# Optimal Uncoordinated Unique IDs 

Peter C. Dillinger* Martín Farach-Colton ${ }^{\dagger}$ Guido Tagliavini ${ }^{\ddagger}$ Stefan Walzer ${ }^{\S}$


#### Abstract

In the Uncoordinated Unique Identifiers Problem (UUIDP) there are $n$ independent instances of an algorithm $\mathcal{A}$ that generates IDs from a universe $\{1, \ldots, m\}$, and there is an adversary that requests IDs from these instances. The goal is to design $\mathcal{A}$ such that it minimizes the probability that the same ID is ever generated twice across all instances, that is, minimizes the collision probability. Crucially, no communication between the instances of $\mathcal{A}$ is possible. Solutions to the UUIDP are often used as mechanisms for surrogate key generation in distributed databases and key-value stores. In spite of its practical relevance, we know of no prior theoretical work on the UUIDP.

In this paper we initiate the systematic study of the UUIDP. We analyze both existing and novel algorithms for this problem, and evaluate their collision probability using worst-case analysis and competitive analysis, against oblivious and adaptive adversaries. In particular, we present an algorithm that is optimal in the worst case against oblivious adversaries, an algorithm that is at most a logarithmic factor away from optimal in the worst case against adaptive adversaries, and an algorithm that is optimal in the competitive sense against both oblivious and adaptive adversaries.


## 1 Introduction

Unique identifiers (IDs) for data are essential to efficient data processing, storage, and retrieval. When natural keys for data do not exist or are impractical, surrogate keys are assigned and used as primary keys. In a distributed context, database systems including Cassandra, Microsoft's Transact-SQL, MongoDB, MySQL, Postgres and RocksDB [10, 7, 14, 13, 8, 25, 22] generate unique surrogate keys without coordination between nodes or use of a central authority.

In the case of RocksDB, IDs are assigned to data files (also known as SSTs) and data blocks within each file, for caching purposes. Although RocksDB is a single-instance key-value store, its users (e.g., Microsoft Bing's web data platform [28], MyRocks [31] and ZippyDB [30]) typically run multiple instances of it and distribute these instances across multiple nodes. Because data are moved between instances (e.g., to balance the load across the nodes), RocksDB's IDs have to be collision-free over all instances, even though these instances are unaware of one another.

There are several benefits to generating IDs without coordination. First, coordination mechanisms are complex, and thus have significant engineering and operational costs. Second, implementations of coordination are often based on features such as MAC addresses (which are subject to spoofing) and clocks (which can become skewed), rendering them notoriously brittle. Finally, coordination in RocksDB is undesirable from a software design point of view, as it would force RocksDB instances to take responsibility for network awareness and communication.

[^0]Generating unique IDs without coordination may seem like fighting a losing battle, since communication is necessary to guarantee ID uniqueness. What we can aim for, however, is to generate IDs randomly, such that the probability that the same ID is ever generated by different instances is near zero.

Formally, in the Uncoordinated Unique Identifiers Problem (UUIDP) an ID-generation algorithm $\mathcal{A}$ plays a game against an adversary. The game proceeds in steps, and at each step the adversary requests an ID from one of $n$ independent instances of $\mathcal{A}$. When an instance of $\mathcal{A}$ receives a request, it must output a value from the space $[m]:=\{1, \ldots, m\}$ of IDs. Importantly, the instances cannot communicate with each other, and they do not know $n$ or the number of IDs that have been requested from other instances. The algorithm loses the game if the adversary causes a collision, that is, if some ID is generated twice. The algorithm wins if by the time the game ends (for instance, after a fixed number of steps has been reached) a collision has not occurred. The goal is to design $\mathcal{A}$ such that the probability that $\mathcal{A}$ wins, called the collision probability, is as small as possible.

Beyond database systems, uncoordinated generation of unique IDs is used in settings such as identification of network connections [9], identification of disk partitions [1], and object identification in video games [12]. Such is the ubiquity of this problem that there exists a standardized solution to it [27, 15], known as globally unique identifiers (GUIDs), and most programming languages have library implementations of GUIDs [4, 2, 6, 11, 5, 3].

GUIDs are generated using some combination of a randomly generated integer and identifying metadata (e.g., creation timestamp or MAC address), with the particular combination depending on the GUID variant. This, however, presumes reliable metadata, which is often times impossible to ensure in practice - an adversary may tamper with the metadata to cause a collision. For this reason, we do not model metadata in the UUIDP. The the random part of GUIDs is modeled as the following simple algorithm for the UUIDP, that we call Random: every time an ID is requested, sample an integer from $[m]$ without replacement.

When a total of $d$ IDs are requested, the collision probability of Random is $O\left(d^{2} / m\right)$. In practice, this means that Random should only be used when far less than $\sqrt{m}$ IDs are needed. Concretely, GUIDs are 128 -bit IDs, so they cannot safely handle more than $\sqrt{m}=2^{64}$ IDs. Unfortunately, with consolidation of cloud computing services and some companies already operating at exabyte scales (i.e., more than $2^{60}$ bytes), we are not so far from a world where we will have enough data objects in a single pool to observe collisions in Random with 128 -bit IDs. Thus, Random is quickly becoming inadequate for large-scale deployments, unless longer IDs are used.

To guarantee collision-freedom of uncoordinated 128-bit IDs in the long-term future, RocksDB recently started using a different algorithm [23, 24], that we call Cluster ${ }^{1}$ This algorithm picks a random integer $x \in[m]$, and generates IDs sequentially, starting at $x$. Although, experimental data showed that the collision probability of Cluster is significantly lower than that of Random [24], it is not known exactly how much better Cluster is, or whether there exists an algorithm that is even better than Cluster.

To the best of our knowledge, no other algorithms for the UUIDP are known. In spite of the widespread use of uncoordinated ID-generation algorithms, there has been no theoretical work on this problem. The overly simple setup of the UUIDP may give the false impression that not much can be said about it. In this paper we show that there is more than meets the eye.

[^1]
### 1.1 Our results

We study the UUIDP in four different settings: using worst-case analysis or competitive analysis; and against an adversary that is oblivious or adaptive.

In the worst-case analysis of an algorithm $\mathcal{A}$, we are interested in minimizing the maximum collision probability of $\mathcal{A}$, over all adversaries that request at most $d$ IDs in total. (Notice that we must bound the number of requests to make the worst case analysis interesting; otherwise the adversary can drive the collision probability up to 1.) In the competitive analysis of $\mathcal{A}$, we are interested, roughly speaking, in minimizing the ratio between the collision probability of $\mathcal{A}$ and the best possible collision probability. We define these notions more formally in Section 2 .

An oblivious adversary chooses how many IDs it will request to each one of the $n$ instances before the game begins. The more powerful adaptive adversary decides from which instance to request an ID on a step-by-step basis, based on the IDs that were produced previously.

We prove the following results, which we will formally state in Section 3;

- There is an algorithm, namely Cluster, that has worst-case collision probability $O(n d / m)$ against oblivious adversaries that request at most $d$ IDs.
- The worst-case collision probability of every algorithm against oblivious adversaries is $\Omega(n d / m)$.
- There is an algorithm that has worst-case collision probability $O((n d / m) \log (1+d / n))$ against adaptive adversaries.
- There is an algorithm that has competitive ratio $O(\log m)$ against adaptive adversaries.
- The competitive ratio of every algorithm against oblivious adversaries is $\Omega(\log m)$.

Naturally, lower bounds against oblivious adversaries also hold against adaptive adversaries, and upper bounds against adaptive adversaries also hold against oblivious adversaries.

Our theorems imply that Cluster is an effective uncoordinated ID-generation mechanism at scales that are orders of magnitude beyond Random's capacity. The $O(n d / m)$ collision probability of Cluster allows it to handle, for example, collections of $d \gg 2^{64}$ objects, using IDs of 128 bits or even less. The matching collision probability lower bound establishes that no other algorithm can do significantly better.

In the adaptive setting, Cluster is far from optimal (and so is Random). Our new algorithms for the adaptive case are either optimal or almost optimal, and may be of interest to database researchers and practitioners working on systems with stringent security requirements.

### 1.2 Related work

A well-studied problem that deals with ID assignments is the renaming problem [18, 17]. In the renaming problem, there is a distributed system with $n$ processors that communicate asynchronously, and can fail. Each processor starts out with a name (i.e., an ID) from an unbounded space. The goal is to design a distributed algorithm that maps the names into a namespace that is as small as possible, such that no processor is assigned the same name, and using as little communication as possible.

The renaming problem departs from the UUIDP in two crucial ways: communication is allowed, and machines only need to produce a single unique name. In the UUIDP, the case where exactly 1 ID is requested to each instance is equivalent to a birthday problem [34, 32].

The existing work on generating unique identifiers in an uncoordinated setting focuses on practical aspects of this problem, like compressing GUIDs [29], designing identifiers with security guarantees [33], and designing identifiers for IoT platforms [16] and mobile environments 26]. To the best of our knowledge, there is no theoretical work on a problem with similar characteristics to the UUIDP.

### 1.3 Organization

In Section 2 we present the mathematical model of the UUIDP, as well as some preliminary definitions that we use in the rest of the paper. In Section 3 we give a full overview of our results and their motivation, including all the algorithms and main theorems. That section presents a complete picture of our work, only leaving out the proofs. The remainder of the paper contains the proofs.

## 2 Model and Definitions

We fix the size $m \in \mathbb{N}$ of the universe $[m]:=\{1, \ldots, m\}$ of IDs. An ID-generation algorithm $\mathcal{A}$ generates the IDs from $[m]$ in some randomized order. Equivalently, we can view $\mathcal{A}$ as a distribution on the set of permutations of $[m]$.

Consider a setting with $n$ instances of $\mathcal{A}$ that use independent randomness and that do not communicate. A vector $D=\left(d_{1}, \ldots, d_{n}\right) \in[m]^{n}$, called the demand profile, indicates that $d_{i}$ IDs are requested from the $i$ th instance. These requests are made one by one, such that the instance does not know in advance how many IDs it will have to produce. We say a collision occurs when the sets of IDs produced by the $n$ instances are not pairwise disjoint. The probability that a collision occurs is denoted by $p_{\mathcal{A}}(D)$.

When $n \geq 2$, we have $p_{\mathcal{A}}(D)>0$ for any $\mathcal{A}$ and $D$. (This follows from the fact that already for $D=(1,1)$ we have $p_{\mathcal{A}}(D)>0$; see Corollary 17 .) In other words, collisions cannot be avoided with certainty for $n \geq 2$. Of course, if $n=1$ we have $p_{\mathcal{A}}(D)=0$; we call these trivial demand profiles.

In our analysis we restrict our attention to natural sets of demand profiles where the number $n$ of instances, the total number $d:=\|D\|_{1}=\sum_{i=1}^{n} d_{i}$ of requests or the maximum number $h$ of requests per instance are restricted, such as

- $\mathcal{D}_{1}(n, d):=\left\{D \in[m]^{n} \mid\|D\|_{1}=d\right\}$,
- $\mathcal{D}_{1}(d):=\bigcup_{i \geq 2} \mathcal{D}(i, d)$,
- $\mathcal{D}_{\infty}^{\leq}(n, h):=\bigcup_{2 \leq i \leq n}[h]^{i}$, and
- $\mathcal{D}_{\infty}^{\leq}(h):=\bigcup_{i \geq 2}[h]^{i}$.
$\mathcal{D}_{1}(n, d)$ and $\mathcal{D}_{1}(d)$ contain demand profiles with $L^{1}$-norm equal to $d$, whereas $\mathcal{D} \leq(n, h)$ and $\mathcal{D}_{\infty}^{\leq}(h)$ contain demand profiles with $L^{\infty}$-norm at most $h$. We will also use $\mathcal{D}_{\infty}^{<}(h)$ when the $L^{\infty}$-norm is strictly less than $h$.

Worst-case vs. competitive analysis. How should we evaluate the performance of an ID generation algorithm $\mathcal{A}$ on a set $\mathcal{D}$ of demand profiles? The simplest option is to consider the worst-case collision probability, namely

$$
\max _{D \in \mathcal{D}} p_{\mathcal{A}}(D)
$$

A natural alternative compares for each $D \in \mathcal{D}$ the actual collision probability $p_{\mathcal{A}}(D)$ to the best possible collision probability for $D$, namely to

$$
p_{*}(D):=\min _{\mathcal{A}^{\prime}} p_{\mathcal{A}^{\prime}}(D),^{23}
$$

When $D$ is a non-trivial demand profile, we have $p_{*}(D)>0$. The competitive ratio of $\mathcal{A}$ for any non-trivial $D$ is defined as

$$
\frac{p_{\mathcal{A}}(D)}{p_{*}(D)} .
$$

The competitive ratio of $\mathcal{A}$ for a set $\mathcal{D}$ of non-trivial demand profiles is then defined as the maximum competitive ratio taken over all $D \in \mathcal{D}$.

Oblivious vs. adaptive adversaries. When $D$ is fixed ahead of time, we say that it is produced by an oblivious adversary. The collision probability in this case is $p_{\mathcal{A}}(D)$.

