Optimal-size problem kernels for d-Hitting Set in linear time and space[⋆]

René van Bevern[∗], Pavel V. Smirnov

Department of Mechanics and Mathematics, Novosibirsk State University, ul. Pirogova 1, 630090 Novosibirsk, Russian Federation

Abstract

The known linear-time kernelizations for d-Hitting Set guarantee linear worst-case running times using a quadratic-size data structure (that is not fully initialized). Getting rid of this data structure, we show that problem kernels of asymptotically optimal size $O(k^d)$ for d-Hitting Set are computable in linear time and space. Additionally, we experimentally compare the linear-time kernelizations for d-Hitting Set to each other and to a classical data reduction algorithm due to Weihe.

Keywords: combinatorial optimization, NP-hard problem, data reduction, parameterized complexity, kernelization

1. Introduction

We study data reduction algorithms for the following combinatorial optimization problem:

Problem 1 (d-Hitting Set, for constant $d \in \mathbb{N}$).
Input: A hypergraph $H = (V, E)$ with vertex set $V = \{1, \ldots, n\}$, edge set $E \subseteq \{e \subseteq V : |e| \le d\}$, and $k \in \mathbb{N}$.
Question: Is there a *hitting set* $S \subseteq V$ of cardinality at most k, that is, $\forall e \in E : e \cap S \ne \emptyset$?

Throughout this work, we denote $n := |V|$ and $m := |E|$ and call $|H| = |V| + |E|$ the size of the hypergraph H.

The d-Hitting Set problem is an NP-complete [22] fundamental combinatorial optimization problem, arising in bioinformatics [25], medicine [24, 34], clustering [8, 21], automatic reasoning [11, 17, 31], feature selection [10, 20], radio frequency allocation [33], software engineering [29], and public transport optimization [9, 35].

Exact algorithms for NP-complete problems usually take time exponential in the input size. Thus, an important preprocessing step is data reduction, which has proven to significantly shrink real-world instances of NP-hard problems [2, 7, 9, 24, 35]. The main notion of data reduction with performance guarantees is *kernelization* [19], here stated for d-Hitting Set:

Definition 2. *A kernelization maps any d-Hitting Set instance (H_{in}, k_{in}) to an instance (H_{out}, k_{out}) in polynomial time such that*

1. *H_{in} has a hitting set of size k_{in} if and only if H_{out} has a hitting set of size k_{out},*

2. *$|H_{out}| + k_{out} \le g(k_{in})$ for a computable function $g \colon \mathbb{N} \to \mathbb{N}$.*

One calls (H_{out}, k_{out}) the problem kernel and $g(k_{in})$ its size. In the kernelizations studied in our work, $k := k_{in} = k_{out}$.

The existence of problem kernels of size $O(k^{d-\varepsilon})$ for any $\varepsilon > 0$ for d-Hitting Set results in a collapse of the polynomial-time hierarchy [14]. There are two known $O(n + m)$-time kernelizations for d-Hitting Set that yield problem kernels of this optimal size $O(k^d)$ [6, 16]. They can be implemented

(a) with expected linear running time and linear space using hash tables with random hash functions, or

(b) with worst-case linear running time using a trie data structure, which may allocate $\Theta(nm)$ cells of random access memory (of which only $O(n + m)$ are initialized) [5].

In experiments, the memory usage of implementation (b) proved to be prohibitively large, yet implementation (b) outperformed implementation (a) when enough memory was available and allocation of zero-initialized memory was cheap [5].

From a theoretical point of view, a natural question is whether one can derandomize implementation (a), so that it runs in worst-case linear time and space, or, equivalently, whether one can lower the memory usage of implementation (b) to linear. From a practical point of view, it is interesting how the two kernelizations compare to each other and to other data reduction algorithms.

Our contributions and organization of this work. Section 2 shows the two known linear-time d-Hitting Set kernelizations. We then describe our new contributions, which are two-fold:

In Section 3, we resolve the apparent paradox that the linear worst-case running time of the known d-Hitting Set kernelizations hinges on quadratic-space data structures. To this end, we show how to implement both of them in $O(n+m)$ time and space.

In Section 4, for the first time, we experimentally compare the two kernelizations to each other and to a well-known data reduction algorithm due to Weihe [35], which runs in superlinear time, does not yield problem kernels, but proved to be very effective on instances of the Station Cover problem. We will see that the kernelizations outperform Weihe's algorithm for small d and that combinations of kernelization and Weihe's algorithm may yield significantly stronger data reduction effects than the individual algorithms.

Algorithm FK: Algorithm of Fafianie and Kratsch [16].

Input: Hypergraph (V_{in}, E_{in}), $k \in \mathbb{N}$.
Output: Problem kernel $((V_{out}, E_{out}), k)$ with $|E_{out}| \leq (k+1)^d$.
// Initially, $\forall s \subseteq V : \text{supersets}[s] = 0$.

1 $E_{out} \leftarrow \emptyset$
2 **foreach** $e \in E_{in}$ **do**
3   **if** $\forall s \subseteq e : \text{supersets}[s] < (k+1)^{d-|s|}$ **then**
4          $E_{out} \leftarrow E_{out} \cup \{e\}$
5   **foreach** $s \subseteq e$ **do** $\text{supersets}[s] \leftarrow \text{supersets}[s] + 1$

6 $V_{out} \leftarrow \bigcup_{e \in E_{out}} e$
7 **return** $((V_{out}, E_{out}), k)$

Algorithm Bev: Algorithm of van Bevern [6].

Input: Hypergraph (V_{in}, E_{in}), $k \in \mathbb{N}$.
Output: Kernel $((V_{out}, E_{out}), k)$ with $|E_{out}| \leq d! d^{d+1}(k+1)^d$.
// Initially, $\forall s \subseteq V, v \in V : \text{petals}[s] = 0, \text{used}[s][v] = \text{false}$.

