

ExaLogLog: Space-Efficient and Practical Approximate Distinct Counting up to the Exa-Scale

Otmar Ertl
 Dynatrace Research
 Linz, Austria
 otmar.ertl@dynatrace.com

ABSTRACT

This work introduces ExaLogLog, a new data structure for approximate distinct counting, which has the same practical properties as the popular HyperLogLog algorithm. It is commutative, idempotent, mergeable, reducible, has a constant-time insert operation, and supports distinct counts up to the exa-scale. At the same time, as theoretically derived and experimentally verified, it requires 43% less space to achieve the same estimation error.

1 INTRODUCTION

Exact counting of distinct elements in a data set or a data stream is known to take linear space [4]. However, the space requirements can be significantly reduced, if approximate results are sufficient. HyperLogLog (HLL) [20] with improved small-range estimation [18, 23, 37, 42, 49] has become the standard algorithm for approximate distinct counting. It achieves a relative standard error of $1.04/\sqrt{m}$ up to distinct counts in the order of $2^{64} \approx 1.8 \cdot 10^{19}$ using only $6m$ bits [23]. Therefore, the query languages of many data stores (see e.g. documentation of Timescale, Redis, Oracle Database, Snowflake, Microsoft SQL Server, Google BigQuery, Vertica, Elasticsearch, Aerospike, Amazon Redshift, KeyDB, or DuckDB) offer special commands for approximate distinct counting that are usually based on HLL. Query optimization [22, 31], caching [46], graph analysis [7, 35], attack detection [9, 11], network volume estimation [6], or metagenomics [5, 8, 17, 28] are further applications of HLL.

HLL is actually very simple as exemplified in Algorithm 1. It typically consists of a densely packed array of 6-bit registers r_0, r_1, \dots, r_{m-1} where the number of registers m is a power of 2, $m = 2^p$ [23]. The choice of the precision parameter p allows trading space for better estimation accuracy. Adding an element requires calculating a 64-bit hash value. p bits are used to choose a register for the update. The number of leading zeros (NLZ) of the remaining $64 - p$ bits are interpreted as a geometrically distributed update value ≥ 1 with success probability $\frac{1}{2}$, that is used to update the selected register. A register always holds the maximum of all its previous update values. Estimating the distinct count from the register values is more challenging, but can also be implemented using a few lines of code [18, 20]. HLL owes its popularity to the following features allowing it to be used in distributed systems [19]:

Speed: Element insertion is a fast and allocation-free operation with a constant time complexity independent of the sketch size. In particular, given the hash value of the element, the update requires only a few CPU instructions.

Idempotency: Further insertions of the same element will never change the state. This is actually a natural property every algorithm for distinct counting should support to prevent duplicates from changing the result.

Mergeability: Partial results calculated over subsets can be easily merged to a final result. This is important when data is distributed, is processed in parallel, or needs to be aggregated.

Reproducibility: The result does not depend on the processing order, which often cannot be guaranteed in practice anyway. Reproducibility is achieved by a commutative insert operation and a commutative and associative merge operation.

Reducibility: The state can be reduced to a smaller state corresponding to a smaller precision parameter. The reduced state is identical to that obtained by direct recording with lower precision. This property allows adjusting the precision without affecting the mergeability with older records.

Estimation: A fast and robust estimation algorithm ensures nearly unbiased estimates with a relative standard error bounded by a constant over the full range of practical distinct counts.

Simplicity: The implementation requires only a few lines of code. The entire state can be stored in a single byte array of fixed length which makes serialization very fast and convenient. Furthermore, add and in-place merge operations do not require any extra memory allocations.

Until recently HLL was the most space-efficient practical data structure having all these desired properties. Space efficiency can be measured in terms of the memory-variance product (MVP) [34], which is the relative variance of the (unbiased) distinct-count estimate \hat{n} multiplied by the storage size in bits

$$\text{MVP} := \text{Var}(\hat{n}/n) \times (\text{storage size in bits}), \quad (1)$$

where n is the true distinct count. If the MVP is asymptotically (for sufficiently large n) a constant specific to the data structure, it can be used for comparison as it eliminates the general inverse dependence of the relative estimation error on the root of the storage size (see Figure 1). Most HLL implementations use 6-bit registers [23] to support distinct counts beyond the billion range, resulting in a theoretical MVP of 6.48 [34]. A recent theoretical work conjectured a general lower bound of 1.98 for the MVP of sketches supporting mergeability and reproducibility [33], which shows the potential for improvement.

Algorithm 1: Inserts an element with 64-bit hash value $\langle h_{63}h_{62} \dots h_0 \rangle_2$ into a HyperLogLog consisting of $m = 2^p$ ($p \geq 2$) 6-bit registers r_0, r_1, \dots, r_{m-1} with initial values $r_i = 0$.

$i \leftarrow \langle h_{63}h_{62} \dots h_{64-p} \rangle_2$	▶ extract register index
$a \leftarrow \langle \underset{p}{0 \dots 0} h_{63-p} h_{62-p} \dots h_0 \rangle_2$	▶ mask register index bits
$k \leftarrow \text{nlz}(a) - p + 1$	▶ update value $k \in [1, 65 - p] \subseteq [1, 63]$
$r_i \leftarrow \max(r_i, k)$	▶ update register

© 2025 Copyright held by the owner/author(s). Published in Proceedings of the 28th International Conference on Extending Database Technology (EDBT), 25th March–28th March, 2025, ISBN 978-3-89318-099-8 on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.




Figure 1: The memory over the relative standard error for different memory-variance products (MVPs) following (1).

1.1 Related Work

In the past, different approaches have been proposed to outperform the space efficiency of HLL. However, many of them sacrificed at least one of the properties listed above [10, 24, 27, 34, 39, 48]. Therefore, in the following, we focus only on algorithms that support at least mergeability, idempotency, and reproducibility, which we consider to be essential properties for practical applications in distributed systems.

Lossless compression can significantly reduce the storage size of HLL [16, 26, 38]. Since the compressed state prevents random access to registers as required for insertion, bulking is needed to realize at least amortized constant update times. The required buffer partially cancels out the memory savings. Recent techniques avoid buffering of insertions. The Apache DataSketches library [1] provides an implementation using 4 bits per register to store the most frequent values relative to a global offset. Out of range values are kept separately in an associative array. Overall, this leads to a smaller MVP but also to a more expensive insert operation. Its runtime is proportional to the memory size in the worst case, because all registers must be updated whenever the global offset is increased. HyperLogLogLog (HLLL) [25] takes this strategy to the extreme with 3-bit registers and achieves a space saving of around 40% at the expense of an insert operation which, except for very large numbers, has been reported to be on average more than an order of magnitude slower compared to HLL [25].

Compression can also be applied to other data structures. Probabilistic counting with stochastic averaging (PCSA) [21] is a predecessor of HLL, also known as FM-sketch. Although less space-efficient when uncompressed, its compressed state leads to a smaller MVP than that of HLL [26, 38]. The compressed probability counting (CPC) sketch as part of the Apache DataSketches library [1] uses this finding. The serialized compressed representation of the CPC sketch achieves a MVP of around 2.31 [2] that is already quite close to the conjectured lower bound of 1.98 [33]. However, this is achieved by a costly consolidation and compression step, as the in-memory representation is more than twice as large, since an amortized constant runtime of insertions can only be achieved by bulking.

Recently, data structures have been proposed that are more space-efficient than HLL while not giving up constant-time insertions. ExtendedHyperLogLog (EHLL) [30] extends the HLL registers from 6 to 7 bits to store not only the maximum update value, but also whether there was an update with a value smaller by one. This additional information can be used to obtain more accurate estimates. In particular, the MVP is reduced by 16% to 5.43. Inspired by this result, we analyzed, if even more additional bits used to store the information about the occurrence of smaller

Table 1: Notations.

Symbol	Description
$\lfloor \dots \rfloor$	floor function, e.g. $\lfloor 3.7 \rfloor = 3$
$\langle \dots \rangle_2$	binary representation, e.g. $\langle 110 \rangle_2 = 6$
$ $	bitwise OR operation, e.g. $\langle 1001 \rangle_2 \langle 1010 \rangle_2 = \langle 1011 \rangle_2$
$\&$	bitwise AND operation, e.g. $\langle 1001 \rangle_2 \& \langle 1010 \rangle_2 = \langle 1000 \rangle_2$
nlz	number of leading zeros if the argument is interpreted as unsigned 64-bit value, e.g. $\text{nlz}(\langle 10110 \rangle_2) = 59$
n	distinct count
\hat{n}	distinct count estimate
b	base, $b > 1$, defines distribution of update values, compare (2)
p	precision parameter
m	number of registers, $m = 2^p$
q	number of bits used for storing an update value
d	number of additional register bits to indicate updates with smaller values
t	parameter of approximated update value distribution, compare (8)
r_i	value of i -th register, $0 \leq i < m$, $0 \leq r_i \leq (65 - p - t)2^{t+d} + 2^d - 1$
ρ_{update}	probability mass function of update values, see (2) and (8)
ρ_{reg}	probability mass function of register values, see Section 3.1
ρ_{token}	probability mass function of hash tokens, see (24)
v	hash token parameter, hash token takes $v + 6$ bits, see Section 4.3
\mathcal{L}	likelihood function, $\mathcal{L} = \mathcal{L}(n r_0 \dots r_{m-1})$, see Sections 3.2 and 4.3
Γ	gamma function, $\Gamma(x) := \int_0^\infty y^{x-1} e^{-y} dy$
ζ	Hurwitz zeta function, $\zeta(x, y) := \sum_{u=0}^\infty (u + y)^{-x} = \frac{1}{\Gamma(x)} \int_0^\infty \frac{z^{x-1} e^{-yz}}{1 - e^{-z}} dz$
μ	probability that the next distinct element causes a state change, see Section 3.3
expm1	$\text{expm1}(x) := e^x - 1$, built-in function available in most standard libraries
log1p	$\text{log1p}(x) := \ln(1 + x)$, built-in function available in most standard libraries

update values could further improve the MVP [19]. The result of our theoretical analysis led to UltraLogLog (ULL), which uses two extra bits and achieves a MVP of 4.63, an improvement of 28% over HLL. Our theoretical analysis also showed that further space savings would be possible with even more additional bits if, at the same time, success probabilities smaller than $\frac{1}{2}$ were used for the geometric distribution of the update values. However, as the generation of such update values is more complicated, this approach has not been pursued so far.