We also consider the case where the demand profile $D$ is produced by an adaptive adversary $Z$ on the fly. In this case, the total number $n$ of instances of $\mathcal{A}$ that are probed is not known in advance. Initially, $D=()$ is empty. When $D=\left(d_{1}, \ldots, d_{i}\right)$, the adversary has three options:

- set $D \leftarrow\left(d_{1}, \ldots, d_{i}, 1\right)$, effectively activating a previously dormant instance of $\mathcal{A}$ from which an ID was never requested;
- set $D \leftarrow\left(d_{1}, \ldots, d_{j}+1, \ldots, d_{i}\right)$ for some $j \in[i]$, effectively requesting another ID from the $j$ th instance of $\mathcal{A}$;
- stop the game, making $D$ the final demand profile.

Crucially, at each step the adaptive adversary learns the ID that is produced and his future decisions may depend on this knowledge. The collision probability for this game is denoted by $p_{\mathcal{A}}(Z)$. Note that the oblivious setting can be thought of as a special case where $Z$ ignores the IDs produced by the instances of $\mathcal{A}$.

Importantly, the adversary knows the algorithm $\mathcal{A}$ it will be playing against. Thus, once $\mathcal{A}$ is fixed, the adversary can tailor its strategy to defeat $\mathcal{A}$.

When evaluating the performance of $\mathcal{A}$ on a set of demand profiles $\mathcal{D}$ in the adaptive setting, we consider the set $\operatorname{Adv}(\mathcal{D})$ of adversaries that always stop the game eventually, and do so with a final demand profile contained in $\mathcal{D}$. We can again choose between a worst-case analysis that ponders

$$
\max _{Z \in \operatorname{Adv}(\mathcal{D})} p_{\mathcal{A}}(Z)
$$

and a competitive analysis that considers the competitive ratio of $\mathcal{A}$ for any $Z$ that produces non-trivial demand profiles, defined as

$$
\frac{p_{\mathcal{A}}(Z)}{\underset{D \sim Z}{\mathbb{E}}\left[p_{*}(D)\right]},
$$

and evaluates the maximum competitive ratio over $Z \in \operatorname{Adv}(\mathcal{D})$. Here we write, in a slight abuse of notation, $D \sim Z$ for the final demand profile $D$ arising randomly from the interaction of $Z$ and $\mathcal{A}$.

In Table 1 we summarize the four settings $\{$ worst-case, competitive $\} \times\{$ oblivious, adaptive $\}$ considered in this paper.

[^2]|  | worst-case | competitive |
| :--- | :---: | :---: |
| oblivious | $\max _{D \in \mathcal{D}} p_{\mathcal{A}}(D)$ | $\max _{D \in \mathcal{D}} \frac{p_{\mathcal{A}}(D)}{p_{*}(D)}$ |
| adaptive | $\max _{Z \in \operatorname{Adv}(\mathcal{D})} p_{\mathcal{A}}(Z)$ | $\max _{Z \in \operatorname{Adv}(\mathcal{D})} \frac{p_{\mathcal{A}}(Z)}{\overline{\mathbb{E}_{D \sim Z}\left[p_{*}(D)\right]}}$ |

Table 1: Four settings for evaluating the performance of an ID generation algorithm $\mathcal{A}$ on a set $\mathcal{D}$ of demand profiles.
$O$-notation In this paper, we use a non-asymptotic version of $O-, \Omega$ - and $\Theta$-notation. Specifically, we say that $f\left(x_{1}, x_{2}, \ldots\right)=O\left(g\left(x_{1}, x_{2}, \ldots\right)\right)$ if there exists a constant $c>0$ such that $f\left(x_{1}, x_{2}, \ldots\right)=c \cdot g\left(x_{1}, x_{2}, \ldots\right)$ for all $x_{1}, x_{2}, \ldots$ In other words, $O$ simply hides constant factors. $\Omega$ and $\Theta$ are defined similarly.

## 3 Overview

In this section we present the roadmap of the rest of the paper, including all the algorithms we analyze, as well as the main theorems.

### 3.1 Basic Results on Random and Cluster

The starting point of this work is an analysis and comparison of the collision probability of Random and Cluster for arbitrary demand profiles. These two algorithms are defined below. In each case we illustrate how the algorithm might behave. A sequence of $m$ squares represents the ID space $[m]$ and a number $i$ in a square indicates that the corresponding ID is the $i$ th ID that is returned.

## Algorithm Random

Return the IDs from $[m]$ in a uniformly random order.

```
Example (m=20,8 requests)
```


## Algorithm Cluster

Pick $x \in[m]$ uniformly at random and return IDs in the order $x, x+1, x+2, \ldots$, all modulo $m$.

```
Example (m=20, 8 requests)
```

                    \(112|3| 4|5| 6|7| 8\)
    [^3]The following theorem gives a tight asymptotic estimate of the collision probability of Cluster.
Theorem 1. Suppose $n \geq 2$. Let $D \in[m]^{n}$. Then,

$$
p_{\mathrm{ClUSTER}}(D)=\Theta\left(\min \left(1, \frac{n\|D\|_{1}}{m}\right)\right) .
$$

The $\min (1, \cdot)$ safeguard in the probability is a recurring pattern in our bounds. If it kicks in we have $p_{\mathcal{A}}(D)=\Theta(1)$, which amounts to total failure of $\mathcal{A}$ on $D$. When we know that we are dealing with small collision probabilities, then the $\min (1, \cdot)$ can be stripped away.

Instead of analyzing Random directly, we study a more general algorithm that we call $\operatorname{Bins}(k)$, which includes Random as a special case for $k=1$, and which also plays a central role in several lower and upper bounds down the road.

## Algorithm Bins( $k$ )

Partition $[m]$ into $\lfloor m / k\rfloor$ bins of $k$ IDs and $m \bmod k$ leftover IDs. Pick a random permutation of the bins. Iterate over the shuffled bins, returning all IDs of a bin in increasing order before moving on to the next bin. Finally, return the leftover IDs in increasing order.

Example ( $m=20, k=3,8$ requests)

## $1456 \mid$ | 78 |123|

Theorem 2. Suppose $n \geq 2$. Let $D \in[m]^{n}$ and $k \in[m]$. Then,

$$
p_{\operatorname{BINS}(k)}(D)=\Theta\left(\min \left(1, \frac{\|D\|_{1}^{2}-\|D\|_{2}^{2}}{k m}+\frac{n\|D\|_{1}}{m}+\frac{n^{2} k}{m}\right)\right) .
$$

A calculation shows that, when $k=1$, the first term in the sum dominates $p_{\operatorname{BiNS}(k)}(D)$. This yields the collision probability of Random.

Corollary 3. Suppose $n \geq 2$. Let $D \in[m]^{n}$. Then,

$$
p_{\text {RANDOM }}(D)=\Theta\left(\min \left(1, \frac{\|D\|_{1}^{2}-\|D\|_{2}^{2}}{m}\right)\right) .
$$

Theorem 2 also implies that $p_{\operatorname{Bins}(k)}(D)=\Omega\left(\min \left(1, \frac{n\|D\|_{1}}{m}\right)\right)$. Combining this with Theorem 1, we conclude that Cluster is asymptotically no worse than $\operatorname{Bins}(k)$ and, consequently, Random.

Corollary 4. Suppose $n \geq 2$. Let $D \in[m]^{n}$ and $k \in[m]$. Then,

$$
\begin{aligned}
& p_{\text {CLuster }}(D)=O\left(p_{\operatorname{Bins}(k)}(D)\right) \text { and, in particular, } \\
& p_{\text {CLuster }}(D)=O\left(p_{\text {Random }}(D)\right) .
\end{aligned}
$$

### 3.2 Worst-Case Analysis Against Oblivious Adversaries

To do a worst-case analysis, we need to restrict the total number $d$ of requests. This is because any adversary that requests $m+1$ or more IDs forces collision probability 1 . We therefore only consider demand profiles from $\mathcal{D}_{1}(n, d)$. Theorem 1 and Theorem 2 imply the following worst-case collision probabilities.

Corollary 5. Suppose $n \geq 2$. Let $d \geq n$. Then,

$$
\begin{aligned}
& \max _{D \in \mathcal{D}_{1}(n, d)} p_{\text {Cluster }}(D)=\Theta(\min (1, n d / m)) \text { and } \\
& \max _{D \in \mathcal{D}_{1}(n, d)} p_{\text {RANDOM }}(D)=\Theta\left(\min \left(1, d^{2} / m\right)\right) .
\end{aligned}
$$

Does Cluster have optimal worst-case collision probability? We answer this question positively, proving a matching lower bound. Further, we show that all but a vanishing fraction of demand profiles from $\mathcal{D}_{1}(n, d)$ force this worst-case collision probability, implying that Cluster is optimal on almost all demand profiles.

Theorem 6. For all but an $\exp (-\Theta(n))$-fraction of $D \in \mathcal{D}_{1}(n, d)$

$$
p_{*}(D)=\Omega\left(\min \left(1, \frac{n d}{m}\right)\right) .
$$

### 3.3 Worst-Case Analysis Against Adaptive Adversaries

Interestingly, when facing adaptive adversaries, the collision probability of Cluster worsens by at least a factor of $n$.

Lemma 7. Suppose $n \geq 2$. Let $d \geq 2 n$. There exists an adaptive adversary $Z \in \operatorname{Adv}\left(\mathcal{D}_{1}(n, d)\right)$ such that

$$
p_{\text {CLUSTER }}(Z)=\Omega\left(\min \left(1, \frac{n^{2} d}{m}\right)\right)
$$

Roughly speaking, this adversary exploits the fact that once the initial ID from each instance is known, the two instances with the two closest IDs can be forced to collide if they are at distance less than $d$.

Notice that Random seems to be robust to adaptivity: because every ID produced is a fresh random number, an adaptive adversary can extract little information about future IDs from past IDs. Of course, the collision probability of Random in the adaptive case is as bad as in the oblivious case.

Is there an algorithm for the adaptive setting that matches the collision probability lower bound $\Omega(\min (1, n d / m))$ from Theorem 6? We give an essentially positive answer, by showing that there exists an algorithm, that we call Cluster*, that is only a small logarithmic factor away from that lower bound.

## Algorithm Cluster*

Let run $(x, r)$ be the sequence $(x, x+1, \ldots, x+(r-1))$ modulo $m$. Repeat the following for $r=1,2,4,8, \ldots$ : Draw $x \in[m]$ uniformly at random, such that run $(x, r)$ does not collide with previously chosen runs. For the next $r$ requests return the IDs from run $(x, r)$.


Intuitively, Cluster* is resilient against adaptive adversaries because an adversary cannot predict too many future IDs from an instance. More specifically, the adversary can only predict a long run of IDs from an instance if it has already requested roughly the same number of IDs from it. Still, Cluster* behaves similarly to Cluster because the exponential growth of the runs implies that the number of runs per instance is small, and that most IDs produced by an instance belong to its largest run.

Theorem 8. Suppose $n \geq 2$. Let $d \geq n$ and $\mathcal{D}:=\mathcal{D}_{1}(d) \cap \mathcal{D}_{\propto}^{\leq}(n, m /(2 \log m))$. Then,

$$
\max _{Z \in \operatorname{Adv}(\mathcal{D})} p_{\text {CLuster }^{*}}(Z)=O\left(\min \left(1, \frac{n d}{m} \log \left(1+\frac{d}{n}\right)\right)\right) .
$$

Due to fragmentation of the ID space, Cluster* may not be able to generate all $m$ IDs according to its rules. This is why we restrict our analysis to demand profiles with at most $m /(2 \log m)$ requests per instance. An instance then allocates at most $\log m$ runs of sizes at most $m /(2 \log m)$, which always fit even for worst-case fragmentation.

### 3.4 Competitive Analysis Against Oblivious Adversaries

The shortcoming of Cluster is that when the requests are skewed towards a small group of instances, the collision probability is far from optimal. For instance, let $D=(d-1,1) \in \mathcal{D}_{1}(2, d)$ be a maximally skewed demand profile with $n=2$. Consider the following algorithm: the first ID is randomly sampled from $[m-(d-2)]$, and all other IDs are deterministically taken from $\{m-d+3, \ldots, m\}$. Then, the collision probability on $D$ is $1 /(m-(d-2))$, which is the probability that first IDs of the two instances collide. This is up to a factor $\Theta(d)$ smaller than $p_{\text {Cluster }}(D)=$ $d / m$.

Therefore, although ClUSTER is optimal in the worst case, there are algorithms that perform far better on some special cases, i.e., Cluster's competitive ratio is far from optimal. To overcome this limitation, we design yet another algorithm, which we call Bins*.

## Algorithm Bins*

Partition the ID space into $O(\log m)$ chunks and partition the $i$ th chunk into bins of $2^{i-1}$ IDs each. Pick a uniformly random bin of size 1 , then of size 2 , then of size 4 , and so on, always returning all IDs of a bin in increasing order before moving on to a bin of twice the size.


Bins* combines the exponential allocations from Cluster* with the spatial partitioning of $\operatorname{Bins}(k)$ for $k \in\{1,2,4, \ldots\}$. Intuitively, the goal of Bins* is that instances with similar loads allocate most of their IDs from the same region of the universe $[m$ ]. In particular, low-demand instances can only collide with few IDs of high-demand instances.

