grid point. If an unit square  $s \in \mathcal{S}_i$  contains exactly one grid point  $g_{p,q}$ , we associate  $s$  with  $g_{p,q}$ . Otherwise, if  $s$  contains multiple grid points, we assign  $s$  to the top-leftmost grid point among others. For the ease of description, we may assume that the dynamic squares are in general positions—that is, each square contains exactly one grid point. Moreover, the preceding assignment does not affect the algorithm description or the analysis. For each  $g_{p,q}$ , we store a Boolean *activity value* 1 or 0 based on its intersection with  $\mathcal{S}_i$  (for any step  $i \in [N]$ ). If  $g_{p,q}$  intersects at least one square of  $\mathcal{S}_i$ , we say that it is *active* and set the value to 1; otherwise, we set the value to 0. Observe that for each grid point  $g_{p,q}$  and each timestep  $i$ , at most one square of  $\mathcal{S}_i$  intersecting  $g_{p,q}$  can be chosen in any MAX-IS. This holds because all squares that intersect the same grid point form a clique in  $G_i$ , and at most one square from a clique can be chosen in any independent set.

For each grid point  $g_{p,q}$ , for some  $p, q \in [\sigma]$ , we store the squares in  $S$  that intersect  $g_{p,q}$  into a list  $L_{pq}$ . Moreover, a counter for the size of  $L_{pq}$  is maintained dynamically for each grid point  $g_{p,q}$  such that we could detect if there exists at least one square intersecting  $g_{p,q}$  efficiently. Note that the set  $L_{pq}$  should allow constant time insertion and deletion (e.g., be stored in sequence containers like lists). We first initialize an independent set  $M_1$  for  $G_1 = (\mathcal{S}_1, E_1)$  with  $|M_1| \geq |OPT_1|/4$ . For each horizontal line  $h_j \in H$ , we compute two independent sets  $M_{h_j}^1$  and  $M_{h_j}^2$ , where  $M_{h_j}^1$  (respectively,  $M_{h_j}^2$ ) contains an arbitrary square intersecting each odd (respectively, even) grid point on  $h_j$ . Since every other grid point is omitted in these sets, any two selected squares are independent. Let  $M(h_j) = \arg \max\{|M_{h_j}^1|, |M_{h_j}^2|\}$  be the larger of the two independent sets. We define  $p(h_j) = |M_{h_j}^1|$  and  $q(h_j) = |M_{h_j}^2|$ , as well as  $c(h_j) = |M(h_j)| = \max\{p(h_j), q(h_j)\}$ .

We construct the independent sets  $M(E_H) = \bigcup_{j=1}^{\lfloor \sigma/2 \rfloor} (M(h_{2j}))$  for  $E_H$  and  $M(O_H) = \bigcup_{j=1}^{\lfloor \sigma/2 \rfloor} (M(h_{2j-1}))$  for  $O_H$ . We return  $M_1 = \arg \max\{|M(E_H)|, |M(O_H)|\}$  as the independent set for  $G_1$ . Figure 2 presents an illustration. The initialization of all  $O(\sigma^2)$  variables and the computation of






Fig. 2. Example instance with bounding square  $\mathcal{B}$  partitioned into a  $5 \times 5$  grid. Red squares represent the computed 4-approximate solution, which here is  $M(O(H))$ .

the first set  $M_1$  take  $O(\sigma^2)$  time. (Alternatively, a hash table would be more space efficient but could not provide the  $O(1)$ -update time guarantee.)

LEMMA 6. *The set  $M_1$  is an independent set of  $G_1 = (\mathcal{S}_1, E_1)$  with  $|M_1| \geq |OPT_1|/4$  and can be computed in  $O(\sigma^2)$  time.*

PROOF. Partition the squares  $S_1$  into 2 sets  $S_E, S_O$ , where  $S_E$  (respectively,  $S_O$ ) consists of all squares intersecting an even (respectively, odd) horizontal line. Let  $M_E$  and  $M_O$  be the MAX-IS of the  $S_E$  and  $S_O$ , respectively. Clearly, the larger one of these two sets contains at least half of as many elements as a MAX-IS of  $S_1$ . We may assume, w.l.o.g., that the  $M_E \geq |OPT_1|/2$ . For each even horizontal line  $h_j$ ,  $M_{h_j}^1$  (respectively,  $M_{h_j}^2$ ) is a MAX-IS of all rectangles stabbed on  $h_j$  and an odd (respectively, even) vertical line. The larger one of these two sets contains at least half as many elements as a MAX-IS of the squares stabbed by  $h_j$ . Overall, this implies that  $|M_1| \geq |OPT_1|/4$ .  $\square$

In the following step, when we move from  $G_i$  to  $G_{i+1}$ , for any  $i \in [N]$ , a square  $s_x$  is inserted into  $\mathcal{S}_i$  or deleted from  $\mathcal{S}_i$ . Let  $g_{pq}$  be the grid point contained in  $s_x$ . We update the list  $L_{pq}$  by either inserting the square  $s_x$  into  $L_{pq}$  or deleting  $s_x$  from  $L_{pq}$ . Moreover, we update the counter recording the size of  $L_{pq}$  and the activity value of the grid point accordingly. Intuitively, we check the activity value of the grid point that  $s_x$  intersects. If the update has no effect on its activity value, we keep  $M_{i+1} = M_i$ . Otherwise, we update the activity value, the corresponding cardinality counters, and report the solution accordingly. All of these operations can be performed in  $O(1)$  time.

A more detailed description of the INSERTION and DELETION operations is given in the following. When we move in the next step from  $G_i$  to  $G_{i+1}$  (for some  $1 \leq i < N$ ), we either insert a new square into  $\mathcal{S}_i$  or delete one square from  $\mathcal{S}_i$ . Let  $s_x$  be the square that is inserted or deleted, and let  $g_{u,v}$  (for some  $u, v \in [\sigma]$ ) be the grid point that intersects  $s_x$ . We next describe how to maintain a 4-approximate MAX-IS with constant update time. We distinguish between the two operations INSERTION and DELETION.

**INSERTION:** If  $g_{u,v}$  is active for  $\mathcal{S}_i$ , there is at least one square intersecting  $g_{u,v}$  that was considered while computing  $M_i$ . Hence, even if we would include  $s_x$  in a modified independent set  $M_{i+1}$ , it

would not make any impact on its cardinality. Hence, we simply set  $M_{i+1} \leftarrow M_i$ . Otherwise, we perform a series of update operations. First, change the activity value of  $g_{u,v}$  from 0 to 1. Second, include  $s_x$  in  $M_{h_u}^1$  (respectively,  $M_{h_u}^2$ ) if  $v$  is odd (respectively, even), and increase the value of  $p(h_u)$  (respectively,  $q(h_u)$ ) by 1. This lets us reevaluate the cardinality  $c(h_u)$  of  $M(h_u)$  in constant time. Third, reevaluate  $M(E_H)$  and  $M(O_H)$  and their cardinalities based on the updated value of  $c(h_u)$ . Note that none of these operations takes more than  $O(1)$  time.

**DELETION:** If there is a square  $s_l$  other than  $s_x$  intersecting  $g_{u,v}$ , then  $g_{u,v}$  stays *active*. We replace  $s_x$  by  $s_l$  in the maintained independent sets  $M_{h_u}^1, M_{h_u}^2, M(E_H)$  and  $M(O_H)$ . Note that this makes no impact on the cardinality of the sets  $M(E_H)$  and  $M(O_H)$ . If there is no other square intersecting  $g_{u,v}$ , we reset the activity value of  $g_{u,v}$  to *false*. Moreover, we delete  $s_x$  from maintained independent sets of line  $h_u$  and reevaluate  $M_E$  and  $M_O$ .

The update procedure described previously ensures that the respective cardinality maximization for the affected stabbing line  $h_j$  and finally  $M_i$  is reevaluated and updated. In this approach, we maintain an  $O(\sigma^2)$  grid. Each of  $n$  rectangles is stored in one of the grid points, thus the storage of rectangles is  $O(n)$ . Thereby, we conclude the following Lemma 7.

**LEMMA 7.** *The set  $M_i$  is an independent set of  $G_i = (\mathcal{S}_i, E_i)$  for each  $i \in [N]$  and  $|M_i| \geq |OPT_i|/4$  and  $O(\sigma^2 + n)$  space.*

**Running time.** We perform either an insertion or a deletion operation at every step  $i \in [N]$ . Both of these operations perform only local operations: (i) compute the grid point intersecting the updates square and check its activity value; (ii) reevaluate the values  $p(h_j)$  and  $q(h_j)$  of the horizontal line  $h_j$  intersecting the square—this may or may not flip the independent set  $M(h_j)$  and its cardinality from  $p(h_j)$  to  $q(h_j)$ , or vice versa; and (iii) finally, if the cardinality of  $M(h_j)$  changes, we reevaluate the sets  $M(E_H)$  and  $M(O_H)$ . All of these operations possibly change one activity value, increase or decrease at most three variables by 1 and perform at most two comparison operation. Therefore, the overall update process takes  $O(1)$  time in each step. Recall that the process to initialize the data structures for the set  $\mathcal{S}_1$  and to compute  $M_1$  for  $G_1$  takes  $O(\sigma^2)$  time.

Lemmas 6 and 7 and the preceding discussion of the  $O(1)$  update time yield the following theorem.

**THEOREM 8.** *We can maintain a 4-approximate MAX-IS in a dynamic unit square intersection graph, deterministically, in  $O(1)$  update time.*

#### 4.2 $2(1 + \frac{1}{k})$ -Approximation Algorithm with $O(k^2 \log n)$ Update Time

Next, we improve the approximation factor from 4 to  $2(1 + \frac{1}{k})$ , for any integer  $k > 1$ , by combining the shifting technique [33] with the insights gained from Section 4.1. This comes at the cost of an increase of the update time to  $O(k^2 \log n)$ , which illustrates the trade-off between solution quality and update time. We reuse the grid partition and some notations from Section 4.1. We first describe how to obtain a solution  $M_1$  for the initial graph  $G_1$  that is of size at least  $|OPT_1|/2(1 + \frac{1}{k})$  and then discuss how to maintain this under dynamic updates.

Let  $h_j \in H$  be a horizontal stabbing line, and let  $\mathcal{S}(h_j) \subseteq \mathcal{S}$  be the set of squares stabbed by  $h_j$ . Since they are all stabbed by  $h_j$ , the intersection graph of  $\mathcal{S}(h_j)$  is equivalent to the unit interval intersection graph obtained by projecting each unit square  $s_x \in \mathcal{S}(h_j)$  to a unit interval  $i_x$  on the line  $h_j$ ; we denote this set of unit intervals as  $I(h_j)$ . First, we sort the intervals in  $I(h_j)$  from left to right. Next, we define  $k + 1$  groups with respect to  $h_j$  that are formed by deleting those squares and their corresponding intervals from  $\mathcal{S}(h_j)$  and  $I(h_j)$ , respectively, that intersect every  $k + 1$ -th grid point on  $h_j$ , starting from some  $g_{j,\alpha}$  with  $\alpha \in [k + 1]$ . Now consider the  $k$  consecutive grid points on  $h_j$  between two deleted grid points in one such group, say,  $\{g_{j,\ell}, \dots, g_{j,\ell+k-1}\}$  for some




Fig. 3. Illustration of a group on line  $h_j$  for  $k = 3$  with the two subgroups  $I_1^3(h_j)$  and  $I_5^3(h_j)$ .

$\ell \in [\sigma]$ . Let  $I_\ell^k(h_j) \subseteq I(h_j)$  be the set of unit intervals intersecting the  $k$  grid points  $g_{j,\ell}$  to  $g_{j,\ell+k-1}$ . We refer to them as *subgroups*. Figure 3 presents an illustration. Observe that the maximum size of an independent set of each subgroup is at most  $k$ , since the width of each subgroup is strictly less than  $k + 1$  and each interval has unit length.

We compute  $M_1$  for  $G_1$  as follows. For each stabbing line  $h_j \in H$ , we form the  $k + 1$  different groups of  $I(h_j)$ . For each group, a MAX-IS is computed optimally and separately inside each subgroup. Since any two subgroups are horizontally separated and thus independent, we can then take the union of the independent sets of the subgroups to get an independent set for the entire group. This is done with the linear-time greedy algorithm to compute maximum independent set(s) for interval graphs [31]. Let  $\{M_{h_j}^1, \dots, M_{h_j}^{k+1}\}$  be  $k + 1$  maximum independent set(s) for the  $k + 1$  different groups, and let  $M(h_j) = \arg \max\{|M_{h_j}^1|, |M_{h_j}^2|, \dots, |M_{h_j}^{k+1}|\}$  be one with maximum size. We store its cardinality as  $c(h_j) = \max\{|M_{h_j}^i| \mid i \in [k + 1]\}$ . Next, we compute an independent set for  $E_H$ , denoted by  $M(E_H)$ , by composing it from the best solutions  $M(h_j)$  from the even stabbing lines—that is,  $M(E_H) = \bigcup_{j=1}^{\lfloor \sigma/2 \rfloor} M(h_{2j})$  and its cardinality  $|M(E_H)| = \sum_{j=1}^{\lfloor \sigma/2 \rfloor} c(h_{2j})$ . Similarly, we compute an independent set for  $O_H$  as  $M(O_H) = \bigcup_{j=1}^{\lfloor \sigma/2 \rfloor} M(h_{2j-1})$  and its cardinality  $|M(O_H)| = \sum_{j=1}^{\lfloor \sigma/2 \rfloor} c(h_{2j-1})$ . Finally, we return  $M_1 = \arg \max\{|M(E_H)|, |M(O_H)|\}$  as the solution for  $G_1$ .