1 $E_{out} \leftarrow \emptyset$
2 **foreach** $e \in E_{in}$ **do**
3   **if** $\forall s \subseteq e : \text{petals}[s] \leq k$ **then**
4          $E_{out} \leftarrow E_{out} \cup \{e\}$
5   **foreach** $s \subseteq e$ **do**
6   **if** $\forall v \in e \setminus s : \text{used}[s][v] = \text{false}$ **then**
7                  $\text{petals}[s] \leftarrow \text{petals}[s] + 1$
8   **foreach** $v \in e \setminus s$ **do** $\text{used}[s][v] \leftarrow \text{true}$

9 $V_{out} \leftarrow \bigcup_{e \in E_{out}} e$
10 **return** $((V_{out}, E_{out}), k)$

Related work. There are several kernelizations for d-Hitting Set [1, 3, 4, 13, 15, 18, 23, 26, 28]. Dell and van Melkebeek [14] showed that the existence of a problem kernel with $O(k^{d-\varepsilon})$ edges for any $\varepsilon > 0$ for d-Hitting Set implies a collapse of the polynomial-time hierarchy. Therefore, we do not expect polynomial-size problem kernels for d-Hitting Set if d is *not* constant. There are problem kernels with $O(k^{d-1})$ *vertices*, however [1, 5, 26].

Lowering the running time and space requirements of d-Hitting Set kernelizations both have been in the focus of research. The first linear-time kernelization is due to van Bevern [5, 6]. The second, due to Fafianie and Kratsch [16], is simpler and has smaller constant factors: the problem kernel of van Bevern [6] has at most $d! \cdot d^{d+1} \cdot (k+1)^d$ edges, whereas the problem kernel of Fafianie and Kratsch [16] has at most $(k+1)^d$ edges. Both kernelizations work in $O(d \cdot n + 2^d d \cdot m)$ time.

Problem kernels of size $O(k^d \log k)$ are computable in $O(k^d \log n)$ space and $O(k^d m)$ time [15], of size $O(k^d)$ in logarithmic space and $O(m^{d+2})$ time [16], and of exponential size even in constant parallel time [4].

2. Known linear-time algorithms

There are two known linear-time kernelizations for d-Hitting Set: FK and Bev. Both iterate over each input edge $e \in E_{in}$ once and decide whether to add e to the output edge set E_{out} as follows.

FK does not add e to E_{out} if e contains a subset s that is contained in $(k+1)^{d-|s|}$ edges of E_{out}. It can be shown that any hitting set of size k for (V_{out}, E_{out}) has to intersect s, and thus e [16]. Thus, for each $s \subseteq e \in E_{in}$, FK maintains the number of supersets of s in E_{out} in a counter $\text{supersets}[s]$, which is updated for each $s \subseteq e$ whenever adding an edge e to E_{out}.

Bev is based on finding *sunflowers*—sets of edges (called *petals*) with mutually equal intersection (called the *core*). If there is a sunflower with $k+1$ petals, then any hitting set of size k has to intersect its core. Thus, Bev does not add e to E_{out} if it finds that e contains the core s of a sunflower with $k+1$ petals in E_{out}. To this end, for each $s \subseteq e \in E_{in}$, Bev maintains the information of *one* sunflower with core s: $\text{petals}[s]$ is its number of petals and $\text{used}[s][v]$ is true if and only if $v \in V$ is contained in one of them. Whenever adding an edge e to E_{out}, Bev adds e to the sunflower with core s for each $s \subseteq e$ (if possible). Depending on the order of the input edges, Bev may not find a

sunflower with $k+1$ petals if it exists, yet if there is a sunflower with $d(k+1)$ petals, it finds one with $k+1$ petals for sure [5].

Access to $\text{supersets}[s]$, $\text{used}[s]$, and $\text{petals}[s]$ for each $s \subseteq e \in E_{in}$ in FK and Bev can be organized in $O(d)$ time using a trie that can be initialized in $O(dn + 2^d d \cdot m)$ time [5, Lemma 5.3]. After initialization, FK and Bev work in $O(2^d \cdot m)$ and $O(2^d d \cdot m)$ time, respectively. The culprit is that the trie can allocate $\Theta(nm)$ random access memory, of which only $O(n + m)$ is initialized [5, Lemma 5.3].

3. Implementing FK and Bev in linear space and time

3.1. *FK in linear space*

To implement FK in linear space, we apply a series of preprocessing steps. First, by iterating over E_{in} once, simultaneously incrementing a counter, we assign to each edge $e \in E_{in}$ a unique index $e_{id} \in \{1, \ldots, m\}$, in $O(m)$ time giving a set

$$E_{id} := \{(e, e_{id}) : e \in E_{in}\}. \tag{1}$$

We will then show how to compute in linear time and space a unique index $s_{id} \in \{1, \ldots, 2^d m\}$ for each $s \subseteq e \in E_{in}$ and a size-m array $A[]$ satisfying

$$A[e_{id}] = \{(s, s_{id}) : s \subseteq e\} \quad \text{for each } (e, e_{id}) \in E_{id}. \tag{2}$$

Then, instead of iterating over each $e \in E_{in}$, each $s \subseteq e$, and looking up $\text{supersets}[s]$ in a trie, it is enough to iterate over each $(e, e_{id}) \in E_{id}$, each $(s, s_{id}) \in A[e_{id}]$, and look up $\text{supersets}'[s_{id}]$ in $O(1)$ time, where $\text{supersets}'[]$ is an ordinary array of length $2^d m$. The modified algorithm is shown in FK'.

Theorem 3. *FK can be run in $O(nd + 2^d \cdot m)$ time and space.*

Proof. To prove the theorem, we show how to compute the array $A[]$ and a unique index s_{id} for each $s \subseteq e \in E_{in}$ in linear space and time. The tricky bit is that s may be a subset of several edges in E_{in}, yet its index s_{id} must be unique. Thus, we use the following canonical encoding of edges: for any subset $s \subseteq V = \{1, \ldots, n\}$ of size at most d, $\llcorner s \lrcorner \in (V \cup \{\square\})^d$ is a d-tuple containing the elements of s in increasing order and padded with $\square$ at

---

Algorithm FK': Linear-space version of FK.

