

h-Index Manipulation by Undoing Merges

René van Bevern^{1,2} and Christian Komusiewicz³
 Hendrik Molter⁴ and Rolf Niedermeier⁴ and Manuel Sorge⁴
 Toby Walsh⁵

Abstract. The h-index is an important bibliographic measure used to assess the performance of researchers. Van Bevern et al. [*Artif. Intel.*, to appear] showed that, despite computational worst-case hardness results, substantial manipulation of the h-index of Google Scholar author profiles is possible by merging articles. Complementing this work, we study the opposite operation, the splitting of articles, which is arguably the more natural operation for manipulation and which is also allowed within Google Scholar. We present numerous results on computational complexity (from linear-time algorithms to parameterized computational hardness results) and empirically indicate that at least small improvements of the h-index by splitting merged articles are easily achievable.

1 INTRODUCTION

Lesk [14] pointed out that the h-index⁶ is the modern equivalent of the old saying “Deans can’t read, they can only count.” He also remarked that the idea of “least publishable units” by dividing one’s reports into multiple (short) papers has been around since the 1970s. A modern version of this scenario is the manipulation of an author’s h-index by splitting a publication in an author’s profile into different versions of an article (which may result in an increased h-index). Google Scholar permits such splitting. We study such manipulation in this work, introducing and discussing several models, performing a thorough analysis of the computational complexity, and providing experimental results.

Our main points of reference are three recent publications dealing with the manipulation of the h-index, particularly motivated by Google Scholar author profile manipulation [2, 13, 16]. Indeed, we will closely follow the notation and concepts introduced by van Bevern et al. [2] and we refer to this work for discussion of related work concerning strategic self-citations to manipulate the h-index [1, 4, 17], other citation indices [6, 16, 18], and manipulation in general [7, 8]. The main difference to these previous publications is that they focus on *merging* articles for increasing the h-index [2, 13, 16] or other indices like the g-index and the i10-index [16], while we focus on *splitting*.

In the manipulation scenario for merging the assumption is that an author has a publication profile, for example in Google Scholar, that

consists of single articles and aims to increase his or her h-index by merging articles. This will result in a new article with a potentially higher number of citations. The merging option is provided by Google Scholar to identify different versions of the same article, for example a journal version and a conference version.

In the case of splitting, we assume that, most of the time, an author will maintain a correct profile in which all necessary merges are performed. Some of these merges may decrease the h-index. For instance, this can be the case when the two most cited papers are the conference and journal version of the same article. A very realistic scenario is that at certain times, for example when being evaluated by their dean, an author may temporarily undo some of these merges to increase artificially his or her h-index. A further point which distinguishes manipulation by splitting from manipulation by merging is that for merging it is easier to detect whether someone cheats too much. This can be done by looking at the titles of merged articles [2]. In contrast, it is much harder to prove that someone is manipulating by splitting; the manipulator can always claim to be too busy or that he or she does not know how to operate the profile.

The main theoretical conclusion from our work is that h-index manipulation by splitting merged articles⁷ is typically computationally easier than manipulation by merging. Hence, undoing all merges and then merging from scratch might be intractable in cases while, on the contrary, computing an optimal splitting is computationally feasible. The only good news (and, in a way, a recommendation) in this sense is that if one would use the citation measure “fusionCite” as defined by van Bevern et al. [2], then manipulation is computationally much harder than for the “unionCite” measure used by Google Scholar. We also experimented with data from Google Scholar profiles [2].

Models for Splitting Articles. We consider the publication profile of an author and denote the articles in this profile by $W \subseteq V$, where V is the set of all articles. Following previous work [2], we call these articles *atomic*. Merging articles yields a partition \mathcal{P} of W in which each part $P \in \mathcal{P}$ with $|P| \geq 2$ is a *merged article*.

Given a partition \mathcal{P} of W , the aim of splitting merged articles is to find a refined partition \mathcal{R} of \mathcal{P} with a large h-index, where the *h-index of a partition* \mathcal{P} is the largest number h such that there are at least h parts $P \in \mathcal{P}$ whose number $\mu(P)$ of citations is at least h . Herein, we have multiple possibilities of defining the number $\mu(P)$ of citations of an article in \mathcal{P} [2]. The first one, $\text{sumCite}(P)$, was introduced by de Keijzer and Apt [13], and is simply the sum of the citations of each atomic article in P . Subsequently, van Bevern et al. [2] introduced the more realistic citation measures unionCite

¹ Novosibirsk State University, Novosibirsk, Russian Federation, rvb@nsu.ru

² Sobolev Institute of Mathematics, Siberian Branch of the Russian Academy of Sciences, Novosibirsk, Russian Federation

³ Institut für Informatik, Friedrich-Schiller-Universität Jena, Germany, christian.komusiewicz@uni-jena.de

⁴ Institut für Softwaretechnik und Theoretische Informatik, TU Berlin, Germany, {h.molter,rolf.niedermeier,manuel.sorge}@tu-berlin.de

⁵ TU Berlin, Germany, University of New South Wales and Data61, Sydney, Australia, toby.walsh@nicta.com.au

⁶ The h-index of a researcher is the maximum number h such that he or she has at least h articles each cited at least h times [10].

⁷ Google Scholar allows authors to group different versions of an article. We call the resulting grouping a *merged article*. Google Scholar author profiles typically contain many merged articles, e.g. an arXiv version with a conference version and with a journal version.




Figure 1: Vertices represent articles, arrows represent citations, numbers are citation counts. The articles on a gray background in (a) have been merged in (b)–(d), and citation counts are given according to the measures sumCite, unionCite, and fusionCite, respectively. The arrows represent the citations counted by the corresponding measure.

(used by Google Scholar), where we take the cardinality of the union of the citations, and fusionCite, where we additionally remove self-citations of merged articles as well as duplicate citations between merged articles. In generic definitions, we denote these measures by μ , see Figure 1 for an illustration and Section 2 for the formal definitions. Note that, to compute these citation measures, we need a *citation graph*, a directed graph whose vertices represent articles and in which an arc from a vertex u to a vertex v means that article u cites article v .

In this work, we introduce three different operations that may be used for undoing merges in a merged article a :

- Atomizing*: splitting a into all its atomic articles,
- Extracting*: splitting off a single atomic article from a , and
- Dividing*: splitting a into two parts arbitrarily.

See Figure 2 for an illustration of the three splitting operations. Note that the atomizing, extracting, and dividing operations are successively more powerful in the sense that successively larger h-indices can be achieved. Google Scholar offers the extraction operation.

The three splitting operations lead to three problem variants, each taking as input a citation graph $D = (V, A)$, a set $W \subseteq V$ of articles belonging to the author, a partition \mathcal{P} of W that defines already-merged articles, and a non-negative integer h denoting the h-index to achieve. For $\mu \in \{\text{sumCite}, \text{unionCite}, \text{fusionCite}\}$, we define the following problems.

ATOMIZING(μ)

Question: Is there a partition \mathcal{R} of W such that

- i) for each $R \in \mathcal{R}$ either $|R| = 1$ or there is a $P \in \mathcal{P}$ such that $R = P$,
- ii) the h-index of \mathcal{R} is at least h with respect to μ ?

EXTRACTING(μ)

Question: Is there a partition \mathcal{R} of W such that

- i) for each $R \in \mathcal{R}$ there is a $P \in \mathcal{P}$ such that $R \subseteq P$,
- ii) for each $P \in \mathcal{P}$ we have $|\{R \in \mathcal{R} \mid R \subset P \text{ and } |R| > 1\}| \leq 1$,
- iii) the h-index of \mathcal{R} is at least h with respect to μ ?

DIVIDING(μ)

Question: Is there a partition \mathcal{R} of W such that

- i) for each $R \in \mathcal{R}$ there is a $P \in \mathcal{P}$ such that $R \subseteq P$,
- ii) the h-index of \mathcal{R} is at least h with respect to μ ?