Theorem 9. The competitive ratio of Bins* for $\mathcal{D}_{\infty}^{<}(m / \log m)$ is $O(\log m)$.
Furthermore, we establish optimality by proving a matching lower bound that even applies to $\mathcal{D}=[\sqrt{m}]^{2}$ (and hence to any $\mathcal{D}^{\prime}$ with $\mathcal{D} \subseteq \mathcal{D}^{\prime}$ ).

Theorem 10. Every algorithm has competitive ratio $\Omega(\log m)$ for $\mathcal{D}=[\sqrt{m}]^{2}$.

### 3.5 Competitive Analysis Against Adaptive Adversaries

Finally, we consider the competitive ratio against adaptive adversaries. We prove the following general reduction from oblivious to adaptive adversaries.

Theorem 11. Let $\mathcal{A}$ be either $\operatorname{Bins}^{*}$ or $\operatorname{Bins}(k)$ for some $k \in[m]$, and let $\mathcal{D}$ be some set of demand profiles. If $\mathcal{A}$ has competitive ratio $c$ for $\mathcal{D}$, then $\mathcal{A}$ has competitive ratio at most $4 c$ for $\operatorname{Adv}(\mathcal{D})$.

This implies that the $O(\log m)$ upper bound on the competitive ratio of Bins* from Theorem 9 also holds in the adaptive case.

Corollary 12. Bins* is has competitive ratio $O(\log m)$ for $\operatorname{Adv}\left(\mathcal{D}_{\infty}^{<}(m / \log m)\right)$.

## 4 Collision Probability On Fixed Demand Profiles

In this section we study the collision probability of $\operatorname{Bins}(k)$ and Cluster, on arbitrary demand profiles. We begin with a technical lemma.

Lemma 13. Let $\mathcal{E}_{1}, \ldots, \mathcal{E}_{\ell}$ be pairwise independent events. Then,

$$
\operatorname{Pr}\left[\bigcup_{i=1}^{\ell} \mathcal{E}_{i}\right]=\Theta\left(\min \left(1, \sum_{i=1}^{\ell} \operatorname{Pr}\left[\mathcal{E}_{i}\right]\right)\right)
$$

Proof. The upper bound follows from the union bound. We now prove the lower bound. We consider three cases:

- Case 1: $\sum_{i=1}^{\ell} \operatorname{Pr}\left[\mathcal{E}_{i}\right] \leq 2 / 3$. By a Bonferroni inequality,

$$
\begin{aligned}
\operatorname{Pr}\left[\bigcup_{i} \mathcal{E}_{i}\right] & \geq \sum_{i} \operatorname{Pr}\left[\mathcal{E}_{i}\right]-\sum_{i} \sum_{j<i} \operatorname{Pr}\left[\mathcal{E}_{i} \cap \mathcal{E}_{j}\right] \\
& =\sum_{i} \operatorname{Pr}\left[\mathcal{E}_{i}\right]-\sum_{i} \sum_{j<i} \operatorname{Pr}\left[\mathcal{E}_{i}\right] \operatorname{Pr}\left[\mathcal{E}_{j}\right] \quad \text { (as the events are pairwise independent) } \\
& =\sum_{i}\left(\operatorname{Pr}\left[\mathcal{E}_{i}\right]\left(1-\sum_{j<i} \operatorname{Pr}\left[\mathcal{E}_{j}\right]\right)\right) \\
& \geq\left(1-\sum_{j} \operatorname{Pr}\left[\mathcal{E}_{j}\right]\right) \sum_{i} \operatorname{Pr}\left[\mathcal{E}_{i}\right] \\
& \geq \frac{1}{3} \sum_{i} \operatorname{Pr}\left[\mathcal{E}_{i}\right]=\Omega\left(\min \left(1, \sum_{i} \operatorname{Pr}\left[\mathcal{E}_{i}\right]\right)\right) .
\end{aligned}
$$

- Case 2: $\operatorname{Pr}\left[\mathcal{E}_{j}\right] \geq 1 / 3$ for some $j$. Then, $\operatorname{Pr}\left[\bigcup_{i} \mathcal{E}_{i}\right] \geq \operatorname{Pr}\left[\mathcal{E}_{j}\right]=\Omega(1)=\Omega\left(\min \left(1, \sum_{i} \operatorname{Pr}\left[\mathcal{E}_{i}\right]\right)\right)$.
- Case 3: otherwise. Then, there exists a $j$ such that $\sum_{i=1}^{j} \operatorname{Pr}\left[\mathcal{E}_{i}\right] \leq 2 / 3$ and $\sum_{i=1}^{j+1} \operatorname{Pr}\left[\mathcal{E}_{i}\right]>2 / 3$. Because $\operatorname{Pr}\left[\mathcal{E}_{j+1}\right]<1 / 3$, we have

$$
1 / 3<\sum_{i=1}^{j} \operatorname{Pr}\left[\mathcal{E}_{i}\right] \leq 2 / 3
$$

Thus, applying the first case to the first $j$ events we get

$$
\operatorname{Pr}\left[\bigcup_{i=1}^{\ell} \mathcal{E}_{i}\right] \geq \operatorname{Pr}\left[\bigcup_{i=1}^{j} \mathcal{E}_{i}\right]=\Omega(1)=\Omega\left(\min \left(1, \sum_{i} \operatorname{Pr}\left[\mathcal{E}_{i}\right]\right)\right) .
$$

Roughly speaking, Lemma 13 says that the union bound almost gives the right answer when the events involved are pairwise independent. The only exception is when the union bound in meaningless, i.e., the sum of the probabilities is greater than 1 , then the probability of the union is close to 1 .

We will extensively use the following well-known asymptotic approximations of the exponential function,

$$
\begin{equation*}
1-\varepsilon \leq e^{-\varepsilon} \leq 1-\varepsilon / 2 \text { for } \varepsilon \in[0,1] \tag{1}
\end{equation*}
$$

We start by establishing the collision probability of Cluster.
Theorem 1. Suppose $n \geq 2$. Let $D \in[m]^{n}$. Then,

$$
p_{\text {Cluster }}(D)=\Theta\left(\min \left(1, \frac{n\|D\|_{1}}{m}\right)\right) .
$$

Originally stated
on page 7

Proof. Let $D=\left(d_{1}, \ldots, d_{n}\right)$. For $i, j \in[n]$ with $i \neq j$, let $\mathcal{C}_{i j}$ be the event that there is a collision between the IDs produced by the instances $i$ and $j$. For every choice of the starting point of instance $j$, there are exactly $d_{i}+d_{j}-1$ choices for the starting point of instance $i$ that cause a collision (this uses $\left.d_{i}, d_{j}>0\right)$. Thus, $\operatorname{Pr}\left[\mathcal{C}_{i j}\right]=\left(d_{i}+d_{j}-1\right) / m$.

The sum of the collision probabilities is

$$
\sum_{i<j} \operatorname{Pr}\left[\mathcal{C}_{i j}\right]=\sum_{i<j} \frac{d_{i}+d_{j}-1}{m}=\frac{1}{m}\left((n-1)\|D\|_{1}-\binom{n}{2}\right)=\Theta\left(\frac{n\|D\|_{1}}{m}\right) .
$$

We claim that the collision events $\mathcal{C}_{i j}$ are pairwise independent. Let $i<j$ and $p<q$ be any two pairs of instances, such that $i \neq p$ or $j \neq q$ (otherwise, they are the same pair). If $|\{i, j, p, q\}|=4$ then $\mathcal{C}_{i j}$ and $\mathcal{C}_{p q}$ are obviously independent. Suppose $|\{i, j, p, q\}|=3$; without loss of generality, assume $i=p$ and $j \neq q$. Then, conditioned on $\mathcal{C}_{i j}$, the probability of $\mathcal{C}_{i q}$ remains unchanged, because conditioning on $\mathcal{C}_{i j}$ neither affects the distribution of the IDs of instance $i$, nor the independence of instances $i$ and $q$.

This means that the hypotheses of Lemma 13 are met for the event $\bigcup_{i<j} \mathcal{C}_{i j}$ that a collision occurs during the game. Hence,

$$
p_{\text {CLUSTER }}(D)=\Theta\left(\min \left(1, \sum_{i<j} \operatorname{Pr}\left[\mathcal{C}_{i j}\right]\right)\right)=\Theta\left(\min \left(1, \frac{n\|D\|_{1}}{m}\right)\right) .
$$

For the analysis of $\operatorname{Bins}(k)$ we need the following auxiliary lemma.
Lemma 14. Let $\mathcal{U}$ be a universe of $u$ elements. Let $S_{1}$ be a set of $s_{1}$ samples from $\mathcal{U}$, drawn without replacement. At every step of the sampling process, each of the remaining elements has equal probability of being sampled. Analogously, let $S_{2}$ be a set of $s_{2}$ samples from $\mathcal{U}$, drawn without replacement, independent of $S_{1}$. Then,

$$
\operatorname{Pr}\left[S_{1} \cap S_{2} \neq \emptyset\right]=\Theta\left(\min \left(1, \frac{s_{1} s_{2}}{u}\right)\right)
$$

Proof. Assume without loss of generality that $1 \leq s_{1} \leq s_{2}$ and that $S_{2}$ is sampled first, after which the set $S_{1}=\left\{x_{1}, \ldots, x_{s_{1}}\right\}$ is sampled element by element. If $s_{2}>u / 4$ then $\operatorname{Pr}\left[S_{1} \cap S_{2} \neq \varnothing\right] \geq$ $\operatorname{Pr}\left[x_{1} \in S_{2}\right] \geq 1 / 4$ and $\min \left(1, \frac{s_{1} s_{2}}{u}\right) \geq \min \left(1, s_{1} / 4\right)=\Theta(1)$ leaving nothing to show. From now on assume $s_{1} \leq s_{2} \leq u / 4$. By the chain rule, we have

$$
\begin{aligned}
\operatorname{Pr}\left[S_{1} \cap S_{2}=\varnothing\right] & =\operatorname{Pr}\left[x_{1} \notin S_{2}, \ldots, x_{s_{1}} \notin S_{2}\right] \\
& =\prod_{i=1}^{s_{1}} \operatorname{Pr}\left[x_{i} \notin S_{2} \mid x_{1} \notin S_{2}, \ldots, x_{i-1} \notin S_{2}\right] \\
& =\prod_{i=1}^{s_{1}}\left(1-\frac{s_{2}}{u-i+1}\right) .
\end{aligned}
$$

We can upper-bound this by

$$
\prod_{i=1}^{s_{1}}\left(1-\frac{s_{2}}{u-i+1}\right) \leq\left(1-\frac{s_{2}}{u}\right)^{s_{1} \mathrm{Eq}} \stackrel{[1]}{\leq}\left(\mathrm{e}^{-s_{2} / u}\right)^{s_{1}}=\mathrm{e}^{-s_{1} s_{2} / u},
$$

and lower bound it similarly by

$$
\prod_{i=1}^{s_{1}}\left(1-\frac{s_{2}}{u-i+1}\right) \geq\left(1-\frac{s_{2}}{u / 2}\right)^{s_{1} \text { Eq. }} \geq{ }^{11}\left(\mathrm{e}^{-4 s_{2} / u}\right)^{s_{1}}=\mathrm{e}^{-4 s_{1} s_{2} / u}
$$

If $s_{1} s_{2} / u \geq 1 / 100$ we obtain, by combining both bounds,

$$
\operatorname{Pr}\left[S_{1} \cap S_{2} \neq \varnothing\right]=1-\mathrm{e}^{-\Theta\left(s_{1} s_{2} / u\right)}=\Theta(1)=\Theta\left(\min \left(1, \frac{s_{1} s_{2}}{u}\right)\right)
$$

as desired. Similarly, if $s_{1} s_{2} / u<1 / 100$, we obtain

$$
\operatorname{Pr}\left[S_{1} \cap S_{2} \neq \varnothing\right]=1-\mathrm{e}^{-\Theta\left(s_{1} s_{2} / u\right)} \stackrel{\mathrm{Eq}}{=}{ }^{\mathbb{1}]} \Theta\left(s_{1} s_{2} / u\right)=\Theta\left(\min \left(1, \frac{s_{1} s_{2}}{u}\right)\right) .
$$

Theorem 2. Suppose $n \geq 2$. Let $D \in[m]^{n}$ and $k \in[m]$. Then,

$$
p_{\operatorname{BiNS}(k)}(D)=\Theta\left(\min \left(1, \frac{\|D\|_{1}^{2}-\|D\|_{2}^{2}}{k m}+\frac{n\|D\|_{1}}{m}+\frac{n^{2} k}{m}\right)\right) .
$$

Originally stated
on page 7

Proof. Let $D=\left(d_{1}, \ldots, d_{n}\right)$. If any $d_{i}$ exceeds $\lfloor m / k\rfloor \cdot k$ then the theorem correctly predicts a collision probability of $\Theta(1)$. Thus, we may assume that no instance runs out of bins. For every $i \in[n]$, let $B_{i}$ be the set of $\left\lceil d_{i} / k\right\rceil$ bins chosen by the $i$ th instance. Notice that $\left(B_{i}\right)_{i}$ is an independent family, because in our model instances behave independently. Let $C_{i j}$ be the event that instances $i$ and $j$ collide, which happens if and only if there is a bin selected by both instances. By Lemma 14 ,

$$
\operatorname{Pr}\left[\mathcal{C}_{i j}\right]=\operatorname{Pr}\left[B_{i} \cap B_{j} \neq \emptyset\right]=\Theta\left(\min \left(1, \frac{\left\lceil d_{i} / k\right\rceil\left\lceil d_{j} / k\right\rceil}{\lfloor m / k\rfloor}\right)\right) .
$$