Input: Hypergraph (V_{in}, E_{in}), $k \in \mathbb{N}$.
Output: Problem kernel $((V_{out}, E_{out}), k)$ with $|E_{out}| \leq (k+1)^d$.
// supersets$'$[] is a zero-initialized size-$2^d m$ array,
// E_{id} and A[] are as in (1) and (2).

1 $E_{out} \leftarrow \emptyset$
2 **foreach** $(e, e_{id}) \in E_{id}$ **do**
3   **if** $\forall (s, s_{id}) \in A[e_{id}]$: supersets$'[s_{id}] < (k+1)^{d-|s|}$ **then**
4          $E_{out} \leftarrow E_{out} \cup \{e\}$
5   **foreach** $(s, s_{id}) \in A[e_{id}]$ **do**
            supersets$'[s_{id}] \leftarrow$ supersets$[s_{id}] + 1$

6 $V_{out} \leftarrow \bigcup_{e \in E_{out}} e$
7 **return** $((V_{out}, E_{out}), k)$

---

the end. For example, for $d = 4$, $\llcorner \{3, 1, 2\} \lrcorner = (1, 2, 3, \square)$. Obviously, $\llcorner e \lrcorner$ for any edge $e \in E_{in}$ is computable in $O(d \log d)$ time. The preprocessing for FK now consists of three steps.

Step 1. Compute a list $L = [(\llcorner s \lrcorner, e_{id}) : s \subseteq e, (e, e_{id}) \in E_{id}]$ of size at most $2^d m$ by first computing $\llcorner e \lrcorner$ for all $e \in E_{in}$ in $O(m \cdot d \log d)$ time and space and then enumerating all substrings of $\llcorner e \lrcorner$ for each $(e, e_{id}) \in E_{id}$ in $O(2^d \cdot m)$ time and space.

Step 2. Sort L by lexicographically non-decreasing $\llcorner s \lrcorner$, where we assume $n < \square$. Since the $\llcorner s \lrcorner$ are d-tuples over $\{1, \ldots, n, \square\}$, this works in $O(d(n + |L|)) = O(nd + 2^d d \cdot m)$ time and space using radix sort [12, Section 8.3]. All pairs $(\llcorner s \lrcorner, e_{id})$ belonging to the same subset s now occur consecutively in L.

Step 3. Initialize a size-m array A[] of empty lists and $s_{id} \leftarrow 1$. Iterate over L as follows. For the current pair $(\llcorner s \lrcorner, e_{id})$, add (s, s_{id}) to $A[e_{id}]$. If there is a next pair $(\llcorner s' \lrcorner, e'_{id})$ on L and $\llcorner s \lrcorner \neq \llcorner s' \lrcorner$, then increment $s_{id} \leftarrow s_{id} + 1$ and continue. Note that $A[e_{id}]$ does not contain duplicates, so that the list is actually a set, as required by (2).

This concludes the computation of the s_{id} and the array A[]. The running time and space bottleneck is step 2. After this preprocessing, FK can be implemented to run in $O(2^d \cdot m)$ time and space as shown in FK'. $\square$

3.2. *Bev in linear space*

To implement Bev in linear time and space, we replace accesses to tries petals$[s]$ and used$[s]$ for each $s \subseteq e \in E_{in}$ by accesses to arrays petals$'[s_{id}]$ and used$'[s_{id}]$, as we did for FK. However, while petals$[s]$ is a counter that translates into a counter petals$'[s_{id}]$, used$[s][]$ is a size-n array indexed by vertices. Holding such an array in used$'[s_{id}][]$ would again use $\Omega(nm)$ space. Instead, we organize used$'[s_{id}][]$ as follows. Let

$$V^s := \bigcup_{e \in E_{in}} (e \setminus s) \quad \text{for each } s \subseteq e \in E_{in}. \quad (3)$$

We will compute unique indices $v_{id}^s \in \{1, \ldots, |V^s|\}$ for the vertices $v \in V^s$ for each $s \subseteq e \in E_{in}$, unique indices $s_{id}^e \in \{1, \ldots, 2^{|e|}\}$ of the subsets $s \subseteq e$ of each $e \in E_{in}$, an array B[] satisfying

$$B[e_{id}] = \{(s, s_{id}, s_{id}^e) : s \subseteq e\} \quad \text{for each } (e, e_{id}) \in E_{id}, \quad (4)$$

Algorithm Bev': Linear-space version of Bev.

Input: Hypergraph (V_{in}, E_{in}), $k \in \mathbb{N}$.
Output: Kernel $((V_{out}, E_{out}), k)$ with $|E_{out}| \leq d! d^{d+1} (k+1)^d$.
// E_{id}, V^s, B[], and C[][] are as in (1), (3), (4), and (5),
// petals$'$[] is a zero-initialized array of size $2^d m$,
// used$'$[] is an array of size $2^d m$,
// used$'[s_{id}]$ is a false-initialized array of size $|V^s|$.