**LEMMA 9.** *The independent set  $M_1$  of  $G_1 = (\mathcal{S}_1, E_1)$  can be computed in  $O(n \log n + kn)$  time and  $|M_1| \geq |OPT_1|/2(1 + \frac{1}{k})$ .*

**PROOF.** Let us begin with the analysis for one horizontal line, say,  $h_j$ . The objective is to show that for  $h_j$ , the size of our solution is least the optimum solution size for  $h_j$  divided by  $(1 + \frac{1}{k})$ . Recall that a group of  $\mathcal{S}(h_j)$  and  $I(h_j)$  is formed by deleting the squares and their corresponding intervals from  $S(h_j)$  and  $I(h_j)$ , respectively, which intersect every  $k + 1$ -th grid point on  $h_j$ , starting at some index  $\alpha \in [k + 1]$ . Now consider a hypothetical MAX-IS  $OPT(h_j)$  on  $h_j$ . By the pigeonhole principle, for at least one of the  $k + 1$  groups of  $\mathcal{S}(h_j)$  we deleted at most  $|OPT(h_j)|/(k + 1)$  squares from  $OPT(h_j)$ . Assume that this group corresponds to the independent set  $M_{h_j}^\ell$  for some  $\ell \in [k + 1]$ , which is maximum within each subgroup. Then we know that  $|M(h_j)| \geq |M_{h_j}^\ell| \geq |OPT(h_j)| - |OPT(h_j)|/(k + 1) = |OPT(h_j)|/(1 + \frac{1}{k})$ . Since this is true for each individual stabbing line  $h_j$  and since any two lines in  $E_H$  (or  $O_H$ ) are independent, this implies that  $|M(E_H)| \geq |OPT(E_H)|/(1 + \frac{1}{k})$  and  $|M(O_H)| \geq |OPT(O_H)|/(1 + \frac{1}{k})$ . Again by pigeonhole principle, if we choose  $M_1$  as the larger of the two independent sets  $M(E_H)$  and  $M(O_H)$ , then we lose by at most another factor of 2 (i.e.,  $|M_1| \geq |OPT_1|/2(1 + \frac{1}{k})$ ).

The algorithm requires  $O(n \log n)$  time to sort all intervals and then computes MAX-IS for the different subgroups with the linear-time greedy algorithm. Since each square belongs to at most  $k$  different subgroups, this takes  $O(kn)$  time in total.  $\square$

Next, we describe a pre-processing step, which is required for the dynamic updates.

**PRE-PROCESSING:** For each horizontal line  $h_j \in H$ , consider a group. For each subgroup  $I_\ell^k(h_j)$  (for some  $\ell \in [k + 1]$ ), we construct a balanced binary tree  $T(I_\ell^k(h_j))$  storing the intervals of  $I_\ell^k(h_j)$  in left-to-right order (indexed by their left endpoints) in the leaves. This process is done for each group of every horizontal line  $h_j \in H$ . This pre-processing step takes  $O(kn \log n)$  time.

When we perform the update step from  $G_i = (\mathcal{S}_i, E_i)$  to  $G_{i+1} = (\mathcal{S}_{i+1}, E_{i+1})$ , either a square is inserted into  $\mathcal{S}_i$  or deleted from  $\mathcal{S}_i$ . Let  $s_x$  and  $i_x$  be this square and its corresponding interval. Let  $g_{u,v}$  (for some  $u, v \in [\sigma]$ ) be the grid point that intersects  $s_x$ .

**INSERTION/DELETION:** The insertion or deletion of  $i_x$  affects all but one of the groups on line  $h_u$ . We describe the procedure for one such group on  $h_u$ ; it is then repeated for the other groups. In each group,  $i_x$  appears in exactly one subgroup and the other subgroups remain unaffected. This subgroup, say  $I_\ell^k(h_u)$ , is determined by the index  $v$  of the grid point  $g_{u,v}$  intersecting  $i_x$ . For each affected subgroups, we do the following update. First, we update the search tree  $T(I_\ell^k(h_u))$  of  $I_\ell^k(h_u)$  by inserting or deleting  $i_x$ , which can be done in  $O(\log n)$  time. Then, we recompute a MAX-IS of the subgroup with the greedy algorithm. Since the intervals of  $I_\ell^k(h_u)$  are sorted, we could locate the left-most interval that is to the right of all chosen intervals in  $O(\log n)$  time. Since a MAX-IS in each subgroup contains at most  $k$  intervals, the recomputation takes  $O(k \log n)$  time in each affected subgroup.

For all groups affected by the insertion or the deletion of  $i_x$ , we first update the corresponding independent sets  $M_{h_u}^p$  for  $p \in [k + 1]$ , whenever some updates of selected intervals are necessary. Then, we select the largest independent set of all  $k + 1$  groups as  $M(h_j)$  and update its new cardinality in  $c(h_j)$ . Finally, we update the independent sets  $M(E_H)$  and  $M(O_H)$  and their cardinalities and return  $M_{i+1} = \arg \max\{|M(E_H)|, |M(O_H)|\}$  as the solution for  $G_{i+1}$ .

Since the intervals computed by the left-to-right greedy algorithm are precisely those intervals that our update procedure selects, we get the following Lemma 10.

**LEMMA 10.** *The set  $M_i$  is an independent set of  $G_i = (\mathcal{S}_i, E_i)$  for each  $i \in [N]$  and  $|M_i| \geq |OPT_i|/2(1 + \frac{1}{k})$ .*

**PROOF.** The fact that  $M_i$  is an independent set follows directly from the construction. Let  $s_x$  be the square added or deleted, and let  $h_u$  be the grid horizontal line stabbed by  $s_x$ . In fact, our update algorithm constructs the same set of independent intervals as the one obtained by running from scratch the greedy MAX-IS algorithm on the set  $I_\ell^k(h_u)$ . The remaining arguments for the claimed approximation ratio of  $M_{i+1}$  are exactly the same as in the proof of Lemma 9.  $\square$

**Running time.** At every step, we perform either an INSERTION or a DELETION operation. Recall from the description of these two operations that an update affects a single stabbing line, say  $h_u$ , for which we have defined  $k + 1$  groups. Of those groups,  $k$  are affected by the update, but only inside a single subgroup. Updating a subgroup can trigger up to  $k$  selection updates, each taking  $O(\log n)$  time. In total, this yields an update time of  $O(k^2 \log n)$ .

This approach requires  $O(\sigma^2)$  space to maintain the grid. Note that each rectangle can be in the stored solutions of at most  $k$  subgroups, thus at most  $O(kn)$  storage is used. With Lemma 10 and the preceding update time discussion, we obtain the following theorem.

**THEOREM 11.** *We can maintain a  $2(1 + \frac{1}{k})$ -approximate MAX-IS in a dynamic unit square intersection graph, deterministically, in  $O(k^2 \log n)$  update time and  $O(\sigma^2 + kn)$  storage.*

### 4.3 2-Approximation Algorithm with $O(\log^2 n + \omega \log n)$ Update Time

We finally design a 2-approximation algorithm for the MAX-IS problem on dynamic axis-aligned unit-height but arbitrary-width rectangles. Note that the coordinates of the input rectangles might

not be integers. Let  $\mathcal{B}$  be the bounding box of the dynamic set of rectangles  $\tilde{\mathcal{R}} = \bigcup_{i \in [N]} \mathcal{R}_i$ . We begin by dividing  $\mathcal{B}$  into horizontal strips of height 1 defined by the set  $H = \{h_1, \dots, h_\sigma\}$  of  $\sigma = O(n)$  horizontal lines. We assume, w.l.o.g., that every rectangle in  $\tilde{\mathcal{R}}$  is stabbed by exactly one line in  $H$ . For a set of rectangles  $\mathcal{R}$ , we denote the subset stabbed by a line  $h_j$  as  $\mathcal{R}(h_j) \subseteq \mathcal{R}$ .

We first describe how to obtain an independent set  $M_1$  for the initial graph  $G_1 = (\mathcal{R}_1, E_1)$  such that  $|M_1| \geq |OPT_1|/2$  by using the following algorithm of Agarwal et al. [3]. For each horizontal line  $h_j \in H$ , we compute a MAX-IS for  $\mathcal{R}_1(h_j)$ . The set  $\mathcal{R}_i(h_j)$  (for any  $i \in [N]$  and  $j \in [\sigma]$ ) can again be seen as an interval graph. For a set of  $n$  intervals, a MAX-IS can be computed by a left-to-right greedy algorithm visiting the intervals in the order of their right endpoints in  $O(n \log n)$  time. So for each horizontal line  $h_j \in H$ , let  $M(h_j)$  be a MAX-IS of  $\mathcal{R}_1(h_j)$ , and let  $c(h_j) = |M(h_j)|$ . Then, we construct the independent set  $M(E_H) = \bigcup_{j=1}^{\lfloor \sigma/2 \rfloor} M(h_{2j})$  for  $E_H$ . Similarly, we construct the independent set  $M(O_H) = \bigcup_{j=1}^{\lfloor \sigma/2 \rfloor} M(h_{2j-1})$  for  $O_H$ . We return  $M_1 = \arg \max\{|M(E_H)|, |M(O_H)|\}$  as the independent set for  $G_1 = (\mathcal{R}_1, E_1)$ . Figure 4 presents an illustration.

**LEMMA 12 (THEOREM 2 [3]).** *The set  $M_1$  is an independent set of  $G_1 = (\mathcal{R}_1, E_1)$  with  $|M_1| \geq |OPT_1|/2$  and can be computed in  $O(n \log n)$  time.*

We describe the following pre-processing step to initialize in  $O(n \log n)$  time the data structures that are required for the subsequent dynamic updates.

**PRE-PROCESSING:** Consider a stabbing line  $h_j$  and the set of rectangles  $\mathcal{R}_i(h_j)$  stabbed on  $h_j$  for some  $i \in [N]$ . We denote the corresponding set of intervals as  $I(h_j)$ .

We build a balanced binary search tree  $T_l(I(h_j))$ , storing the intervals in  $I(h_j)$  in left-to-right order based on their left endpoints.

This is called the *left tree* of  $I(h_j)$ . For each internal node in the left tree, we associate it with an augmented balanced binary search tree storing the intervals in its subtree based on their right endpoints. Thus, we could get the interval with leftmost right endpoint in each subtree in constant time.

Such range tree-like data structure can be constructed in  $O(n \log n)$  and can be dynamically maintained in  $O(\log^2 n)$  time [50]. Additionally, we compute a MAX-IS of  $I(h_j)$  and store it in left-to-right order in a balanced binary search tree  $T_s(I(h_j))$ , denoted by the *solution tree* of  $h_j$ . Let  $\omega_j$  be the cardinality of a MAX-IS of  $I(h_j)$  for  $j \in [\sigma]$ , and let  $\omega = \max_j \omega_j$  for  $j \in [\sigma]$  be the maximum of these cardinalities over all stabbing lines.

When we move from  $G_i$  to  $G_{i+1}$  (for some  $1 \leq i < N$ ), either we insert a new rectangle into  $\mathcal{R}_i$  or delete one rectangle from  $\mathcal{R}_i$ . Let  $r_x$  be the rectangle that is inserted or deleted, let  $i_x$  be its corresponding interval, and let  $h_j$  (for some  $j \in [\sigma]$ ) be the horizontal line that intersects  $r_x$ . By maintaining the MAX-IS for  $I(h_j)$ , and then reevaluating the solution set, a 2-approximate MAX-IS can be maintained. Note that this recomputation costs  $O(\omega \log n + \log^2 n)$ , including updating the left tree of  $h_j$  and recomputing the MAX-IS for  $h_j$  by the greedy approach described in the pre-processing phase. In what follows, we describe how to maintain the MAX-IS for  $h_j$  dynamically in  $O(\omega \log n + \log^2 n)$  time.

Note that our update operation is faster in practice while having the same asymptotic worst-case update time  $O(n \log n)$  as recomputing the MAX-IS of  $h_j$ .

Given an interval  $i$ , we denote the left endpoint and right endpoint of  $i$  as  $l(i)$  and  $r(i)$ , respectively.

**INSERTION/DELETION:** Because the greedy algorithm for constructing the MAX-IS visits the intervals in left-to-right order based on the right endpoints, it would make the same decisions for all intervals with their right endpoint before the right endpoint of  $i_x$ . Let  $i_y$  be the rightmost

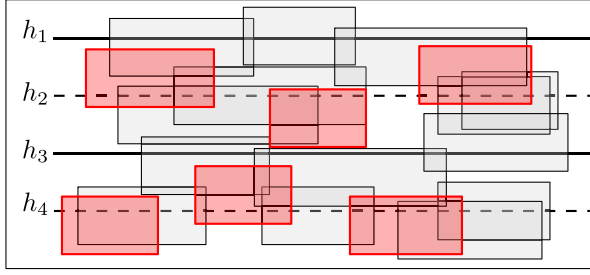


Fig. 4. Example instance with four horizontal lines. Red rectangles represent the computed 2-approximate solution, which here is  $M(E(H))$ .

interval in the current solution such that the right endpoint of  $i_y$  is before the right endpoint of  $i_x$  (i.e.,  $r(i_y) < r(i_x)$ ). We may assume that such interval  $i_y$  always exist by adding a dummy interval  $(0, \epsilon)$  in each solution tree  $T_s(I(h_j))$  for an arbitrary small value  $\epsilon$ . We could find  $i_y$  for  $i_x$  by querying the solution tree  $T_s(I(h_j))$  in  $O(\log \omega_j)$  time. Now we need to identify the next selected interval right of  $i_y$  that would have been found by the greedy algorithm. We use the left tree  $T_l(I(h_j))$  to search in  $O(\log n)$  time for the interval  $i'_z$  with leftmost right endpoint, whose left endpoint is right of the right endpoint  $r(i_y)$  of  $i_y$ . More precisely, we search for  $r(i_y)$  in  $T_l(I(h_j))$ , and whenever the search path branches into the left subtree, we compare whether the leftmost right endpoint stored in the root of the right subtree is left of the right endpoint of the current candidate interval. If so, we use this interval as the new candidate interval. Once a leaf is reached, the leftmost found candidate interval is the desired interval  $i'_z$ . This interval  $i'_z$  is precisely the first interval considered by the greedy algorithm after  $I_y$  and thus must be the next selected interval. We repeat the update process for  $i'_z$  as if it would have been the newly inserted interval until either  $i'_z$  is also selected in the previous solution or we reach the end of  $I(h_j)$ . We now reevaluate the new MAX-IS  $M(h_j)$  and its cardinality, which possibly affects  $M(E_H)$  or  $M(O_H)$ . We obtain the new independent set  $M_{i+1} = \arg \max\{|M(E_H)|, |M(O_H)|\}$  for  $G_{i+1} = (\mathcal{R}_{i+1}, E_{i+1})$ .

**Running time.** An update in the left tree (interval insertion/deletion) costs  $O(\log^2 n)$  time.

To update the solution of  $h_j$ , we perform at most  $\omega_j$  searches in  $T_l(I(h_j))$ , each of which takes  $O(\log n)$  time. Finally, we need to delete  $O(\omega_j)$  old selected intervals from and insert  $O(\omega_j)$  new selected intervals into the solution tree  $T_s(I(h_j))$ , each of which takes  $O(\log \omega_j)$  time. We now reevaluate the new MAX-IS  $M(h_j)$  and its cardinality  $c(h_j)$ , which possibly affects  $M(E_H)$  or  $M(O_H)$ . We obtain the new independent set  $M_{i+1} = \arg \max\{|M(E_H)|, |M(O_H)|\}$  for  $G_{i+1} = (\mathcal{R}_{i+1}, E_{i+1})$ . Overall, the total update time is  $O(\omega_j \log n + \log^2 n)$ .