Conservative Splitting. We study for each of the problem variants an additional upper bound on the number of merged articles that are split. We call these variants *conservative*: if an insincere author would like to manipulate his or her profile temporarily, then he or she would prefer a manipulation that can be easily undone. To formally define CONSERVATIVE ATOMIZING, CONSERVATIVE EXTRACTING, and CONSERVATIVE DIVIDING, we add the following restriction to the partition \mathcal{R} : “the number $|\mathcal{P} \setminus \mathcal{R}|$ of changed articles is at most k ”.

A further motivation for the conservative variants is that, in a Google Scholar profile, an author can click on a merged article and




Figure 2: Vertices represent articles, arrows represent citations, numbers are citation counts. The articles on a gray background have been merged in the initial profile (a) and correspond to remaining merged articles after applying one operation in (c) and (d). Each (merged) article has the same citation count, regardless of the used measure sumCite, unionCite, and fusionCite.

tick a box for each atomic article that he or she wants to extract. Since Google Scholar uses the unionCite measure [2], CONSERVATIVE EXTRACTING(unionCite) thus corresponds closely to manipulating the Google Scholar h-index via few of the splitting operations available to the user.

Cautious Splitting. For each splitting operation, we also study an upper bound k on the number of operations. Following our previous work [2], we call this variant *cautious*. In the case of atomizing, conservativity and caution coincide since exactly one operation is performed per changed article. Thus, we obtain two cautious problem variants: CAUTIOUS EXTRACTING and CAUTIOUS DIVIDING. For both we add the following restriction to the partition \mathcal{R} : “the number $|\mathcal{R}| - |\mathcal{P}|$ of extractions (or divisions, respectively) is at most k ”. In both variants we consider k to be part of the input.

Our results. Our theoretical (complexity classification) results are summarized in Table 1. The measures sumCite and unionCite behave basically the same. In particular, in case of atomizing and extracting, manipulation is doable in linear time, while fusionCite mostly leads to (parameterized) intractability, that is, to high worst-case computational complexity. Moreover, the dividing operation (the most general one) seems to lead to computationally much harder problems than atomizing and extracting. As indicated in Table 1, the computational complexity of two specific problems remains open.

We performed experiments with real-world data [2] and the mentioned linear-time algorithms, in particular for the case directly relevant to Google Scholar, that is, using the extraction operation and the unionCite measure. Our general findings are that increases of the h-index by one or two typically are easily achievable with few operations. The good news is that dramatic manipulation opportunities due to splitting are rare. They cannot be excluded, however, and they could be easily executed when relying on standard operations and measures (as used in Google Scholar). Working with fusionCite instead of the other two could substantially hamper manipulation.

2 PRELIMINARIES

Throughout this work, we use $n := |V|$ for the number of input articles and $m := |A|$ for the overall number of arcs in the input citation graph $D = (V, E)$. Let $\text{deg}^{\text{in}}(v)$ denote the indegree of an article v in a citation graph $D = (V, A)$, that is, v ’s number of citations. Furthermore, let $N_D^{\text{in}}(v) := \{u \mid (u, v) \in A\}$ denote the set of articles that cite v and $N_{D-W}^{\text{in}}(v) := \{u \mid (u, v) \in A \wedge u \notin W\}$ be the set of articles outside W that cite v . For each part $P \in \mathcal{P}$, the

Table 1: Computational complexity of the various variants of manipulating the h-index by splitting operations. For all FPT and W[1]-hardness results we also show NP-hardness.

†: wrt. parameter h , the h-index to achieve.

◊: wrt. parameter k , the number of operations.

★: wrt. parameter $h + k + s$, where s is the largest number of articles merged into one.

Problem	sumCite / unionCite	fusionCite
Atomizing	Linear (Theorem 1)	FPT† (Theorems 5, 6)
Conservative A.	Linear (Theorem 1)	W[1]-h★ (Theorem 7)
Extracting	Linear (Theorem 2)	?
Conservative E.	Linear (Theorem 2)	W[1]-h★ (Corollary 1)
Cautious E.	Linear (Theorem 2)	W[1]-h★ (Corollary 1)
Dividing	FPT† (Theorem 3)	NP-h (Proposition 1)
Conservative D.	FPT† (Theorem 3)	W[1]-h★ (Corollary 1)
Cautious D.	W[1]-h◊ (Theorem 4)	W[1]-h★ (Corollary 1)

following three measures for the number $\mu(P)$ of citations of P have been introduced [2]. They are illustrated in Figure 1. The measure

$$\text{sumCite}(P) := \sum_{v \in P} \text{deg}^{\text{in}}(v)$$

defines the number of citations of a merged article P as the sum of the citations of the atomic articles it contains. This measure was proposed by de Keijzer and Apt [13]. In contrast, the measure

$$\text{unionCite}(P) := \left| \bigcup_{v \in P} N_D^{\text{in}}(v) \right|$$

defines the number of citations of a merged article P as the number of distinct atomic articles citing at least one atomic article in P . Google Scholar uses the unionCite measure [2]. The measure

$$\text{fusionCite}(P) := \left| \bigcup_{v \in P} N_{D-W}^{\text{in}}(v) \right| + \sum_{P' \in \mathcal{P} \setminus \{P\}} \begin{cases} 1 & \text{if } \exists v \in P' \exists w \in P : (v, w) \in A, \\ 0 & \text{otherwise} \end{cases}$$

is perhaps the most natural one: at most one citation of a part $P' \in \mathcal{P}$ to a part $P \in \mathcal{P}$ is counted, that is, we additionally remove duplicate citations between merged articles and self-citations of merged articles.

Our theoretical analysis is in the framework of parameterized complexity [3, 5, 9, 15]. That is, for those problems that are NP-hard, we study the influence of a *parameter*, an integer associated with the input, on the computational complexity. For a problem P , we seek to decide P using a *fixed-parameter algorithm*, an algorithm with running time $f(p) \cdot |q|^{O(1)}$, where q is the input and $f(p)$ a computable function depending only on the parameter p . If such an algorithm exists, then P is *fixed-parameter tractable* (FPT) with respect to p . W[1]-hard parameterized problems presumably do not admit FPT algorithms. For instance, to find an order- k clique in an undirected graph is known to be W[1]-hard for the parameter k . W[1]-hardness of a problem P parameterized by p can be shown via a *parameterized reduction* from a known W[1]-hard problem Q parameterized by q . That is, a reduction that runs in $f(q) \cdot n^{O(1)}$ time on input of size n with parameter q and produces instances that satisfy $p \leq f(q)$ for some function f .

3 SUM CITE AND UNION CITE

In this section, we study the sumCite and unionCite measures. We provide linear-time algorithms for atomizing and extracting and analyze the parameterized complexity of dividing with respect to the

Algorithm 1: Atomizing

Input: A citation graph $D = (V, A)$, a set $W \subseteq V$ of articles, a partition \mathcal{P} of W , a nonnegative integer h and a measure μ .

Output: A partition \mathcal{R} of W .

```

1  $\mathcal{R} \leftarrow \emptyset$ 
2 foreach  $P \in \mathcal{P}$  do
3    $\mathcal{A} \leftarrow \text{Atomize}(P)$ 
4   if  $\exists A \in \mathcal{A} : \mu(A) \geq h$  then  $\mathcal{R} \leftarrow \mathcal{R} \cup \mathcal{A}$ 
5   else  $\mathcal{R} \leftarrow \mathcal{R} \cup \{P\}$ 
6 return  $\mathcal{R}$ 

```

Algorithm 2: Conservative Atomizing

Input: A citation graph $D = (V, A)$, a set $W \subseteq V$ of articles, a partition \mathcal{P} of W , nonnegative integers h and k , and a measure μ .