1 $E_{out} \leftarrow \emptyset$
2 **foreach** $(e, e_{id}) \in E_{id}$ **do**
3   **if** $\forall (s, s_{id}, s_{id}^e) \in B[e_{id}]$: petals$'[s_{id}] \leq k$ **then**
4          $E_{out} \leftarrow E_{out} \cup \{e\}$
5   **foreach** $(s, s_{id}, s_{id}^e) \in B[e_{id}]$ **do**
6   **if** $\forall v_{id}^s \in C[e_{id}][s_{id}^e]$: used$'[s_{id}][v_{id}^s] =$ false **then**
7   petals$'[s_{id}] \leftarrow$ petals$'[s_{id}] + 1$
8   **foreach** $v_{id}^s \in C[e_{id}][s_{id}^e]$ **do**
9   used$'[s_{id}][v_{id}^s] \leftarrow$ true

10 $V_{out} \leftarrow \bigcup_{e \in E_{out}} e$
11 **return** $((V_{out}, E_{out}), k)$

---

and an array C[] of arrays satisfying

$$C[e_{id}][s_{id}^e] = \{v_{id}^s : v \in e \setminus s\} \quad (5)$$
$$\text{for each } (e, e_{id}) \in E_{id}, (s, s_{id}, s_{id}^e) \in B[e_{id}].$$

Bev can then be implemented using arrays petals$'$[] and used$'$[] of size $2^d m$ each, where for each s_{id}, used$'[s_{id}][]$ is an array of size $|V^s|$: instead of iterating over each $e \in E_{in}$, each $s \subseteq e$, each $v \in e \setminus s$, and looking up petals$[s]$ and used$[s][v]$ in tries, Bev can iterate over each $(e, e_{id}) \in E_{id}$, each $(s, s_{id}, s_{id}^e) \in B[e_{id}]$, each $v_{id}^s \in C[e_{id}][s_{id}^e]$, and look up petals$'[s_{id}]$ and used$'[s_{id}][v_{id}^s]$. These are simple array accesses, each working in constant time. The modified algorithm is shown in Bev'.

Theorem 4. *Bev can be run in* $O(nd + 2^d d \cdot m)$ *time and space.*

Proof. We describe how to compute the indices v_{id}^s, s_{id}^e, and the arrays B[] and C[] in linear time and space. First, the indices s_{id} and array A[] in (2) are computed as in Theorem 3 in $O(nd + 2^d d \cdot m)$ time and space. For Bev, we use three additional preprocessing steps.

Step 1. Initialize a size-m array B[]. For each $(e, e_{id}) \in E_{id}$, compute $B[e_{id}]$ from $A[e_{id}]$ by iterating over each $(s, s_{id}) \in A[e_{id}]$, simultaneously incrementing a counter s_{id}^e from 0 to $2^{|e|}$. This works in time $O(2^d \cdot m)$ and space.

Step 2. Iterating over each $(e, e_{id}) \in E_{id}$ and each $(s, s_{id}, s_{id}^e) \in B[e_{id}]$, in $O(2^d d \cdot m)$ time, generate a list

$$L := [(s_{id}, v, s_{id}^e, e_{id}) \mid v \in e \setminus s, s \subseteq e, e \in E_{in}].$$

Sort the list by lexicographically non-decreasing (s_{id}, v). Since these are pairs of numbers in $\{1, \ldots, 2^d m\} \cup \{1, \ldots, n\}$, this works in $O(n + 2^d m + |L|) = O(n + 2^d m)$ time using radix sort [12, Section 8.3]. Thereafter, all quadruples belonging to the same s_{id} occur consecutively in L. Also, for each fixed s_{id}, all quadruples belonging to s_{id} and the same v occur consecutively in L.

Step 3. Initialize a size-m array C[], and for each $(e, e_{id}) \in E_{id}$, a size-$2^{|e|}$ array $C[e_{id}][]$ of empty lists. Iterate over

3

Algorithm Wei: Algorithm due to Weihe [35]

Input: Hypergraph $(V_{\text{in}}, E_{\text{in}})$.

Output: Hypergraph $(V_{\text{out}}, E_{\text{out}})$ that has a hitting set of size k if and only if $(V_{\text{in}}, E_{\text{in}})$ has.

Exhaustively apply the following two data reduction rules:

1. If, for some vertex v, all edges containing v also contain some vertex $u \neq v$, then delete v.

2. If there are two edges $e \subseteq e'$, then delete e'.

Return the result $(V_{\text{out}}, E_{\text{out}})$.

---

each $(s_{\text{id}}, v, s^e_{\text{id}}, e_{\text{id}}) \in L$. If there is no predecessor on L or the predecessor $(s'_{\text{id}}, v', s^e_{\text{id}}{}', e'_{\text{id}})$ satisfies $s_{\text{id}} \neq s'_{\text{id}}$, then initialize $v^s_{\text{id}} \leftarrow 1$. If $s_{\text{id}} = s'_{\text{id}}$ but $v \neq v'$, then increment $v^s_{\text{id}} \leftarrow v^s_{\text{id}} + 1$. Add v^s_{id} to $C[e_{\text{id}}][s^e_{\text{id}}]$ and continue.

Note that, as a by-product, Step 3 also computes $|V^s|$, which is just the largest v^s_{id}. The running time for the preprocessing is dominated by $O(nd + 2^d d \cdot m)$ for computing array $A[]$ as in Theorem 3. The space used additionally to Theorem 3 is the array $B[]$ of overall size $O(2^d d \cdot m)$, and the size-m array $C[]$. Each entry of $C[]$ is an array of size at most 2^d, whose entries are lists of length at most d. Thus, $C[]$ takes at most $O(2^d d \cdot m)$ total space.

After preprocessing, Bev can be implemented to run in $O(2^d d \cdot m)$ time as shown in Bev': it uses arrays petals$'[]$ and used$'[]$ with $2^d \cdot m$ entries each. For each $s \subseteq e \in E_{\text{in}}$, used$'[s_{\text{id}}][]$ is an array indexed by $\{1, \ldots, |V^s|\}$, where

$$\sum_{s \subseteq e \in E_{\text{in}}} |V^s| \leq \left| \{(s, e, v) \mid s \subseteq e \in E_{\text{in}}, v \in e \setminus s\} \right| \leq 2^d d \cdot m.$$

Thus, the total size of used$'[][]$ is $O(2^d d \cdot m)$. $\qquad \square$

4. Experiments

In this section, we compare the kernelizations FK and Bev and the well-known data reduction algorithm Wei. Wei does not yield problem kernels (as it does not give size bounds), does not work in linear time, yet works independently of k.