The collision events $\mathcal{C}_{i j}$ are pairwise independent, by a similar argument as in the proof of Theorem 1. Thus, by Lemma 13 ,

$$
\begin{align*}
p_{\operatorname{BINS}(k)}(D) & =\Theta\left(\min \left(1, \sum_{i<j} \operatorname{Pr}\left[\mathcal{C}_{i j}\right]\right)\right) \\
& =\Theta\left(\min \left(1, \sum_{i<j} \min \left(1, \frac{\left\lceil d_{i} / k\right\rceil\left\lceil d_{j} / k\right\rceil}{\lfloor m / k\rfloor}\right)\right)\right) \\
& =\Theta\left(\min \left(1, \sum_{i<j} \frac{\left\lceil d_{i} / k\right\rceil\left\lceil d_{j} / k\right\rceil}{\lfloor m / k\rfloor}\right)\right) \tag{2}
\end{align*}
$$

Now we wish to compute the inner summation. Notice that

$$
\begin{aligned}
\frac{\left\lceil d_{i} / k\right\rceil\left\lceil d_{j} / k\right\rceil}{\lfloor m / k\rfloor} & =\Theta\left(\frac{\left(1+d_{i} / k\right)\left(1+d_{j} / k\right)}{m / k}\right) \\
& =\Theta\left(\frac{k+\left(d_{i}+d_{j}\right)+d_{i} d_{j} / k}{m}\right) .
\end{aligned}
$$

Straightforward calculations show that

$$
\sum_{i<j}\left(d_{i}+d_{j}\right)=\Theta\left(n\|D\|_{1}\right) \text { and } \sum_{i<j} d_{i} d_{j}=\Theta\left(\|D\|_{1}^{2}-\|D\|_{2}^{2}\right) .
$$

Thus,

$$
\sum_{i<j} \frac{\left\lceil d_{i} / k\right\rceil\left\lceil d_{j} / k\right\rceil}{\lfloor m / k\rfloor}=\Theta\left(\frac{n^{2} k+n\|D\|_{1}+\left(\|D\|_{1}^{2}-\|D\|_{2}^{2}\right) / k}{m}\right) .
$$

Plugging this equality into Equation (2) completes the proof.

## 5 Worst-Case Lower Bound Against Oblivious Adversaries

The goal of this section is to prove Theorem 6. The proof is in two steps: The first one is to establish a lower bound on the collision probability for the case of uniform demand profiles, that is, demand profiles of the form $(h, h, \ldots, h)$ for some $h \in \mathbb{N}$. Then, we extend this lower bound to most demand profiles, using the fact that most demand profiles "contain" a uniform demand profile with almost the same total number of requests.

### 5.1 Uniform Demand Profiles

We will use the following technical lemma.
Lemma 15. Assume $n \geq 2$ balls are thrown independently into $\ell \geq n$ bins where the bins are chosen with probabilities $p_{1}, \ldots, p_{\ell}$. The probability that all balls land in distinct bins is maximized if and only if $p_{1}=\ldots=p_{\ell}=1 / \ell$.

Proof. For $\ell=n=2$ the success probability $2 p_{1}\left(1-p_{1}\right)$ is maximized only for $p_{1}=p_{2}=1 / 2$. Now consider $\ell \geq 3$ and assume, for the sake of contradiction, that the maximum is achieved for a non-uniform distribution in which, say, $p_{1} \neq p_{2}$. We may imagine that the throwing of balls happens in two stages. In the first stage, the balls are thrown into $\ell-1$ bins with probabilities $p_{1}+p_{2}, p_{3}, \ldots, p_{\ell}$. In the second stage, the balls from the first bin are reclaimed and thrown into two new sub-bins with probabilities $p_{1} /\left(p_{1}+p_{2}\right)$ and $p_{2} /\left(p_{1}+p_{2}\right)$.

Now note that if we vary $p_{1}$ and $p_{2}$ while keeping $p_{1}+p_{2}$ fixed, the distribution of the balls in the first stage of the experiment is unaffected. And when we equalize the probabilities $p_{1}$ and $p_{2}$ of the sub-bins, the success probability in the second stage is improved, by the argument for the case $\ell=2$. (Of course, when the number of balls that land in the first bin in the first phase is not equal to 2 , the distribution in the sub-bins does not matter.) Thus, $p_{1}=p_{2}$ yields higher success probability, which is a contradiction.

Lemma 16. For all $D=(h, \ldots, h) \in \mathbb{N}^{n}$,

$$
p_{*}(D)=p_{\operatorname{BINS}(h)}(D) .
$$

Proof. In the scenario at hand every instance of an ID generation algorithm $\mathcal{A}$ has to produce exactly $h$ IDs, allowing us to think of $\mathcal{A}$ in simpler terms, namely as a distribution $\mu_{\mathcal{A}}$ on $\binom{[m]}{h}$, i.e., a distribution on sets of $h$ IDs. Conversely, any such distribution characterizes an algorithm. Assume $\mathcal{A}$ is optimal for $D$, i.e., $p_{\mathcal{A}}(D)=p_{*}(D)$ and that $\mathcal{A}$ minimizes the support size $\left|\operatorname{supp}\left(\mu_{\mathcal{A}}\right)\right|$ among the optimal algorithms, where

$$
\operatorname{supp}(\mu):=\left\{\left.X \in\binom{[m]}{h} \right\rvert\, \mu(X)>0\right\} .
$$

Our plan is to show that $\mathcal{A}$ behaves like $\operatorname{Bins}(h)$, up to relabeling of IDs.
We call a sequence $X_{1}, \ldots, X_{n} \in\binom{[m]}{h}$ of sets of IDs an outcome. When using $\mathcal{A}$ on demand profile $D$ the $n$ instances of $\mathcal{A}$ produce this outcome with probability $\mu_{\mathcal{A}}\left(X_{1}\right) \ldots \mu_{\mathcal{A}}\left(X_{n}\right)$. Let

$$
\operatorname{disj}\left(X_{1}, \ldots, X_{n}\right)= \begin{cases}1 & \text { if } X_{1}, \ldots, X_{n} \text { are pairwise disjoint } \\ 0 & \text { otherwise }\end{cases}
$$

The probability that no collision occurs is $q\left(\mu_{\mathcal{A}}\right)$ where $q: \mathbb{R}\binom{[m]}{h} \rightarrow \mathbb{R}$ is the function

$$
\begin{equation*}
q(\mu):=\sum_{X_{1}, \ldots, X_{n} \in\binom{[m]}{h}} \mu\left(X_{1}\right) \ldots \mu\left(X_{n}\right) \cdot \operatorname{disj}\left(X_{1}, \ldots, X_{n}\right) . \tag{3}
\end{equation*}
$$

Claim 1. $q(\mu)$ is an affine function ${ }^{4}$ on any two entries of $\mu$ corresponding to intersecting sets. More specifically, if $P, Q \in\binom{[m]}{h}$, with $P \cap Q \neq \varnothing$, and we fix every entry of $\mu$ except $\mu(P)$ and $\mu(Q)$, the resulting 2-dimensional function is affine.

[^4]Proof. This is because $q(\mu)$ is structurally a polynomial with a variable for every entry $\mu(X)$ of the input $\mu$, and the $\operatorname{disj}(\cdot)$-predicate filters out all monomials that contain $\mu(X) \mu(Y)$ with $X \cap Y \neq \varnothing$ (including, in particular, $\mu(X)^{2}$ ).

Claim 2. If $P, Q \in \operatorname{supp}\left(\mu_{\mathcal{A}}\right), P \neq Q$, then $P \cap Q=\varnothing$.
Proof. To reach a contradiction, assume $P \cap Q \neq \varnothing$. Then, Claim 1 implies that $q(\mu)$ is affine on $P$ and $Q$.

Consider a continuum $\left(\mu_{t}\right)_{t \in[0,1]}$ of variants of $\mu_{\mathcal{A}}$ where the probability mass on $\{P, Q\}$ linearly shifts from $Q$ to $P$, namely

$$
\mu_{t}(X):= \begin{cases}\left(\mu_{\mathcal{A}}(P)+\mu_{\mathcal{A}}(Q)\right) \cdot t & \text { if } X=P \\ \left(\mu_{\mathcal{A}}(P)+\mu_{\mathcal{A}}(Q)\right) \cdot(1-t) & \text { if } X=Q \\ \mu_{\mathcal{A}}(X) & \text { otherwise } .\end{cases}
$$

Since $\mu_{t}$ only linearly varies the entries corresponding to $P$ and $Q$, and $q$ is affine on those entries, the map $t \mapsto q\left(\mu_{t}\right)$ is affine. Note that $\mu_{\mathcal{A}}=\mu_{t}$ for some $t \in(0,1)$. Thus, we either have $q\left(\mu_{0}\right) \geq q\left(\mu_{\mathcal{A}}\right)$ or $q\left(\mu_{1}\right) \geq q\left(\mu_{\mathcal{A}}\right)$. Assume without loss of generality the former. Then $\mu_{0}$ is no worse than $\mu_{\mathcal{A}}$ in terms of success probability and has smaller support than $\mu_{\mathcal{A}}$, contradicting the choice of $\mathcal{A}$.

Hence, $\operatorname{supp}\left(\mu_{\mathcal{A}}\right)$ is some collection of pairwise disjoint sets from $\binom{[m]}{h}$. The sets in $\operatorname{supp}\left(\mu_{\mathcal{A}}\right)$ correspond to the disjoint bins used by $\operatorname{Bins}(h)$ up to relabeling of IDs. There are just two missing details. First, that $\mu_{\mathcal{A}}$ must assign the same probability to each bin, by Lemma $155^{5}$ Second, that $\operatorname{supp}\left(\mu_{\mathcal{A}}\right)$ is a maximal set of disjoint sets; otherwise, we can view $\mathcal{A}$ as a variant of $\operatorname{Bins}(h)$ that assigns probability 0 to some bins, which is again sub-optimal by Lemma 15. Overall we obtain

$$
p_{*}(D)=p_{\mathcal{A}}(D)=p_{\operatorname{BiNs}(h)}(D)
$$

as desired.
Because $\operatorname{Bins}(h)$ is optimal on demand profiles $(h, \ldots, h) \in \mathbb{N}^{n}$, it plays a central role in some of our analyses. In particular, the collision probability on uniform demand profiles, namely

$$
\begin{equation*}
p_{\operatorname{BiNS}(h)}((h, \ldots, h))=\Theta\left(\min \left(1, \frac{n^{2} h}{m}\right)\right) \tag{4}
\end{equation*}
$$

will come in handy. This probability follows directly from Theorem 2. Lemma 16 and Equation (4) imply the following.

Corollary 17. Let $D=(h, \ldots, h) \in \mathbb{N}^{n}$. Then

$$
p_{*}(D)=\Omega\left(\min \left(1, \frac{n^{2} h}{m}\right)\right) .
$$

### 5.2 Proof of Theorem 6

The intuition is that most demand profiles are sufficiently similar to the uniform demand profile $\left(\frac{d}{n}, \ldots, \frac{d}{n}\right)$, which we have analyzed in the preceding subsection. Let us call $D \in \mathcal{D}_{1}(n, d) \varepsilon$-good if at least $\varepsilon n$ entries of $D$ exceed $\varepsilon d / n$. Call $D \varepsilon$-bad otherwise. Let $\mathcal{B}_{\varepsilon}$ be the set of $\varepsilon$-bad demand profiles in $\mathcal{D}_{1}(n, d)$.