Output: A partition \mathcal{R} of W .

```

1  $\mathcal{R} \leftarrow \mathcal{P}$ 
2 foreach  $P \in \mathcal{P}$  do
3    $\ell_P \leftarrow 0$ 
4    $\mathcal{A} \leftarrow \text{Atomize}(P)$ 
5    $\ell_P \leftarrow \ell_P + |\{A \in \mathcal{A} \mid \mu(A) \geq h\}|$ 
6   if  $\mu(P) \geq h$  then  $\ell_P \leftarrow \ell_P - 1$ 
7 for  $i \leftarrow 1$  to  $k$  do
8    $P^* \leftarrow \arg \max_{P \in \mathcal{P}} \{\ell_P\}$ 
9   if  $\ell_{P^*} > 0$  then
10     $\mathcal{A} \leftarrow \text{Atomize}(P^*)$ 
11     $\mathcal{R} \leftarrow (\mathcal{R} \setminus \{P^*\}) \cup \mathcal{A}$ 
12     $\ell_{P^*} \leftarrow -1$ 
13 return  $\mathcal{R}$ 

```

number k of splits and the h-index h to achieve. In our results for sumCite and unionCite, we often tacitly use the observation that local changes to the merged articles do not influence the citations of other merged articles.

Manipulation by Atomizing. Recall that the atomizing operation splits a merged article into singletons and that, for the atomizing operation, the notions of *conservative* (touching few articles) and *cautious* (making few operations) manipulation coincide and are thus both captured by CONSERVATIVE ATOMIZING. Both ATOMIZING and CONSERVATIVE ATOMIZING are solvable in linear time. Intuitively, it suffices to find the merged articles which, when atomized, increase the number of articles with at least h citations the most. This leads to Algorithms 1 and 2 for ATOMIZING and CONSERVATIVE ATOMIZING. Herein, the `Atomize()` operation takes a set S as input and returns $\{\{s\} \mid s \in S\}$. The algorithms yield the following theorem.

Theorem 1. ATOMIZING(μ) and CONSERVATIVE ATOMIZING(μ) are solvable in linear time for $\mu \in \{\text{sumCite}, \text{unionCite}\}$.

Proof. We first consider ATOMIZING(μ). Let \mathcal{R} be a partition created from a partition \mathcal{P} by atomizing a part $P^* \in \mathcal{P}$. Observe that for all $P \in \mathcal{P}$ and $R \in \mathcal{R}$ we have that $P = R$ implies $\mu(P) = \mu(R)$, for $\mu \in \{\text{sumCite}, \text{unionCite}\}$. Intuitively this means that atomizing a single part $P^* \in \mathcal{P}$ does not alter the μ -value of any other part of the partition.

Algorithm 1 computes a partition \mathcal{R} that has a maximal number of parts R with $\mu(R) \geq h$ that can be created by applying atom-

Algorithm 3: Extracting

Input: A citation graph $D = (V, A)$, a set $W \subseteq V$ of articles, a partition \mathcal{P} of W , a nonnegative integer h and a measure μ .

Output: A partition \mathcal{R} of W .

```

1  $\mathcal{R} \leftarrow \emptyset$ 
2 foreach  $P \in \mathcal{P}$  do
3   foreach  $v \in P$  do
4     if  $\mu(\{v\}) \geq h$  then
5        $\mathcal{R} \leftarrow \mathcal{R} \cup \{\{v\}\}$ 
6        $P \leftarrow P \setminus \{v\}$ 
7   if  $P \neq \emptyset$  then  $\mathcal{R} \leftarrow \mathcal{R} \cup \{P\}$ 
8 return  $\mathcal{R}$ 

```

izing operations to \mathcal{P} : It applies the atomizing operation to each part $P \in \mathcal{P}$ if there is at least one singleton A in the atomization of P with $\mu(A) \geq h$. By the above observation, this cannot decrease the total number of parts in the partition that have a μ -value of at least h . Furthermore, we have that for all $R \in \mathcal{R}$, we cannot potentially increase the number of parts with μ -value at least h by atomizing R . Thus, we get the maximal number of parts R with $\mu(R) \geq h$ that can be created by applying atomizing operations to \mathcal{P} .

Obviously, if \mathcal{R} has at least h parts R with $\mu(R) \geq h$, we face a yes-instance. Conversely, if the input is a yes-instance, then there is a number of atomizing operations that can be applied to \mathcal{P} such that the resulting partition \mathcal{R} has at least h parts R with $\mu(R) \geq h$.

It is easy to see that the algorithm runs in linear time and finds a yes-instance if it exists. If the output partition \mathcal{R} does not have at least h parts R with $\mu(R) \geq h$, then the input is a no-instance.

The pseudocode for solving CONSERVATIVE ATOMIZING(μ) is given in Algorithm 2. First, in Lines 2–6, for each part P , Algorithm 2 records how many singletons A with $\mu(A) \geq h$ are created when atomizing P . Then, in Lines 7–12, it repeatedly atomizes the part yielding the most such singletons. This procedure creates the maximum number of parts that have a μ -value of at least h , since the μ -value cannot be increased by exchanging one of these atomizing operations by another.

Obviously, if \mathcal{R} has at least h parts R with $\mu(R) \geq h$, we face a yes-instance. Conversely, if the input is a yes-instance, then there are k atomizing operations that can be applied to \mathcal{P} to yield an h -index of at least h . Since Algorithm 2 takes successively those operations that yield the most new parts with h citations, the resulting partition \mathcal{R} has at least h parts R with $\mu(R) \geq h$. It is not hard to verify that the algorithm has linear running time. \square

Manipulation by Extracting. Recall that the extracting operation removes a single article from a merged article. All variants of the extraction problem are solvable in linear time. Intuitively, in the cautious case, it suffices to find k extracting operations that each increase the number of articles with h citations. In the conservative case, we determine for each merged article a set of extraction operations that increases the number of articles with h citations the most. Then we use the extraction operations for those k merged articles that yield the k largest increases in the number of articles with h citations. This leads to Algorithms 3, 4, and 5 for EXTRACTING, CAUTIOUS EXTRACTING, and CONSERVATIVE EXTRACTING, respectively, which yield the following theorem.

Theorem 2. EXTRACTING(μ), CONSERVATIVE EXTRACTING(μ)

Algorithm 4: Cautious Extracting

Input: A citation graph $D = (V, A)$, a set $W \subseteq V$ of articles, a partition \mathcal{P} of W , nonnegative integers h and k , and a measure μ .

Output: A partition \mathcal{R} of W .

```

1  $\mathcal{R} \leftarrow \emptyset$ 
2 foreach  $P \in \mathcal{P}$  do
3   foreach  $v \in P$  do
4     if  $k > 0$  and  $\mu(\{v\}) \geq h$  and  $\mu(P \setminus \{v\}) \geq h$  then
5        $\mathcal{R} \leftarrow \mathcal{R} \cup \{\{v\}\}$ 
6        $P \leftarrow P \setminus \{v\}$ 
7        $k \leftarrow k - 1$ 
8   if  $P \neq \emptyset$  then  $\mathcal{R} \leftarrow \mathcal{R} \cup \{P\}$ 
9 return  $\mathcal{R}$ 

```

Algorithm 5: Conservative Extracting

Input: A citation graph $D = (V, A)$, a set $W \subseteq V$ of articles, a partition \mathcal{P} of W , nonnegative integers h and k , and a measure μ .