Section 4.1, describes our experimental setup, Section 4.2 presents time and memory measurements. We analyze the effect of data reduction on instances arising in data clustering (Section 4.3) and public transportation optimization (Section 4.4).

4.1. Experimental setup

FK, Bev and Wei were implemented in C++.[1] The experiments were conducted on a 3.60 GHz processor with 16 GB of RAM. The running time is measured with the standard C++ library ctime. The memory consumption is measured using the valgrind memory measurement tool. The data reduction effect is measured by comparing the number $|E_{\text{in}}|$ of input edges to the number $|E_{\text{out}}|$ of output edges.

Implementation of FK and Bev. We implemented three variants of data structures for supersets[], petals[], and used[]:

- using arrays with constant worst-case look-up time after linear-time precomputations, as described in Section 3,

- using hash tables (the unordered_map type in C++), with $O(d)$ expected look-up time (we account $O(d)$ time for computing the hash value of a set of size d),

- using balanced search trees (the map type in C++), with $O(d \log 2^d m) \subseteq O(d^2 \log n)$ worst-case look-up time (a factor d is for lexicographically comparing sets of size d).

All implementations use linear space. Additionally to the data reduction described in Section 3, FK and Bev also delete all edges that are supersets of other edges (it is easy to add this data reduction without increasing their running time or space usage).

Implementation of Wei. Our implementation of Wei uses hash tables and runs in $O(z \cdot (n^2 + nm))$ time, where z is the number of vertices deleted by Wei and we omit factors depending on hash-table look-up and d.

Combinations of algorithms. We analyze the data reduction effect of the individual algorithms as well as of their combinations, applying one data reduction algorithm to the output of previous data reduction algorithms. Since the data reduction effect of the algorithms may depend on the processing order of the edges [5, Fig. 5.3], each algorithm is applied to a *random permutation* of edges. This excludes the possibility that instance generators or previous data reduction algorithms generate particularly "friendly" input orders. The order of combining the algorithms is determined by their running times: since the running time of Wei is non-linear, it is applied last, so that it is run on a problem kernel with size independent of $n + m$. Bev is slower than FK, so it is applied after FK.

Computing k. The algorithms FK and Bev require an upper bound k on the minimum hitting set size as input. We compute k using a greedy approach: repeatedly pick a vertex with a maximum number of incident edges, add it to the hitting set, and remove all incident edges, until all edges are hit.

4.2. Time and memory measurements

In the following, we present measurements of the running time and memory usage.

Data generation. We randomly generate one instance for each combination of $d \in \{1, \ldots, 5\}$ and $m \in \{i \cdot 10^5 \mid i \in \{1, \ldots, 10\}\}$, consisting of $n = 100$ vertices and m edges of size d, each chosen with equal probability. We observed these instances to be reluctant to data reduction: the result of applying FK, Bev, and Wei to these hypergraphs was exactly the same and was limited to deleting edges that are supersets of other edges. The greedily computed upper bound k was too high for FK and Bev to apply any k-dependent data reduction. Thus, the random data gives a pessimistic estimate of the running time of FK and Bev: both algorithms iterate over each subset of each input edge and
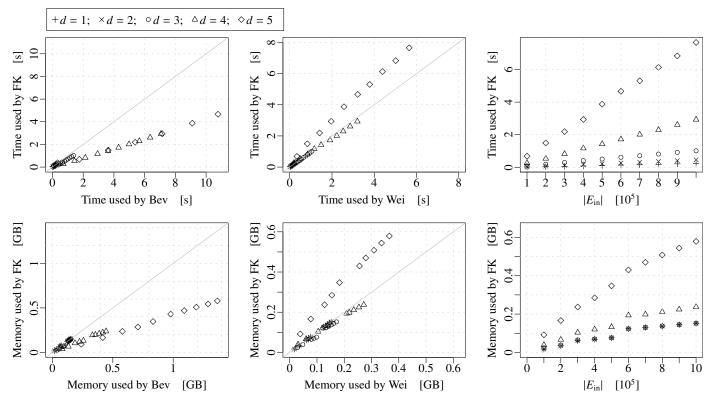

Figure 1: Comparison of running time and memory usage of FK, Bev, and Wei on the data set described in Section 4.2. Each point represents an instance. The leftmost column shows that FK outperforms Bev. The rightmost column shows time and memory usage of FK. The middle column compares FK to Wei: on this particular data set, Wei does run in linear time and outperforms FK for $d = 5$.

add almost all of them to the output hypergraph, updating their data structures. However, the random data gives a too optimistic estimate of the running time of Wei, since the number n of vertices is constant and it does not delete any vertex. Thus, on this data set, Wei also exhibits linear running-time behavior.

Results. On *each* generated instance, an implementation of FK and Bev using hash tables ran about three times faster than the implementation using arrays and uses about four times less memory. This is not surprising since the intricate precomputation described in Section 3, iterates over the input hypergraph three to four times, building four helper arrays and lists whose size each is linear in that of the input hypergraph. Thus, it seems that using tries (with O(nm) pre-initialized memory) remains the only way to outperform hash tables [5]. The balanced tree variant on *each* instance uses roughly the same amount of memory as the hash table variant and is only about 1.5 times faster than the array variant. Thus, hash tables are the most reasonable choice to implement FK and Bev in linear space. In the following, we only compare the hash table implementations of FK, Bev, and Wei. The resource consumption of these implementations is shown in Fig. 1.

4.3. Cluster Vertex Deletion

In this section, we analyze the data reduction effect of FK, Bev, and Wei on 3-Hitting Set instances arising from the Cluster Vertex Deletion problem: the task is to delete at most k vertices from a graph so that each connected component in the remaining graph is a clique [21]. A Cluster Vertex Deletion instance (G, k)

with $G = (V, E)$ can be reduced to a 3-Hitting Set instance (H, k) with $H = (V, \{e \subseteq V : G[e] \text{ is a path on three vertices}\})$ [21].