LEMMA 13. *The set  $M_i$  is an independent set of  $G_i = (\mathcal{R}_i, E_i)$  for each  $i \in [N]$  and  $|M_i| \geq |OPT_i|/2$ .*

PROOF. We prove the lemma by induction. From Lemma 12, we know that  $M_1$  satisfies the claim, and in particular, each set  $M(h)$  for  $h \in H$  is a MAX-IS of the interval set  $I(h)$ . So let us consider the set  $M_i$  for  $i \geq 2$  and assume that  $M_{i-1}$  satisfies the claim by the induction hypothesis. Let  $r_x$  and  $i_x$  be the updated rectangle and its interval, and assume that it belongs to the stabbing line  $h_j$ . Then we know that for each  $h_k \in H$  with  $k \neq j$ , the set  $M(h_k)$  is not affected by the update to  $r_x$  and thus is a MAX-IS by the induction hypothesis. It remains to show that the update operations described earlier restore a MAX-IS  $M(h_j)$  for the set  $I(h_j)$ . But in fact the updates are designed in such a way that the resulting set of selected intervals is identical to the set of intervals that would be found by the greedy MAX-IS algorithm for  $I(h_j)$ . Therefore,  $M(h_j)$  is a MAX-IS for  $I(h_j)$ , and by the pigeonhole principle,  $|M_i| \geq |OPT_i|/2$ .  $\square$



Fig. 5. Example instance of a unit-height rectangle with width-height aspect ratio 5.

**Running time.** Each update of a rectangle  $r_x$  (and its interval  $i_x$ ) triggers either an INSERTION or a DELETION operation on the unique stabbing line of  $r_x$ . As we have argued in the description of these two update operations, the insertion or deletion of  $i_x$  requires one  $O(\log^2 n)$ -time update in the left tree data structure. If  $i_x$  is a selected independent interval, the update further triggers a sequence of at most  $\omega_j$  selection updates, each of which requires  $O(\log n)$  time. Hence, the update time is bounded by  $O(\log^2 n + \omega_j \log n) = O(\log^2 n + \omega \log n)$ . Recall that  $\omega_j$  and  $\omega$  are output-sensitive parameters describing the maximum size of an independent set of  $I(h)$  for a specific stabbing line  $h = h_j$  or any stabbing line  $h$ .

In this approach, we have to maintain  $O(\sigma)$  stabbing lines. For each stabbing line  $h_l$ , let  $n_l$  be the number of rectangles stabbed by  $h_l$ . For each stabbing  $h_l$ , we maintain a dynamic left tree for the  $n_l$  corresponding intervals, which requires  $O(n_l \log n_l)$  space [50]. Overall, the total space required is  $O(n \log n + \sigma)$ .

**THEOREM 14.** *We can maintain a 2-approximate MAX-IS in a dynamic unit-height arbitrary-width rectangle intersection graph, deterministically, in  $O(\log^2 n + \omega \log n)$  time and in  $O(n \log n + \sigma)$  space, where  $\omega$  is the maximum size of an independent set of the unit-height rectangles stabbed by any horizontal line.*

*Remark.* We note that Gavruskin et al. [28] gave a dynamic algorithm for maintaining a MAX-IS on *proper* interval graphs. Their algorithm runs in amortized time  $O(\log^2 n)$  for insertion and deletion, and  $O(\log n)$  for element-wise decision queries. The complexity to report a MAX-IS  $J$  is  $\Theta(|J|)$ . Whether the same result holds for general interval graphs was posed as an open problem [28]. Our algorithm in fact solves the MAX-IS problem on arbitrary dynamic interval graphs, which is of independent interest. Moreover, it explicitly maintains an exact MAX-IS at every step. Recently, Bhore et al. [11] showed that for intervals, a  $(1 + \epsilon)$ -approximate MAX-IS can be maintained with logarithmic worst-case update time, where  $\epsilon > 0$  is any positive constant.

## 5 EXPERIMENTS

We implemented all of our MAX-IS approximation algorithms presented in Sections 3 and 4 to empirically evaluate their trade-offs in terms of *solution quality* (i.e., the cardinality of the computed independent sets) and *update time* measured on a set of suitable synthetic and real-world map-labeling benchmark instances of two types of dynamic rectangle sets: (i) unit squares and (ii) rectangles of uniform height and bounded width-height integer aspect ratio (Figure 5). We believe that these two models are representative models in map labeling applications. The goal is to identify those algorithms that best balance the two performance criteria.

Moreover, for smaller benchmark instances with up to 2,000 squares, we compute exact MAX-IS solutions using a MAXSAT model by Klute et al. [36] that we solve with MaxHS 3.0 (see <http://www.maxhs.org>). These exact solutions allow us to evaluate the optimality gaps of the different algorithms in light of their worst-case approximation guarantees. Finally, we investigate the speedups gained by using our dynamic update algorithms compared to the baseline of recomputing new solutions from scratch with their respective static algorithm after each update.

## 5.1 Experimental Setup

**Implemented Algorithms.** We implemented the following six algorithms (and their greedy augmentation variants) in C++. We implemented two MIS algorithms for unit squares, *MIS-graph* and *MIS-ORS*, where an MIS is maintained dynamically. In both of these approaches, we use a dynamic orthogonal range searching data structure to check the intersections of squares. The main difference is that in *MIS-graph*, we maintain the geometric intersection graph explicitly and thus the update time is affected by the degree of the corresponding vertex. Note that in our implementation, the MIS algorithms *MIS-graph* and *MIS-ORS* compute the same MIS at each round and provide a 4-approximation. Moreover, we extended the approach *MIS-graph* for unit-height rectangles.

**MIS-graph for unit squares.** This is a naive graph-based dynamic MIS algorithm for unit squares, explicitly maintaining the square intersection graph and an MIS [4, Section 3]. To evaluate and compare the performance of our algorithm *MIS-ORS* (Section 3) for the MIS problem, we implemented this alternative dynamic algorithm as the baseline approach. This algorithm maintains the current instance in a dynamic geometric data structure and maintains the square intersection graph explicitly. We use standard adjacency lists to represent the intersection graph, implemented as unordered sets in C++.

In the initialization step, we store the center points of all unit squares in the dynamic point range query structure<sup>2</sup> implemented in CGAL (version 5.2.1). By performing neighbor searches in this range search structure, we build the initial geometric intersection graph. Precisely, for each square  $s_x$ , we query the orthogonal range tree with the range  $s_x^2$ , where  $s_x^2$  is the square of side length 2 concentric with  $s_x$ . This range contains the center points of all squares that intersect  $s_x$ . Now, to obtain an MIS at the first step, we add the first (unmarked) vertex  $v$  to the solution and mark  $N(v)$  in the corresponding intersection graph. This process is repeated iteratively until there is no unmarked vertex left in the intersection graph. Clearly, by following this greedy method, we obtain an MIS.

Moreover, for each vertex  $v$ , we maintain an augmenting *counter* that stores the number of vertices from its neighborhood  $N(v)$  that are contained in the current MIS. Note that in our implementation, the approach greedily checks the vertices in their ordering as given in the input file.

This approach handles the updates in a straightforward manner. To add a new square, we first insert its center point into the dynamic orthogonal range query data structure and add a new vertex in the intersection graph. When a new vertex is inserted, its corresponding square may introduce new intersections. Therefore, when adding a vertex, we also determine the edges that are required to be added to the intersection graph. Notice that unlike the canonical vertex update operation defined in the literature, where the adjacencies of the new vertex are part of the dynamic update, here we actually need to figure out the neighborhood of a vertex. Let  $s_x$  be the square to add, and let  $v_x$  be its corresponding vertex, which is added to the intersection graph. To find all squares that overlap  $s_x$ , we query the orthogonal range tree with  $s_x^2$ . This output-sensitive operation takes  $O(\log n + \deg(v))$  time, where  $\deg(v)$  is the size of neighborhood of this newly added vertex  $v$ . If the newly inserted square has no intersection with any square from the current solution, then we simply add its vertex to the solution; otherwise, we ignore it. Finally, we update the counters. If a vertex is deleted, we update the orthogonal range query structure and the intersection graph by deleting its corresponding center point and vertex, respectively. If the deleted vertex was in the solution, then we decrease the counters of its neighbors by 1. Once the counter of a vertex is updated to

<sup>2</sup>2D Range and Neighbor Search in CGAL: [https://doc.cgal.org/latest/Point\\_set\\_2/index.html](https://doc.cgal.org/latest/Point_set_2/index.html).



0, we add this vertex into the solution. Both the insertion (after computing  $N(v)$ ) and deletion operation for a vertex  $v$  take  $O(\deg(v))$  time each to update the intersection graph and the MIS solution. By maintaining the conflict graph in an adjacency list, the space requirement of this implemented approach is  $O(n^2)$ .

**MIS-graph for unit-height rectangles.** We extend the approach *MIS-graph* for axis-parallel rectangles of unit height and with bounded integer width-height aspect ratio. Let  $r$  be an axis-parallel rectangle with unit height and aspect ratio  $w$  for an integer  $w$ . The rectangle  $r$  can be partitioned into  $w$  unit squares. Each corner of these partitioning unit squares is denoted as a *witness point* of  $r$  in the following. In this extended *MIS-graph* approach, instead of maintaining the center points of squares in the orthogonal range query data structure as in *MIS-graph*, all of the witness points of all rectangles are stored in the range query data structure. Furthermore, we use a hash table, which assigns each witness point to its corresponding rectangle. We assume a general position property of the dynamic set of rectangles in the input: each pair of rectangles intersects at most twice at their boundaries. This property is denoted as corner intersection property in the following.<sup>3</sup> Thus, given a rectangle  $r$ , we could find all rectangles that overlap with  $r$  by making a range query of  $r$  in the dynamic orthogonal range tree. Note that in our input instances, the width-height aspect ratio is from  $\{1, \dots, b\}$  for a positive constant integer  $b$ . Thus, the initialization takes  $(bn \log bn)$  time. Let  $r_x$  be the rectangle to add or delete, and let  $v_x$  be its corresponding vertex in the conflict graph. Both the insertion and deletion (after updating the range tree and conflict graph) take  $O(\deg(v_x))$ , where  $\deg(v_x)$  is the size of the neighborhood of the vertex  $v_x$  in the conflict graph.

**MIS-ORS.** The dynamic MIS algorithm is based on orthogonal range searching (Section 3); this algorithm provides a 4-approximation. In the implementation, we used the dynamic orthogonal range searching data structure implemented in CGAL (version 5.2.1), which is based on a dynamic Delaunay triangulation [40, Chapter 10.6]. More precisely, given a rectilinear area  $R$ , a range query with  $R$  is implemented in CGAL by making a circular range query with its circumscribed circle and then check if the reported points are in  $R$ . That means that finding and reporting one point of a rectilinear area take in the worst-case  $O(n)$ . Hence, this implementation does not provide the polylogarithmic worst-case update time of Theorem 4. However, we did not observe such behavior in most of our experiments. Note that the initial solution is computed greedily in the ordering of the instance file, and in each update round the MIS is maintained dynamically. Overall, the solution computed in each round by this approach is identical as the solution computed by *MIS-graph*.

We implement MIS algorithms *grid* (Section 4.1), *grid-k* (Section 4.2), and *line* (Section 4.3) and their greedy augmentation variants in C++. Since all of these algorithms are based on partitioning the set of squares and considering only sufficiently segregated subsets, they produce a lot of white space in practice.

For instance, they ignore the squares stabbed by either all of the even or all of the odd stabbing lines completely to create isolated subinstances. In practice, it is therefore interesting to augment the computed approximate MAX-IS by greedily adding independent but initially discarded squares. We also implemented the greedy variants of these algorithms, which are denoted as *g-grid*, *g-grid-k*, and *g-line*.

**grid.** Recall that in the grid-based 4-approximation algorithm for unit squares (Section 4.1), we either omit all squares intersecting even horizontal lines or all squares intersecting odd horizontal lines. Then, for each horizontal line  $l_i$ , we maintain a Boolean value based on

<sup>3</sup>In computational geometry, a set of geometric objects with such property is denoted as a family of pseudo-disks.




Fig. 6. Example instance within a  $4 \times 4$  grid. The solution obtained by the *g-grid* approach is the union of the solution  $M^1(h_1) \cup M^2(h_3)$  (red) by the *grid* approach, the line augmentation set of  $M^1(h_1)$  (violet), the line augmentation set of  $M^3(h_3)$  ( $= \emptyset$ ), and the combination augmentation set of  $M^1(h_1) \cup M^2(h_3)$  (green).

whether we omit all rectangles intersecting the odd vertical lines or all rectangles intersecting the even vertical lines.

In our implementation, to obtain the initial solution, we iterate over all grid points that are not omitted and collect the first square stored in their square lists. Let this be our initial solution.

***g-grid***. This is the grid-based approach with greedy augmentation. We describe the greedy augmenting procedure. Recall that for each horizontal line  $h_j$ , we maintain two candidate independent sets  $M_{h_j}^1$  and  $M_{h_j}^2$  of rectangles of odd grid points and rectangles of even grid points of  $h_j$ , respectively. Given a horizontal grid line  $h_j$ , we check all rectangles intersecting even (respectively, odd) grid points on  $h_j$  and compute an augmentation for  $M_{h_j}^1$  (respectively,  $M_{h_j}^2$ ), which is denoted as a *line augmentation set*. For three consecutive horizontal grid lines  $h_j, h_{j+1}$  and  $h_{j+2}$  in the grid partitioning, we consider the four possible combinations of one candidate set of  $h_j$  and one candidate set of  $h_{j+2}$  with their corresponding line augmentation sets. For each combination, we compute a greedy augmentation of rectangles stabbed by line  $h_{j+1}$ , denoted as a *combination augmentation set*. To achieve a constant-time update, each of these computed augmentation sets is stored in a vector of size  $\sigma$  such that the rectangle intersecting the  $k$ -th vertical grid line is stored as the  $k$ -th element of the vector. Overall, for each horizontal grid line  $h_j$ , we maintain six greedy augmentation sets: four combination augmentation sets for the four possible combinations of  $h_{j-1}$  and  $h_{j+1}$  as well as two line augmentation set for  $M_{h_j}^1$  and  $M_{h_j}^2$ , respectively. Thus, the initial solution obtained by this approach contains three parts: the initial solution obtained by the *grid* approach, the line augmentation sets for the chosen candidate sets of the odd/even (based on the choice by *grid* approach) horizontal lines, and the combination augmentation sets from the omitted horizontal lines based on the choices of its two neighboring lines (Figure 6).