Output: A partition \mathcal{R} of W .

```

1 foreach  $P \in \mathcal{P}$  do
2    $\ell_P \leftarrow 0$ 
3    $\mathcal{R}_P \leftarrow \emptyset$ 
4   foreach  $v \in P$  do
5     if  $\mu(\{v\}) \geq h$  and  $\mu(P \setminus \{v\}) \geq h$  then
6        $\mathcal{R}_P \leftarrow \mathcal{R}_P \cup \{\{v\}\}$ 
7        $P \leftarrow P \setminus \{v\}$ 
8        $\ell_P \leftarrow \ell_P + 1$ 
9   if  $P \neq \emptyset$  then  $\mathcal{R}_P \leftarrow \mathcal{R}_P \cup \{P\}$ 
10  $\mathcal{P}^* \leftarrow$  the  $k$  elements of  $P \in \mathcal{P}$  with largest  $\ell_P$ -values
11  $\mathcal{R} \leftarrow \bigcup_{P \in \mathcal{P}^*} \mathcal{R}_P \cup (\mathcal{P} \setminus \mathcal{P}^*)$ 
12 return  $\mathcal{R}$ 

```

and CAUTIOUS EXTRACTING(μ) are solvable in linear time for $\mu \in \{\text{sumCite}, \text{unionCite}\}$.

Proof. We first consider EXTRACTING(μ). Let \mathcal{R} be a partition produced from \mathcal{P} by extracting an article from a part $P^* \in \mathcal{P}$. Recall that this does not alter the μ -value of any other part, i.e., for all $P \in \mathcal{P}$ and $R \in \mathcal{R}$, we have that $P = R$ implies $\mu(P) = \mu(R)$ for $\mu \in \{\text{sumCite}, \text{unionCite}\}$.

Consider Algorithm 3. It is easy to see that the algorithm only performs extracting operations and that the running time is polynomial. So we have to argue that whenever there is a partition \mathcal{R} that can be produced by extracting operations from \mathcal{P} such that the h -index is at least h , then the algorithm finds a solution.

We show this by arguing that the algorithm produces the maximum number of articles with at least h citations possible. Extracting an article that has strictly less than h citations cannot produce an h -index of at least h unless we already have an h -index of at least h , because the number of articles with h or more citations does not increase. Extracting an article with h or more citations cannot decrease the number of articles with h or more citations. Hence, if there are no articles with at least h citations that we can extract, we cannot create more articles with h or more citations. Therefore, we have produced the maximum number of articles with h or more citations when the algorithm stops.

The pseudocode for solving `CAUTIOUS EXTRACTING(μ)` is given in Algorithm 4. We perform up to k extracting operations (Line 6). Each of them increases the number of articles that have h or more citations by one. As Algorithm 4 checks each atomic article in each merged article, it finds k extraction operations that increase the number of articles with h or more citations if they exist. Thus, it produces the maximum-possible number of articles that have h or more citations and that can be created by k extracting operations.

To achieve linear running time, we need to efficiently compute $\mu(P \setminus \{v\})$ in Line 4. This can be done by representing articles as integers and using an n -element array A which stores throughout the loop in Line 3, for each article $v \in N_D^{\text{in}}[P]$, the number $A[w]$ of articles in P that are cited by w . Using this array, one can compute $\mu(P \setminus \{v\})$ in $O(\deg^{\text{in}}(v))$ time in Line 4, amounting to overall linear time. The time needed to maintain array A is also linear: We initialize it once in the beginning with all zeros. Then, before entering the loop in Line 3, we can in $O(|N_D^{\text{in}}(P)|)$ total time store for each article $v \in N_D^{\text{in}}[P]$, the number $A[w]$ of articles in P that are cited by w . To update the array within the loop in Line 3, we need $O(\deg^{\text{in}}(v))$ time if Line 6 applies. In total, this is linear time.

Finally, the pseudocode for solving `CONSERVATIVE EXTRACTING(μ)` is given in Algorithm 5. For each merged article $P \in \mathcal{P}$, Algorithm 5 computes a set \mathcal{R}_P and the number ℓ_P of additional articles v with $\mu(v) \geq h$ that can be created by extracting. Then it chooses a set \mathcal{P}^* of k merged articles $P \in \mathcal{P}$ with maximum ℓ_P and, from each $P \in \mathcal{P}^*$, extracts the articles in \mathcal{R}_P .

This procedure creates the maximum number of articles that have a μ -value of at least h while only performing extraction operations on at most k merges.

Obviously, if the solution \mathcal{R} has at least h parts R with $\mu(R) \geq h$, then we face a yes-instance. Conversely, if the input is a yes-instance, then there are k merged articles that we can apply extraction operations to, such that the resulting partition \mathcal{R} has at least h parts R with $\mu(R) \geq h$. Since the algorithm produces the maximal number of parts R with $\mu(R) \geq h$, it achieves an h -index of at least h .

The linear running time follows by implementing the check in line 5 in $O(\deg^{\text{in}}(v))$ time as described for Algorithm 4 and by using counting sort to find the k parts to extract from in line 10. \square

Manipulation by Dividing. Recall that the dividing operation splits a merged article into two arbitrary parts. First we consider the basic and the conservative case and show that they are FPT when parameterized by the h -index. Then we show that the cautious variant is W[1]-hard when parameterized by k . `DIVIDING(μ)` is closely related to `H-INDEX MANIPULATION(μ)` [2, 13] which is, given a citation graph $D = (V, A)$, a subset of articles $W \subseteq V$, and a non-negative integer h , to decide whether there is a partition \mathcal{P} of W such that \mathcal{P} has h -index h with respect to μ . De Keijzer and Apt [13] showed that `H-INDEX MANIPULATION(sumCite)` is NP-hard, even if merges are unconstrained. The hardness of `H-INDEX MANIPULATION` for $\mu \in \{\text{unionCite}, \text{fusionCite}\}$ follows. We can reduce `H-INDEX MANIPULATION` to `CONSERVATIVE DIVIDING` by defining the partition $\mathcal{P} = \{W\}$, hence we get the following.

Proposition 1. `DIVIDING` and `CONSERVATIVE DIVIDING` are NP-hard for $\mu \in \{\text{sumCite}, \text{unionCite}, \text{fusionCite}\}$.

As to computational tractability, `DIVIDING` and `CONSERVATIVE DIVIDING` are FPT when parameterized by h —the h -index to achieve.

Theorem 3. `DIVIDING` and `CONSERVATIVE DIVIDING(μ)` can be solved in $2^{O(h^4 \log h)} \cdot n^{O(1)}$ time, where h is the h -index to achieve and $\mu \in \{\text{sumCite}, \text{unionCite}\}$.

Algorithm 6: Conservative Dividing

Input: A citation graph $D = (V, A)$, a set $W \subseteq V$ of articles, a partition \mathcal{P} of W , nonnegative integers h and k , and a measure μ .

Output: `true` if k dividing operations can be applied to \mathcal{P} to yield h -index h and `false` otherwise.

```

1 foreach  $P \in \mathcal{P}$  do
2    $D' \leftarrow$  The graph obtained from  $D$  by removing all
   citations  $(u, v)$  such that  $v \notin P$  and adding
    $h + 1$  articles  $r_1, \dots, r_{h+1}$ 
3    $W' \leftarrow P, \ell_P \leftarrow 0$ 
4   for  $i \leftarrow 0$  to  $h$  do
5     if Merge ( $D', W', h, \mu$ ) then
6        $\ell_P \leftarrow h - i$ 
7       Break
8     Add  $r_i$  to  $W'$  and add each citation  $(r_i, r_j)$ ,
        $j \in \{1, \dots, h + 1\} \setminus \{i\}$  to  $D'$ 
9 return  $\exists \mathcal{P}' \subseteq \mathcal{P}$  s.t.  $|\mathcal{P}'| \leq k$  and  $\sum_{P \in \mathcal{P}'} \ell_P \geq h$ 

```

Proof. The pseudocode is given in Algorithm 6. Herein, `Merge` (D, W, h, μ) decides `H-INDEX MANIPULATION(μ)`, that is, it returns `true` if there is a partition \mathcal{Q} of W such that \mathcal{Q} has h -index h and `false` otherwise. It follows from van Bevern et al. [2, Theorem 7] that `Merge` can be carried out in $2^{O(h^4 \log h)} \cdot n^{O(1)}$ time.