Data acquisition. We used Cluster Vertex Deletion instances arising when clustering real-world protein similarity graphs initially used by Rahmann et al. [30].[2] In fact, they used these graphs as instances for the weighted Cluster Editing problem, where one adds and deletes edges instead of deleting vertices. As suggested by Rahmann et al. [30], we create an edge between two proteins if their similarity score is positive. Since our problem is unweighted, we ignore edge weights. The data set contains hypergraphs up to 10^8 edges. However, due to the high running time of Wei, we only used those with up to 10^6 edges in our comparison (discarding 12 hypergraphs).

Results. FK processed each instance in under 1.1 seconds, Bev in under 2.5 seconds, Wei in under 7.3 seconds.

Fig. 2 shows that the data reduction effect of FK and Bev is about the same—stronger than that of Wei. Indeed, it shows that FK and Bev work well on the very same instances, which is surprising, since they are based on different data reduction criteria and it seems that Bev should be able to apply data reduction earlier than FK.

Fig. 3 shows the absolute and relative data reduction effect of FK, Bev, Wei, and their combinations. Combining any of

---

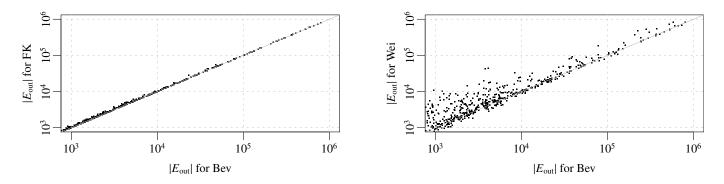[2]Available at `https://bio.informatik.uni-jena.de/data/` as `biological_bielefeld.zip`

Figure 2: Comparison between Bev, FK, and Wei on the 3-Hitting Set instances described in Section 4.3. Each point represents an instance. The left-hand plot shows that FK and Bev are about on par. Thus, the right-hand plot shows that FK and Bev outperform Wei.
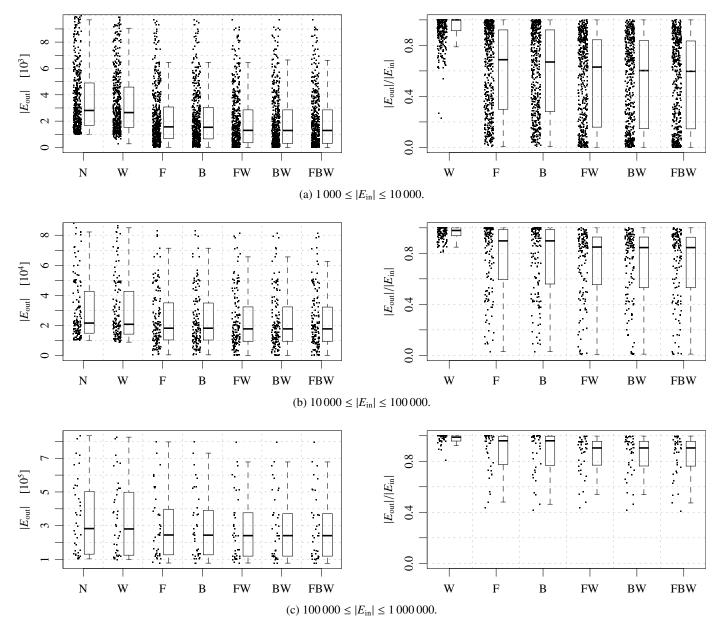

(a) $1\,000 \leq |E_{in}| \leq 10\,000$.

(b) $10\,000 \leq |E_{in}| \leq 100\,000$.

(c) $100\,000 \leq |E_{in}| \leq 1\,000\,000$.

Figure 3: Compression of the 3-Hitting Set instances described in Section 4.3 for different orders of magnitude of $|E_{in}|$. Each dot represents an instance. The boxes show the first quartile, the median, and the third quartile. The whiskers extend up to 1.5 times the interquartile range. The columns are: N — no data reduction (repeats the distribution of $|E_{in}|$); W — Wei; F — FK; B — Bev; FW — FK followed by Wei; BW — Bev followed by Wei; FBW — FK, Bev, and Wei applied in this order.
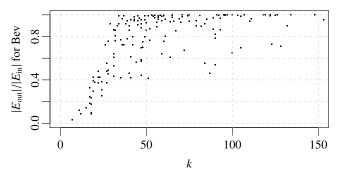
6

Figure 4: Data reduction effect on the d-Hitting Set instances described in Section 4.3, in dependence of k.
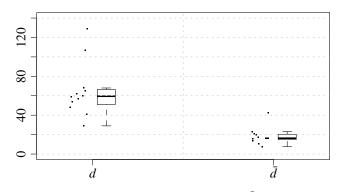

Figure 5: Maximum and average edge cardinalities, d and $\bar{d}$, respectively, of the instances in Section 4.4. The boxes show the first quartile, the median, and the third quartile. The whiskers extend up to 1.5 times the interquartile range.
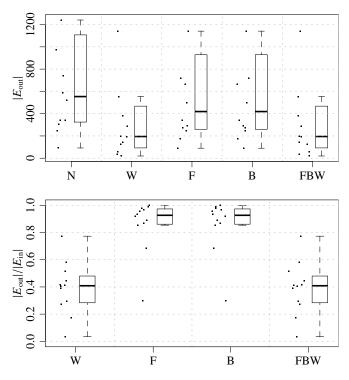

Figure 6: Compression of the instances described in Section 4.4. Each dot represents an instance. The boxes show the first quartile, the median, and the third quartile. The whiskers extend up to 1.5 times the interquartile range. The columns are as follows: N — no data reduction (repeats the distribution of $|E_{\text{in}}|$); W — Wei; F — FK; B — Bev; FBW — FK, Bev, and Wei applied in this order.