Whenever a square  $s$  is inserted on the horizontal grid line  $h_r$ , the two candidate sets  $M_{h_r}^1$  and  $M_{h_r}^2$  of  $h_r$  might be updated by adding  $s$ . Then we update all greedy augmentation sets involving  $h_r$  by removing the rectangles intersecting  $s$ . Note that we only need to check the rectangles on a neighboring grid point of the grid point of  $s$ . Thus, to update one greedy

augmentation set takes constant time. When removing a square  $s$  from  $h_r$ , we remove  $s$  from all greedy augmentation sets of  $h_r$ . Note that this greedy augmentation procedure does not affect the approximation bound (i.e., 4-approximation) since the solution of the *grid* approach is unaffected. Moreover, this procedure takes constant update time and  $O(\sigma^2)$  space, which is the same as for the *grid* approach.

**grid- $k$ .** This is the shifting-based  $2(1 + \frac{1}{k})$ -approximation algorithm (Section 4.2). In the experiments, we use  $k = 2$  (i.e., a 3-approximation) and  $k = 4$  (i.e., a 2.5-approximation).

**g-grid- $k$ .** We describe the greedy augmented version of the *grid- $k$*  approach and generalize the idea of the *g-grid* approach. Recall that for each horizontal grid line  $h_j$ , there is one candidate independent set for each of  $k + 1$  group. In the initialization phase, for each candidate set of  $h_j$ , we compute its augmented set consisting of rectangles of  $h_j$ , which is denoted as the *line augmentation set* of the candidate set. For every three consecutive horizontal grid lines  $h_j, h_{j+1}, h_{j+2}$ , we consider  $k^2$  unions of one candidate set of  $h_j$  and one candidate set of  $h_{j+1}$  with their corresponding line augmentation sets and compute for each union an augmented set of rectangles on  $h_{j+1}$  greedily. For each computed augmented set, we store it in a vector of size  $\sigma$  such that the index of each stored rectangle is identical to the index of its vertical grid point. Thus, the initial solution obtained by this approach contains three parts: the initial solution obtained by the *grid- $k$*  approach, greedy augmented rectangles for the chosen candidate sets of the odd/even (based on the choice by *grid- $k$*  approach) horizontal lines, and the greedy augmented rectangles from the omitted horizontal lines based on the choices of its two neighboring lines. Whenever a square is inserted or deleted in an update phase, the involved  $O(k^2)$  greedy augmentation sets need to be updated accordingly. Overall, this greedy augmented version of the approach *grid- $k$*  retains the same approximation ratio  $2(1 + \frac{1}{k})$ , the same update time  $O(k^2 \log n)$ , and the same space requirement  $O(\sigma^2 + kn)$  as the *grid- $k$*  approach.

**line.** This is the stabbing line-based 2-approximation algorithm (Section 4.3).

**g-line.** We describe the greedy augmentation procedure of the *line* approach. Recall that *line* maintains a candidate set, which is a MAX-IS, for each stabbing line, and we omit either all rectangles stabbed by even lines or by odd lines. In the initialization phase of this greedy-augmented approach, we first compute the candidate set for each horizontal line as in the *line* approach. For every three consecutive horizontal grid lines  $h_j, h_{j+1}, h_{j+2}$ , we consider the union of the candidate set of  $h_j$  and the candidate set of  $h_{j+1}$  and compute for this union an augmentation set of rectangles on  $h_{j+1}$  greedily. Each greedy augmentation set is sorted from left to right and stored in an ordered set. The solution obtained by this approach consists of the solution obtained by *line* and the greedy augmented sets of the omitted lines. When a square is inserted into or removed from a stabbing line  $h_j$ , we first update the candidate set of  $h_j$  and then update the greedy augmentation sets of the stabbing line above and the stabbing line below  $h_j$ , if needed. More precisely, when the candidate set of  $h_j$  is updated, we mark the leftmost and the rightmost rectangles  $r_s, r_t$ , which are the newly added squares in this candidate set. Then, when we update a greedy augmentation set, we first find the leftmost rectangle in this greedy augmentation set that is to the left of the left endpoint of  $r_s$  and recompute the greedy augmentation set from this rectangle. This recomputing procedure terminates when the chosen rectangle is in the greedy augmentation set from previous round and is right to the right endpoint of  $r_t$ . This means that to update the greedy augmentation set of one stabbing line  $h_j$ , it takes in the worst-case  $O(n_j)$  time, where  $n_j$  is the number of rectangles stabbed by  $h_j$ . Thus, this greedy augmentation takes worst-case  $O(n)$  time for an update. However, we note that such worst-case behavior was not observed in the experiments.




Fig. 7. Example distributions of small real-world instances.

Table 1. Specification of the Six OSM Unit Square Instances

	Post-CH	Viewpoint-AT	Hotels-CH	Hotels-AT	Peaks-CH	Hamlets-CH
Features ( $n$ )	646	652	1,788	2,209	4,320	4,326
Overlaps ( $m$ )	5,376	5,418	28,124	68,985	107,372	159,270
Density ( $m/n$ )	8.32	8.31	15.73	31.23	24.85	36.92

**System Specifications.** The experiments were run on a server equipped with two Intel Xeon E5-2640 v4 processors (2.4-GHz 10-core) and 160 GB of RAM. The machine ran the 64-bit version of Ubuntu Bionic (18.04.2 LTS). The code was compiled using g++ 7.5.0 with optimization level O3.

**Benchmark Data of Unit Squares.** We created three types of benchmark instances. The two synthetic datasets consist of  $n$   $30 \times 30$ -pixel squares placed inside a bounding rectangle  $\mathcal{B}$  of size  $1,080 \times 720$  pixels, which also creates different densities. The real-world instances use the same square size but geographic feature distributions. For the updates, we consider three models: *insertion-only*, *deletion-only*, and *mixed*, where the latter selects insertion or deletion uniformly at random. The new squares to insert are generated uniformly.

**Gaussian.** In the Gaussian model, we generate  $n$  squares randomly in  $\mathcal{B}$  according to an overlay of three Gaussian distributions, where 70% of the squares are from the first distribution, 20% from the second one, and 10% from the third one. First, the three means of the distributions are sampled uniformly at random in  $\mathcal{B}$ ; for each Gaussian, the standard deviation is set to 100 in both dimensions. Next, in each distribution, we sample a point and check whether the unit square centered at this point is inside the bounding box  $\mathcal{B}$ . If so, we add this point to the input; otherwise, we discard it. This process is repeated until the required number of rectangles of each distribution is collected.

**Uniform.** In the uniform model, we generate  $n$  squares in  $\mathcal{B}$  uniformly at random.

**Real world.** We created six real-world datasets by extracting point features from OpenStreetMap (OSM) (Table 1 presents their detailed properties). Note that in our experiment, the input set is from a real-world instance and the new squares to insert are generated by an uniform sampling. In Figure 7, we illustrate the distribution of point features in our small real-world instances.

**Benchmark Data of Unit-Height Rectangles.** Similarly to the unit squares data, we created three types of benchmark instances with unit-height rectangles. Since the approach *MIS-graph*

Table 2. Specification of the Six OSM Unit-Height Rectangle Instances

	Post-CH	Viewpoint-AT	Hotels-CH	Hotels-AT	Peaks-CH	Hamlets-CH
Features ( $n$ )	612	918	1,476	1,755	2,899	2,930
Overlaps ( $m$ )	4,688	6,710	19,886	46,621	54,912	64,098
Density ( $m/n$ )	7.66	7.3	13.47	18.56	18.94	21.87
Label lengths (range/aver.)	3–50/18	3–63/16	2–46/15	2–57/17	3–63/12	3–29/10

requires the corner intersection property of the input set, we created instances such that the boundary of every pair of rectangles intersects at most twice. The synthetic datasets consist of unit-height rectangles with height of 10 pixels placed inside a bounding rectangle  $\mathcal{B}$  of size  $1,080 \times 720$  pixels, which also creates different densities. The real-world instances use the geographic feature distributions and the original label texts extracted from the OSM files. Each label of length  $l$  is represented as a  $10l \times 10$ -pixel rectangle placed inside the bounding rectangle  $\mathcal{B}$ . For synthetic datasets, we sample a label length for each rectangle between 2 and 21 based on the distribution of word lengths in English.<sup>4</sup>

**Gaussian.** In the Gaussian model, we generate center points of  $n$  unit-height rectangles randomly in  $\mathcal{B}$  according to an overlay of three Gaussian distributions, where 70% of the center points are from the first distribution, 20% from the second one, and 10% from the third one. The means that they are sampled uniformly at random in  $\mathcal{B}$  and the standard deviation is 100 in both dimensions. Once each center point is sampled, we sample a label length for it based on the distribution of word lengths in English. Whenever a new rectangle candidate is generated, we check if it is inside the bounding box  $\mathcal{B}$ . Furthermore, we verify if the rectangles set keeps the corner intersection property after adding this newly generated rectangle. This process is repeated for each distribution until the number of collected rectangles reaches the corresponding required number for this distribution.

**Uniform.** In the uniform model, we generate  $n$  unit-height rectangles in  $\mathcal{B}$  uniformly at random. To guarantee the corner intersection property, we check each newly generated rectangle.

**Real world.** We created six real-world datasets by extracting point features from OSM (see Table 1 for their detailed properties). To guarantee the corner intersection property of the instances, we filter the the instances in a post-processing step (see Table 2 for their detailed properties).

The source code of benchmark instance generator is available at <https://dyna-mis.github.io/dynaMIS/>.

## 5.2 Experimental Results for Unit Squares

**Time-quality trade-offs.** For our first set of experiments we compare the five implemented algorithms, including their greedy variants, in terms of update time and size of the computed independent sets. Figure 8 shows scatter plots of runtime vs. solution size on uniform and Gaussian benchmarks, where algorithms with dots in the top-left corner perform well in both measures.

We first consider the results for the uniform instances with  $n = 10,000$  squares in the top row of Figure 8. Each algorithm performed  $N = 400$  updates, either insertions (Figure 8(a)) or deletions (Figure 8(b)), and each update is shown as one point in the respective color. Both plots show that the two MIS algorithms compute the best solutions with almost the same size and well ahead of the rest. The MIS algorithm *MIS-ORS* is clearly faster than *MIS-graph* on both insertions and deletions. The approximation algorithms *grid*, *grid-2*, *grid-4*, and *line* (without the greedy optimizations) show

<sup>4</sup>Taken from <http://www.ravi.io/language-word-lengths>.

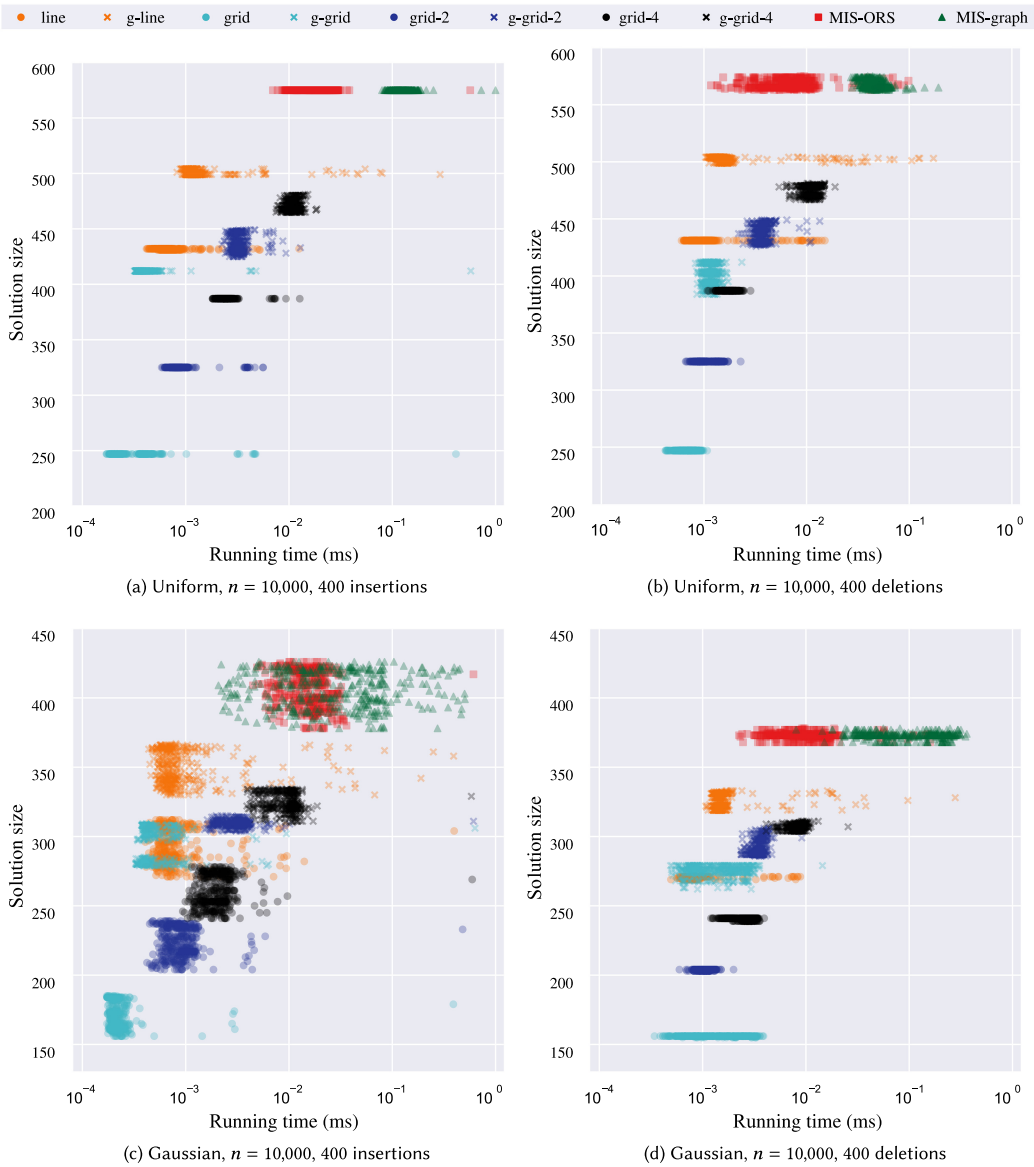


Fig. 8. Time-quality scatter plots for synthetic benchmark instances. The x-axis (log-scale) shows runtime, and the y-axis shows the solution size. We use semi-transparent markers in the scatter plots.

their predicted relative behavior: the better the solution quality, the worse the update times. Algorithms *line* and *g-line* show a wide range of update times, spanning almost two orders of magnitude. Adding the greedy optimization drastically improves the solution quality in all cases but typically at the cost of higher runtimes. For *g-grid-k*, the algorithms get slower by an order of magnitude and increase the solution size by 30% to 50%. For *g-grid*, the additional runtime is not as significant (but deletions are slower than insertions), and the solution size almost doubles. Finally, for *g-line*, the additional runtime is not as significant and reaches the best quality among the approximation algorithms with about 90% of the MIS solutions, but faster by one order of magnitude.