Another recent data structure, SpikeSketch [15], has chosen a different approach and combines update values following a geometric distribution with a success probability $\frac{3}{4}$ and a special lossy encoding scheme to improve space efficiency. It preserves mergeability and idempotency and also has a constant update time. Unfortunately, our experiments presented later could not confirm the claimed MVP of 4.08 over the entire range of distinct counts. We found values that are significantly larger for distinct counts below 10^4 caused by lossy compression and also the stepwise smoothing.

This work focuses on distinct counting of single, possibly distributed, data flows. Data sketches for (non-distinct) counting [29, 36, 40] or distinct counting of multiple flows [42, 47] were thus out of scope.

1.2 Summary of Contributions

We introduce ExaLogLog (ELL), which is based on a recently proposed and theoretically analyzed data structure [19] that generalizes HLL, EHLL, ULL, and PCSA. However, we replace the geometric distribution of the update values by a distribution for which it is easier to map a 64-bit hash value to a corresponding random value. When optimally configured, ELL achieves a MVP of 3.67 as theoretically predicted and experimentally confirmed. Compared to HLL with 6-bit registers [23], which is still the standard algorithm in most applications, ELL supports the same operating range up to the exa-scale, but requires up to 43% less space. In contrast to compressed probability counting (CPC) [1, 26] or HyperLogLogLog (HLLL) [25], the insert operation takes constant time independent of the sketch size and therefore also independent of the configured precision.

We will show that the new update value distribution leads to a simple maximum likelihood (ML) equation for the ELL state, which has only a small number of terms regardless of the chosen precision. We propose a robust and efficient algorithm based on Newton's method to solve this equation and, hence, to quickly find the distinct count estimate.

Moreover, we considered martingale estimation, which is the standard estimation algorithm, if the data is not distributed and merging is not needed [12, 41]. For this case, the theory also predicts significant reductions of the MVP by up to 33 % compared to HLL, which was again confirmed experimentally.

Finally, we present a strategy to realize a sparse mode for ELL by transforming the 64-bit hash values to smaller hash tokens which can be collected and equivalently used for insertions. We also demonstrate how the distinct count can be directly estimated from those hash tokens using ML estimation.

A reference implementation of ELL written in Java can be found at <https://github.com/dynatrace-research/exaloglog-paper>. This repository also contains all the source code and instructions for reproducing the results and figures presented in this work. Table 1 summarizes the notations used in the following.

2 DATA STRUCTURE

Recently, we have introduced and theoretically analyzed a data structure that can be seen as generalization of HLL, EHLL and PCSA, and that finally led us to ULL, another special case [19]. The generalized data structure consists of $m = 2^p$ registers. The update operation uses p bits of a uniformly distributed hash value for selecting a register and another geometrically distributed hash value with probability mass function (PMF)

$$\rho_{\text{update}}(k) = (b-1)b^{-k} \quad k \geq 1, b > 1 \quad (2)$$

for updating the selected register. Each register consists of $q+d$ bits. The first q bits store the maximum u of all update values the register was updated with, and the remaining d bits indicate the occurrences of update values from the range $[u-d, u-1]$. By definition, this data structure is idempotent, as duplicate insertions will never modify the state. Since the state may only change for distinct insertions, it enables distinct count estimation.

2.1 Previous Theoretical Results

We have derived and presented various expressions for the MVP of this generalized data structure [19]. If the registers are densely stored in a bit array, the theoretical MVP for an efficient unbiased estimator meeting the Cramér-Rao bound is

$$\text{MVP} \approx \frac{(q+d) \ln b}{\zeta\left(2, 1 + \frac{b-d}{b-1}\right)}, \quad (3)$$

where ζ is the Hurvitz zeta function as defined in Table 1. For ULL, as a special case with $b = 2$ and $d = 2$, we have shown that this MVP can also be practically achieved using ML estimation. We have also derived a first-order bias correction [14] that can be applied to the ML estimate \hat{n}_{ML} according to

$$\hat{n} = \frac{\hat{n}_{\text{ML}}}{1 + \frac{c}{m}} \quad \text{with } c := (\ln b) \left(1 + 2 \frac{b-d}{b-1}\right) \frac{\zeta\left(3, 1 + \frac{b-d}{b-1}\right)}{\left(\zeta\left(2, 1 + \frac{b-d}{b-1}\right)\right)^2}. \quad (4)$$

c is a constant dependent on the parameters. The bias correction factor $(1 + \frac{c}{m})^{-1}$ approaches 1 as $m \rightarrow \infty$.

If the state size is measured in terms of the Shannon entropy, corresponding to optimal compression of the state, a different




Figure 2: Comparing the PMFs (2) and (8) for $b = 2^{2^{-t}}$.

expression was obtained for the MVP

$$\text{MVP} \approx \frac{\left(1 + \frac{b-d}{b-1}\right)^{-1} + \int_0^1 z \frac{b-d}{b-1} \frac{(1-z) \ln(1-z)}{z \ln z} dz}{\zeta\left(2, 1 + \frac{b-d}{b-1}\right) \ln 2}. \quad (5)$$

This number is also known as Fisher-Shannon (FISH) number, for which a theoretical lower bound of 1.98 was postulated [33].

In a non-distributed setting, where merging of sketches is not needed, martingale estimation can be used [12, 41]. This estimation method is known to be efficient [34] and leads to a smaller MVP than (3). The corresponding asymptotic MVP, if again the registers are densely stored in a bit array, is given by

$$\text{MVP} \approx \frac{(q+d) \ln b}{2} \left(1 + \frac{b-d}{b-1}\right). \quad (6)$$

In contrast, if the registers are optimally compressed, the MVP would be

$$\text{MVP} \approx \frac{1 + \left(1 + \frac{b-d}{b-1}\right) \int_0^1 z \frac{b-d}{b-1} \frac{(1-z) \ln(1-z)}{z \ln z} dz}{2 \ln 2}. \quad (7)$$

This expression has a lower bound of 1.63 which corresponds to the theoretical limit [34].

2.2 Approximated Update Value Distribution

Standard hash functions usually give a uniformly distributed 64-bit hash value. Getting a hash value that is geometrically distributed according to (2) is not straightforward except for $b = 2$, where the number of leading zeros (NLZ) of a uniformly distributed hash value can be taken. For $b \neq 2$, the hash value could be used to seed a pseudo-random number generator that generates values according to the desired distribution. However, this involves floating-point operations, which might be slow on devices without floating-point units. Alternatively, hash ranges that map to the same update value can be precomputed, which allows finding the update value via linear or binary search.

Both approaches will be slower and not branch-free, unlike for $b = 2$. They may also introduce inaccuracies for huge distinct counts, as the probabilities of update values will always be multiples of 2^{-64} when using 64-bit hashes.

Therefore, we propose to use a different update value distribution with PMF

$$\rho_{\text{update}}(k) = \frac{1}{2^{t+1+(k-1)/2^t}} \quad \text{with } k \geq 1, t \geq 0. \quad (8)$$

This distribution approximates geometric distributions (2) with base value $b = 2^{2^{-t}}$ quite well as shown in Figure 2. The reason is that chunks of 2^t subsequent update values have always the same total probability. In other words, $\sum_{k=c2^t+1}^{c2^t+2^t} \rho_{\text{update}}(k) = \frac{1}{2^{c+1}}$ holds for (2) and (8) and for all $c \geq 0$.