[^5]Lemma 18. There is a constant $\varepsilon \in\left(0, \frac{1}{2}\right]$ such that

$$
\left|B_{\varepsilon}\right| /\left|\mathcal{D}_{1}(n, d)\right|=\exp (-\Theta(n)) .
$$

Proof. In this proof we use standard bounds on binomial coefficients, namely

$$
\left(\frac{n}{k}\right)^{k} \leq\binom{ n}{k} \leq\left(\frac{n \mathrm{e}}{k}\right)^{k}
$$

We will also use the following basic combinatorial fact: Let $f(a, b)$ be the number of ways to throw $a$ indistinguishable balls into $b \leq a$ numbered bins, such that every bin contains at least one ball. Then,

$$
f(a, b)=\binom{a-1}{b-1} .
$$

The total number of demand profiles in $\mathcal{D}_{1}(n, d)$ is equal to $f(d, n)$, hence

$$
\left|\mathcal{D}_{1}(n, d)\right|=\binom{d-1}{n-1} \geq\left(\frac{d-1}{n-1}\right)^{n-1} \geq\left(\frac{d}{n}\right)^{n-1}
$$

Each $D \in \mathcal{B}_{\varepsilon}$ has at most $\varepsilon n$ entries exceeding $\varepsilon d / n$. To upper bound $\left|\mathcal{B}_{\varepsilon}\right|$ we first select $\varepsilon n$ out of $n$ indices for "large" entries that (potentially) exceed $\varepsilon d / n$. We then specify the remaining $n-\varepsilon n$ "small" entries explicitly, each between 1 and $\varepsilon d / n$. Finally, we specify how the remaining demand, which is at most $d$, is partitioned among the $\varepsilon n$ large entries. We get

$$
\begin{aligned}
\left|\mathcal{B}_{\varepsilon}\right| & \leq\binom{ n}{\varepsilon n}\left(\frac{\varepsilon d}{n}\right)^{n-\varepsilon n} f(d, \epsilon n) \\
& \leq 2^{n}\left(\frac{\varepsilon d}{n}\right)^{n-1}\left(\frac{n}{\varepsilon d}\right)^{\varepsilon n-1}\left(\frac{d \mathrm{e}}{\varepsilon n-1}\right)^{\varepsilon n-1} \\
& \leq 2^{n}\left(\frac{\varepsilon d}{n}\right)^{n-1}\left(\frac{2 \mathrm{e}}{\varepsilon^{2}}\right)^{\varepsilon n-1} .
\end{aligned}
$$

For the fraction of $\varepsilon$-bad demand profiles we obtain, assuming $\varepsilon>0$ is small enough

$$
\frac{\left|\mathcal{B}_{\varepsilon}\right|}{\left|\mathcal{D}_{1}(n, d)\right|} \leq 2^{n} \varepsilon^{n-1}\left(\frac{2 \mathrm{e}}{\varepsilon^{2}}\right)^{\varepsilon n-1}=\exp (-\Theta(n))
$$

We are now ready to present the proof of the main theorem of this section.
Theorem 6. For all but an $\exp (-\Theta(n))$-fraction of $D \in \mathcal{D}_{1}(n, d)$

$$
p_{*}(D)=\Omega\left(\min \left(1, \frac{n d}{m}\right)\right)
$$

Proof. By Lemma 18 it suffices to show that the lower bound on $p_{*}(D)$ holds for the $\varepsilon$-good demand profiles, for a suitable constant $\varepsilon$. For any $\varepsilon$-good $D \in \mathcal{D}_{1}(n, d)$, we can construct a new demand profile $D^{\prime}=\left(\frac{\varepsilon d}{n}, \ldots, \frac{\varepsilon d}{n}\right)$ with $n^{\prime}:=\varepsilon n$ entries, by decreasing some entries and removing some other entries from $D$. Then,

$$
\begin{aligned}
& p_{*}(D) \geq p_{*}\left(D^{\prime}\right) \stackrel{\text { Lem. }}{=} 16 \\
& p_{\operatorname{Bins}(\varepsilon d / n)}\left(D^{\prime}\right) \\
& \stackrel{\text { Eq. }}{=} \Omega\left(\min \left(1, \frac{\varepsilon n^{\prime} d}{m}\right)\right) \\
&=\Omega\left(\min \left(1, \frac{n d}{m}\right)\right) .
\end{aligned}
$$

## 6 Worst-Case Upper Bound Against Adaptive Adversaries

In this section we set out to analyze the worst-case collision probability against adaptive adversaries. Our first result shows that Cluster is vulnerable to adaptive adversaries, which can worsen the collision probability (compared to oblivious adversaries) by at least a factor of $n$.

Lemma 7. Suppose $n \geq 2$. Let $d \geq 2 n$. There exists an adaptive adversary $Z \in \operatorname{Adv}\left(\mathcal{D}_{1}(n, d)\right)$

Originally

Proof. Consider an adversary $Z$ that behaves as follows:

1. Request an ID from each of the $n$ instances.
2. Pick the two closest IDs; say they were produced by instances $i$ and $j$. Without loss of generality, assume instance $i$ produced the smaller ID of the two.
3. Request $d-n$ IDs from instance $i$.

Note that $Z \in \operatorname{Adv}\left(\mathcal{D}_{1}(n, d)\right)$. For each $i \in[n]$, let $x_{i}$ be the first ID produced by instance $i$. Observe that $Z$ causes a collision if and only if $x_{i}$ and $x_{j}$ are at distance at most $d-n-1$ modulo $m$, for some $i \neq j$. Let $\mathcal{E}_{i}$ be the event that $x_{i}$ is at distance at least $d-n-1$ from each of $x_{1}, \ldots, x_{i-1}$. Then,

$$
\begin{aligned}
1-p_{\text {Cluster }}(Z) & =\operatorname{Pr}\left[\mathcal{E}_{1}, \ldots, \mathcal{E}_{n}\right] \\
& =\prod_{i=1}^{n} \operatorname{Pr}\left[\mathcal{E}_{i} \mid \mathcal{E}_{1}, \ldots, \mathcal{E}_{i-1}\right] \\
& \leq \prod_{i=1}^{n}\left(1-\frac{(i-1)(d-n)}{m}\right) \\
& \text { Eq. } \mathbb{1} 1 \\
& \leq \prod_{i=1}^{n} \exp \left(-\frac{(i-1)(d-n)}{m}\right) \\
& =\exp \left(-\frac{n(n-1)(d-n)}{2 m}\right) \\
& d \geq 2 n \\
& \leq \exp \left(-\frac{n(n-1) d}{4 m}\right) .
\end{aligned}
$$

If $4 m<n(n-1) d$, then $1-p_{\text {Cluster }}(Z) \leq e^{-1}$. Thus, $p_{\text {Cluster }}(Z) \geq \Omega(1)=\Omega\left(\min \left(1, n^{2} d / m\right)\right)$, as desired. Otherwise, $4 m \geq n(n-1) d$, so we can apply Equation (1), to get

$$
1-p_{\text {CLUSTER }}(Z) \leq 1-\frac{n(n-1) d}{8 m}
$$

Thus, $p_{\text {Cluster }}(Z)=\Omega\left(n^{2} d / m\right)=\Omega\left(\min \left(1, n^{2} d / m\right)\right)$, concluding the proof.

### 6.1 A Nearly Optimal Algorithm

In the rest of this section we show that Cluster* is optimal in the worst case against adaptive adversaries, up to a logarithmic factor.

Theorem 8. Suppose $n \geq 2$. Let $d \geq n$ and $\mathcal{D}:=\mathcal{D}_{1}(d) \cap \mathcal{D}_{\infty}^{\leq}(n, m /(2 \log m))$. Then,

$$
\max _{Z \in \operatorname{Adv}(\mathcal{D})} p_{\text {Cluster }^{*}}(Z)=O\left(\min \left(1, \frac{n d}{m} \log \left(1+\frac{d}{n}\right)\right)\right)
$$

Proof. We may assume

$$
\begin{equation*}
m=\omega(n d \log (1+d / n)) ; \tag{5}
\end{equation*}
$$

otherwise the upper bound holds trivially. Let $Z \in \operatorname{Adv}(\mathcal{D})$. As soon as an instance selects a run and returns the first ID in it, we say the run has been opened. At any point during the game, we say an ID is active if it belongs to an opened run of some instance. Notice that the number of active IDs is at most $2 d$ at all times, because the number of IDs that are active but have not been requested yet is at most $d$.

Since we aim for an upper bound on the collision probability, we may count any intersection between the open runs of two instances as a collision, regardless of which IDs from these runs have already been returned by the corresponding instances.

Let $\mathcal{C}_{i}$ be the event that at least $i$ runs are open and that the $i$ th run that is opened collides with an already opened run. Let $I_{i} \in[n]$ be the instance that opens a run and let $T_{i}:=\mid\{j: j<$ $\left.i, I_{j}=I_{i}\right\} \mid$ be the number of runs previously opened by the same instance (with $\left(I_{i}, T_{i}\right)=(\perp,-\infty)$ if fewer than $i$ runs are ever opened). Let $\lambda$ be an upper bound on the number of runs that can be opened in total.

Claim 3. We have for any $t_{i} \in\{-\infty, 0,1,2, \ldots\}$

$$
\operatorname{Pr}\left[\mathcal{C}_{i} \mid \overline{\mathcal{C}_{1}}, \ldots, \overline{\mathcal{C}_{i-1}}, T_{i}=t_{i}\right] \leq O\left(\frac{d+\lambda 2^{t_{i}}}{m}\right) .
$$

Proof. Since $\operatorname{Pr}\left[C_{i} \mid T_{i}=-\infty\right]=0$ we may assume $t_{i} \geq 0$. When the $i$ th run is opened, instance $I_{i}$ has at least $m-\left(d+t_{i}\left(2^{t_{i}}-1\right)\right)$ valid choices for the first ID of the run. This is because the first ID cannot be any of the $d$ or less IDs within the runs $I_{i}$ has previously opened, nor one of the $2^{t_{i}}-1$ IDs directly to the left of the $t_{i}$ runs that $I_{i}$ has previously opened. Notice that Equation (5) implies that $m=\omega(d \log d)$. Combining this with the fact that $t_{i} \leq \log d$, we conclude that there are $\Omega(m)$ valid choices.

On the other hand, the number of choices that cause a collision are at most $2 d+\lambda\left(2^{t_{i}}-1\right)$. This is because a collision occurs when the first ID is either one of the $2 d$ or less active IDs, or is one of the $2^{t_{i}}-1$ IDs to the left of one the $\lambda$ or less runs from other instances.

Dividing the choices causing a collision by the number of valid choices yields the stated upper bound.

Claim 3 implies

$$
\operatorname{Pr}\left[\mathcal{C}_{i} \mid \overline{\mathcal{C}_{1}}, \ldots, \overline{\mathcal{C}_{i-1}}, T_{i}\right] \leq O\left(\frac{d}{m}+\frac{\lambda 2^{T_{i}}}{m}\right)
$$

Summing over $i$ gives,

$$
\begin{align*}
\sum_{i=1}^{\lambda} \operatorname{Pr}\left[\mathcal{C}_{i} \mid \overline{\mathcal{C}_{1}}, \ldots, \overline{\mathcal{C}_{i-1}}, T_{i}\right] & \leq O\left(\sum_{i=1}^{\lambda}\left(\frac{d}{m}+\frac{\lambda 2^{T_{i}}}{m}\right)\right) \\
& =O\left(\frac{d \lambda}{m}+\frac{\lambda}{m} \sum_{i=1}^{\lambda} 2^{T_{i}}\right) \\
& =O\left(\frac{d \lambda}{m}\right) \tag{6}
\end{align*}
$$

where the last step uses that $\sum_{i=1}^{\lambda} 2^{T_{i}}$ counts the number of active IDs in the end, which is at most $2 d$. We have,

$$
\begin{aligned}
& p_{\text {Cluster }^{*}}(Z) \leq \operatorname{Pr}\left[\mathcal{C}_{1} \cup\left(\overline{\mathcal{C}_{1}} \cap \mathcal{C}_{2}\right) \cup\left(\overline{\mathcal{C}_{1}} \cap \overline{\mathcal{C}_{2}} \cap \mathcal{C}_{3}\right) \cup \ldots\right] \\
&=\operatorname{Pr}\left[\mathcal{C}_{1}\right]+\operatorname{Pr}\left[\overline{\mathcal{C}_{1}} \cap \mathcal{C}_{2}\right]+\operatorname{Pr}\left[\overline{\mathcal{C}_{1}} \cap \overline{\mathcal{C}_{2}} \cap \mathcal{C}_{3}\right]+\ldots \\
& \leq \sum_{i=1}^{\lambda} \operatorname{Pr}\left[\mathcal{C}_{i} \mid \overline{\mathcal{C}_{1}}, \ldots, \overline{\mathcal{C}_{i-1}}\right] \\
&=\sum_{i=1}^{\lambda} \underset{T_{i}}{\mathbb{E}}\left[\operatorname{Pr}\left[\mathcal{C}_{i} \mid \overline{\mathcal{C}_{1}}, \ldots, \overline{\mathcal{C}_{i-1}}, T_{i}\right]\right] \quad \text { (by the law of total expectation) } \\
&=\underset{\left(T_{i}\right)_{i}}{\mathbb{E}}\left[\sum_{i=1}^{\lambda} \operatorname{Pr}\left[\mathcal{C}_{i} \mid \overline{\mathcal{C}_{1}}, \ldots, \overline{\mathcal{C}_{i-1}}, T_{i}\right]\right] \quad \text { (by linearity of expectation) } \\
& \leq O\left(\frac{d \lambda}{m}\right) . \\
& \text { (by Equation (6)) }
\end{aligned}
$$

To conclude the proof, we must show that $\lambda=O(n \log (1+d / n))$. Let $d_{i}$ be the (random) number of IDs requested from instance $i \in[n]$. Note that $d_{i}=0$ is permitted if only $n^{\prime}<n$ instances receive at least one request. Instance $i$ creates exactly $\left\lceil\log \left(1+d_{i}\right)\right\rceil$ runs. Thus, the total number of runs is

$$
\begin{aligned}
\sum_{i=1}^{n}\left\lceil\log \left(1+d_{i}\right)\right\rceil & \leq O\left(\sum_{i=1}^{n} \log \left(1+d_{i}\right)\right) \\
& \leq O\left(n \log \left(1+\frac{\sum_{i=1}^{n} d_{i}}{n}\right)\right)
\end{aligned}
$$

(by the concavity of $x \mapsto \log (1+x)$, and Jensen's inequality)

$$
\leq O(n \log (1+d / n)) .
$$

## 7 Competitive Ratio Upper Bound Against Oblivious Adversaries

In this section we take a closer look at Bins*, proving that it has competitive ratio $O(\log m)$ for a large class of demand profiles. This result will be strengthened to adaptive adversaries in Section 9 .