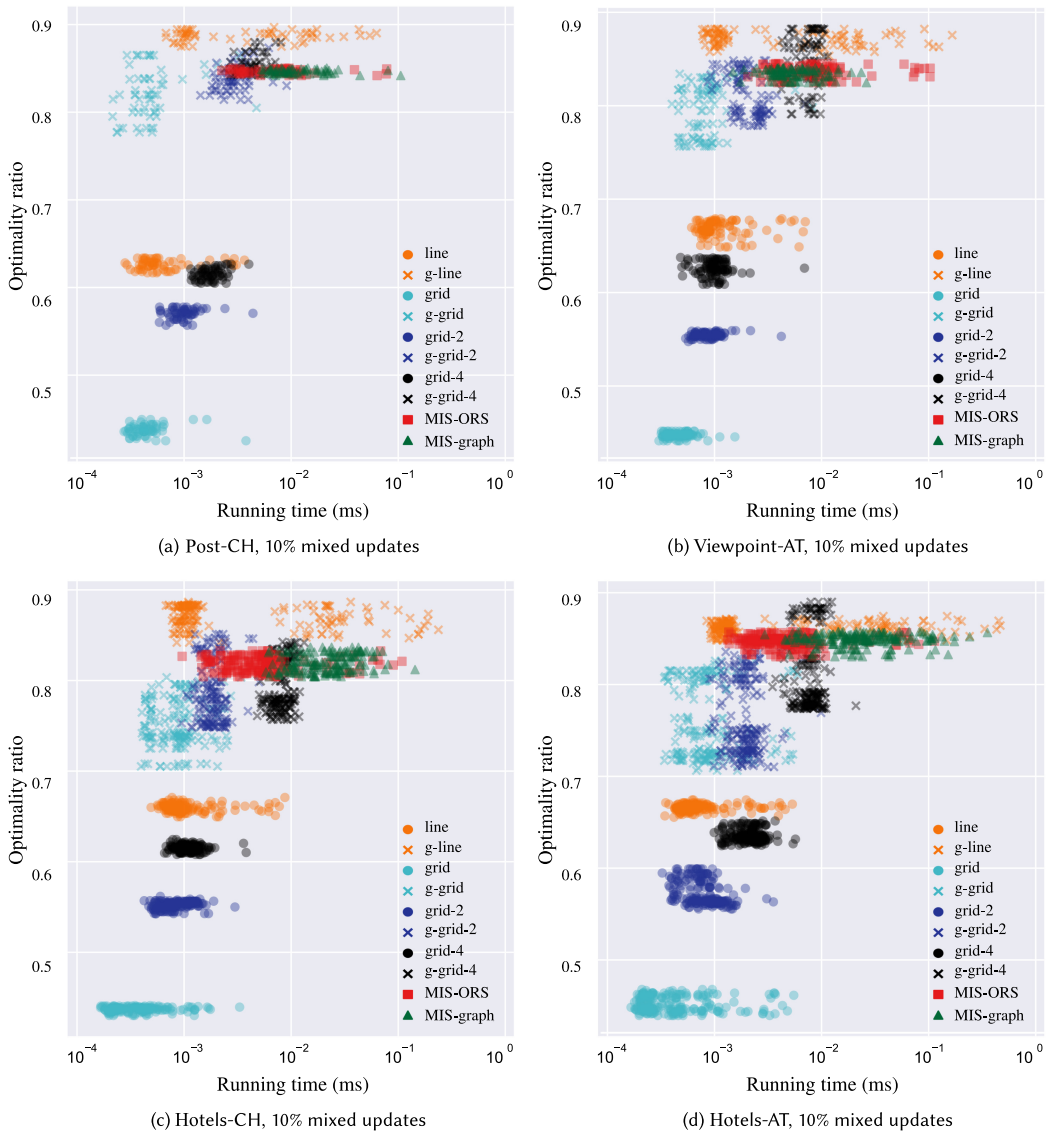


Fig. 9. Time-quality scatter plots for the small OSM instances. The x-axis (log-scale) shows runtime. The y-axis shows the quality ratio compared to an optimal MAX-IS solution.

For the results of the Gaussian instances with  $n = 10,000$  squares and  $N = 400$  updates plotted in Figure 8(c) (insertions) and 8(d) (deletions), we observe the same ranking between the different algorithms. However, due to the non-uniform distribution of squares, the solution sizes are more varying, especially for the insertions. For the deletions, it is interesting to see that *grid* and *MIS-graph* have more strongly varying runtimes, which is in contrast to the deletions in the uniform instance, possibly due to the dependence on the vertex degree. The best solutions are computed by *MIS-ORS* and *MIS-graph*. Regarding the runtime, *MIS-ORS* has more homogeneous update times ranging between the extrema of *MIS-graph*, whereas they are comparably fast for insertions on average.

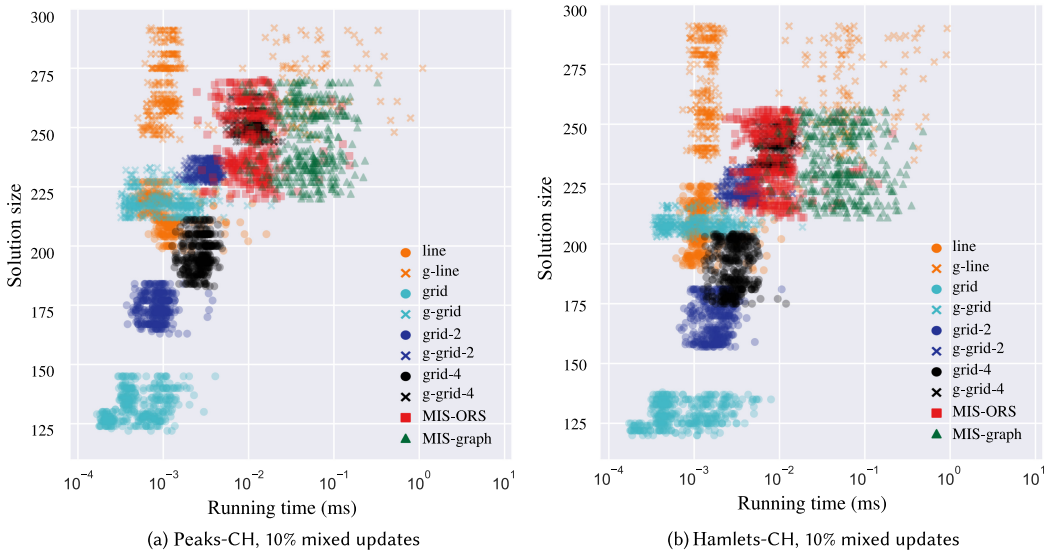


Fig. 10. Time-quality scatter plots for the large OSM instances. The x-axis (log-scale) shows runtime. The y-axis shows the quality ratio compared to the optimal Max-IS solution size.

Algorithm *g-line* again reaches nearly 90% of the quality of the MIS algorithms, with a speedup almost one order of magnitude.

**Optimality gaps.** Next, let us look at the results of the real-world instances in Figure 9 and in Figure 10. The first four instances in Figure 9 were small enough so that we could compute each MAX-IS exactly with MaxHS and compare the solutions of the approximation algorithms with the optimum on the y-axis. The largest two instances in Figure 10 plot the solution size on the y-axis. First, let us consider Figure 9(c) as a representative, which is based on a data set of 1,788 hotels and hostels in Switzerland with mixed updates of 10% of the squares ( $N = 179$ ). Generally speaking, the results of the different algorithms are much more overlapping in terms of quality than for the synthetic instances. The plot shows that the MIS algorithms reach consistently between 80% and 85% of the optimum but are sometimes outperformed by the greedy-augmented approaches. Interestingly, *g-line*, the best of the approximation algorithms with greedy augmentation, contributes consistently the best solutions. Regarding the runtime, *MIS-ORS* has generally faster update time than the *MIS-graph* approach. The original approximations are well above their respective worst-case ratios but stay between 45% and 65% of the optimum. The greedy extensions push this toward larger solutions, at the cost of higher runtimes. However, *g-line* seems to provide a very good balance between quality and speed.

Let us next consider the largest OSM instance in Figure 10(b). It again reflects the same findings as obtained from the smaller instances. The instance consists of  $n = 4,326$  hamlets in Switzerland with 10% mixed updates ( $N = 433$ ) and is denser by a factor of about 2.3 than hotels-CH (see Table 1). There is quite some overlap of the different algorithms in terms of the solution size, yet the algorithms form the same general ranking pattern as observed before. The approach *g-line* contributes the best solutions in most of the rounds. Moreover, regarding the running time, *g-line* is again about nearly an order of magnitude faster than the MIS algorithms, except for a few slower outliers. Comparing the two MIS approaches, *MIS-ORS* is significantly faster than *MIS-graph*.



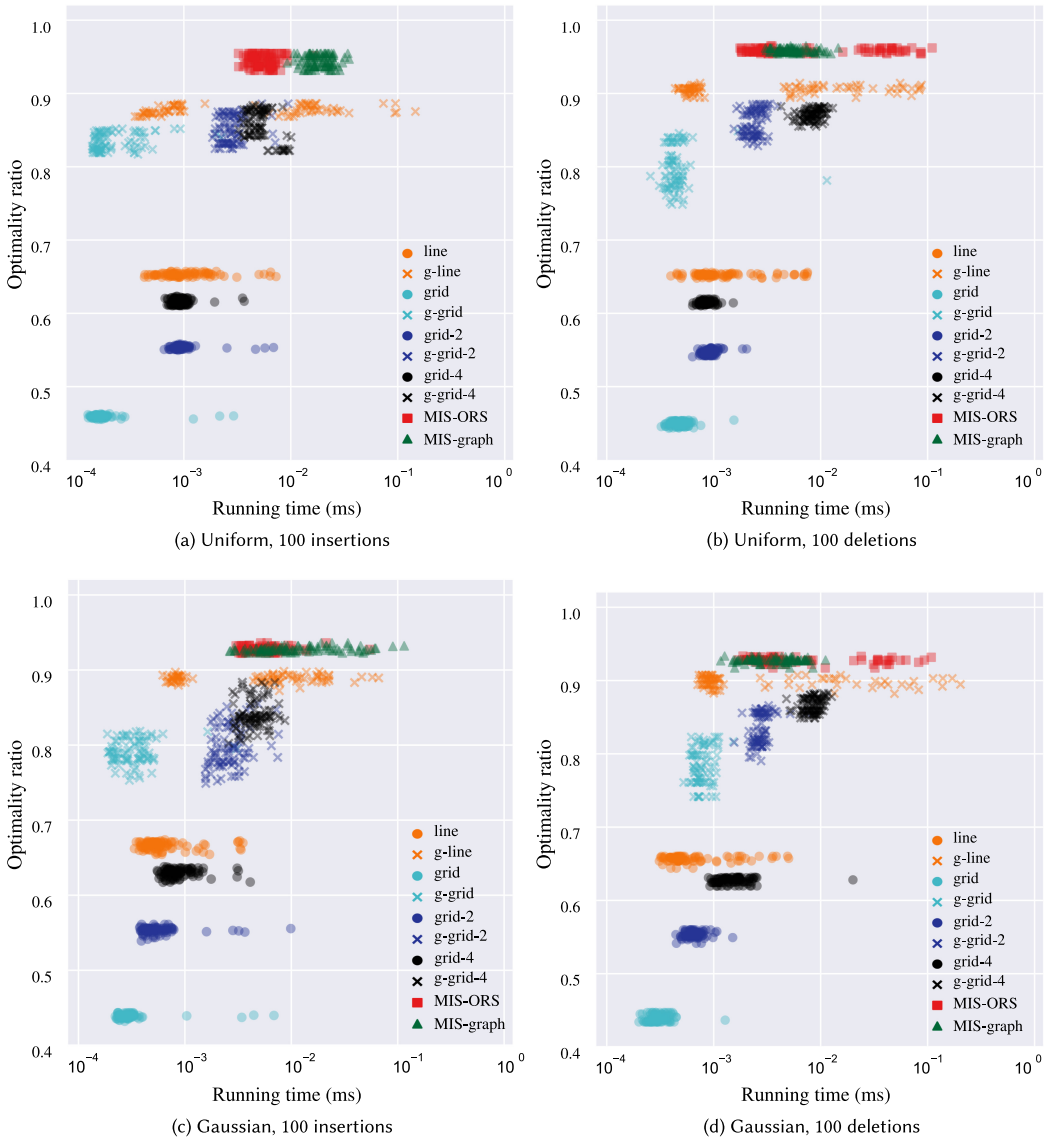


Fig. 11. Time-quality scatter plots for uniform and Gaussian instances with  $n = 1,000$  squares. The  $x$ -axis (log-scale) shows runtime. The  $y$ -axis shows the quality ratio compared to an optimal MAX-IS solution.

Finally, Figure 11 shows the optimality ratios of the algorithms for small uniform and Gaussian instances with  $n = 1,000$  squares. They confirm our earlier observations but also show that for these small instances, *MIS-graph* and *MIS-ORS* are comparable in terms of running time. This is because the graph size and vertex degrees do not yet influence the running time of *MIS-graph* strongly. Yet, as the next experiment shows, this changes drastically, as the instance size grows.