A big advantage of (8) over (2) is that update values can be easily and accurately generated from a 64-bit hash value using a few CPU instructions by taking t bits of the hash value and

Algorithm 2: Inserts an element with 64-bit hash value $\langle h_{63}h_{62} \dots h_0 \rangle_2$ into an ExaLogLog with $(6 + t + d)$ -bit registers r_0, r_1, \dots, r_{m-1} ($m = 2^p$) and initial values $r_i = 0$.

```

 $i \leftarrow \langle h_{p+t-1}h_{p+t-2} \dots h_t \rangle_2$  ▶ extract register index
 $a \leftarrow \langle h_{63}h_{62} \dots h_{p+t-1} \dots 1 \rangle_2$ 
 $k \leftarrow \text{nlz}(a)2^t + \langle h_{t-1}h_{t-2} \dots h_0 \rangle_2 + 1$  ▶ compute update value (9),
 $k \in [1, (65 - p - t)2^t]$ 

 $u \leftarrow \lfloor r_i / 2^d \rfloor$  ▶ get max. update value from  $r_i$ , right-shift by  $d$  bits
 $\Delta \leftarrow k - u$ 
if  $\Delta > 0$  then
   $r_i \leftarrow k \cdot 2^d + \lfloor (2^d + (r_i \bmod 2^d)) / 2^\Delta \rfloor$ 
else if  $\Delta < 0 \wedge d + \Delta \geq 0$  then
   $r_i \leftarrow r_i \mid 2^{d+\Delta}$  ▶  $|$  denotes bitwise OR

```

determining the NLZ of the remaining bits according to

$$\text{update value} = \text{NLZ} \times 2^t + (\text{value of the } t \text{ bits}) + 1. \quad (9)$$

Furthermore, as we will see in Section 3.2, this distribution simplifies maximum likelihood estimation a lot compared to a geometric distribution with $b \neq 2$ thanks to its power-of-two probabilities.

2.3 ExaLogLog

Our new data structure ExaLogLog (ELL) is based on the generalized data structure as we have recently proposed in [19] and also briefly discussed in Section 2, however, using the approximate distribution (8) instead of the geometric distribution (2) for the update values. Hence, the parameter b is replaced by t were the two distributions are similar for $b = 2^{2^{-t}}$. For $t = 0 \Leftrightarrow b = 2$ the two distributions are even identical.

As the generalized data structure, ELL consists of $m = 2^p$ registers. Every time an element is inserted, a 64-bit hash value is computed of which p bits are used to select one of the registers. The remaining $64 - p$ bits are used to generate an update value according to (8). Each register consists of $q + d$ bits. The first q bits store the maximum update value u seen so far for that register. The remaining d bits memorize the occurrences of update values in the range $[u - d, u - 1]$.


To live up its name, ELL should, like HLL with 6-bit registers [23], support distinct counts up to the exa-scale, which is sufficient for any conceivable practical use case. As the maximum supported distinct count is roughly given by b^{2^q} [19], $q = 6 + t$ must be chosen in order to get $b^{2^q} = (2^{2^{-t}})^{2^{6+t}} = 2^{64} \approx 1.8 \times 10^{19}$.

The update procedure for inserting an element is summarized by Algorithm 2 and exemplified in Figure 3. It splits the hash value into three parts. The first $64 - p - t$ bits are used to determine the number of leading zeros (NLZ) which is therefore in the range $[0, 64 - p - t]$. The next p bits address the register, and the remaining last t bits, in combination with the NLZ of the first part, determine the update value according to (9). It is possible to use the bits in a different order. In particular, the three parts consisting of $64 - p - t$, p , and t bits could be permuted. However, only if the $64 - p - t$ bits for the NLZ and the p bits for the address are adjacent and in this order, ELL will be reducible, as described later in Section 4.2.

Since ELL uses 64-bit hashes, update values computed according to (9) are limited by $(64 - p - t)2^t + (2^t - 1) + 1 = (65 - p - t)2^t$. For $p \geq 2$, all possible update values thus fit into $6 + t$ bits as $(65 - p - t)2^t \leq 63 \cdot 2^t \leq 2^{6+t} - 1$ holds in any case. Adapting distribution (8) to incorporate the limitation of update values to the range $[1, (65 - p - t)2^t]$ gives

$$\rho_{\text{update}}(k) = \frac{1}{2^{\phi(k)}} \quad \text{with } k \in [1, (65 - p - t)2^t], t \geq 0, (10)$$

Initial state:



Element insertion:

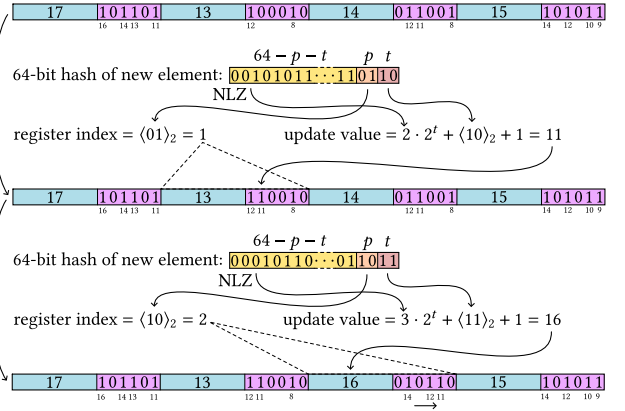


Figure 3: Two element insertions into an ExaLogLog sketch with parameters $p = 2, t = 2, d = 6$ which has $2^p = 4$ registers with a size of $6 + t + d = 14$ bits.

where we introduced the function

$$\phi(k) := \min(t + 1 + \lfloor (k - 1) / 2^t \rfloor, 64 - p). \quad (11)$$

2.4 Choice of Parameters

Under the assumption that the formulas for the MVP presented in Section 2.1 are still a good approximation after the exchange of the update value distribution, we can use them to search for parameters that lead to a good space efficiency. As our experimental results will show later, the predicted MVPs are indeed very accurate despite the deviation from a geometric distribution which eventually justifies the made assumption.

We evaluated the MVPs given in Section 2.1 with $q = 6 + t$ and $b = 2^{2^{-t}}$ for $t \in \{0, 1, 2, 3\}$ and $d \geq 0$. Figure 4 shows the MVP according to (3) for an unbiased and efficient estimator. Figure 5 shows the MVP given by (6) when using the martingale estimator that can only be used for non-distributed setups. Figures 6 and 7 show the corresponding MVPs when the registers are assumed to be optimally compressed as given by (5) and (7).

These figures allow identifying useful configurations. The optimal setting in Figure 4 is $t = 2$ and $d = 20$, resulting in a total register size of $q + d = 6 + t + d = 28$ bits and a theoretical MVP of 3.67 which is 43% less than that of HLL with 6-bit registers. Since two registers can be packed into exactly 7 bytes, register access is not too complicated.

An interesting configuration is also $t = 2$ and $d = 24$ despite the larger theoretical MVP with a value of 3.78. The registers with a size of 32 bits allow very fast register access when stored in a 32-bit integer array. The 32-bit register alignment makes this configuration even convenient for concurrent updates using compare-and-swap instructions available on modern central processing units (CPUs). Furthermore, as Figures 6 and 7 indicate, this configuration is probably more efficient than $d = 20$ or $d = 16$ when using compression.

When setting $t = 1$, the choice $d = 9$ with a MVP of 3.90 is also worth mentioning. Although less space-efficient than the mentioned configurations with $t = 2$, this setting also results in




Figure 4: The MVP according to (3) with $b = 2^{2^{-t}}$ and $q = 6 + t$ when storing the registers in a bit array and using an efficient unbiased estimator. Arrows indicate minima.




Figure 5: The MVP according to (6) with $b = 2^{2^{-t}}$ and $q = 6 + t$ when storing the registers in a bit array and using the martingale estimator. Arrows indicate minima.




Figure 6: The MVP according to (5) with $b = 2^{2^{-t}}$ when assuming optimal compression of the state and using an efficient unbiased estimator.




Figure 7: The MVP according to (7) with $b = 2^{2^{-t}}$ when assuming optimal compression of the state and using the martingale estimator.

byte-aligned registers, as their size is exactly 16 bits. Variants with $t \geq 3$ lead to larger MVPs than for $d = 2$. However, as they also lead to large d and therefore to quite large register sizes, we do not consider them to be useful in practice.

In the case of martingale estimation, which can be used for non-distributed setups, the optimum is achieved for $t = 2$ and $d = 16$ (cf. Figure 5), which leads to a MVP of 2.77, which is 33% below that of HLL with 6-bit registers. As the register size is 24 bits and therefore fits exactly into 3 bytes, register access is also relatively simple.

As the theoretical MVPs only depend on t and d , we will use the notation $\text{ELL}(t, d)$ to describe a specific class of ELL sketches with equal MVP. The third parameter, the precision parameter p , controls the trade-off between space and estimation error.

2.5 Relationship to Other Data Structures

In Section 2.3 we already mentioned that ELL with $t = 0$ is the same as the generalized data structure with $b = 2$ we have recently proposed [19]. As this data structure is in turn a generalization of HLL with $d = 0$, EHLL with $d = 1$, and ULL with $d = 2$ [19], these therefore correspond to $\text{ELL}(0, 0)$, $\text{ELL}(0, 1)$, and $\text{ELL}(0, 2)$, respectively. There is also a direct relationship to PCSA [21] and CPC [1, 26], as $\text{ELL}(0, 64)$ stores exactly the same information, albeit encoded differently. In addition, HyperMinHash [51] corresponds to $\text{ELL}(t, 0)$, whose registers only store the maxima of update values. HyperMinHash uses an update value distribution equivalent to (8) but defines the ordering of register and update values differently.