FK and Bev with Wei significantly improves the data reduction effect compared to the single algorithms, whereas combining all three algorithms shows no improvement. Also, the data reduction effect on larger instances can be seen to be lower than that on smaller instances. As Fig. 4 shows, this effect is mainly determined by k: not much data reduction is happening for $k \geq 50$, and we observed that almost all instances with $|E_{\text{in}}| \geq 10^5$ have $k \geq 50$ for our greedily computed value of k.

4.4. Station Cover

Wei is well known for its data reduction effect on the Station Cover problem in real-world transportation networks [35]. In this section, we compare FK, Bev, and Wei on the corresponding d-Hitting Set instances.

Data acquisition. We applied FK, Bev, and Wei to twelve d-Hitting Set instances modeling the Station Cover problem in European transportation networks (of cities, rural areas, and countries) that were kindly made available to us by Bläsius et al. [9]

Results. As shown by Fig. 5, the hypergraphs in this data set have much larger edges than the ones arising from Cluster Vertex Deletion in Section 4.3. Thus, as shown in Fig. 6, Wei obviously outperforms FK and Bev. Moreover, FK and Bev work equally bad and even combining all three algorithms did not yield any additional data reduction effect compared to Wei.

Indeed, due to high values of d, Wei outperforms FK and Bev even with respect to running time. FK and Bev are not even

applicable to the shown instances right away: it is infeasible to iterate over all 2^d subsets of an edge e of size d when d is large (say, 10). In such cases, we do not iterate over all subsets of e, but only over intersections of e with other edges and its subsets of size bounded by $d' \leq d$ (we used $d' = 1$). That is, we do not iterate over 2^d subsets, but only over m subsets of e. Using this implementation trick, the running times of FK and Bev were under six seconds on each instance, the running time of Wei was under one second on each instance.

5. Conclusion

We presented the first linear-time and linear-space kernelizations for d-Hitting Set, and thus resolved the apparent paradox that worst-case linear running times of d-Hitting Set kernelizations hinge on quadratic-size and partly initialized memory [5, 16].

We also conducted the first experimental evaluation of FK, significantly extended previous experimental results for Bev, and compared them to the well-known Wei data reduction algorithm. The experiments show that Wei is outperformed by FK and Bev on hypergraphs of small edge cardinality when one has good upper bounds on the hitting set size. Otherwise, Wei outperforms FK and Bev. The data reduction effect of Wei can be strengthened by applying FK and Bev in advance. Thus, the algorithms complement each other. We have also seen that, the data reduction of FK and Bev is about equal, yet FK is significantly faster.

Given that FK often kernelizes instances within a fraction of a second and yields good data reduction results when k is small,

it seems to be a good candidate for effectively applying the technique of *interleaving* kernelization and branching [27]: in a branch-and-bound algorithm for computing minimum hitting sets, FK can be applied with k set to the size of the smallest hitting set found so far, giving a good upper bound. Herein, there is no need to apply FK to the input hypergraph each time, but to the hypergraphs already kernelized on higher levels of the search tree.

References

[1] F. N. Abu-Khzam, A kernelization algorithm for *d*-Hitting Set, Journal of Computer and System Sciences 76 (2010) 524–531, doi:10.1016/j.jcss.2009.09.002.

[2] J. Alber, N. Betzler, R. Niedermeier, Experiments on data reduction for optimal domination in networks, Annals of Operations Research 146 (2006) 105–117, doi:10.1007/s10479-006-0045-4.

[3] M. Bannach, Z. Heinrich, R. Reischuk, T. Tantau, Dynamic kernels for Hitting Set and Set Packing, ECCC preprint TR19-146, 2019, URL https://eccc.weizmann.ac.il/report/2019/146/.

[4] M. Bannach, T. Tantau, Computing hitting set kernels by AC_0-circuits, Theory of Computing Systems 62 (2020) 374–399, doi:10.1007/s00224-019-09941-z.

[5] R. van Bevern, Fixed-Parameter Linear-Time Algorithms for NP-hard Graph and Hypergraph Problems Arising in Industrial Applications, Universitätsverlag der TU Berlin, Berlin, Germany, volume 1 of Foundations of Computing, chapter 5, 2014, pp. 123–150, doi:10.14279/depositonce-4131.

[6] R. van Bevern, Towards optimal and expressive kernelization for *d*-Hitting Set, Algorithmica 70 (2014) 129–147, doi:10.1007/s00453-013-9774-3.

[7] R. van Bevern, T. Fluschnik, O. Yu. Tsidulko, On approximate data reduction for the Rural Postman Problem: Theory and experiments, Networks (accepted for publication, 2020), URL https://arxiv.org/abs/1812.10131.

[8] R. van Bevern, H. Moser, R. Niedermeier, Approximation and tidying—a problem kernel for *s*-Plex Cluster Vertex Deletion, Algorithmica 62 (2012) 930–950, doi:10.1007/s00453-011-9492-7.

[9] T. Bläsius, P. Fischbeck, T. Friedrich, M. Schirneck, Understanding the effectiveness of data reduction in public transportation networks, in: K. Avrachenkov, P. Prałat, N. Ye (Eds.), WAW 2019, Springer, volume 11631 of Lecture Notes in Computer Science, 2019, pp. 87–101, doi:10.1007/978-3-030-25070-6_7.

[10] T. Bläsius, T. Friedrich, J. Lischeid, K. Meeks, M. Schirneck, Efficiently enumerating hitting sets of hypergraphs arising in data profiling, in: S. Kobourov, H. Meyerhenke (Eds.), 2019 Proceedings of the Meeting on Algorithm Engineering and Experiments (ALENEX), SIAM, 2019, pp. 130–143, doi:10.1137/1.9781611975499.11.

[11] G. Brewka, M. Thimm, M. Ulbricht, Strong inconsistency, Artificial Intelligence 267 (2019) 78–117, doi:10.1016/j.artint.2018.11.002.