**Runtimes.** In our last experiment, we explore in more detail the scalability of the algorithms for larger instances, both relative to each other and in comparison to the recomputation times of their corresponding static algorithms. We generated 10 random instances with  $n = 1,000k$

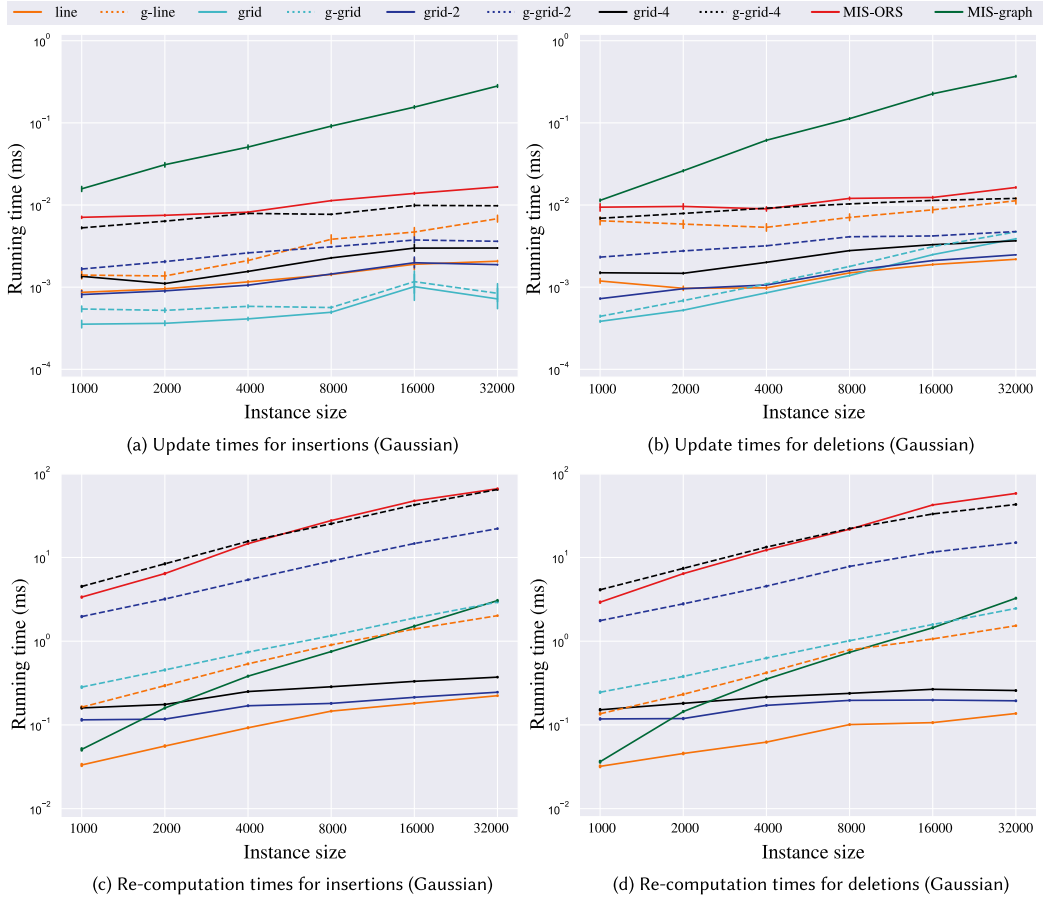


Fig. 12. Log-log runtime plots (notice the different  $y$ -offsets) for dynamic updates and recomputation on Gaussian instances of size  $n = 1,000$  to  $32,000$ , averaged over  $n/10$  updates. Error bars indicate the standard deviation.

squares for each  $k \in \{1, 2, 4, 8, 16, 32\}$  and measured the average update times over  $n/10$  insertions or deletions. The results for the Gaussian and uniform model are plotted in Figure 12 and in Figure 13. Considering the update times, we confirm the observations from the scatter plots in terms of the performance ranking. The running time of most algorithms grows only very slowly as the input size grows larger with the notable exception of *MIS-graph*, but that was to be expected.

In the comparison with their non-dynamic versions (i.e., recomputing solutions after each update), the dynamic algorithms indeed show a significant speedup in practice, already for small instance sizes of  $n = 1,000$ , and even more so as  $n$  grows (notice the different  $y$ -offsets). For some algorithms, including *MIS-ORS* and *g-grid-4*, this can be as high as three to four orders of magnitude for  $n = 32,000$ . This clearly confirms that the investigation of algorithms for dynamic MIS and MAX-IS problems for rectangles is well justified also from a practical point of view.

**Discussion.** Our experimental evaluation provides several interesting insights into the practical performance of the different algorithms. For the synthetic instances, both MIS-based algorithms *MIS-graph* and *MIS-ORS* generally showed the best solution quality in the field, reaching 90% of the

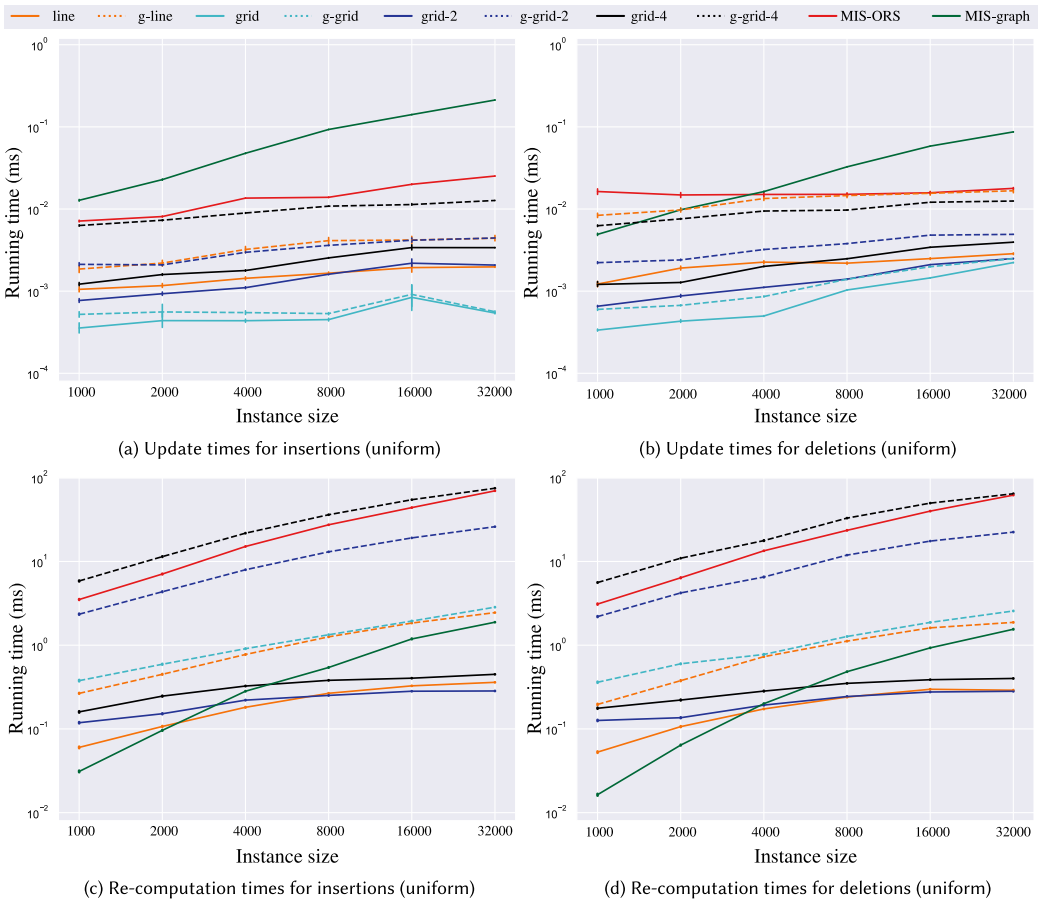


Fig. 13. Log-log runtime plots (notice the different y-offsets) for dynamic updates and recomputation on uniform instances of size  $n = 1,000$  to  $32,000$ , averaged over  $n/10$  updates. Error bars indicate the standard deviation.

exact MAX-IS size, where we could compare against optimal solutions. This is in strong contrast to their factor-4 worst-case approximation guarantee of only 25%.

Our algorithm *MIS-ORS* avoids storing the intersection graph explicitly. Instead, we only store the relevant geometric information in a dynamic data structure and derive edges on demand. Therefore, it overcomes the natural barrier of  $\Omega(\Delta)$  vertex update in a dynamic graph, where  $\Delta$  is the maximum degree in the graph. Instead, it has to find the intersections using the complex range query, which takes  $O(\log n \log \log n)$  time. We observe that *MIS-ORS* provides faster update times than *MIS-graph* in general and is more scalable. Recall that in our implementation, we used the dynamic range searching data structure from CGAL, which does not provide the theoretical worst-case update time of  $O(\log n \log \log n)$  from Theorem 4. Exploring how *MIS-ORS* can benefit from such a state-of-the-art dynamic data structure in practice remains to be investigated in future work. Notwithstanding, it remains to state that even with the suboptimal data structure, *MIS-ORS* was able to compute its solutions for up to 32,000 squares in less than 1 ms.

An expected observation is that although consistently exceeding their theoretical guarantees, the approximation algorithms do not perform too well in practice due to their pigeonhole choice

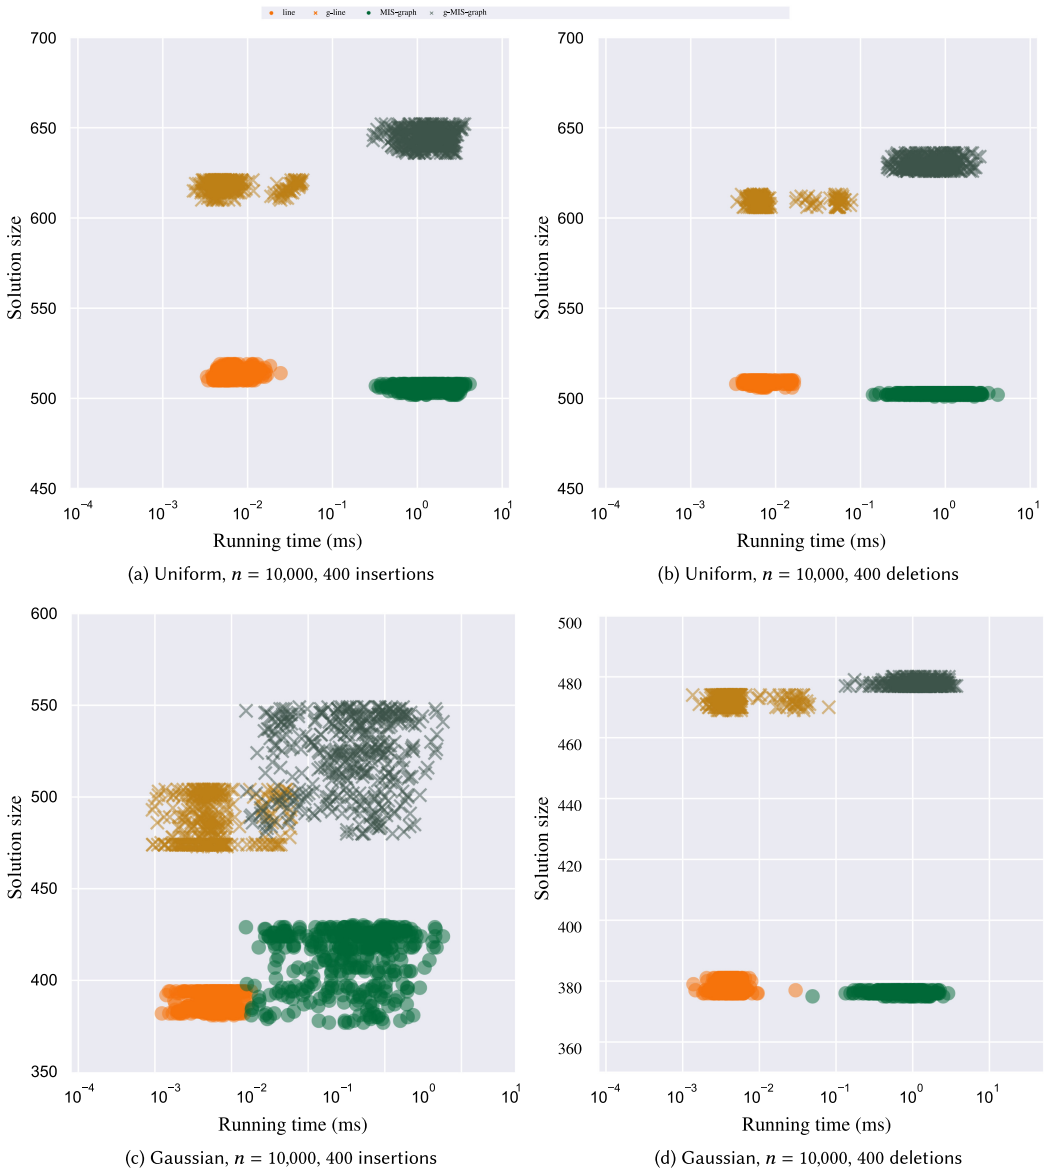


Fig. 14. Time-quality scatter plots for synthetic benchmark instances of uniform-height rectangles. The  $x$ -axis (log-scale) shows runtime, and the  $y$ -axis shows the solution size.

of too strictly separated substances. However, a simple greedy augmentation of the approximate solutions can boost the solution size significantly, and for some algorithms even to almost that of the MIS algorithms. Of course, at the same time, this increases the runtime of the algorithms. We want to point out *g-line*, the greedy-augmented version of the 2-approximation algorithm *line*, as it computes very good solutions, even comparable or better than *MIS-ORS* and *MIS-graph* for the real-world instances, and at 90% of the MIS solutions for the synthetic instances. At the same time, *g-line* is still significantly faster than *MIS-ORS* and *MIS-graph* and thus turns out to be a

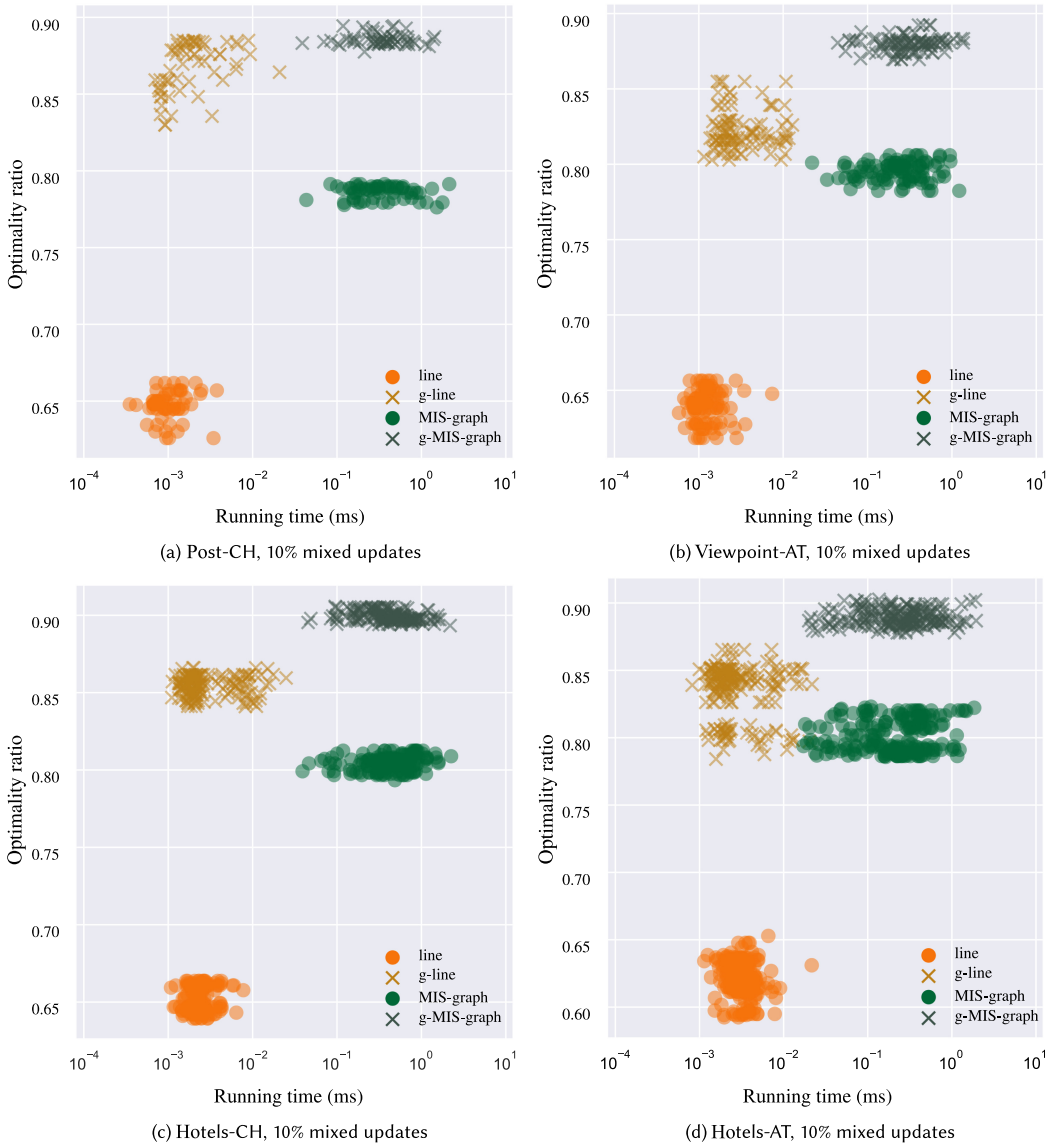


Fig. 15. Time-quality scatter plots for the small OSM uniform-height rectangle instances. The x-axis (log-scale) shows runtime. The y-axis shows the quality ratio compared to an optimal MAX-IS solution.

well-balanced compromise between time and quality. It is our recommended method if *MIS-ORS* or *MIS-graph* are too slow for an application.