### 7.1 Missing Details in Bins*

Let us start by filling in some missing details from our description of Bins* from Section 3. Recall that Bins* partitions the IDs space $[m]$ into $O(\log m)$ chunks, and these are in turn broken into bins. Specifically, the number of chunks is $C:=\lceil\log m-\log \log m\rceil$, and each one has $2^{C-1}$ IDs. This works because

$$
\begin{aligned}
C \cdot 2^{C-1} & =\lceil\log m-\log \log m\rceil \cdot 2^{\lceil\log m-\log \log m\rceil-1} \\
& \leq \log m \cdot 2^{\log m-\log \log m}=m .
\end{aligned}
$$

For each $i \in[C]$, the $i$ th chunk is further partitioned into $2^{C-i}$ bins of size $2^{i-1}$ each.
Note that Bins* does not specify how to proceed after the selected bin from the last chunk has been exhausted. In that case at least $2^{C}$ IDs have been requested. Since $2^{C} \geq m / \log m$, the underlying demand profile would not be in $\mathcal{D}_{\infty}^{<}(m / \log m)$, so Theorem 9 makes no claim in this case.

### 7.2 Bins* $^{*}$ has Competitive Ratio $O(\log m)$ for Oblivious Adversaries

We begin by limiting the set of demand profiles we have to consider. For a (non-trivial) demand profile $D$ let $D^{-}$be the rounded demand profile arising from $D$ by

- first, rounding each entry in $D$ down to the next power of 2 , and
- second, if there is a unique largest entry, we reduce it to the second largest entry; when a unique largest entry exists, we call its associated instance heavy.

For instance if $D=(9,5,4,42)$ then $D^{-}=(8,4,4,8)$. For a set $\mathcal{D}$ of demand profiles we define $\mathcal{D}^{-}:=\left\{D^{-} \mid D \in \mathcal{D}\right\}$.

Lemma 19. Let $D \in \mathcal{D}_{\infty}^{<}(m / \log m)$. Then

$$
p_{\mathrm{BINS}^{*}}(D)=p_{\mathrm{BiNS}^{*}}\left(D^{-}\right)
$$

Proof. When using Bins* a collision occurs if and only if two instances allocate the same bin. It therefore only matters which bins are allocated by the instances. An instance allocates its $i$ th bin (of size $2^{i-1}$ ) when the previously allocated bins of sizes $1,2, \ldots, 2^{i-2}$ are all full and an additional ID is requested, i.e., when the $2^{i-1}$ th ID is requested. Thus, rounding down to powers of 2 does not change the collision probability.

If after rounding the is a heavy instance with $2^{i-1}$ requests for some $i$, while there are at most $2^{j-1}$ requests for each of the others instances for some $j<i$, then the heavy instance is the only one allocating a bin in each of the chunks $j+1, \ldots, i$. These bins cannot cause a collision, so the heavy instance might as well have only $2^{j-1}$ requests.

We define the rank distribution of a rounded demand profile $D^{-}$to be a vector $\left(s_{1}, \ldots, s_{k}\right)$ where $s_{i}$ is the number of times $2^{i-1}$ occurs in $D^{-}$and where $2^{k-1}$ is the largest entry of $D^{-}$.

Lemma 20. Let $D^{-}$be a rounded demand profile with rank distribution $\left(s_{1}, \ldots, s_{k}\right)$. Then

$$
p_{*}\left(D^{-}\right)=\Omega\left(\min \left(1, \frac{1}{m} \sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right)\right) .
$$

Proof. Assume we use an optimal algorithm $\mathcal{A}$ for $D^{-}$, i.e., one with $p_{\mathcal{A}}\left(D^{-}\right)=p_{*}\left(D^{-}\right)$. For $i \in[k]$, let $\mathcal{E}_{i}$ be the event that there is a collision between two instances with demand $2^{i-1}$. We have $\operatorname{Pr}\left[\mathcal{E}_{i}\right]=p_{\mathcal{A}}\left(D_{i}\right)$ where $D_{i}=\left(2^{i-1}, \ldots, 2^{i-1}\right)$ has length $s_{i} \geq 0$. Recall that $\operatorname{Bins}\left(2^{i-1}\right)$ is the optimal algorithm for $D_{i}$, by Lemma 16 . We compute

$$
\begin{aligned}
& \operatorname{Pr}\left[\mathcal{E}_{i}\right]=p_{\mathcal{A}}\left(D_{i}\right) \\
& \geq p_{*}\left(D_{i}\right) \\
& =p_{\operatorname{BiNS}\left(2^{i-1}\right)}\left(D_{i}\right) \\
& \stackrel{\text { Eq.(4] }}{=} \Theta\left(\min \left(1,\binom{s_{i}}{2} \frac{2^{i-1}}{m}\right)\right) .
\end{aligned}
$$

Notice that because we use $\binom{s_{i}}{2}$ instead of $s_{i}^{2}$, we correctly handle the cases $s_{i} \in\{0,1\}$, using the convention $\binom{0}{2}=\binom{1}{2}=0$. Using that the events $\mathcal{E}_{1}, \ldots, \mathcal{E}_{k}$ relate to disjoint sets of instances and
are hence independent, we get

$$
\begin{aligned}
p_{*}\left(D^{-}\right) & =p_{\mathcal{A}}\left(D^{-}\right) \\
& \geq \operatorname{Pr}\left[\bigcup_{i=1}^{k} \mathcal{E}_{i}\right] \\
& \stackrel{\text { Lem. }}{=} \Theta\left(\min \left(1, \sum_{i=1}^{k} \operatorname{Pr}\left[\mathcal{E}_{i}\right]\right)\right) \\
& =\Theta\left(\min \left(1, \frac{1}{m} \sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right)\right) .
\end{aligned}
$$

We will use the following simple inequalities involving binomial coefficients.
Lemma 21. For any $x, y \geq 0$
(i) $\binom{x+y}{2} \leq 3\binom{x}{2}+2 x+\frac{3}{2}\binom{y}{2}+\frac{y}{2}$,
(ii) $\binom{x+y}{2} \leq 4\binom{x}{2}+\frac{5}{3}\binom{y}{2}+O(1)$.

Proof. (ii) follows from (i) because $2 x=o\left(\binom{x}{2}\right)$ and $\frac{y}{2}=o\left(\binom{y}{2}\right)$. To prove (i) we proceed by induction. For $x=0$ or $y=0$ the claim is trivial. For $x, y \geq 1$ we distinguish two cases. In both we first apply the identity $\binom{a}{2}=\binom{a-1}{2}+\binom{a-1}{1}$ and then the induction hypothesis.

- Case 1: $x \leq \frac{y}{2}$.

$$
\begin{aligned}
\binom{x+y}{2} & =\binom{x+(y-1)}{2}+\binom{x+y-1}{1} \\
& \text { Ind. } 3\binom{x}{2}+2 x+\frac{3}{2}\binom{y-1}{2}+\frac{y-1}{2}+x+y-1 \\
& \leq 3\binom{x}{2}+2 x+\frac{3}{2}\binom{y-1}{2}+\frac{y}{2}+\frac{y}{2}+y-\frac{3}{2} \\
& \leq 3\binom{x}{2}+2 x+\frac{3}{2}\binom{y-1}{2}+\frac{y}{2}+\frac{3}{2}\binom{y-1}{1} \\
& \leq 3\binom{x}{2}+2 x+\frac{3}{2}\binom{y}{2}+\frac{y}{2} .
\end{aligned}
$$

- Case 2: $y \leq 2 x$.

$$
\begin{aligned}
\binom{x+y}{2} & =\binom{(x-1)+y}{2}+\binom{x+y-1}{1} \\
& \text { Ind. } \\
& \leq 3\binom{x-1}{2}+2(x-1)+\frac{3}{2}\binom{y}{2}+\frac{y}{2}+x+y-1 \\
& \leq 3\binom{x-1}{2}+2 x+\frac{3}{2}\binom{y}{2}+\frac{y}{2}+x+2 x-3 \\
& \leq 3\binom{x-1}{2}+2 x+\frac{3}{2}\binom{y}{2}+\frac{y}{2}+3\binom{x-1}{1} \\
& \leq 3\binom{x}{2}+2 x+\frac{3}{2}\binom{y}{2}+\frac{y}{2} .
\end{aligned}
$$

Lemma 22. Let $D^{-}$be a rounded demand profile with rank distribution $\left(s_{1}, \ldots, s_{k}\right)$. Then

$$
p_{\mathrm{BINS}^{*}}\left(D^{-}\right)=O\left(\frac{\log m}{m} \sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right)
$$

Proof. A collision within chunk $i$ can only be caused by instances that allocate a bin in chunk $i$. Let $s \geq i:=s_{i}+\ldots+s_{k}$ be the number of such instances. Any fixed one of such such instances collides with probability $O\left(\frac{\log m}{m} 2^{i}\right)$, because there are $2^{C-i}=\Omega\left(\frac{m}{\log m} 2^{-i}\right)$ bins within chunk $i$. By union bound over all pairs of instances and all chunks we obtain

$$
p_{\mathrm{BINS}^{*}}\left(D^{-}\right) \leq \frac{\log m}{m} \sum_{i=1}^{k}\binom{s \geq i}{2} 2^{i} .
$$

We define $X:=\sum_{i=1}^{k}\binom{s \geq i}{2} 2^{i}$ and compute, with the understanding that $s \geq k+1=0$,

$$
\begin{aligned}
X & =\sum_{i=1}^{k}\binom{s \geq i}{2} 2^{i}=\sum_{i=1}^{k}\binom{s_{i}+s_{\geq i+1}}{2} 2^{i} \\
& \left.\left.=4\left(\sum_{i=1}^{\text {Lem. }}{ }^{211} \sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right)+\frac{5}{6}\binom{s_{i}}{2}+\frac{5}{3}\binom{s \geq i+1}{2}+O(1)\right) \cdot 2^{i}\binom{s \geq i+1}{2} 2^{i+1}\right)+O(1) \cdot 2^{k+1} \\
& \leq O\left(\sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right)+\frac{5}{6} \cdot X .
\end{aligned}
$$

In the last step we used that $s_{k} \geq 2$ so that the $O(1) \cdot 2^{k+1}$ term is dominated by the last term of the first sum. After subtracting $\frac{5}{6} X$ we see that $X=O\left(\sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right)$, which implies the claim.

Theorem 9. The competitive ratio of $\operatorname{Bins}^{*}$ for $\mathcal{D}_{\infty}^{<}(m / \log m)$ is $O(\log m)$.

Originally

$$
\begin{aligned}
& p_{\text {Bins }^{*}}(D) \stackrel{\text { Lem. }}{=} \stackrel{19}{ } p_{\text {Bins }^{*}}\left(D^{-}\right) \\
& \stackrel{\text { Lem. }}{=}{ }^{[22]} O\left(\min \left(1, \frac{\log m}{m} \sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right)\right) \\
& \leq O(\log m) \min \left(1, \frac{1}{m} \sum_{i=1}^{k}\binom{s_{i}}{2} 2^{i}\right) \\
& \stackrel{\text { Lem. 20 }}{\leq} O(\log m) p_{*}\left(D^{-}\right) \leq O(\log m) p_{*}(D) \text {. }
\end{aligned}
$$

## 8 Competitive Ratio Lower Bound Against Oblivious Adversaries

In the following we consider the performance of an algorithm $\mathcal{A}$ on demand profiles of the form $D=(i, j)$ for $i, j \in[m]$, which will serve as "hard instances" for our lower bounds down the line. Denote $\mathcal{A}(i)$ the distribution of the set of the first $i$ IDs returned by an instance of $\mathcal{A}$. Moreover, let $q_{c, i}:=\operatorname{Pr}_{X \sim \mathcal{A}(i)}[c \in X]$ for $c, i \in[m]$.
Lemma 23. Let $\mathcal{A}$ be any algorithm and let $i, j \in[m]$. We have

$$
p_{\mathcal{A}}((i, j)) \geq \frac{1}{\min (i, j)} \sum_{c \in[m]} q_{c, i} q_{c, j} .
$$

Proof. We will use that a bounded random variable $Z \in\{0,1, \ldots, B\}$ satisfies $\operatorname{Pr}[Z>0]=\operatorname{Pr}[Z \geq$ $1] \geq \mathbb{E}[Z] / B$. We have

$$
\begin{aligned}
p_{\mathcal{A}}((i, j)) & =\operatorname{Pr}_{\substack{X \sim \mathcal{A}(i) \\
Y \sim \mathcal{A}(j)}}[|X \cap Y|>0] \\
& \geq \frac{1}{\min (i, j)} \underset{\substack{X \sim \mathcal{A}(i) \\
Y \sim \mathcal{A}(j)}}{\mathbb{E}}[|X \cap Y|] \\
& =\frac{1}{\min (i, j)} \sum_{c \in[m]} \operatorname{Pr}_{\substack{X \sim \mathcal{A}(i)}}^{Y \sim \mathcal{A}(j)}[c \in X \cap Y] \\
& =\frac{1}{\min (i, j)} \sum_{c \in[m]} q_{c, i} q_{c, j} .
\end{aligned}
$$

The following lemma determines the best achievable collision probability for demand profiles $D=(i, j)$, up to constants.