Algorithm 6 first finds, using `Merge`, the maximum number ℓ_P of (merged) articles with at least h citations that we can create in each part $P \in \mathcal{P}$. For this, we first prepare an instance (D', W', h, μ) of `H-INDEX MANIPULATION(μ)` in Lines 2 and 3. In the resulting instance, we ask whether there is a partition of P with h -index h . If this is the case, then we set ℓ_P to h and, otherwise, we add one artificial article with h citations to W' in Line 8. Then we use `Merge` again and we iterate this process until `Merge` returns `true`, or we find that there is not even one merged article contained in P with h citations. Clearly, this process correctly computes ℓ_P . Thus, the algorithm is correct. The running time is clearly dominated by the calls to `Merge`. Since `Merge` runs in $2^{O(h^4 \log h)} \cdot n^{O(1)}$ time [2, Theorem 7], the running time bound follows. \square

We note that `Merge` can be modified so that it outputs the desired partition. Hence, we can modify Algorithm 6 to output the actual solution. Furthermore, for $k = n$, Algorithm 6 solves the non-conservative variant, which is therefore also fixed-parameter tractable parameterized by h .

In contrast, for the cautious variant we show W[1]-hardness when parameterized by k , the number of allowed operations.

Theorem 4. `CAUTIOUS DIVIDING(μ)` is NP-hard and W[1]-hard when parameterized by k for $\mu \in \{\text{sumCite}, \text{unionCite}, \text{fusionCite}\}$, even if the citation graph is acyclic.

Proof. We reduce from the `UNARY BIN PACKING` problem: given a set S of n items with integer sizes $s_i, i \in \{1, \dots, n\}$, ℓ bins and a maximum bin capacity B , can we distribute all items into the ℓ bins? Herein, all sizes are encoded in unary. `UNARY BIN PACKING` parameterized by ℓ is W[1]-hard [11].

Given an instance (S, ℓ, B) of `UNARY BIN PACKING`, we produce an instance $(D, W, \mathcal{P}, h, \ell - 1)$ of `CAUTIOUS DIVIDING(sumCite)`. Let $s^* = \sum_i s_i$ be the sum of all item sizes. We assume that $B < s^*$ and $\ell \cdot B \geq s^*$ as, otherwise, the problem is trivial, since all items fit into one bin or they collectively cannot fit into all bins, respectively.

Furthermore, we assume that $\ell < B$ since, otherwise, the instance size is upper bounded by a function of ℓ and, hence, is trivially FPT with respect to ℓ . We construct the instance of CAUTIOUS DIVIDING(sumCite) in polynomial time as follows.

- Add s^* articles x_1, \dots, x_{s^*} to D . These are only used to increase the citation count of other articles.
- Add one article a_i to D and W for each s_i .
- For each article a_i , add citations (x_j, a_i) for all $1 \leq j \leq s_i$ to G . Note that, after adding these citations, each article a_i has citation count s_i .
- Add $\Delta := \ell \cdot B - s^*$ articles u_1, \dots, u_Δ to D and W .
- For each article u_i with $i \in \{1, \dots, \Delta\}$, add a citation (x_1, u_i) to D . Note that each article u_i has citation count 1.
- Add $B - \ell$ articles $h_1, \dots, h_{B-\ell}$ to D and W .
- For each article h_i with $i \in \{1, \dots, B - \ell\}$, add citations (x_j, h_i) for all $1 \leq j \leq B$ to D . Note that each article h_i has citation count B .
- Add $P^* = \{a_1, \dots, a_n, u_1, \dots, u_\Delta\}$ to \mathcal{P} , for each article h_i with $i \in \{1, \dots, B - \ell\}$, add $\{h_i\}$ to \mathcal{P} , and set $h = B$.

Now we show that (S, ℓ, B) is a yes-instance if and only if $(D, W, \mathcal{P}, h, \ell - 1)$ is a yes-instance.

(\Rightarrow) Assume that (S, ℓ, B) is a yes-instance and let S_1, \dots, S_ℓ be a partition of S such that items in S_i are placed in bin i . Now we split P^* into ℓ parts R_1, \dots, R_ℓ in the following way. Note that for each S_i , we have that $\sum_{s_j \in S_i} s_j = B - \delta_i$ for some $\delta_i \geq 0$. Furthermore, $\sum_i \delta_i = \Delta$. Recall that there are Δ articles u_1, \dots, u_Δ in P^* . Let $\delta_{<i} = \sum_{j < i} \delta_j$ and $U_i = \{u_{\delta_{<i}+1}, \dots, u_{\delta_{<i}+\delta_i}\}$, with $\delta_0 = 0$ and if $\delta_i > 0$, let $U_i = \emptyset$ for $\delta_i = 0$. We set $R_i = \{a_j \mid s_j \in S_i\} \cup U_i$. Then for each R_i , we have that $\text{sumCite}(R_i) = \text{sumCite}(\{a_j \mid s_j \in S_i\}) + \text{sumCite}(U_i)$, which simplifies to $\text{sumCite}(R_i) = \sum_{s_j \in S_i} s_j + \delta_i = B$. For each i , $1 \leq i \leq n$, we have $\text{sumCite}(\{h_i\}) = B$. Hence, $\mathcal{R} = \{R_1, \dots, R_\ell, \{h_1\}, \dots, \{h_{B-\ell}\}\}$ has h-index B .

(\Leftarrow) Assume that $(D, W, \mathcal{P}, h, \ell - 1)$ is a yes-instance and let \mathcal{R} be a partition with h-index h . Recall that \mathcal{P} consists of P^* and $B - \ell$ singletons $\{h_1\}, \dots, \{h_{B-\ell}\}$, which are hence also contained in \mathcal{R} . Furthermore, $\text{sumCite}(\{h_i\}) = B$ for each h_i and, by the definition of the h-index, there are ℓ parts R_1, \dots, R_ℓ with $R_i \subset P^*$ and $\text{sumCite}(R_i) \geq B$ for each i . Since, by definition, $\text{sumCite}(P^*) = \ell \cdot B$ and $\text{sumCite}(P^*) = \sum_{1 \leq i \leq \ell} \text{sumCite}(R_i)$ we have that $\text{sumCite}(R_i) = B$ for all i . It follows that $\text{sumCite}(R_i \setminus \{u_1, \dots, u_\Delta\}) \leq B$ for all i . This implies that packing into bin i each item in $\{s_j \mid a_j \in R_i\}$ solves the instance (S, ℓ, B) .

Note that this proof can be modified to cover also the unionCite and the fusionCite case by adding $\ell \cdot s^*$ extra x -articles and ensuring that no two articles in W are cited by the same x -article. \square

4 FUSION CITE

We now consider the fusionCite measure, which makes manipulation considerably harder than the other measures. In particular, we obtain that even in the most basic case, the manipulation problem is NP-hard.

Theorem 5. *ATOMIZING(fusionCite) is NP-hard, even if the citation graph is acyclic.*

Proof. We reduce from the NP-hard 3-SAT problem: given a 3-CNF formula F with n variables and m clauses, decide whether F allows for a satisfying truth assignment to its variables. Without loss

of generality, we assume $n + m > 3$. Given a formula F with variables x_1, \dots, x_n and clauses c_1, \dots, c_m such that $n + m > 3$, we produce an instance $(D, W, \mathcal{P}, m + n)$ of ATOMIZING(fusionCite) in polynomial time as follows.