### 5.3 Experimental Results for Unit-Height Rectangles

In this section, we compare our approaches *line* and *MIS-graph* for unit-height rectangles with the same sets of experiments as in the previous section. Recall that the *line* approach performs a greedy MAX-IS approach for interval graphs on each stabbing line where the intervals are sorted with increasing right endpoints. For fairness, we also provide the “greedy” version of *MIS-graph*

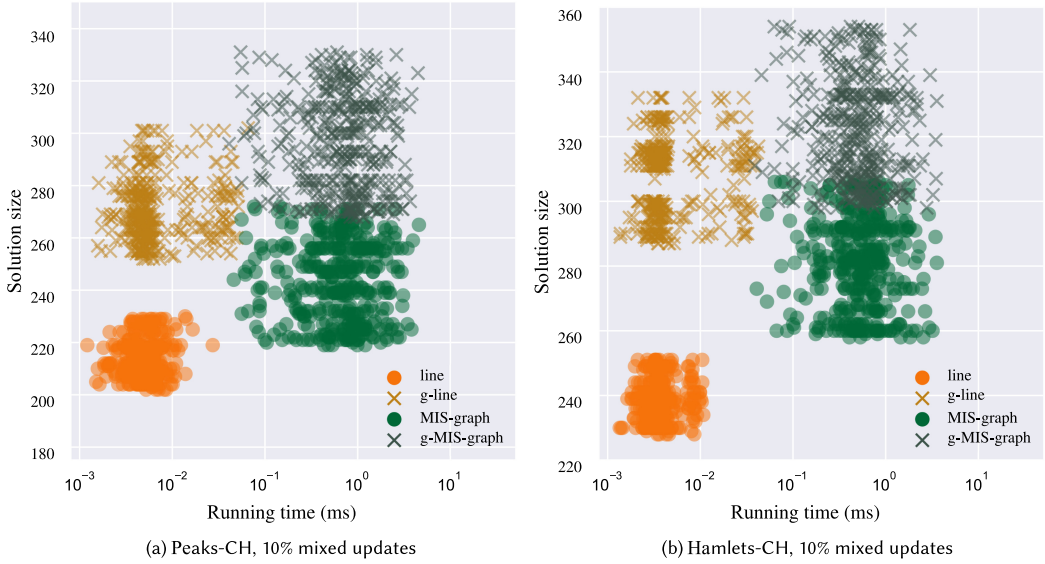


Fig. 16. Time-quality scatter plots for the large OSM uniform-height rectangle instances. The  $x$ -axis (log-scale) shows runtime. The  $y$ -axis shows the quality ratio compared to the optimal MAX-IS solution size.

approach,  $g$ -MIS-graph. In the initialization phase,  $g$ -MIS-graph sorts the vertices of the conflict graph by vertex degrees incrementally and builds the MIS greedily by iterating the vertices in this new order. We expect that this greedy variant would provide larger solutions than MIS-graph and the same update time as MIS-graph.

**Time-quality trade-offs.** For our first set of experiments, we compare the *line* approach with MIS-graph, including their greedy variants, in terms of update time and size of the computed solution. Figure 14 shows scatter plots of runtime vs. solution size on uniform and Gaussian rectangle benchmarks. For each instance with  $n = 10,000$ , each algorithm performed 400 updates, either insertions (Figures 14(a) and 14(c)) or deletions (Figures 14(b) and 14(d)). All plots show that  $g$ -MIS-graph, the greedy variant of MIS-graph, computes the best solutions and well ahead of the rest. Both greedy variants increase the solution size significantly without significant additional runtime. Note that  $g$ -line is nearly two orders of magnitude faster than graph-based approaches with about 90% of the MIS solution obtained by  $g$ -MIS-graph. It is interesting to see that  $g$ -line gets larger solutions than MIS-graph, which is in contrast to our experimental results for unit squares.

**Optimality gaps.** Next, we explore our approaches on real-world instances and show the results in Figures 15 and 16. For small instances, we compute MAX-IS exactly with MaxHS at each round and compare the solution of our presented dynamic approaches with the optimum. These plots show that  $g$ -MIS-graph reaches consistently about 90% of the optimum. In contrast with the graph-based approaches, algorithms *line* and  $g$ -line show a wider range of optimization ratios. The original approximations are well above their respective worst-case ratios, namely around 65% and 80%, respectively. The greedy variants of these two approaches push this toward larger solutions with nearly no additional running time. Note that  $g$ -line reaches between 80% and 85% of the optimum but faster is by one to two orders of magnitude compared to the graph-based approaches.

Consider the large OSM instances in Figure 16. Here, we also observe a similar pattern as the one from the smaller instance, except that they show larger variances of solution sizes for each

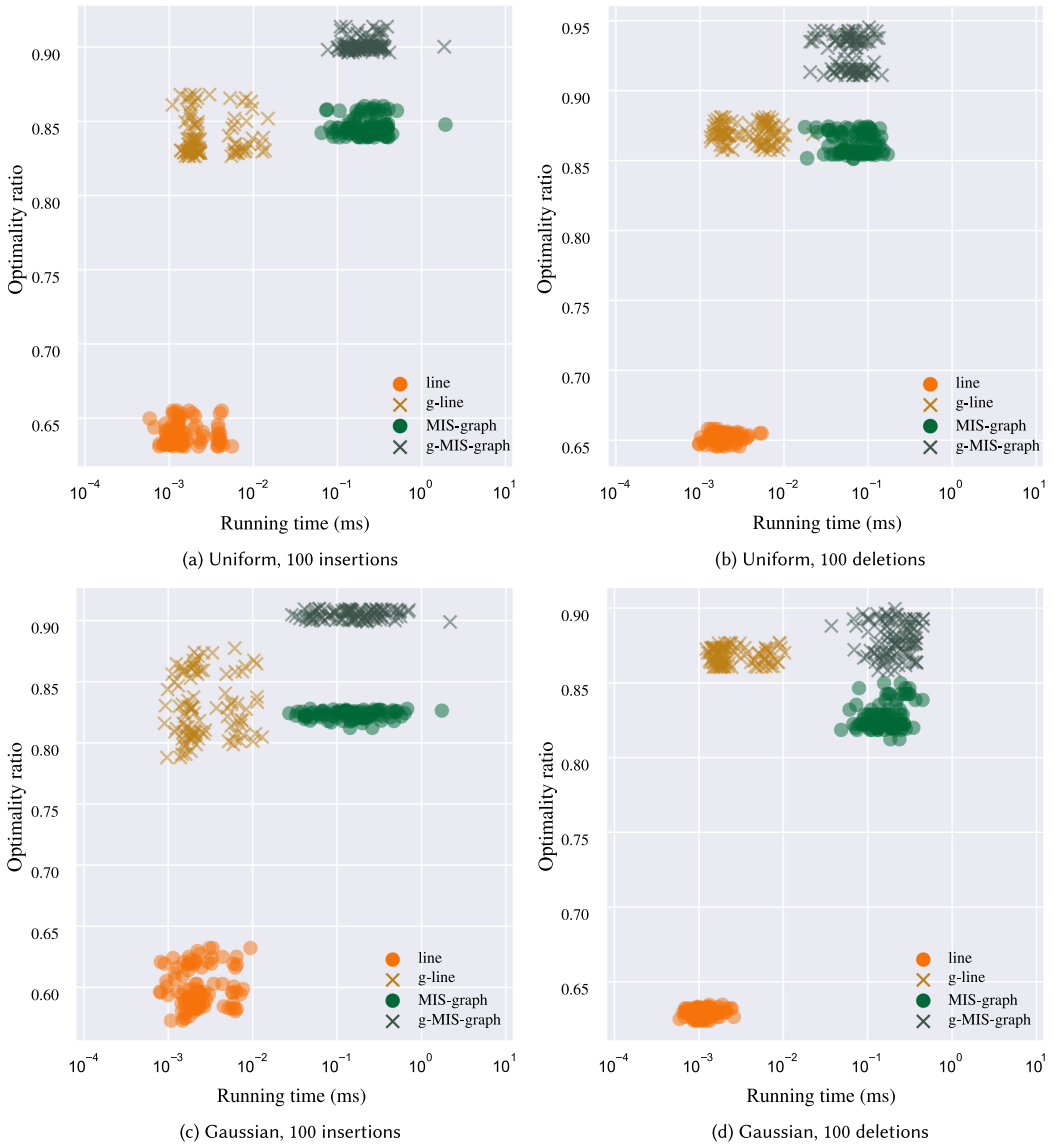


Fig. 17. Time-quality scatter plots for uniform and Gaussian instances with  $n = 1,000$  unit-height rectangles. The x-axis (log-scale) shows runtime. The y-axis shows the quality ratio compared to an optimal MAX-IS solution.

algorithm and *MIS-graph* and *g-MIS-graph* get closer to each other in terms of solution size. One possible reason could be that with the change of input, the ordering of vertices with incremental vertex degrees is not maintained anymore.

Finally, we also compare the solutions of the approximation approaches with the optimum on the small uniform and Gaussian instances with  $n = 1,000$  unit-height rectangles (Figure 17). This confirms our observations from the small OSM instances that *g-line* provides a very good balance between quality and computation time.




Fig. 18. Log-log runtime plots (notice the different  $y$ -offsets) for dynamic updates and recomputation on Gaussian uniform-height rectangle instances of size  $n = 1,000$  to  $32,000$ , averaged over  $n/10$  updates. Error bars indicate the standard deviation.

**Runtimes.** In the last experiment, we explore the scalability of our presented approaches, both relative to each other and in comparison to the recomputing times. For each  $k \in \{1, 2, 4, 8, 16, 32\}$ , we generated 10 random instances with  $1,000k$  squares and measured the average update time over  $100k$  insertion or deletions. The results are plotted in Figures 18 (for Gaussian instances) and 19 (for uniform instances). Note here that since the update procedure of  $g$ -MIS-graph is exactly the same as the MIS-graph approach, we do not include  $g$ -MIS-graph explicitly in this set of runtime experiments.

Considering the update time, the plots confirm the observations we had before. The  $g$ -line approach is nearly as fast as  $line$ . The runtime of MIS-graph shows steeper increase as the instance size increases than the runtime of  $line$  and  $g$ -line.

In the comparison with non-dynamic versions (i.e., recomputing the solution after each update), the dynamic approaches show a speedup by at least one order of magnitude.

**Discussion.** Our experimental results for unit-height and arbitrary-width rectangles confirm several of our findings obtained in the experiment for unit squares. Moreover, all approaches are well above their respective approximation ratios. A simple greedy variant of the MIS-graph






Fig. 19. Log-log runtime plots (notice the different y-offsets) for dynamic updates and recomputation on uniform-height rectangle instances of size  $n = 1,000$  to  $32,000$ , averaged over  $n/10$  updates. Error bars indicate the standard deviation.

approach can significantly boost the solution size and provide the best solution quality. Therefore, if the solution size is then priority, then the  $g$ -MIS-graph approach can be chosen. Our greedy augmented version of  $line$  reaches a good balance of solution quality and update time since it is significantly faster than graph-based approaches and computes very good solutions in practice.

## 6 CONCLUSION

We investigated the MIS and MAX-IS problems on dynamic sets of uniform rectangles and uniform-height rectangles from an algorithm engineering perspective, providing both theoretical results for maintaining an MIS or an approximate MAX-IS and reporting insights from an experimental study. Open problems for future work include (i) finding MAX-IS sublinear-update-time approximation algorithms for dynamic unit squares with approximation ratio better than 2, (ii) studying similar questions for dynamic disk graphs, and (iii) implementing improvements such as a faster dynamic range searching data structure to speed up our algorithm MIS-ORS. Moreover, it would be interesting to design dynamic approximation schemes for MAX-IS that maintain stability in a solution.