Lemma 24. Let $1 \leq i \leq j \leq m / 2$. Then, $p_{*}((i, j))=\Theta\left(\frac{i}{m}\right)$.
Proof. Using the optimality of $\operatorname{Bins}(i)$ on $D=(i, i)$ we get

$$
p_{*}((i, j)) \geq p_{*}((i, i)) \stackrel{\text { Lem. }}{=}{ }^{[16} p_{\operatorname{BINS}(i)}((i, i)) \stackrel{\text { Eq. [4] }}{=} \Omega\left(\frac{i}{m}\right) .
$$

For the upper bound we construct an algorithm $\mathcal{A}$ for the fixed demand profile $(i, j)$. The algorithm sets aside $j-i$ hard-wired IDs. The first $i$ requests are handled using $\operatorname{Bins}(i)$ on the rest of the IDs. All other requests (which are at most $j-i$ ) are served from the hard-wired IDs.

A collision on the demand profile $(i, j)$ occurs if and only if there is a collision between the first $i$ IDs of each instance, produced via $\operatorname{BINs}(i)$ on the reduced ID space with $m-j+i$ IDs. This probability is $p_{\operatorname{BiNS}(i)}((i, i))=\Theta(1 /(m-j+i))=\Theta(i / m)$.

Let $k=\left\lfloor\frac{1}{2} \log (m)\right\rfloor$. Consider the following distribution $\Phi$ on demand profiles:

$$
\begin{equation*}
\underset{D \sim \Phi}{\operatorname{Pr}}\left[D=\left(2^{i}, 2^{j}\right)\right]=\frac{1}{W} \cdot 2^{-\max (i, j)} \text { for } 0 \leq i, j \leq k \tag{7}
\end{equation*}
$$

where $W$ is a normalization factor. Note that

$$
W:=\sum_{i=0}^{k} \sum_{j=0}^{k} 2^{-\max (i, j)} \leq 2 \sum_{i \geq 0} \sum_{j \geq i} 2^{-j}=2 \sum_{i \geq 0} 2^{-i+1}=8=O(1) .
$$

Lemma 25. For any algorithm $\mathcal{A}$ and $\Phi$ as above we have

$$
\underset{D \sim \Phi}{\mathbb{E}}\left[p_{\mathcal{A}}(D)\right]=\Omega\left(\frac{\log ^{2} m}{m}\right)
$$

Proof. We reuse the numbers $q_{c, i}$ from Lemma 23. We have

$$
\begin{aligned}
& \underset{D \sim \Phi}{\mathbb{E}}\left[p_{\mathcal{A}}(D)\right]=\frac{1}{W} \sum_{i=0}^{k} \sum_{j=0}^{k} 2^{-\max (i, j)} p_{\mathcal{A}}\left(\left(2^{i}, 2^{j}\right)\right) \\
& \stackrel{\text { Lem. [23 }}{\geq} \frac{1}{W} \sum_{i=0}^{k} \sum_{j=0}^{k} 2^{-\max (i, j)} \frac{\sum_{c \in[m]} q_{c, 2^{i}} q_{c, 2^{j}}}{\min \left(2^{i}, 2^{j}\right)} \\
& \geq \frac{1}{W} \sum_{i=0}^{k} \sum_{j=0}^{k} 2^{-(i+j)} \sum_{c \in[m]} q_{c, 2^{i}} q_{c, 2^{j}} \\
& =\frac{1}{W} \sum_{c \in[m]} \sum_{i=0}^{k} 2^{-i} q_{c, 2^{i}} \sum_{j=0}^{k} 2^{-j} q_{c, 2^{j}} \\
& =\frac{1}{W} \sum_{c \in[m]}\left(\sum_{i=0}^{k} 2^{-i} q_{c, 2^{i}}\right)^{2} \geq \frac{1}{W m}\left(\sum_{c \in[m]} \sum_{i=0}^{k} 2^{-i} q_{c, 2^{i}}\right)^{2} \\
& \geq \frac{1}{W m}\left(\sum_{i=0}^{k} 2^{-i} 2^{i}\right)^{2}=\frac{1}{W m} \Omega\left(\log ^{2} m\right)=\Omega\left(\frac{\log ^{2} m}{m}\right) .
\end{aligned}
$$

Theorem 10. Every algorithm has competitive ratio $\Omega(\log m)$ for $\mathcal{D}=[\sqrt{m}]^{2}$.
Proof. Let $\mathcal{A}$ be an algorithm with competitive ratio $c$ for $[\sqrt{m}]^{2}$. Note that $[\sqrt{m}]^{2}$ includes the support of the distribution $\Phi$ from Equation (7). We have

Originally

$$
\begin{aligned}
\underset{D \sim \Phi}{\mathbb{E}}\left[p_{\mathcal{A}}(D)\right] & =\frac{1}{W} \sum_{i=0}^{k} \sum_{j=0}^{k} 2^{-\max (i, j)} p_{\mathcal{A}}\left(\left(2^{i}, 2^{j}\right)\right) \\
& \leq \frac{c}{W} \sum_{i=0}^{k} \sum_{j=0}^{k} 2^{-\max (i, j)} p_{*}\left(\left(2^{i}, 2^{j}\right)\right) \\
& \leq \frac{2 c}{W} \sum_{i=0}^{k} \sum_{j=i}^{k} 2^{-j} p_{*}\left(\left(2^{i}, 2^{j}\right)\right) \\
& \stackrel{\text { Lem. } 24}{\leq 2 c} \frac{2 c}{W} \sum_{i=0}^{k} \sum_{j=i}^{k} 2^{-j} \cdot O\left(\frac{2^{i}}{m}\right) \\
& \leq O\left(\frac{c}{m} \sum_{i=0}^{k} 2^{i} \sum_{j \geq i} 2^{-j}\right) \\
& =O\left(\frac{c}{m} \sum_{i=0}^{k} 2^{i} \cdot 2^{-i+1}\right)=O\left(\frac{c \log m}{m}\right) .
\end{aligned}
$$

Since $\mathbb{E}_{D \sim \Phi}\left[p_{\mathcal{A}}(D)\right]=\Omega\left(\frac{\log ^{2} m}{m}\right)$ by Lemma 25 , we conclude that $c=\Omega(\log m)$, as claimed.

## 9 Competitive Ratio Upper Bound Against Adaptive Adversaries

Assume we are using an ID generation algorithm $\mathcal{A}$, the current demand profile is $D=\left(d_{1}, \ldots, d_{n}\right)$ and no collision has occurred yet. In Lemma 7 we have seen an example where an adaptive adversary
of Cluster observes and exploits useful information about the state of the game namely, which two instances of Cluster produced initial IDs that are close to one another. In contrast, Random, $\operatorname{Bins}(k)$ and Bins* exhibit a symmetry between the bins that makes every game state with the same current demand profile (and no collision so far) equivalent. This means that an adaptive adversary cannot make any meaningful observations that could guide its choices, except for its ability to stop immediately when a collision occurs. This suggests that the competitive ratio for these algorithms does not increase (much) if adaptive adversaries are considered. In this section, we will prove the following formalization of this intuition:

Theorem 11. Let $\mathcal{A}$ be either $\operatorname{Bins}^{*}$ or $\operatorname{Bins}(k)$ for some $k \in[m]$, and let $\mathcal{D}$ be some set of demand profiles. If $\mathcal{A}$ has competitive ratio $c$ for $\mathcal{D}$, then $\mathcal{A}$ has competitive ratio at most $4 c$ for $\operatorname{Adv}(\mathcal{D})$.

The following notions will be useful. A sequence $S=\left(D_{0}, D_{1}, \ldots, D_{k}\right)$ of demand profiles is called a $\mathcal{D}$-demand sequence if $D_{0}=(), D_{k} \in \mathcal{D}$ and $D_{i+1}$ arises from $D_{i}$ either by appending a 1 or by incrementing one of the entries of $D_{i}$. For such an $S$ there is a simple adaptive adversary fol $(S)$ that follows the demand sequence $S$ as long as no collision has occurred. Should a collision occur when the current demand profile is $D_{i}, \operatorname{fol}(S)$ transitions to a demand profile $\tilde{D}_{i} \in \mathcal{D}$ reachable from $D_{i}$ that minimizes $p_{*}\left(\tilde{D}_{i}\right)$, and then immediately stops the game $\sqrt{6}$ We call fol $(S)$ a semi-adaptive adversary, as its only adaptive decisions depend on whether or not a collision has occurred so far.

In the following we take the adversarial perspective trying to maximize the competitive ratio of $Z \in \operatorname{Adv}(\mathcal{D})$ against a fixed $\mathcal{A}$.

Proof of Theorem 11. We fix the algorithm $\mathcal{A}$ to be either $\operatorname{Bins*}$ or $\operatorname{Bins}(k)$ and an arbitrary adaptive adversary $Z \in \operatorname{Adv}(\mathcal{D})$ with competitive ratio $c^{\prime}$. We have to show $c^{\prime} \leq 4 c$.

Below we will prove the following three claims:
Claim 4. There exists an adaptive adversary $Z^{\prime} \in \operatorname{Adv}(\mathcal{D})$ that achieves competitive ratio $c^{\prime}$, and that only learns after every step whether a collision has occurred or not. In particular, $Z^{\prime}$ does not learn the IDs produced by the instances of $\mathcal{A}$.

Claim 5. There exists a $\mathcal{D}$-demand sequence $S$ such that the semi-adaptive adversary $\operatorname{fol}(S)$ achieves competitive ratio at least $c^{\prime}$.
Claim 6. There exists a demand profile $D \in \mathcal{D}$ (i.e., an oblivious adversary) that achieves competitive ratio at least $c^{\prime} / 4$.

The last claim immediately implies what we need: since $\mathcal{A}$ has competitive ratio $c$ for $\mathcal{D}$, the competitive ratio $c^{\prime} / 4$ for $D \in \mathcal{D}$ from the claim satisfies $c^{\prime} / 4 \leq c$.

We now prove the claims in order.
Proof of Claim 4 . Consider the case of $\mathcal{A}=\operatorname{Bins}(k)$. For $\mathcal{A}=$ Bins* a similar argument works.
Since $\mathcal{A}$ repeatedly selects bins of size $k$ and then produces the IDs within the selected bins in increasing order, the $i$ th ID selected by an instance of $\mathcal{A}$ only carries information if $i \equiv 1(\bmod k)$, and this information is precisely the identity of the bin that has been selected. Let $Z^{\prime \prime}$ be a variant of $Z$ that is only told these bin identities in a correspondingly modified game. Clearly $Z^{\prime \prime}$ has the same competitive ratio as $Z$.

Now consider a further modification of the game where a uniformly random permutation $\pi$ of the bins is picked beforehand and where any bin identity that would be given to $Z^{\prime \prime}$ is first

[^6]permuted by $\pi$. Since $\operatorname{Bins}(k)$ is symmetric under any permutation of the bins and since bin permutations do not affect whether or not a collision has occurred, the random process as observed by $Z^{\prime \prime}$ does not change and the competitive ratio is still unaffected. Furthermore, the only relevant information obtained by $Z^{\prime \prime}$ is whether or not a collision has occured: Conditioned on a collision having occurred, no further information on the produced IDs makes a difference. Conditioned on no collision having occurred, the sequence of distinct bin identities that is given to $Z^{\prime \prime}$ is stochastically independently of the behavior of the instances of $\mathcal{A}$.

We can therefore construct a version $Z^{\prime}$ of $Z^{\prime \prime}$ that only learns if a collision has occurred so far and produces a random sequence of distinct bin identities internally, if needed.

Proof of Claim 5. The adversary $Z^{\prime}$ from Claim 4 may use randomization in every step. However, it learns no information during the game (except whether a collision has occurred) and can select the demand sequence $S$ that it follows (as long as no collision has occurred) at the start of the game, possibly at random according to some distribution $\Delta$. If a collision does occur, then the course of action that maximizes the competitive ratio is to stop as soon as possible with a final demand profile $D \in \mathcal{D}$ minimizing $p_{*}(D)$. In other words, we may assume that $Z^{\prime}$ behaves like fol $(S)$ for $S \sim \Delta$. Then,

$$
\begin{aligned}
c^{\prime} & =\frac{p_{\mathcal{A}}\left(Z^{\prime}\right)}{\underset{D \sim Z^{\prime}}{\mathbb{E}}\left[p_{*}(D)\right]} \\
& =\frac{\underset{S \sim \Delta}{\mathbb{E}}\left[p_{\mathcal{A}}(\operatorname{fol}(S))\right]}{\underset{S \sim \Delta}{\mathbb{E}}\left[\underset{D \sim \operatorname{fol}(S)}{\mathbb{E}}\left[p_{*}(D)\right]\right]} \\
& \leq \max _{S} \frac{p_{\mathcal{A}}(\operatorname{fol}(S))}{\underset{D \sim \operatorname{fol}(S)}{\mathbb{E}}\left[p_{*}(D)\right]} .
\end{aligned}
$$

In the last step we used that $\frac{a}{b} \leq \frac{c}{d}$ implies $\frac{a+c}{b+d} \leq \frac{c}{d}$ for any $a, b, c, d \in \mathbb{R}^{+}$and, more generally,

$$
\frac{\sum_{i \in[n]} a_{i}}{\sum_{i \in[n]} b_{i}} \leq \max _{i \in[n]} \frac{a_{i}}{b_{i}} \quad \text { for } a_{i}, b_{i} \in \mathbb{R}^{+} \text {and } i \in[n] .
$$

The demand sequence $S$ maximizing the quotient yields a semi-adaptive adversary fol $(S)$ with competitive ratio at least $c^{\prime}$, as desired.