For each variable x_i of F , add to D and W sets $\mathcal{X}_i^F := \{X_{i,1}^F, \dots, X_{i,2(n+m)}^F\}$ and $\mathcal{X}_i^T := \{X_{i,1}^T, \dots, X_{i,2(n+m)}^T\}$ of *variable articles*. Add \mathcal{X}_i^F and \mathcal{X}_i^T to \mathcal{P} and, for $1 \leq \ell \leq n + m$, add citations $(X_{i,\ell}^F, X_{i,2\ell}^T)$ and $(X_{i,\ell}^T, X_{i,2\ell}^F)$ to D . Next, for each clause c_i of F , add a *clause article* C_i to D , to W , and add $\{C_i\}$ to \mathcal{P} . Finally, if a positive literal x_i occurs in a clause c_j , then add citations $(X_{i,\ell}^T, C_j)$ to D for $1 \leq \ell \leq n + m$. If a negative literal $\neg x_i$ occurs in a clause c_j , then add citations $(X_{i,\ell}^F, C_j)$ to D for $1 \leq \ell \leq n + m$. This concludes the construction. Observe that D is acyclic since all citations go from variable articles to clause articles or to variable articles with a higher index. It remains to show that F is satisfiable if and only if $(D, W, \mathcal{P}, m + n)$ is a yes-instance.

(\Rightarrow) If F is satisfiable, then a solution \mathcal{R} for $(D, W, \mathcal{P}, m + n)$ looks as follows: for each $i \in \{1, \dots, n\}$, if x_i is true, then we put $\mathcal{X}_i^F \in \mathcal{R}$ and we put $\mathcal{X}_i^T \notin \mathcal{R}$ otherwise. All other articles of D are added to \mathcal{R} as singletons. We count the citations that every part of \mathcal{R} gets from other parts of \mathcal{R} . If x_i is true, then \mathcal{X}_i^F gets $m + n$ citations from $\{X_{i,\ell}^T\}$ for $1 \leq \ell \leq n + m$. Moreover, for the clause c_j containing the literal x_i , $\{C_j\}$ gets $n + m$ citations from $\{X_{i,\ell}^T\}$ for $1 \leq \ell \leq n + m$. Similarly, if x_i is false, then $\{X_{i,\ell}^T\}$ gets $m + n$ citations and so does every $\{C_j\}$ for each clause c_j containing the literal $\neg x_i$. Since every clause is satisfied and every variable is either true or false, it follows that each of the m clause articles gets $m + n$ citations and that, for each of the n variables x_i , either \mathcal{X}_i^F or \mathcal{X}_i^T gets $m + n$ citations. It follows that $m + n$ parts of \mathcal{R} get at least $m + n$ citations and thus, that \mathcal{R} has h-index at least $m + n$.

(\Leftarrow) Let \mathcal{R} be a solution for $(D, W, \mathcal{P}, m + n)$. We first show that, for each variable x_i , we have either $\mathcal{X}_i^T \in \mathcal{R}$ or $\mathcal{X}_i^F \in \mathcal{R}$. To this end, it is important to note two facts:

1. For each variable x_i , every variable article in $\mathcal{X}_i^T \cup \mathcal{X}_i^F$ has at most one incoming arc in D . Thus, no (singleton) variable article in \mathcal{R} can get $m + n$ citations.
2. If, for some variable x_i , the part $\mathcal{X}_i^T \in \mathcal{R}$ gets $m + n$ citations, then $\mathcal{X}_i^F \notin \mathcal{R}$ and vice versa.

Thus, since there are at most m clause articles and \mathcal{R} contains $m + n$ parts with $m + n$ citations, \mathcal{R} contains exactly one of the parts $\mathcal{X}_i^T, \mathcal{X}_i^F$ of each variable x_i . It follows that, in \mathcal{R} , all singleton clause articles have to receive $m + n$ citations. Thus, for each clause c_j , there is a literal x_i in c_j or a literal $\neg x_i$ in c_j such that $\mathcal{X}_i^T \in \mathcal{R}$ or $\mathcal{X}_i^F \in \mathcal{R}$, respectively. It follows that setting each x_i to true if and only if $\mathcal{X}_i^T \in \mathcal{R}$ gives a satisfying truth assignment to the variables of F . \square

This NP-hardness result motivates the search for fixed-parameter tractability.

Theorem 6. *ATOMIZING(fusionCite) can be solved in $O(4^{h^2}(n + m))$ time, where h is the h-index to achieve.*

Proof. We use the following procedure to solve an instance (D, W, \mathcal{P}, h) of ATOMIZING(fusionCite).

Let $\mathcal{P}_{\geq h}$ be the set of merged articles $P \in \mathcal{P}$ with $\text{fusionCite}(P) \geq h$. If $|\mathcal{P}_{\geq h}| \geq h$, then we face a yes-instance and output “yes”. To see that we can do this in linear time, note that, given \mathcal{P} , we can compute $\text{fusionCite}(P)$ in linear time for each $P \in \mathcal{P}$. Below we assume that $|\mathcal{P}_{\geq h}| < h$.

First, we atomize all $P \in \mathcal{P}$ that cannot have h or more citations, that is, for which, even if we atomize all merged articles except

for P , we have $\text{fusionCite}(P) < h$. Formally, we atomize P if $\sum_{v \in P} |N_{D-P}^{\text{in}}(v)| < h$. Let \mathcal{P}' be the partition obtained from \mathcal{P} after these atomizing operations; note that \mathcal{P}' can be computed in linear time.

The basic idea is now to look at all remaining merged articles that receive at least h citations from atomic articles; they form the set $\mathcal{P}_{<h}$ below. They are cited by at most $h - 1$ other merged articles. Hence, if the size of $\mathcal{P}_{<h}$ exceeds some function $f(h)$, then, among the contained merged articles, we find a large number of merged articles that do not cite each other. If we have such a set, then we can atomize all other articles, obtaining h-index h . If the size of $\mathcal{P}_{<h}$ is smaller than $f(h)$, then we can determine by brute force whether there is a solution.

Consider all merged articles $P \in \mathcal{P}'$ that have less than h citations but can obtain h or more citations by applying atomizing operations to merged articles in \mathcal{P}' . Let us call the set of these merged articles $\mathcal{P}_{<h}$. Formally, $P \in \mathcal{P}_{<h}$ if $\sum_{v \in P} |N_{D-P}^{\text{in}}(v)| \geq h$ and $\text{fusionCite}(P) < h$. Again, $\mathcal{P}_{<h}$ can be computed in linear time. Note that $\mathcal{P}' \setminus (\mathcal{P}_{\geq h} \cup \mathcal{P}_{<h})$ consists only of singletons.

Now, we observe the following. If there is a set $\mathcal{P}^* \subseteq \mathcal{P}_{<h}$ of at least h merged articles such that, for all $P_i, P_j \in \mathcal{P}^*$, neither P_i cites P_j nor P_j cites P_i , then we can atomize all merged articles in $\mathcal{P}' \setminus \mathcal{P}^*$ to reach an h-index of at least h . We finish the proof by showing that we can conclude the existence of the set \mathcal{P}^* if $\mathcal{P}_{<h}$ is sufficiently large and solve the problem using brute force otherwise.

Consider the undirected graph G that has a vertex v_P for each $P \in \mathcal{P}_{<h}$ and an edge between v_{P_i} and v_{P_j} if P_i cites P_j or P_j cites P_i . Note that $\{v_P \mid P \in \mathcal{P}^*\}$ forms an independent set in G . Furthermore, let I be an independent set in G that has size at least h . Let $\mathcal{P}^{**} = \{P \in \mathcal{P}_{<h} \mid v_P \in I\}$. Then, we can atomize all merged articles in $\mathcal{P}' \setminus \mathcal{P}^{**}$ to reach an h-index of at least h .