3 STATISTICAL INFERENCE

To simplify the statistical model, we use the common Poisson approximation [18, 20, 45] that the number of inserted distinct elements is not fixed, but follows a Poisson distribution with mean n . Consequently, since the updates are evenly distributed over all m registers, the number of updates with value k per register is again Poisson distributed with mean $\frac{n}{m}\rho_{\text{update}}(k)$ when using ρ_{update} defined in (10). The probability that a register was updated with value k at least once, denoted by event A_k , is therefore

$$\Pr(A_k) = 1 - e^{-\frac{n}{m}\rho_{\text{update}}(k)}. \quad (12)$$

The probability that $u \in [1, (65 - p - t)2^t]$ was the largest update value, which implies that there were no updates with values greater than u , is given by

$$\begin{aligned} \Pr(A_u \wedge \bigwedge_{k=u+1}^{(65-p-t)2^t} \bar{A}_k) &= \Pr(A_u) \prod_{k=u+1}^{(65-p-t)2^t} (1 - \Pr(A_k)) \\ &= \left(1 - e^{-\frac{n}{m}\rho_{\text{update}}(u)}\right) \exp\left(-\frac{n}{m} \sum_{k=u+1}^{(65-p-t)2^t} \rho_{\text{update}}(k)\right) \\ &= \left(1 - e^{-\frac{n}{m}\rho_{\text{update}}(u)}\right) e^{-\frac{n}{m}\omega(u)}, \end{aligned} \quad (13)$$

where ω is defined as (see Lemma B.1)

$$\omega(u) := \sum_{k=u+1}^{(65-p-t)2^t} \rho_{\text{update}}(k) = \frac{2^t(1-t+\phi(u))-u}{2\phi(u)}. \quad (14)$$

3.1 Probability Mass Function for Registers

As described in Section 2, a register stores the maximum update value u and also if there were updates with values in the range $[u - d, u - 1]$. Multiplying the probability (13) that u was the maximum update value with $\Pr(A_k)$ or $(1 - \Pr(A_k))$ (compare (12)) for all $k \in [u - d, u - 1]$ dependent on whether an update

Algorithm 3: Computation of the coefficients of log-likelihood function (15) for an ExaLogLog sketch with registers r_0, r_1, \dots, r_{m-1} .

```

 $\alpha' \leftarrow 0$  ▷  $\alpha' = \alpha \cdot 2^{64-p}$  is an integer
 $(\beta_{t+1}, \beta_{t+2}, \dots, \beta_{64-p}) \leftarrow (0, 0, \dots, 0)$ 
for  $i \leftarrow 0$  to  $m-1$  do ▷ iterate over all  $m$  registers
   $u \leftarrow \lfloor r_i / 2^d \rfloor$  ▷ get max. update value from  $r_i$ , right-shift by  $d$  bits
   $\alpha' \leftarrow \alpha' + 2^{64-p} \omega(u)$  ▷ see (14) for a definition of function  $\omega$ 
  if  $u \geq 1$  then
     $j \leftarrow \phi(u)$  ▷ see (11) for a definition of function  $\phi$ 
     $\beta_j \leftarrow \beta_j + 1$ 
    if  $u \geq 2$  then
      for  $k \leftarrow \max(1, u-d)$  to  $u-1$  do
         $j \leftarrow \phi(k)$ 
        if  $(r_i \& 2^{d-u+k}) = 0$  then ▷  $\&$  denotes bitwise AND
           $\alpha' \leftarrow \alpha' + 2^{64-p-j}$ 
        else
           $\beta_j \leftarrow \beta_j + 1$ 
   $\alpha \leftarrow \alpha' / 2^{64-p}$ 

```

with value k has occurred or not, gives the probability mass function (PMF) for a single register r :

Case $r = 0$:

$$\rho_{\text{reg}}(r|n) = e^{-\frac{n}{m}\omega(0)} = e^{-\frac{n}{m}}.$$

Case $r = u2^d + \langle l_1 \dots l_{u-1} \rangle_2 2^{d+1-u}$ with $1 \leq u \leq d$:

$$\rho_{\text{reg}}(r|n) = (1 - e^{-\frac{n}{m}\rho_{\text{update}}(u)})e^{-\frac{n}{m}\omega(u)} \cdot \prod_{j=1}^{u-1} (1 - e^{-\frac{n}{m}\rho_{\text{update}}(u-j)})^{l_j} (e^{-\frac{n}{m}\rho_{\text{update}}(u-j)})^{1-l_j}.$$

Case $r = u2^d + \langle l_1 \dots l_d \rangle_2$ with $d+1 \leq u \leq (65-p-t)2^t$:

$$\rho_{\text{reg}}(r|n) = (1 - e^{-\frac{n}{m}\rho_{\text{update}}(u)})e^{-\frac{n}{m}\omega(u)} \cdot \prod_{j=1}^d (1 - e^{-\frac{n}{m}\rho_{\text{update}}(u-j)})^{l_j} (e^{-\frac{n}{m}\rho_{\text{update}}(u-j)})^{1-l_j}.$$

3.2 Maximum-Likelihood Estimation

Since the registers are independent due to the Poisson approximation, the log-likelihood function for register states r_0, r_1, \dots, r_{m-1} can be written as

$$\ln \mathcal{L} = \ln \mathcal{L}(n|r_0, \dots, r_{m-1}) = \sum_{i=0}^{m-1} \ln \rho_{\text{reg}}(r_i|n).$$

According to (10), ρ_{update} is always a power of two from the set $\{\frac{1}{2^{t+1}}, \frac{1}{2^{t+2}}, \dots, \frac{1}{2^{64-p}}\}$. Therefore, the log-likelihood function always has the shape

$$\ln \mathcal{L} = -\frac{n}{m}\alpha + \sum_{u=t+1}^{64-p} \beta_u \ln(1 - e^{-\frac{n}{m2^u}}), \quad (15)$$

where the sum has at most $64-p-t$ nonzero terms. The coefficients α and β_u depend just on the register states r_0, r_1, \dots, r_{m-1} and can be computed according to Algorithm 3. Since j is the result of $\phi(k)$ (10), it is bounded by $64-p$ and all contributions to α are integer multiples of $\frac{1}{2^{64-p}}$. Therefore, the summation can be performed with integer arithmetic only, if $\alpha' = \alpha \cdot 2^{64-p}$ is computed instead.

In the following we consider the case where α and at least one of the coefficients β_u are positive. β_u all zero requires all registers to be in the initial state leading to a ML estimate of zero. $\alpha = 0$ can only occur, if all registers are saturated. In this case, which only occurs with a noteworthy probability for distinct counts that are entirely unrealistic, the ML estimate would be infinite. When introducing

$$x := \exp\left(\frac{n}{m2^{u_{\max}}}\right) - 1, \quad (16)$$

$$u_{\min} := \min_{t+1 \leq u \leq 64-p} \{u \mid \beta_u > 0\},$$

$$u_{\max} := \max_{t+1 \leq u \leq 64-p} \{u \mid \beta_u > 0\},$$

$$\varphi(x) := \beta_{u_{\max}} + \sum_{j=1}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} \prod_{l=0}^{j-1} \frac{2}{(1+x)^{2^l+1}}, \quad (17)$$

the ML equation can be equivalently written as $f(x) = 0$ where f is defined as

$$f(x) := -x(1+x) \frac{\partial}{\partial x} \ln \mathcal{L} = \alpha 2^{u_{\max}} x - \sum_{j=0}^{u_{\max}-u_{\min}} \frac{\beta_{u_{\max}-j} 2^j x}{(1+x)^{2^j-1}} = \alpha 2^{u_{\max}} x - \varphi(x). \quad (18)$$

This function is strictly increasing and concave for $x \geq 0$ as shown in Lemma B.2. As a result, it has a well-defined root, because $f(0) = -\sum_{j=u_{\min}}^{u_{\max}} \beta_j < 0$ and $f(\infty) \rightarrow \infty > 0$. If \hat{x} denotes the root of f , hence $f(\hat{x}) = 0$, the ML estimate \hat{n}_{ML} is given according to (16) by

$$\hat{n}_{\text{ML}} = m 2^{u_{\max}} \ln(1 + \hat{x}). \quad (19)$$

As an optional last step, the ML estimate can be corrected according to (4) with $b = 2^{2^{-t}}$ to reduce the bias.

The evaluation of f is cheap, because the number of terms is limited by $u_{\max} - u_{\min} + 1 \leq 64 - p - t$. Moreover, as all occurring exponents are powers of two, they can be computed recursively by squaring using simple multiplications. As the denominator in (17) is always greater than or equal to 2, the evaluation is also numerically safe. To reduce the numerical error, we can replace the denominator by $2 + y_l$ with

$$y_l := (1+x)^{2^l} - 1 \geq 0 \quad (20)$$

and use the recursion

$$y_{l+1} = y_l \cdot (2 + y_l). \quad (21)$$

The product appearing in (17)

$$\lambda_j := \prod_{l=0}^{j-1} \frac{2}{(1+x)^{2^l+1}}$$

can then be also computed recursively following

$$\lambda_{j+1} = \lambda_j \frac{2}{2+y_j}. \quad (22)$$

The simplicity of the ML equation is a result of distribution (8). For comparison, geometrically distributed update values following (2) would have led to significantly more terms. In addition, the computations of the resulting power expressions with real exponents would have been much more expensive. Appendix A describes a robust and fast-converging algorithm for finding the root of the ML equation.