## REFERENCES

- [1] Amir Abboud, Raghavendra Addanki, Fabrizio Grandoni, Debmalya Panigrahi, and Barna Saha. 2019. Dynamic set cover: Improved algorithms and lower bounds. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23–26, 2019*, Moses Charikar and Edith Cohen (Eds.). ACM, New York, NY, 114–125. <https://doi.org/10.1145/3313276.3316376>
- [2] Anna Adamaszek and Andreas Wiese. 2013. Approximation schemes for maximum weight independent set of rectangles. In *Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26–29 October, 2013, Berkeley, CA, USA*. IEEE, Los Alamitos, CA, 400–409. <https://doi.org/10.1109/FOCS.2013.50>
- [3] Pankaj K. Agarwal, Marc Van Kreveld, and Subhash Suri. 1998. Label placement by maximum independent set in rectangles. *Comput. Geom. Theory Appl.* 11, 3–4 (1998), 209–218. [https://doi.org/10.1016/S0925-7721\(98\)00028-5](https://doi.org/10.1016/S0925-7721(98)00028-5)
- [4] Sepehr Assadi, Krzysztof Onak, Baruch Schieber, and Shay Solomon. 2019. Fully dynamic maximal independent set with sublinear in  $n$  update time. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, CA, USA, January 6–9, 2019*, Timothy M. Chan (Ed.). SIAM, 1919–1936. <https://doi.org/10.1137/1.9781611975482.116>
- [5] Ken Been, Eli Daiches, and Chee-Keng Yap. 2006. Dynamic map labeling. *IEEE Trans. Vis. Comput. Graph.* 12, 5 (2006), 773–780. <https://doi.org/10.1109/TVCG.2006.136>
- [6] Ken Been, Martin Nöllenburg, Sheung-Hung Poon, and Alexander Wolff. 2010. Optimizing active ranges for consistent dynamic map labeling. *Comput. Geom. Theory Appl.* 43, 3 (2010), 312–328. <https://doi.org/10.1016/j.comgeo.2009.03.006>
- [7] Soheil Behnezhad, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Cliff Stein, and Madhu Sudan. 2019. Fully dynamic maximal independent set with polylogarithmic update time. In *Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, MD, USA, November 9–12, 2019*, David Zuckerman (Ed.). IEEE, Los Alamitos, CA, 382–405. <https://doi.org/10.1109/FOCS.2019.00032>
- [8] Aaron Bernstein, Sebastian Forster, and Monika Henzinger. 2019. A deamortization approach for dynamic spanner and dynamic maximal matching. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, CA, USA, January 6–9, 2019*, Timothy M. Chan (Ed.). SIAM, 1899–1918. <https://doi.org/10.1137/1.9781611975482.115>
- [9] Sayan Bhattacharya, Deeparnab Chakrabarty, and Monika Henzinger. 2017. Deterministic fully dynamic approximate vertex cover and fractional matching in  $O(1)$  amortized update time. In *Integer Programming and Combinatorial Optimization*. Lecture Notes in Computer Science, Vol. 10328. Springer, 86–98. [https://doi.org/10.1007/978-3-319-59250-3\\_8](https://doi.org/10.1007/978-3-319-59250-3_8)
- [10] Sayan Bhattacharya, Deeparnab Chakrabarty, Monika Henzinger, and Danupon Nanongkai. 2018. Dynamic algorithms for graph coloring. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7–10, 2018*, Artur Czumaj (Ed.). SIAM, 1–20. <https://doi.org/10.1137/1.9781611975031.1>
- [11] Sujoy Bhore, Jean Cardinal, John Iacono, and Grigoriou Koumoutsos. 2020. Dynamic geometric independent set. *CoRR abs/2007.08643* (2020). <https://arxiv.org/abs/2007.08643>.
- [12] Parinya Chalermsook and Julia Chuzhoy. 2009. Maximum independent set of rectangles. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2009, New York, NY, USA, January 4–6, 2009*, Claire Mathieu (Ed.). SIAM, 892–901. <https://doi.org/10.1137/1.9781611973068.97>
- [13] Timothy M. Chan and Sariel Har-Peled. 2012. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete Comput. Geom.* 48, 2 (2012), 373–392. <https://doi.org/10.1007/s00454-012-9417-5>
- [14] Timothy M. Chan and Konstantinos Tsakalidis. 2017. Dynamic orthogonal range searching on the RAM, revisited. In *Computational Geometry (SoCG'17) (LIPIcs)*, Boris Aronov and Matthew J. Katz (Eds.), Vol. 77. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Article 28, 13 pages. <https://doi.org/10.4230/LIPIcs.SocG.2017.28>
- [15] Shiri Chechik and Tianyi Zhang. 2019. Fully dynamic maximal independent set in expected poly-log update time. In *Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, MD, USA, November 9–12, 2019*, David Zuckerman (Ed.). IEEE, Los Alamitos, CA, 370–381. <https://doi.org/10.1109/FOCS.2019.00031>
- [16] Jon Christensen, Joe Marks, and Stuart M. Shieber. 1995. An empirical study of algorithms for point-feature label placement. *ACM Trans. Graph.* 14, 3 (1995), 203–232. <https://doi.org/10.1145/212332.212334>
- [17] Julia Chuzhoy and Alina Ene. 2016. On approximating maximum independent set of rectangles. In *Proceedings of the IEEE 57th Annual Symposium on Foundations of Computer Science, FOCS 2016, 9–11 October 2016, Hyatt Regency, New Brunswick, NJ, USA*, Irit Dinur (Ed.). IEEE, Los Alamitos, CA, 820–829. <https://doi.org/10.1109/FOCS.2016.92>
- [18] Graham Cormode, Jacques Dark, and Christian Konrad. 2019. Independent sets in vertex-arrival streams. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP'19) (LIPIcs)*, Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi (Eds.), Vol. 132. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Article 45, 14 pages. <https://doi.org/10.4230/LIPIcs.ICALP.2019.45>

- [19] Mark de Berg and Dirk H. P. Gerrits. 2013. Labeling moving points with a trade-off between label speed and label overlap. In *Algorithms*. Lecture Notes in Computer Science, Vol. 8125. Springer, 373–384. [https://doi.org/10.1007/978-3-642-40450-4\\_32](https://doi.org/10.1007/978-3-642-40450-4_32)
- [20] Hugo A. D. do Nascimento and Peter Eades. 2008. User hints for map labeling. *J. Vis. Lang. Comput.* 19, 1 (2008), 39–74. <https://doi.org/10.1016/j.jvlc.2006.03.004>
- [21] David Eppstein, Zvi Galil, and Giuseppe F. Italiano. 1999. Dynamic graph algorithms. In *Algorithms and Theory of Computation Handbook*, Mikhail J. Atallah (Ed.). CRC Press, Boca Raton, FL, 1–28.
- [22] Thomas Erlebach, Klaus Jansen, and Eike Seidel. 2005. Polynomial-time approximation schemes for geometric intersection graphs. *SIAM J. Comput.* 34, 6 (2005), 1302–1323. <https://doi.org/10.1137/s0097539702402676>
- [23] Michael Formann and Frank Wagner. 1991. A packing problem with applications to lettering of maps. In *Proceedings of the 7th Annual Symposium on Computational Geometry, North Conway, NH, USA, June 10–12, 1991*, Robert L. Scot Drysdale (Ed.). ACM, New York, NY, 281–288. <https://doi.org/10.1145/109648.109680>
- [24] Robert J. Fowler, Mike Paterson, and Steven L. Tanimoto. 1981. Optimal packing and covering in the plane are NP-complete. *Inf. Process. Lett.* 12, 3 (1981), 133–137. [https://doi.org/10.1016/0020-0190\(81\)90111-3](https://doi.org/10.1016/0020-0190(81)90111-3)
- [25] Edith Gabriel. 2015. Spatio-temporal point pattern analysis and modeling. In *Encyclopedia of GIS*, Shashi Shekhar, Hui Xiong, and Xun Zhou (Eds.). Springer, 1–8. [https://doi.org/10.1007/978-3-319-23519-6\\_1646-1](https://doi.org/10.1007/978-3-319-23519-6_1646-1)
- [26] Waldo Gálvez, Arindam Khan, Mathieu Mari, Tobias Mömke, Madhusudhan Reddy Pittu, and Andreas Wiese. 2021. A 3-Approximation Algorithm for Maximum Independent Set of Rectangles. *CoRR* abs/2106.00623 (2021). <https://dblp.org/rec/conf/soda/GalvezKMMPW22.html?view=bibtex>.
- [27] Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. 2019. Online matching with general arrivals. In *Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, MD, USA, November 9–12, 2019*, David Zuckerman (Ed.). IEEE, Los Alamitos, CA, 26–37. <https://doi.org/10.1109/FOCS.2019.00011>
- [28] Alexander Gavruskin, Bakhadyr Khossainov, Mikhail Kokho, and Jiamou Liu. 2015. Dynamic algorithms for monotonic interval scheduling problem. *Theor. Comput. Sci.* 562 (2015), 227–242. <https://doi.org/10.1016/j.tcs.2014.09.046>
- [29] Andreas Gamsa, Martin Nöllenburg, and Ignaz Rutter. 2016. Evaluation of labeling strategies for rotating maps. *ACM J. Exp. Algorithmics* 21, 1 (2016), Article 1.4, 21 pages. <https://doi.org/10.1145/2851493>
- [30] Andreas Gamsa, Martin Nöllenburg, and Ignaz Rutter. 2016. Consistent labeling of rotating maps. *J. Comput. Geom.* 7, 1 (2016), 308–331. <https://doi.org/10.20382/jocg.v7i1a15>
- [31] U. I. Gupta, D. T. Lee, and Joseph Y.-T. Leung. 1982. Efficient algorithms for interval graphs and circular-arc graphs. *Networks* 12, 4 (1982), 459–467. <https://doi.org/10.1002/net.3230120410>
- [32] Monika Henzinger, Stefan Neumann, and Andreas Wiese. 2020. Dynamic approximate maximum independent set of intervals, hypercubes and hyperrectangles. In *Proceedings of the Symposium on Computational Geometry (SoCG'20) (LIPIcs)*, Sergio Cabello and Danny Z. Chen (Eds.), Vol. 164. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Article 51, 14 pages. <https://doi.org/10.4230/LIPIcs.SoCG.2020.51>
- [33] Dorit S. Hochbaum and Wolfgang Maass. 1985. Approximation schemes for covering and packing problems in image processing and VLSI. *J. ACM* 32, 1 (1985), 130–136. <https://doi.org/10.1145/2455.214106>
- [34] John E. Hopcroft and Richard M. Karp. 1973. An  $n^{5/2}$  algorithm for maximum matchings in bipartite graphs. *SIAM J. Comput.* 2, 4 (1973), 225–231. <https://doi.org/10.1137/0202019>
- [35] Richard M. Karp. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, R. E. Miller, J. W. Thatcher, and J. D. Bohlinger (Eds.). Springer, 85–103. [https://doi.org/10.1007/978-1-4684-2001-2\\_9](https://doi.org/10.1007/978-1-4684-2001-2_9)
- [36] Fabian Klute, Guangping Li, Raphael Löffler, Martin Nöllenburg, and Manuela Schmidt. 2019. Exploring semi-automatic map labeling. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2019, Chicago, IL, USA, November 5–8, 2019*, Farnoush Banaei Kashani, Goce Trajcevski, Ralf Hartmut Güting, Lars Kulik, and Shawn D. Newsam (Eds.). ACM, New York, NY, 13–22. <https://doi.org/10.1145/3347146.3359359>
- [37] Nathan Linial. 1987. Distributive graph algorithms-global solutions from local data. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science (FOCS'87)*. IEEE, Los Alamitos, CA, 331–335. <https://doi.org/10.1109/SFCS.1987.20>
- [38] Alan M. MacEachren, Anuj R. Jaiswal, Anthony C. Robinson, Scott Pezanowski, Alexander Savelyev, Prasenjit Mitra, Xiao Zhang, and Justine I. Blanford. 2011. SensePlace2: GeoTwitter analytics support for situational awareness. In *Proceedings of the Conference on Visual Analytics Science and Technology (VAST'11)*. IEEE, Los Alamitos, CA, 181–190.
- [39] Kurt Mehlhorn and Stefan Näher. 1990. Dynamic fractional cascading. *Algorithmica* 5, 1–4 (1990), 215–241. <https://doi.org/10.1007/BF01840386>
- [40] Kurt Mehlhorn and Stefan Näher. 1999. *The LEDA Platform of Combinatorial and Geometric Computing*. Cambridge University Press, Cambridge, MA. <https://doi.org/10.1145/204865.204889>

- [41] Joseph S. B. Mitchell. 2021. Approximating maximum independent set for rectangles in the plane. *CoRR* abs/2101.00326 (2021). <https://arxiv.org/abs/2101.00326>.
- [42] Huy N. Nguyen and Krzysztof Onak. 2008. Constant-time approximation algorithms via local improvements. In *Proceedings of the 49th Annual IEEE Symposium on Foundations of Computer Science (FOCS'08)*. IEEE, Los Alamitos, CA, 327–336. <https://doi.org/10.1109/FOCS.2008.81>
- [43] Panos M. Pardalos and Jue Xue. 1994. The maximum clique problem. *J. Glob. Optim.* 4, 3 (1994), 301–328. <https://doi.org/10.1007/BF01098364>
- [44] Maxim A. Rylov and Andreas W. Reimer. 2014. A comprehensive multi-criteria model for high cartographic quality point-feature label placement. *Cartogr. Int. J. Geogr. Inf. Geovisualization* 49, 1 (2014), 52–68. <https://doi.org/10.3138/carto.49.1.2137>
- [45] Pedro V. Sander, Diego Nehab, Eden Chlamtac, and Hugues Hoppe. 2008. Efficient traversal of mesh edges using adjacency primitives. *ACM Trans. Graph.* 27, 5 (2008), 144. <https://doi.org/10.1145/1409060.1409097>
- [46] Dennis Thom, Harald Bosch, Steffen Koch, Michael Wörner, and Thomas Ertl. 2012. Spatiotemporal anomaly detection through visual analysis of geolocated Twitter messages. In *Proceedings of the 2012 Pacific Visualization Symposium (PacificVis'12)*. IEEE, Los Alamitos, CA, 41–48.
- [47] René van Bevern, Matthias Mnich, Rolf Niedermeier, and Mathias Weller. 2015. Interval scheduling and colorful independent sets. *J. Sched.* 18, 5 (2015), 449–469. <https://doi.org/10.1007/s10951-014-0398-5>
- [48] Marc J. van Kreveld, Tycho Strijk, and Alexander Wolff. 1998. Point set labeling with sliding labels. In *Proceedings of the 14th Annual Symposium on Computational Geometry, Minneapolis, MN, USA, June 7–10, 1998*, Ravi Janardan (Ed.). ACM, New York, NY, 337–346. <https://doi.org/10.1145/276884.276922>
- [49] Frank Wagner and Alexander Wolff. 1997. A practical map labeling algorithm. *Comput. Geom.* 7 (1997), 387–404. [https://doi.org/10.1016/S0925-7721\(96\)00007-7](https://doi.org/10.1016/S0925-7721(96)00007-7)
- [50] Dan E. Willard and George S. Lueker. 1985. Adding range restriction capability to dynamic data structures. *J. ACM* 32, 3 (1985), 597–617. <https://doi.org/10.1145/3828.3839>
- [51] David Zuckerman. 2007. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory Comput.* 3, 1 (2007), 103–128. <https://doi.org/10.4086/toc.2007.v003a006>

Received October 2020; revised August 2021; accepted January 2022