We claim that the number of edges in G is at most $(h - 1) \cdot |\mathcal{P}_{<h}|$. This is because the edge set of G can be enumerated by enumerating for every vertex v_P the edges incident with v_P that result from a citation of P from another $P' \in \mathcal{P}_{<h}$. The citations for each P are less than h as, otherwise, we would have that $P \in \mathcal{P}_{\geq h}$. Now, we can make use of Turán's Theorem, which can be stated as follows: If a graph with ℓ vertices has at most $\ell k/2$ edges, then it admits an independent set of size at least $\ell/(k + 1)$ [12, Exercise 4.8]. Hence, if $|\mathcal{P}_{<h}| \geq 2h^2 - h$, then we face a yes-instance and we can find a solution by taking an arbitrary subset $\mathcal{P}'_{<h}$ of $\mathcal{P}_{<h}$ with $|\mathcal{P}'_{<h}| = 2h^2 - h$, by atomizing every merged article outside of $\mathcal{P}'_{<h}$, and by guessing which merged articles we need to atomize inside of $\mathcal{P}'_{<h}$. If $|\mathcal{P}_{<h}| < 2h^2 - h$, then we guess which merged articles in $\mathcal{P}_{<h} \cup \mathcal{P}_{\geq h}$ we need to atomize to obtain a solution if it exists. In both cases, for each guess we need linear time to determine whether we have found a solution, giving the overall running time of $O(4^{h^2} \cdot (m + n))$. \square

For the conservative variant, however, we cannot achieve FPT, even if we add the number of atomization operations and the maximum size of a merged article to the parameter.

Theorem 7. *CONSERVATIVE ATOMIZING(fusionCite) is NP-hard and W[1]-hard when parameterized by $h + k + s$, where $s := \max_{P \in \mathcal{P}} |P|$, even if the citation graph is acyclic.*

Proof. We reduce from the CLIQUE problem: given a graph G and an integer k , decide whether G contains a clique on at least k vertices. CLIQUE parameterized by k is known to be W[1]-hard.

Given an instance (G, k) of CLIQUE, we produce an instance $(D, W, \mathcal{P}, h, k)$ of CONSERVATIVE ATOMIZING(fusionCite) in polynomial time as follows. Without loss of generality, we assume $k \geq 4$ so that $\binom{k}{2} \geq 4$. For each vertex v of G , introduce a

set R_v of $\lceil \binom{k}{2} / 2 \rceil$ vertices to D and W and add R_v as a part to \mathcal{P} . For an edge $\{v, w\}$ of G , add to D and W a vertex $e_{\{v,w\}}$ and add $\{e_{\{v,w\}}\}$ to \mathcal{P} . Moreover, add a citation from each vertex in $R_v \cup R_w$ to $e_{\{v,w\}}$. Finally, set $h := \binom{k}{2}$. Each of h, k and s in our constructed instance of CONSERVATIVE ATOMIZING(fusionCite) depends only on k in the input CLIQUE instance. It remains to show that (G, k) is a yes-instance for CLIQUE if and only if $(D, W, \mathcal{P}, h, k)$ is.

(\Rightarrow) Assume that (G, k) is a yes-instance and let S be a clique in G . Then, atomizing R_v for each $v \in S$ yields $\binom{k}{2}$ articles with at least $\binom{k}{2}$ citations in D : for each of the $\binom{k}{2}$ pairs of vertices $v, w \in S$, the vertex $e_{\{v,w\}}$ gets $\lceil \binom{k}{2} / 2 \rceil$ citations from the vertices in R_v and the same number of citations from the vertices in R_w and, thus, at least $\binom{k}{2}$ citations in total.

(\Leftarrow) Assume that $(D, W, \mathcal{P}, h, k)$ is a yes-instance and let \mathcal{R} be a solution. We construct a subgraph $S = (V_S, E_S)$ of G that is a clique of size k . Let $V_S := \{v \in V(G) \mid R_v \in \mathcal{P} \setminus \mathcal{R}\}$ and $E_S := \{\{v, w\} \in E(G) \mid \{v, w\} \subseteq V_S\}$, that is, $S = G[V_S]$. Obviously, $|V_S| \leq k$. It remains to show $|E_S| \geq \binom{k}{2}$, which implies both that $|V_S| = k$ and that S is a clique. To this end, observe that the only vertices with incoming citations in D are the vertices $e_{\{v,w\}}$ for the edges $\{v, w\}$ of G . The only citations of a vertex $e_{\{v,w\}}$ are from the parts R_v and R_w in \mathcal{P} . That is, with respect to the partition \mathcal{P} , each vertex $e_{\{v,w\}}$ has two citations. Since the h-index h to reach is $\binom{k}{2}$, at least $\binom{k}{2}$ vertices $e_{\{v,w\}}$ have to receive $\binom{k}{2} \geq 4$ citations, which is only possible by atomizing both R_v and R_w . That is, for at least $\binom{k}{2}$ vertices $e_{\{v,w\}}$, we have $\{R_v, R_w\} \subseteq \mathcal{P} \setminus \mathcal{R}$ and, thus, $v, w \subseteq V_S$ and $\{v, w\} \in E_S$. It follows that $|E_S| \geq \binom{k}{2}$. \square

The reduction given above easily yields the same hardness result for most other problem variants: a vertex $e_{\{v,w\}}$ receives a sufficient number of citations only if R_v and R_w are atomized. Hence, even if we allow extractions or divisions on R_v , it helps only if we extract or split off all articles in R_v . The only difference is that the number of allowed operations is set to $k \cdot (\lceil \binom{k}{2} / 2 \rceil - 1)$ for these two problem variants. By the same argument, we obtain hardness for the conservative variants.

Corollary 1. For $\mu = \text{fusionCite}$, CONSERVATIVE EXTRACTING(μ), CAUTIOUS EXTRACTING(μ), CONSERVATIVE DIVIDING(μ), and CAUTIOUS DIVIDING(μ) are NP-hard and W[1]-hard when parameterized by $h + k + s$, where $s := \max_{P \in \mathcal{P}} |P|$, even if the citation graph is acyclic.

5 COMPUTATIONAL EXPERIMENTS

To assess how much the h-index of a researcher can be manipulated by splitting articles, we performed computational experiments with data extracted from Google Scholar.

Description of the Data. We use three data sets provided by van Bevern et al. [2]. One data set consists of 22 selected authors of IJCAI'13. The selection of these authors was biased to obtain profiles of authors in their early career. More precisely, the selected authors have a Google Scholar profile, an h-index between 8 and 20, between 100 and 1000 citations, and have been active between 5 and 10 years. The other two data sets contain Google Scholar data of 'AI's 10 to Watch', a list of young accomplished researchers in AI compiled by *IEEE Intelligent Systems*. One data set contains five profiles from the 2011 edition, the other eight profiles from the 2013 edition of the list.

Generation of Profiles with Merged Articles. In our setting, the input consists of a profile which already contains some merged articles.

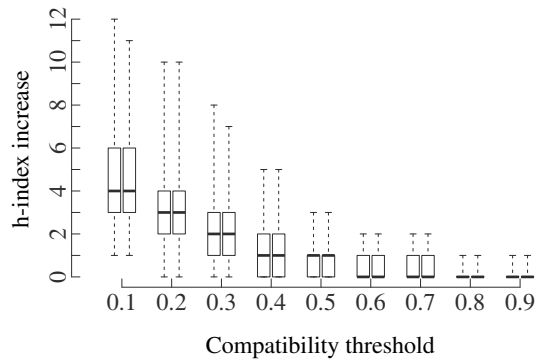


Figure 3: For each compatibility threshold, the left box shows maximum h-index increases for sumCite, the right box for unionCite.

To obtain such merged profiles, we used the compatibility graphs for each profile provided by van Bevern et al. [2], which they generated as follows. For each article u let $T(u)$ denote the set of words in its title. There is an edge between articles u and v if $|T(u) \cap T(v)| \geq t \cdot |T(u) \cup T(v)|$, where $t \in [0, 1]$ is the *compatibility threshold*. For $t = 0$, the compatibility graph is a clique; for $t = 1$ only articles with the same words in the title are adjacent. For $t \leq 0.3$, very dissimilar articles are still considered compatible [2]. Hence, we focus on $t \geq 0.4$ below.