3.3 Martingale Estimation

A simple, efficient, and unbiased way to estimate the distinct count is to start from zero and increment the estimate, whenever the state of the sketch is modified, according to the inverse of the probability that such a modification occurs with the insertion of any unseen element. This online approach is known as martingale or historic inverse probability (HIP) estimation [12, 41] and even leads to smaller estimation errors as already mentioned in Sections 2.1 and 2.4. However, martingale estimation is limited to cases where the data is not distributed and merging of sketches is not needed.

In addition to the estimate, the martingale estimator also keeps track of the current state change probability μ . Initially, $\mu = 1$ as the first update will certainly change the state. Whenever a register is modified, the probability of state changes for further

Algorithm 4: Incrementally updates the martingale estimate $\hat{n}_{\text{martingale}}$ and the state change probability μ whenever a register is altered from r to r' ($r < r'$). Initially, $\hat{n}_{\text{martingale}} = 0$ and $\mu = 1$.

```

 $\hat{n}_{\text{martingale}} \leftarrow \hat{n}_{\text{martingale}} + \frac{1}{\mu}$  ▷ update estimate
 $\mu \leftarrow \mu - (h(r) - h(r'))$  ▷ update state change probability,
 $h(r) > h(r')$ , compare (23)

```

Algorithm 5: Merges two corresponding registers r and r' of Exa-LogLog sketches with identical parameters t , d , and p .

```

function MergeRegister( $r, r', d$ )
   $u \leftarrow \lfloor r/2^d \rfloor$  ▷ get max. update value from  $r$ , right-shift by  $d$  bits
   $u' \leftarrow \lfloor r'/2^d \rfloor$  ▷ get max. update value from  $r'$ , right-shift by  $d$  bits
  if  $u > u' \wedge u' > 0$  then
    return  $r \mid \lfloor (2^d + (r' \bmod 2^d))/2^{u-u'} \rfloor$  ▷  $\mid$  denotes bitwise OR
  else if  $u' > u \wedge u > 0$  then
    return  $r' \mid \lfloor (2^d + (r \bmod 2^d))/2^{u'-u} \rfloor$ 
  else
    return  $r \mid r'$ 

```

elements decreases. The probability, that a new unseen element changes the ELL state, is given by

$$\mu(r_0, \dots, r_{m-1}) = \sum_{i=0}^{m-1} h(r_i). \quad (23)$$

$h(r_i)$ is the probability that register r_i is changed with the next new element. The function h is strictly monotonically decreasing and defined for a register value $r = u2^d + \langle l_1 \dots l_d \rangle_2$ as

$$h(r) = \frac{1}{m} (\omega(u) + \sum_{k=\max(1, u-d)}^{u-1} (1 - l_{u-k}) \cdot \rho_{\text{update}}(k))$$

when using ω from (14) and ρ_{update} from (10).

The martingale estimator is incremented with every state change by $\frac{1}{\mu}$ prior the update as demonstrated by Algorithm 4. μ itself can also be incrementally adjusted, such that the whole update takes constant time. The martingale estimator is unbiased and optimal, if mergeability is not needed [34].

4 PRACTICAL IMPLEMENTATION

Like other probabilistic data structures such as HLL, ELL also relies on high-quality 64-bit hash values for the elements. Known good hash functions are WyHash [50], Komihash [44], or PolymurHash [32]. Insert operations according to Algorithm 2 obviously take constant time and are very fast, because all statements can be translated into inexpensive CPU instructions. Expressions of kind $\lfloor x/2^y \rfloor$ can be realized by a right-shift operation by y bits, and $x \bmod 2^y$ is the same as masking the lower y bits. Furthermore, specific instructions exist on modern CPUs to obtain the number of leading zeros (NLZ).

4.1 Mergeability

If two ELL data structures are equally configured, thus they have equal p , t , and d values, they can be easily merged by pairwise merging of individual registers. A register stores the maximum update value u in its upper $6 + t$ and the occurrences of update values in the range $[u - d, u - 1]$ in its lower d bits (compare Figure 3). Since the merged state is the result of the union of all updates, the merged register must finally store the common maximum update value, and the d bits must indicate the combined occurrences of the next d smaller update values relative to this common maximum. Algorithm 5 demonstrates how efficient register merging can be realized using bitwise operations. The result of the merge operation is equivalent to inserting the union of all individual original elements, previously inserted into one

Algorithm 6: Reduces an ExaLogLog (ELL) sketch with registers r_0, r_1, \dots, r_{m-1} and parameters t, d, p to an ELL sketch with registers $r'_0, r'_1, \dots, r'_{m'-1}$ and parameters t, d', p' , where $m = 2^p$, $m' = 2^{p'}$, $d \geq d' \geq 0$, and $p \geq p' \geq 0$.

```

 $a \leftarrow (64 - t - p) \cdot 2^t + 1$ 
for  $i \leftarrow 0$  to  $m' - 1$  do
   $r' \leftarrow 0$ 
  for  $j \leftarrow 0$  to  $2^{p-p'} - 1$  do
     $r \leftarrow \lfloor r_{i+j \cdot m'} / 2^{d-d'} \rfloor$  ▷ right-shift by  $d - d'$  bits
     $u \leftarrow \lfloor r/2^{d'} \rfloor$  ▷ right-shift by  $d'$  bits
    if  $u \geq a$  then ▷ satisfied if NLZ was  $64 - t - p$  in (9) for  $u$ 
      ▷  $r$  must be adapted, if  $u$  was different at precision  $p'$ 
       $s \leftarrow ((p - p') - (64 - \text{nlz}(j))) \cdot 2^t$  ▷ assuming  $j$  has 64 bits
      ▷  $\text{nlz}(j) \in [0, 64]$ 
      if  $s > 0$  then
         $v \leftarrow d' + a - u$ 
        if  $v > 0$  then  $r \leftarrow \lfloor r/2^v \rfloor \cdot 2^v + \lfloor (r \bmod 2^v)/2^s \rfloor$ 
         $r \leftarrow r + s \cdot 2^{d'}$ 
     $r' \leftarrow \text{MergeRegister}(r, r', d')$  ▷ see Algorithm 5
   $r'_i \leftarrow r'$ 

```

of the two data structures, directly into an empty data structure using Algorithm 2.

ELL sketches are also mergeable if not all the parameters are equal as long as the sketches share the same t -parameter. If the sketch parameters are (t, d, p) and (t, d', p') , respectively, they both can be reduced to an ELL sketch with parameters $(t, \min(d, d'), \min(p, p'))$ first. This is useful for migration scenarios, if the precision p or parameter d must be changed while mergeability with older records should be preserved.

4.2 Reducibility

The reduction of the d -parameter is straightforward. Decrementing it from d to d' with $d \geq d'$ only requires right-shifting all registers by $d - d'$ bits. The reduction of the precision from p to p' is more complex, as $2^{p-p'}$ registers need to be combined into one. However, due to the way in which the hash bits are consumed in Algorithm 2, this is also possible in a lossless way, meaning that the result is the same as if all elements would be recorded directly using a sketch with the reduced parameters, as demonstrated by Algorithm 6.

4.3 Sparse Mode

ExaLogLog uses a fixed array of registers, which guarantees constant-time insertions. However, if space efficiency is more important, allocating that array from the beginning does not make sense if keeping the raw input data takes less space. Therefore, many data sketches start with a *sparse* mode with a linearly scaling memory footprint and only switch to the *dense* representation at the break-even point.

For ELL, a sparse representation could be realized by just storing the 64-bit input hash values in a list. To save space, we can reduce those hash values to $(v + 6)$ -bit values, which we call hash tokens, by keeping only information needed for insertions into ELL sketches with $p + t \leq v$. A hash token stores the least significant v bits of the original hash value and, in addition, the NLZ of the remaining $(64 - v)$ most significant bits of the hash value. If $v \geq 1$ the NLZ fits into 6 bits and a 64-bit hash value $\langle h_{63}h_{62} \dots h_0 \rangle_2$ can be mapped to a $(v + 6)$ -bit hash token w according to

$$w = \langle h_{v-1} \dots h_0 000000 \rangle_2 + \text{nlz}(\langle h_{63}h_{62} \dots h_v \underline{11} \dots 1 \rangle_2).$$

Algorithm 7: Computation of the coefficients for the log-likelihood function (26) from a set T of distinct $(v + 6)$ -bit hash tokens.

```

 $\alpha' \leftarrow 2^{64}$       ▶ start from 0 when using an unsigned 64-bit integer
 $(\beta_{v+1}, \beta_{t+2}, \dots, \beta_{64}) \leftarrow (0, 0, \dots, 0)$ 
for  $w \in T$  do      ▶ iterate over all collected distinct tokens
     $j \leftarrow \min(v + 1 + (w \bmod 64), 64)$ 
     $\beta_j \leftarrow \beta_j + 1$ 
     $\alpha' \leftarrow \alpha' - 2^{64-j}$ 
 $\alpha \leftarrow \alpha' / 2^{64};$ 

```

While in sparse mode, it is sufficient to keep only distinct hash tokens. When switching to dense mode, the hash tokens can be transformed back to representative 64-bit hash values following

$$\langle h'_{63} h'_{62} \dots h'_0 \rangle_2 = 2^{64 - \langle s_5 s_4 s_3 s_2 s_1 s_0 \rangle_2} - 2^v + \langle s_{v+5} s_{v+4} \dots s_6 \rangle_2$$

where $\langle s_{v+5} s_{v+4} \dots s_0 \rangle_2$ is the binary representation of the token. The reconstructed hash values can be equivalently used for the insertion using Algorithm 2 as the original hash value.