[12] T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, Introduction to Algorithms, MIT Press, Cambridge, Massachusetts, USA, 2nd edition, 2001.

[13] P. Damaschke, Parameterized enumeration, transversals, and imperfect phylogeny reconstruction, Theoretical Computer Science 351 (2006) 337–350, doi:10.1016/j.tcs.2005.10.004.

[14] H. Dell, D. van Melkebeek, Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses, Journal of the ACM 61 (2014) 23, doi:10.1145/2629620.

[15] S. Fafianie, S. Kratsch, Streaming kernelization, in: E. Csuhaj-Varjú, M. Dietzfelbinger, Z. Ésik (Eds.), MFCS 2014, Springer, volume 8635 of Lecture Notes in Computer Science, 2014, pp. 275–286, doi:10.1007/978-3-662-44465-8_24.

[16] S. Fafianie, S. Kratsch, A shortcut to (sun)flowers: Kernels in logarithmic space or linear time, in: G. F. Italiano, G. Pighizzini, D. T. Sannella (Eds.), MFCS 2015, Springer, volume 9235 of Lecture Notes in Computer Science, 2015, pp. 299–310, doi:10.1007/978-3-662-48054-0_25.

[17] K. Fazekas, F. Bacchus, A. Biere, Implicit hitting set algorithms for maximum satisfiability modulo theories, in: D. Galmiche, S. Schulz, R. Sebastiani (Eds.), IJCAR 2018, Springer, volume 10900 of Lecture Notes in Computer Science, 2018, pp. 134–151, doi:10.1007/978-3-319-94205-6_10.

[18] J. Flum, M. Grohe, Parameterized Complexity Theory, Texts in Theoretical Computer Science, An EATCS Series, Springer, 2006, doi:10.1007/3-540-29953-X.

[19] F. V. Fomin, D. Lokshtanov, S. Saurabh, M. Zehavi, Kernelization, Cambridge University Press, Cambridge, England, 2019, doi:10.1017/9781107415157.

[20] V. Froese, R. van Bevern, R. Niedermeier, M. Sorge, Exploiting hidden structure in selecting dimensions that distinguish vectors, Journal of Computer and System Sciences 82 (2016) 521–535, doi:10.1016/j.jcss.2015.11.011.

[21] F. Hüffner, C. Komusiewicz, H. Moser, R. Niedermeier, Fixed-parameter algorithms for cluster vertex deletion, Theory of Computing Systems 47 (2010) 196–217, doi:10.1007/s00224-008-9150-x.

[22] R. M. Karp, Reducibility among combinatorial problems, in: R. E. Miller, J. W. Thatcher, J. D. Bohlinger (Eds.), Complexity of Computer Computations, The IBM Research Symposia Series, Springer, 1972, pp. 85–103, doi:10.1007/978-1-4684-2001-2_9.

[23] S. Kratsch, Polynomial kernelizations for MIN $F^+\Pi_1$ and MAX NP, Algorithmica 63 (2012) 532–550, doi:10.1007/s00453-011-9559-5.

[24] D. Mellor, E. Prieto, L. Mathieson, P. Moscato, A kernelisation approach for multiple *d*-hitting set and its application in optimal multi-drug therapeutic combinations, PLOS ONE 5 (2010) e0013055, doi:10.1371/journal.pone.0013055.

[25] E. Moreno-Centeno, R. M. Karp, The implicit hitting set approach to solve combinatorial optimization problems with an application to multigenome alignment, Operations Research 61 (2013) 453–468, doi:10.1287/opre.1120.1139.

[26] H. Moser, Finding Optimal Solutions for Covering and Matching Problems, Cuvillier, Göttingen, Germany, 2010, URL https://www.hmoser.info/diss_moser.pdf.

[27] R. Niedermeier, P. Rossmanith, A general method to speed up fixed-parameter-tractable algorithms, Information Processing Letters 73 (2000) 125–129, doi:10.1016/S0020-0190(00)00004-1.

[28] R. Niedermeier, P. Rossmanith, An efficient fixed-parameter algorithm for 3-Hitting Set, Journal of Discrete Algorithms 1 (2003) 89–102, doi:10.1016/S1570-8667(03)00009-1.

[29] R. O'Callahan, J.-D. Choi, Hybrid dynamic data race detection, in: R. Eigenmann, M. Rinard (Eds.), PPoPP'03, ACM, 2003, pp. 167–178, doi:10.1145/781498.781528.

[30] S. Rahmann, T. Wittkop, J. Baumbach, M. Martin, A. Truss, S. Böcker, Exact and heuristic algorithms for weighted cluster editing, in: P. Markstein, Y. Xu (Eds.), Computational Systems Bioinformatics, World Scientific Publishing, volume 6, 2007, pp. 391–401, doi:10.1142/9781860948732_0040.

[31] R. Reiter, A theory of diagnosis from first principles, Artificial Intelligence 32 (1987) 57–95, doi:10.1016/0004-3702(87)90062-2.

[32] P. V. Smirnov, Reduktsiya dannykh dlya zadachi o vershinnom pokrytii gipergrafa za lineinoe vremya s lineinoi pamyat'yu, Bachelor's thesis, Novosibirsk State University, Novosibirsk, Russian Federation, 2018, URL http://rvb.su/pdf/Smi18.pdf.

[33] M. Sorge, H. Moser, R. Niedermeier, M. Weller, Exploiting a hypergraph model for finding Golomb rulers, Acta Informatica 51 (2014) 449–471, doi:10.1007/s00236-014-0202-1.

[34] A. Vazquez, Optimal drug combinations and minimal hitting sets, BMC Systems Biology 3 (2009) 81, doi:10.1186/1752-0509-3-81.

[35] K. Weihe, Covering trains by stations or the power of data reduction, in: R. Battiti, A. A. Bertossi (Eds.), Proceedings of Algorithms and Experiments (ALEX 1998), 1998, pp. 1–8.