For each profile and corresponding compatibility graph G , we obtained a profile with merged articles as follows. While the compatibility graph G contains an edge, compute a maximal clique C by a greedy algorithm, add C as a merged article to the profile, remove C from G , and continue. If the compatibility graph has no edge, then add all remaining articles as atomic articles of the profile.

Experimental Results. We implemented Algorithms 2, 4, and 5—the exact, linear-time algorithms from Section 3 for CONSERVATIVE ATOMIZING, CONSERVATIVE EXTRACTING, and CAUTIOUS EXTRACTING, respectively, each for both the sumCite and unionCite measures. Using them, we computed the maximum-possible h-index increases under the respective restrictions. The implementation is in Python 2.7.10 under Ubuntu Linux 15.10. Using an Intel Xeon E3-1231 CPU with 3.4 GHz and 32 GB RAM, the instances could be solved within three minutes altogether.

Figure 3 shows h-index increases for the IJCAI’13 authors for extracting articles: the lower edge of a box is the first quartile, the upper edge the third quartile, and the thick bar is the median; whiskers extend to the maximum and minimum values. Note that the h-index increase achievable by extracting articles is always at least as large as the one for atomizing articles. For the IJCAI’13 authors, atomizing articles yields essentially the same curve/same results as in Figure 3. Qualitatively, the results for AI’s 10 to Watch 2013 are the same, whereas AI’s 10 to Watch 2011 can achieve larger h-index increases for compatibility threshold 0.1. Hence, supposing that compatibility thresholds of at least 0.4 yield realistic profiles, we can conclude that 25 % of the authors could improve their h-index by unmerging articles by at least two and some outliers by five.

The results concerning the influence of restrictions on the number of operations and number of touched merged articles are as follows. For atomizing articles, most authors can increase their h-indices in increments of one for each atomizing operation up to their individual maximum. There is, however, one IJCAI’13 author who can achieve an increment of five with one atomizing operation ($t = 0.4$, sumCite). For extracting articles, clearly, each operation can increase the h-index by at most one. The results for CAUTIOUS EXTRACTING




Figure 4: For each upper bound k on the number of allowed article extractions, the left box shows the maximum h-index increases for sumCite, the right box for unionCite.

over the IJCAI’13 authors are shown in Figure 4 for compatibility threshold 0.4. Interestingly, in the experiments for CONSERVATIVE EXTRACTING and $t = 0.4$, all selected IJCAI’13 authors can achieve their maximum h-index increases by extracting articles out of at most two merged articles. In general, for threshold at least 0.4, they need to touch at most three merged articles to achieve the maximum h-index increase. This is also true for AI’s 10 to Watch 2013, whereas AI’s 10 to Watch 2011 can improve further by manipulating four merged articles (for $t = 0.5$).

Summarizing, our findings indicate that realistic profiles can be manipulated by splitting articles to yield h-index increases of at most two for the majority of authors. This can mean saving at least a year of work, since the average increase of the h-indices per year is 1.22 in the considered IJCAI data set. Furthermore, our findings indicate that the increase can be obtained by tampering with a small number of merged articles.

6 CONCLUSION

Regarding theory, we leave three main open questions concerning the computational complexity of EXTRACTING(fusionCite), the parameterized complexity of DIVIDING(fusionCite), as well as the parameterized complexity of CAUTIOUS DIVIDING(sumCite / unionCite) with respect to h (see Table 1), as the most immediate challenges for future work. Also, finding hardness reductions that produce more realistic instances would be desirable. From the experimental side, evaluating the potentially possible h-index increase by splitting on real merged profiles would be interesting as well as experiments using fusionCite as a measure. Moreover, it makes sense to consider the manipulation of the h-index also in context with the simultaneous manipulation of other indices (e.g., Google’s i10-index, see also Pavlou and Elkind [16]) and to look for Pareto-optimal solutions. We suspect that our algorithms easily adapt to other indices. In addition, it is natural to consider combining merging and splitting in manipulation of author profiles. Finally, from a practical point of view, our experimental results indicate that author profiles with surprisingly large h-index may be worth inspecting concerning potential manipulation.

ACKNOWLEDGEMENTS

René van Bevern is supported by the Russian Science Foundation, grant 16-11-10041, under support of which in particular the linear running times and Theorems 5 and 7 were proven. Hendrik Molter and Manuel Sorge are supported by the DFG, project DAPA (NI 369/12).

REFERENCES

- [1] Christoph Bartneck and Servaas Kokkelmans, ‘Detecting *h*-index manipulation through self-citation analysis’, *Scientometrics*, **87**(1), 85–98, (2011).
- [2] René van Bevern, Christian Komusiewicz, Rolf Niedermeier, Manuel Sorge, and Toby Walsh, ‘H-index manipulation by merging articles: Models, theory, and experiments’, *Artificial Intelligence*, (2016). Accepted for publication, available on arXiv:1412.5498.
- [3] Marek Cygan, Fedor V Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipeczuk, Michal Pilipeczuk, and Saket Saurabh, *Parameterized Algorithms*, Springer, 2015.
- [4] Emilio Delgado López-Cózar, Nicolás Robinson-García, and Daniel Torres-Salinas, ‘The Google Scholar experiment: How to index false papers and manipulate bibliometric indicators’, *Journal of the Association for Information Science and Technology*, **65**(3), 446–454, (2014).
- [5] Rod G. Downey and Michael R. Fellows, *Fundamentals of Parameterized Complexity*, Springer, 2013.
- [6] Leo Egghe, ‘Theory and practise of the *g*-index’, *Scientometrics*, **69**(1), 131–152, (2006).
- [7] Piotr Faliszewski, Edith Hemaspaandra, and Lane A. Hemaspaandra, ‘Using complexity to protect elections’, *Communications of the ACM*, **53**(11), 74–82, (2010).
- [8] Piotr Faliszewski and Ariel D. Procaccia, ‘AI’s war on manipulation: Are we winning?’, *AI Magazine*, **31**(4), 53–64, (2010).
- [9] Jörg Flum and Martin Grohe, *Parameterized Complexity Theory*, Springer, 2006.
- [10] J. E. Hirsch, ‘An index to quantify an individual’s scientific research output’, *Proceedings of the National Academy of Sciences of the United States of America*, **102**(46), 16569–16572, (2005).
- [11] Klaus Jansen, Stefan Kratsch, Dániel Marx, and Ildikó Schlotter, ‘Bin packing with fixed number of bins revisited’, *Journal of Computer and System Sciences*, **79**(1), 39–49, (2013).
- [12] Stasys Jukna, *Extremal Combinatorics - with Applications in Computer Science*, Texts in theoretical computer science, Springer, 2001.
- [13] Bart de Keijzer and Krzysztof R. Apt, ‘The H-index can be easily manipulated’, *Bulletin of the EATCS*, **110**, 79–85, (2013).
- [14] Michael Lesk, ‘How many scientific papers are not original?’, *Proceedings of the National Academy of Sciences of the United States of America*, **112**(1), 6–7, (2015).
- [15] Rolf Niedermeier, *Invitation to Fixed-Parameter Algorithms*, number 31 in Oxford Lecture Series in Mathematics and Its Applications, Oxford University Press, 2006.
- [16] Chrystalla Pavlou and Edith Elkind, ‘Manipulating citation indices in a social context’, in *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS ’16)*, pp. 32–40. IFAAMAS, (2016).
- [17] Péter Vinkler, ‘Would it be possible to increase the Hirsch-index, π -index or CDS-index by increasing the number of publications or citations only by unity?’, *Journal of Informetrics*, **7**(1), 72–83, (2013).
- [18] Gerhard J. Woeginger, ‘An axiomatic analysis of Egghe’s *g*-index’, *Journal of Informetrics*, **2**(4), 364–368, (2008).