It is also possible, to estimate the distinct count directly from a given set of distinct hash tokens T . Since the first v bits are uniformly distributed and the NLZ, stored in the remaining 6 bits, are distributed according to a truncated geometric distribution with maximum value $64 - v$, the probability mass function (PMF) of hash tokens is given by

$$\rho_{\text{token}}(w) = \begin{cases} \frac{1}{2^{\min(v+1+(w \bmod 64), 64)}} & w \bmod 64 \leq 64 - v, \\ 0 & \text{else,} \end{cases} \quad (24)$$

with $w \in [0, 2^{v+6})$ and $v \geq 1$. As for any PMF, summing up the probabilities for all possible values yields 1

$$\sum_{w=0}^{2^{v+6}-1} \rho_{\text{token}}(w) = 1. \quad (25)$$

As in Section 3, we use again the Poisson approximation, which allows to write the probability that some hash token w is in the set of collected hash tokens T as $\Pr(w \in T) = 1 - e^{-n\rho_{\text{token}}(w)}$. Therefore, the log-likelihood function is

$$\begin{aligned} \ln \mathcal{L} &= \ln \mathcal{L}(n|T) \\ &= \sum_{w \notin T} \ln(e^{-n\rho_{\text{token}}(w)}) + \sum_{w \in T} \ln(1 - e^{-n\rho_{\text{token}}(w)}) \\ &= -n \sum_{w \notin T} \rho_{\text{token}}(w) + \sum_{w \in T} \ln(1 - e^{-n\rho_{\text{token}}(w)}) \\ &= -n(1 - \sum_{w \in T} \rho_{\text{token}}(w)) + \sum_{w \in T} \ln(1 - e^{-n\rho_{\text{token}}(w)}) \\ &= -n\alpha + \sum_{u=v+1}^{64} \beta_u \ln(1 - e^{-\frac{n}{2^u}}) \end{aligned} \quad (26)$$

where we used (25). The coefficients α and β_u can be obtained using Algorithm 7. The log-likelihood function (26) has the same shape as that for the ELL registers (15) when setting $m = 1 \Leftrightarrow p = 0$ and $t = v$. Therefore, the ML estimate can be found again with the same root-finding algorithm described in Appendix A.

5 EXPERIMENTS

We provide a Java reference implementation of ExaLogLog (ELL) together with instructions and source code to reproduce all the presented results and figures at <https://github.com/dynatrace-research/exaloglog-paper>. The repository also includes numerous unit tests that cover 100% of the code. In particular, merging of ELL sketches based on Algorithm 5 was tested by creating many pairs of random ELL sketches for which we compared the merged ELL sketch with a sketch into which the unified stream of elements was inserted. Similarly, Algorithm 6 was tested by inserting the identical elements into two ELL sketches with different configurations and checking whether the state was the same after reduction to the same parameters.

5.1 Estimation Error

We need to verify that the estimation error of ELL really matches the theoretically predicted error, despite modeling (8) by (2) and the number of distinct element insertions by a Poisson distribution. According to (1) we would expect a root-mean-square error (RMSE) error of $\sqrt{\text{MVP}/((q+d)m)}$ with $q = 6 + t$, $b = 2^{2-t}$, and the MVP either given by (3) or (6), depending on whether the ML or the martingale estimator is used.

An accurate evaluation of the error requires thousands of estimates of different data sets with identical true distinct count n . It is infeasible to use real data sets, if we want to repeat that for many different and also large n . Therefore, we use a more efficient approach [19]. Extensive empirical tests [43] have shown that the output of modern hash functions such as WyHash [50], Komihash [44], or PolymurHash [32] can be considered like uniform random values. Otherwise, field-tested probabilistic data structures like HLL would not work. This fact allows us to perform the experiments without real or artificially generated data.

Insertion of a new element can be simulated by simply generating a 64-bit random value to be used directly as the hash value of the inserted element in Algorithm 2. Duplicate insertions of the same element can be ignored as they cannot change the state of ELL by definition. Processing a random data set with true distinct count n is thus equivalent to using n random values instead. Accidental collisions of random values can be ignored because they occur with the same probability as hash collisions for real data. To simulate the estimation error for a predefined distinct count value n , the estimate is computed after updating the ELL sketch using Algorithm 2 with n random values and finally compared against the true distinct count n . By repeating this process with many different random sequences, in our experiments 100 000, the bias and the RMSE can be empirically determined.

As this approach becomes computationally infeasible for distinct counts beyond 1 million, we switch to a different strategy [19]. After the first million insertions, for which a random value was generated each time, we just generate the waiting time (the number of distinct count increments) until a register is processed with a certain update value the next time. The probability that a register is updated with any possible update value $k \in [1, (65 - p - t)2^t]$ is given by $\rho_{\text{update}}(k)/m$ with $\rho_{\text{update}}(k)$ from (10). Therefore, the number of distinct count increments until a register is updated with a specific value k the next time is geometrically distributed with corresponding success probability. In this way, we determine the next update time for each register and for each possible update value. Since the same update value can only modify a register once, we do not need to consider further updates which might occur with the same value for the same register. Knowing these $m \times ((65 - p - t)2^t)$ distinct counts leading to possible state changes in advance, enables us to make large distinct count increments, resulting in a huge speedup. This eventually allowed us to simulate the estimation error for distinct counts up to values of 10^{21} and also to test the presented estimators up to the exa-scale.

Figure 8 shows the empirical relative bias and RMSE together with the theoretical RMSE given by $\sqrt{\text{MVP}/((q+d)m)}$ for the ML and the martingale estimator for configurations $(t, d) \in \{(1, 9), (2, 16), (2, 20), (2, 24)\}$ and precisions $p \in \{4, 6, 8, 10\}$. For intermediate distinct counts, perfect agreement with theory is observed. For small distinct counts the estimation error is even much smaller. Interestingly, the estimation error also decreases slightly at the end of the operating range, which is in the order

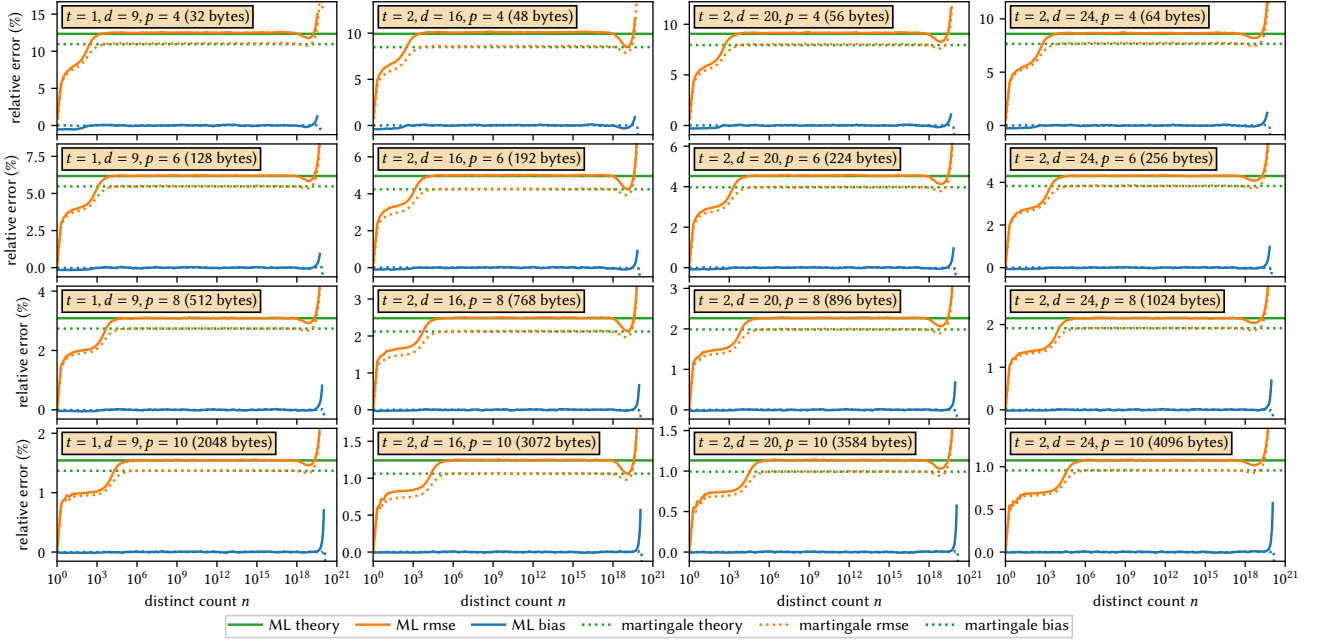


Figure 8: The relative bias and the RMSE for the ML and the martingale estimator for different ELL configurations obtained from 100 000 simulation runs. The theoretically predicted errors perfectly match the experimental results. Individual insertions were simulated up to a distinct count of 10^6 before switching to the fast simulation strategy.