Proof of Claim 6. Let $S=\left(D_{0}, \ldots, D_{k}\right)$ be the $\mathcal{D}$-demand sequence from Claim 5. Note that $p_{\mathcal{A}}(\operatorname{fol}(S))=p_{\mathcal{A}}\left(D_{k}\right)$, that the probability that $\operatorname{fol}(S)$ reaches the last demand profile without causing a collision is $\operatorname{Pr}_{D \sim \operatorname{fol}(S)}\left[D=D_{k}\right]=1-p_{\mathcal{A}}\left(D_{k-1}\right)$, and that $p_{\mathcal{A}}\left(D_{i}\right)$ is increasing in $i$. There are two cases:

- Case 1: $p_{\mathcal{A}}\left(D_{k}\right) \leq \frac{1}{2}$. Using that fol $(S)$ has competitive ratio at least $c^{\prime}$ we get

$$
\begin{aligned}
c^{\prime} & \leq \frac{p_{\mathcal{A}}(\mathrm{fol}(S))}{\underset{D \sim \operatorname{fol}(S)}{\mathbb{E}}\left[p_{*}(D)\right]} \\
& \leq \frac{p_{\mathcal{A}}\left(D_{k}\right)}{\operatorname{Pr}_{D \sim \mathrm{fol}(S)}\left[D=D_{k}\right] \cdot p_{*}\left(D_{k}\right)} \\
& \leq \frac{p_{\mathcal{A}}\left(D_{k}\right)}{\left(1-p_{\mathcal{A}}\left(D_{k-1}\right)\right) \cdot p_{*}\left(D_{k}\right)} \\
& \leq \frac{p_{\mathcal{A}}\left(D_{k}\right)}{\left(1-p_{\mathcal{A}}\left(D_{k}\right)\right) \cdot p_{*}\left(D_{k}\right)} \\
& \leq \frac{p_{\mathcal{A}}\left(D_{k}\right)}{\frac{1}{2} \cdot p_{*}\left(D_{k}\right)}=2 \cdot \frac{p_{\mathcal{A}}\left(D_{k}\right)}{p_{*}\left(D_{k}\right)} .
\end{aligned}
$$

Thus, $D_{k}$ has competitive ratio at least $c^{\prime} / 2$.

- Case 2: $p_{\mathcal{A}}\left(D_{k}\right) \geq \frac{1}{2}$. Since $p_{\mathcal{A}}\left(D_{0}\right)=0$, there exists an index $i$ such that $p_{\mathcal{A}}\left(D_{i-1}\right) \leq \frac{1}{2}$ and $p_{\mathcal{A}}\left(D_{i}\right) \geq \frac{1}{2}$. Let $J \in[k]$ be the random index of the step in which the first collision occurs, if any, and $J=k$ otherwise. Then,

$$
\begin{aligned}
c^{\prime} & \leq \frac{p_{\mathcal{A}}(\operatorname{fol}(S))}{\underset{D \sim \operatorname{fol}(S)}{\mathbb{E}}\left[p_{*}(D)\right]} \\
& \leq \frac{2 \cdot p_{\mathcal{A}}\left(D_{i}\right)}{\sum_{j=1}^{k} \operatorname{Pr}[J=j] \cdot p_{*}\left(\tilde{D}_{j}\right)} \\
& \leq \frac{2 \cdot p_{\mathcal{A}}\left(D_{i}\right)}{\operatorname{Pr}[J \geq i] \cdot p_{*}\left(D_{i}\right)} \\
& \leq \frac{2 \cdot p_{\mathcal{A}}\left(D_{i}\right)}{\left(1-p_{\mathcal{A}}\left(D_{i-1}\right)\right) \cdot p_{*}\left(D_{i}\right)} \\
& \leq \frac{2 \cdot p_{\mathcal{A}}\left(D_{i}\right)}{\frac{1}{2} \cdot p_{*}\left(D_{i}\right)}=4 \cdot \frac{p_{\mathcal{A}}\left(D_{i}\right)}{p_{*}\left(D_{i}\right)} .
\end{aligned}
$$

Thus, $D_{i}$ has competitive ratio at least $c^{\prime} / 4$.

## Acknowledgments

We thank the anonymous reviewers of PODS ' 23 for their valuable comments that greatly improved the presentation of our paper. We gratefully acknowledge support from NSF grants CNS-2118620 and CCF-2106999, as well as Deutsche Forschungsgemeinschaft (DFG) grant 465963632.

## References

[1] blkid(8). Linux man page;https://linux.die.net/man/8/blkid. Last accessed: 2023-03-20.
[2] Boost's uuids. https://www.boost.org/doc/libs/1_65_0/libs/uuid/uuid.html. Last accessed: 2022-11-28.
[3] C\#'s uuids. https://learn.microsoft.com/en-us/dotnet/api/system.guid?view= net-7.0. Last accessed: 2022-11-28.
[4] Go's uuids. https://pkg.go.dev/github.com/google/uuid. Last accessed: 2022-11-28.
[5] Java's uuids. https://docs.oracle.com/javase/7/docs/api/java/util/UUID.html. Last accessed: 2022-11-28.
[6] Linux's uuids. https://linux.die.net/man/3/libuuid. Last accessed: 2022-11-28.
[7] Mysql blog archive: Storing uuid values in mysql tables. https://dev.mysql.com/ blog-archive/storing-uuid-values-in-mysql-tables/. Last accessed: 2022-12-04.
[8] Newid function. Transact-SQL Documentation; https://learn.microsoft.com/en-us/sql/ t-sql/functions/newid-transact-sql?view=sql-server-ver16. Last accessed: 2022-1204.
[9] nmcli(1). Linux man page; https://linux.die.net/man/1/nmcli. Last accessed: 2023-0320.
[10] Objectid. MongoDB Manual; https://www.mongodb.com/docs/manual/reference/method/ ObjectId/. Last accessed: 2022-09-04.
[11] Python's uuids. https://docs.python.org/3/library/uuid.html. Last accessed: 2022-1128.
[12] Universally unique identifier. Minecraft Wiki; https://minecraft.fandom.com/wiki/ Universally_unique_identifier, Last accessed: 2022-12-04.
[13] Uuid function. Cassandra Documentation; https://cassandra.apache.org/doc/3.11/ cassandra/cql/functions.html\#uuid. Last accessed: 2022-12-04.
[14] Uuid type. PostgreSQL Documentation; https://www.postgresql.org/docs/current/ datatype-uuid.html. Last accessed: 2022-12-04.
[15] Information technology - open systems interconnection - procedures for the operation of osi registration authorities: Generation and registration of universally unique identifiers (uuids) and their use as asn. 1 object identifier components. ITU-T Recommendation X.667, International Telecommunication Union, September 2004.
[16] Haris Aftab, Komal Gilani, Ji-Eun Lee, Lewis Nkenyereye, SeungMyeong Jeong, and JaeSeung Song. Analysis of identifiers in iot platforms. Digital Communications and Networks, 6(3):333340, 2020.
[17] Dan Alistarh, James Aspnes, Seth Gilbert, and Rachid Guerraoui. The complexity of renaming. In Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science (FOCS), pages 718-727, USA, 2011. IEEE Computer Society.
[18] Hagit Attiya, Amotz Bar-Noy, Danny Dolev, David Peleg, and Rüdiger Reischuk. Renaming in an asynchronous environment. Journal of the ACM, 37(3):524-548, July 1990.
[19] Muhammad A. Awad, Saman Ashkiani, Rob Johnson, Martin Farach-Colton, and John D. Owens. Engineering a high-performance GPU b-tree. In PPoPP, pages 145-157. ACM, 2019.
[20] Michael A. Bender, Martin Farach-Colton, Mayank Goswami, Rob Johnson, Samuel McCauley, and Shikha Singh. Bloom filters, adaptivity, and the dictionary problem. In FOCS, pages 182193. IEEE Computer Society, 2018.
[21] Michael A. Bender, Martin Farach-Colton, William Jannen, Rob Johnson, Bradley C. Kuszmaul, Donald E. Porter, Jun Yuan, and Yang Zhan. An introduction to be-trees and writeoptimization. login Usenix Mag., 40(5), 2015.
[22] Zhichao Cao, Siying Dong, Sagar Vemuri, and David HC Du. Characterizing, modeling, and benchmarking rocksdb key-value workloads at facebook. In 18th USENIX Conference on File and Storage Technologies (FAST 20), pages 209-223, 2020.
[23] Peter Dillinger. Experimental support for sst unique ids. RocksDB repository; https:// github.com/facebook/rocksdb/pull/8990. Last accessed: 2022-11-28.
[24] Peter Dillinger. New stable, fixed-length cache keys. RocksDB repository; https://github. com/facebook/rocksdb/pull/9126. Last accessed: 2022-11-28.
[25] Siying Dong, Mark Callaghan, Leonidas Galanis, Dhruba Borthakur, Tony Savor, and Michael Strum. Optimizing space amplification in rocksdb. In CIDR, volume 3, page 3, 2017.
[26] Paulo Jesus, Carlos Baquero, and Paulo Sérgio Almeida. Id generation in mobile environments. In Proceedings of the Conference on Mobile and Ubiquitous Systems (CSMU), 2006.
[27] Paul J. Leach, Michael Mealling, and Richard Salz. A universally unique identifier (uuid) urn namespace. RFC 4122, Internet Engineering Task Force, July 2005.
[28] Pu Li, Zengzhong Li, Max Sigalov, Dafan Liu, and Knut Magne Risvik. Rocksdb in microsoft bing. Microsoft Bing Blogs; https://blogs.bing.com/Engineering-Blog/october-2021/ RocksDB-in-Microsoft-Bing. Last accessed: 2023-03-22.
[29] Christof Lutteroth and Gerald Weber. Efficient use of guids. In Proceedings of the Ninth International Conference on Parallel and Distributed Computing, Applications and Technologies (PDCAT), pages 115-120, 2008.
[30] Sarang Masti. How we built a general purpose key value store for facebook with zippydb. Engineering at Meta; https://engineering.fb.com/2021/08/06/core-data/zippydb/. Last accessed: 2023-03-22.
[31] Yoshinori Matsunobu. Myrocks: A space- and write-optimized mysql database. Engineering at Meta; https://engineering.fb.com/2016/08/31/core-data/ myrocks-a-space-and-write-optimized-mysql-database/. Last accessed: 2023-0322.
[32] E. H. Mckinney. Generalized birthday problem. The American Mathematical Monthly, 73(4):385-387, 1966.
[33] Martin Schaffer, Peter Schartner, and Stefan Rass. Universally unique identifiers: How to ensure uniqueness while protecting the issuer's privacy. In Selim Aissi and Hamid R. Arabnia, editors, Proceedings of the 2007 International Conference on Security 83 Management (SAM), pages 198-204. CSREA Press, June 2007.
[34] Richard von Mises. Selected Papers of Richard Von Mises: Probability and statistics, general. Number Volume 2. American Mathematical Society, 1964.


[^0]:    *Meta Platforms, Inc. peterd@meta.com
    ${ }^{\dagger}$ Rutgers University. martin@farach-colton.com
    ${ }^{\ddagger}$ Rutgers University. guido.tag@rutgers.edu
    ${ }^{\S}$ University of Cologne. walzer@cs.uni-koeln. de

[^1]:    ${ }^{1}$ Notably, MongoDB also uses a variant of Cluster to identify records 10 .

[^2]:    ${ }^{2}$ The minimum is well-defined because it is taken over a compact space of distributions over permutations, and $\mathcal{A}^{\prime} \mapsto p_{\mathcal{A}^{\prime}}(D)$ is continuous.
    ${ }^{3}$ Our competitive analysis does not compare against an offline setting where the $i$ th instance of $\mathcal{A}$ knows in advance how many requests $d_{i}$ it will receive. This would make $\mathcal{A}$ overly powerful, as it could sometimes avoid collisions with

[^3]:    certainty (e.g., when the $d_{i}$ 's are all different).

[^4]:    ${ }^{4}$ A function $f: \mathbb{R}^{n} \rightarrow \mathbb{R}$ is affine if $f\left(x_{1}, \ldots, x_{n}\right)=a_{1} x_{1}+\cdots+a_{n} x_{n}+b$ for some constants $a_{1}, \ldots, a_{n}, b \in \mathbb{R}$.

[^5]:    ${ }^{5} \mathrm{An}$ exception is the case $n h>m$. Then collisions are guaranteed and any algorithm is optimal.

[^6]:    ${ }^{6}$ If $\mathcal{D}$ is downward closed, then we always have $D_{i} \in \mathcal{D}$ and hence $\tilde{D}_{i}=D_{i}$, meaning fol $(S)$ always stops immediately upon a collision.