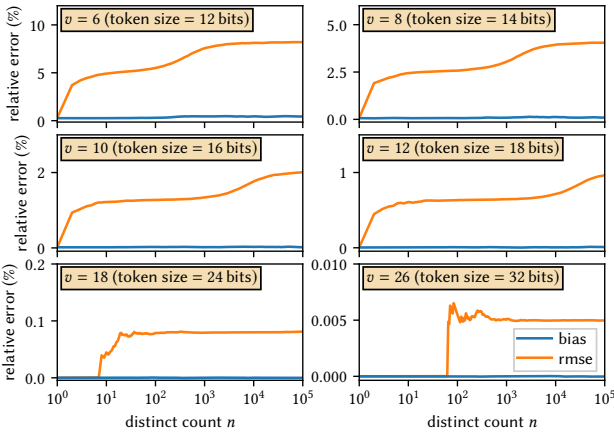


Figure 9: The relative bias and the RMSE when estimating the distinct count from a set of collected distinct hash tokens with different sizes.

of $2^{64} \approx 1.9 \times 10^{19}$ and thus lies in the exa-scale range. The estimators are essentially unbiased. The tiny bias which appears for small p for the ML estimator can be ignored in practice as it is much smaller than the theoretical RMSE.

We also verified estimation from sets of hash tokens as proposed in Section 4.3. Again, we performed 100 000 simulation runs, where we simulated 64-bit hash values by taking 64-bit random values and transforming them into corresponding hash tokens using different parameters $v \in \{6, 8, 10, 12, 18, 26\}$. We considered distinct counts up to 10^5 , which is typically far beyond the break-even point where a transition to the dense representation takes place. As shown in Figure 9, the estimates are unbiased, and the relative estimation error is slightly smaller than the estimation error of an ELL sketch for which $p + t = v$. The reason is that

the set of hash tokens contains information that is equivalent to an ELL sketch with $d \rightarrow \infty$. From a practical perspective, a hash token size of 4 bytes ($v = 26$) is particularly interesting, because it is big enough to support any practical ELL configurations. Furthermore, as the tokens can be stored in a plain 32-bit integer array, off-the-shelf sorting algorithms can be used for deduplication.

5.2 Space Efficiency Comparison

Our experiments have shown that the estimation error matches the theoretical predicted estimation error. Since the space requirement of ELL is constant $(q + d)m$ bits, the theoretically predicted MVP, as discussed in Section 2.4, can be achieved if memory overhead for the Java object or auxiliary fields can be ignored. However, for a fairer comparison with other practical algorithms, we considered the empirical MVP based on the total space allocated by the whole data structure.

We performed 1 million simulation runs. In each cycle, the distinct count was estimated and the allocated amount of memory as well as the serialization size were measured after adding up to 10^6 distinct random elements. This allowed us to compute the RMSE and, together with the average space requirements, the empirical MVPs according to (1).

Table 2 compares our ELL reference implementation to other state-of-the-art algorithms. All algorithms were configured to give roughly 2% estimation error. Since the reference implementation of SpikeSketch [15] is not very space-efficient and also does not support serialization, we used the size of the plain register array without any additional overheads as lower bound size estimates. For algorithms that allocate variable space, such as HLL with 4-bit registers, CPC, or HLLL, the standard deviation of the size is also shown.

The serialization size is always smaller than the in-memory size to which object overhead or auxiliary fields such as buffers

Table 2: Comparison of mergeable approximate distinct counting algorithms when estimating the distinct count after inserting $n = 10^6$ distinct elements. The parameters were chosen to obtain roughly 2 % root-mean-square error (RMSE). The actual RMSE was empirically determined from 1 million simulation runs. The memory-variance product (MVP) is estimated as $MVP = (\text{average memory/serialization size in bits}) \times (\text{RMSE})^2$ and is a fair measure of the space efficiency. The table is sorted by the in-memory MVP.

Algorithm	References	Source code (https://github.com/...)	RMSE	Size in bytes		MVP (space efficiency)		Constant-time insert operation
				memory	serialized	memory	serialized	
HyperLogLog (HLL, 8-bit registers, $p = 11$)	[1]	apache/datasketches-java	2.29 %	2296	2088	9.66	8.78	✓
HyperLogLog (HLL, 6-bit registers, $p = 11$)	[1, 23]	apache/datasketches-java	2.29 %	1792	1577	7.54	6.63	✓
HyperLogLog (HLL, ML estimator, $p = 11$)	[18]	dynatrace-oss/hash4j	2.29 %	1576	1536	6.63	6.46	✓
HyperLogLog (HLL, 4-bit registers, $p = 11$)	[1]	apache/datasketches-java	2.29 %	1331±56	1067±4	5.60	4.49	-
Compressed probability counting (CPC, $p = 10$)	[1, 26]	apache/datasketches-java	2.16 %	1416±34	656±11	5.30	2.46*	-
UltraLogLog (ULL, ML estimator, $p = 10$)	[19]	dynatrace-oss/hash4j	2.38 %	1056	1024	4.78	4.64	✓
HyperLogLog (HLLL, $p = 11$)	[25]	mkarppa/hyperloglog	2.30 %	1100±13	898±16	4.64	3.79	-
SpikeSketch (128 buckets)	[15]	duyang92/SpikeSketch	2.26%**	≥ 1024***	≥ 1024***	≥ 4.19***	≥ 4.19***	✓
ExaLogLog (ELL, $t = 2, d = 24, p = 8$)	this work	dynatrace-research/exaloglog-paper	2.15 %	1064	1024	3.93	3.79	✓
ExaLogLog (ELL, $t = 2, d = 20, p = 8$)	this work	dynatrace-research/exaloglog-paper	2.27 %	936	896	3.86	3.69	✓
Conjectured lower bound	[33]	-	-	-	-	1.98	1.98	not known

* achieved by expensive compression during serialization

** error can be much larger for smaller distinct counts

*** lower bound values based on theoretical considerations completely ignoring auxiliary data fields (empirical values are meaningless as the reference implementation is not optimized)

also contribute. The difference is particularly large for CPC, whose serialization method also applies a specialized and relatively expensive (see Section 5.3) compression step [26]. A fair comparison would require the development of specific compression techniques for all other data structures which is out of the scope of this work. However, the theoretical MVPs for optimal compression shown in Figure 6 indicate that the size of ELL could be further reduced. For ULL, which is a special case of ELL, we have already shown that MVPs below 3 can be achieved. Its 1-byte register array seems to be very convenient for standard compression algorithms [19].

Figure 10 also shows the memory consumption and the corresponding MVP for other distinct counts. ELL requires constant space and never allocates additional data. The data structures from the Apache DataSketches library have implemented a sparse mode that allows them to be more space-efficient for small distinct counts. However, a sparse mode could also be easily implemented for ELL as discussed in Section 4.3.

The MVP of SpikeSketch stands out at lower distinct counts. The high values are a result of the lossy compression and stepwise smoothing. The latter reduces the update probability even of empty SpikeSketches by a factor of 64 %. As a consequence, the estimation error is 100 % with a 36 % probability independent of the number of buckets for the extreme case of $n = 1$. We, therefore, do not consider SpikeSketch to be suitable for practical use. HLLL also shows a spike around $n = 5 \times 10^3$, which is a result of using the original HLL estimator [20] that is known to have some issues [18, 23].

5.3 Performance Comparison

To compare the performance of ELL with the other algorithms listed in Table 2, we used an Amazon EC2 c5.metal instance running Ubuntu Server 24.04 LTS. Turbo Boost was disabled by setting the processor P-state to 1 [3]. We used Java implementations for all algorithms except for HLLL and SpikeSketch whose reference implementations are written in C++.

Figure 11 shows the results of our benchmarks. First, element insertion was tested by adding up to 10^6 random 16-byte arrays that were generated and stored in memory in advance. As Apache DataSketches uses the 128-bit version of Murmur3 as built-in hash function without the flexibility of defining a different one, we used it also for all other algorithms to make a fair comparison.

The corresponding graph shows the average time per inserted element, which also includes the initial allocation of the data structure. For this reason, the measured times for small n tend to be higher. All insertion times, except those for HLLL and SpikeSketch are between 20 and 50 ns. For ELL, we investigated insertion with and without martingale estimator.

The fastest estimation times are achieved by algorithms, including those from Apache DataSketches, that maintain a martingale estimator or keep track of other redundant statistics during insertion. The ELL ML estimator compares well to those algorithms that insert elements without such additional bookkeeping.

To analyze serialization, we measured the time to write the state into a newly allocated byte array. For ELL, this means just copying the byte array holding the register values, which is very fast. The results also show that the serialization of CPC is more than an order of magnitude slower due to the expensive compression, as discussed before. We do not have data for HLLL and SpikeSketch as their reference implementations come without a serialization method.

Finally, we also investigated the time to merge two data structures, both filled with n random elements. The results show that ELL is very fast. Algorithms that have implemented a sparse mode are, as expected, faster for small n as less data has to be processed. The comparison with the algorithms from Apache DataSketches is not entirely fair, as they rebuild internal statistics for estimation during the merging process. For this reason, we have also considered merging followed by estimation, which is a common operation sequence in practice. ELL also performs quite well in this case. We have no data for SpikeSketch as its reference implementation does not include a working merge operation.

It is worth noting that our ELL reference implementation is generic and supports arbitrary values of t and d . Hardcoding these values could potentially further improve its performance.

6 FUTURE WORK

A topic for future research is the compressibility of ELL. According to Figures 6 and 7, much lower MVPs could be achieved with optimal compression. For ULL, a special case of ELL with $t = 0$ and $d = 2$ (compare Section 2.5), a reduction close to the theoretical limit can be achieved with standard compression algorithms [19] such as Zstandard [13]. Unlike ULL, whose register size is exactly one byte, we assume that standard algorithms will work




Figure 10: The average memory footprint and the empirical MVP for $n \in \{10, 20, 50, 100, 200, 500, \dots, 10^6\}$ obtained from 1 million simulation runs. The legend is given in Figure 11.

worse in the general case. Since the shape of the register distribution is known (see Section 3.1), some sort of entropy coding could be a way to approach the theoretical limit.

As discussed in Section 2.5, HyperMinHash is a special case of ELL, and PCSA and CPC contain the same information as an ELL(0, 64) sketch. Therefore, our proposed ML estimation approach, in which the ML equation is first reduced to the simple form (15) with a relatively small number of terms, should also work for them. We assume that this method could lead to slightly lower estimation errors than current approaches, as ML estimation is generally known to be asymptotically efficient.

7 CONCLUSION

We have introduced a new algorithm for distinct counting called ExaLogLog (ELL), which includes already-known algorithms as special cases. With the right parameters, the space efficiency can be improved significantly. In particular, a configuration was presented that reduces the MVP by 43% compared to the widely used HyperLogLog algorithm. ELL also supports practical properties such as mergeability, idempotency, reproducibility, and reducibility. We have also shown that maximum likelihood (ML) estimation using the developed robust numerical solver and martingale estimation are feasible. The observed estimation errors are fully consistent with the theory. In contrast to other recent approaches, insertions always require constant time, regardless of the chosen accuracy. We also proposed a sparse mode based on hash tokens, which allows the allocation of the register array to be postponed. All this makes ELL very attractive for wider use in practice.

A NUMERICAL ROOT-FINDING

As the function f defined in (18) is strictly increasing and concave (see Lemma B.2), the root of the ML equation $f(x) = 0$ can be robustly found using Newton’s method when starting from a point x_0 satisfying $f(x_0) \leq 0$. According to Lemma B.3 such a point is

$$x_0 = \exp\left(\ln\left(1 + \frac{\sigma_1}{\alpha 2^{u_{\max}}}\right) \frac{\sigma_0}{\sigma_1}\right) - 1$$

with $\sigma_0 := \sum_{j=u_{\min}}^{u_{\max}} \beta_j$ and $\sigma_1 := \sum_{j=u_{\min}}^{u_{\max}} \beta_j 2^{u_{\max}-j}$. (27)

As a result, the sequence of points obtained by Newton’s method

$$x_{t+1} = x_t - \frac{f(x_t)}{f'(x_t)} = x_t \left(1 - \frac{f(x_t)}{x_t f'(x_t)}\right)$$

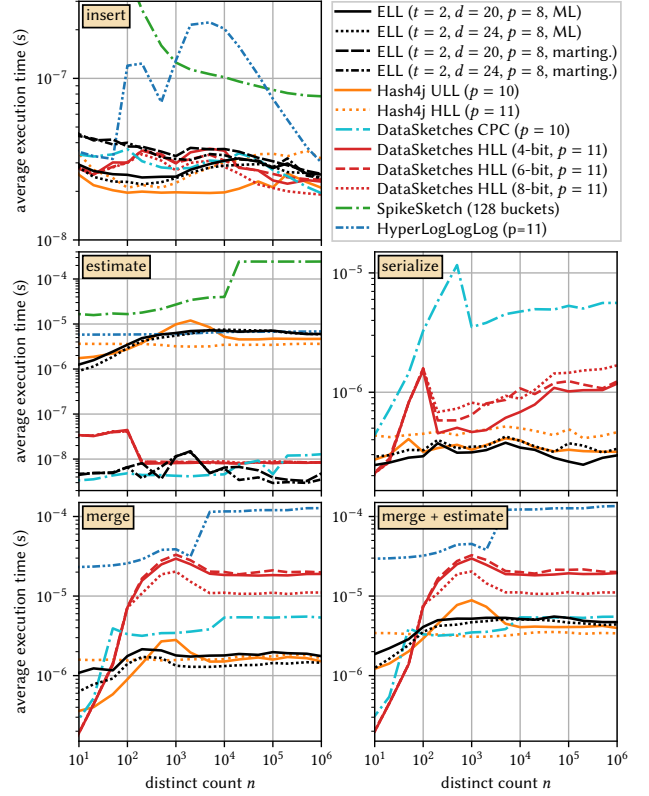


Figure 11: The average execution time for insert, estimate, serialize, merge, and combined merge and estimate operations for $n \in \{10, 20, 50, 100, 200, 500, \dots, 10^6\}$.

is always increasing and approaches \hat{x} . The recursion can be transformed using

$$\begin{aligned} x f'(x) - \alpha 2^{u_{\max}} x &= -x \frac{\partial}{\partial x} \sum_{j=1}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} \frac{2^j x}{(1+x)^{2^j-1}} \\ &= \sum_{j=1}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} \left(\frac{4^j x^2 (1+x)^{2^j-1}}{((1+x)^{2^j-1})^2} - \frac{2^j x}{(1+x)^{2^j-1}} \right) \\ &= \sum_{j=1}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} \frac{2^j x}{(1+x)^{2^j-1}} \left(\frac{2^j x (1+x)^{2^j-1}}{(1+x)^{2^j-1}} - 1 \right) = \psi(x) \end{aligned}$$

with

$$\psi(x) := \sum_{j=1}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} \left(\prod_{l=0}^{j-1} \frac{2}{(1+x)^{2^l+1}} \right) \cdot \left(\left(\prod_{l=0}^{j-1} \frac{2(1+x)^{2^l}}{(1+x)^{2^l+1}} \right) - 1 \right) \quad (28)$$

into

$$x_{t+1} = x_t \left(1 + \frac{\varphi(x_t) - \alpha 2^{u_{\max}} x_t}{\psi(x_t) + \alpha 2^{u_{\max}} x_t}\right). \quad (29)$$

Function ψ shares the same product as φ (17), which therefore needs to be computed only once. The last factor in (28) is always nonnegative as the product consists of factors that are all greater than or equal to 1, because $1 \leq \frac{2(1+x)^{2^l}}{(1+x)^{2^l+1}} < 2$. Furthermore, when using standard floating-point types, overflows will not occur as the product has at most $u_{\max} - u_{\min} \leq 63 - p - t$ factors and all of them are smaller than 2. To reduce numerical errors and minimize the number of operations, we compute

$$\eta_j := \left(\prod_{l=0}^{j-1} \frac{2(1+x)^{2^l}}{(1+x)^{2^l+1}} \right) - 1$$

Algorithm 8: Numerical computation of the distinct count estimate by solving the ML equation using Newton's method. α and $\beta_{t+1}, \dots, \beta_{64-p}$ are the coefficients of the log-likelihood function (15).

```

 $\sigma_0 \leftarrow 0, \sigma_1 \leftarrow 0, u_{\min} \leftarrow -1, u_{\max} \leftarrow 0$ 
for  $j \leftarrow t+1$  to  $64-p$  do
  if  $\beta_j > 0$  then
    if  $u_{\min} < 0$  then  $u_{\min} \leftarrow j$ 
     $u_{\max} \leftarrow j$ 
     $\sigma_0 \leftarrow \sigma_0 + \beta_j, \sigma_1 \leftarrow \sigma_1 + \beta_j \cdot 2^{-j}$  ▷ see (27)
if  $u_{\min} < 0$  then return 0 ▷ all  $\beta_j$  are zero
 $\sigma_1 \leftarrow \sigma_1 \cdot 2^{u_{\max}}$ 
 $x \leftarrow \sigma_1 / (\alpha \cdot 2^{u_{\max}})$ 
if  $u_{\min} < u_{\max}$  then ▷  $x$  is already the root of  $f$ , if  $u_{\min} = u_{\max}$ 
   $x \leftarrow \text{expm1}(\log 1p(x)) \cdot (\sigma_0 / \sigma_1)$  ▷ starting point, see (27)
  loop ▷ main loop of Newton iteration
     $\lambda \leftarrow 1, \eta \leftarrow 0, y \leftarrow x, u \leftarrow u_{\max}$ 
     $\varphi \leftarrow \beta_u, \psi \leftarrow 0$ 
    loop ▷ loop for summing up  $\varphi$  (17) and  $\psi$  (28)
       $u \leftarrow u - 1$ 
       $z \leftarrow 2 / (2 + y)$  ▷  $z \in [0, 1]$ 
       $\lambda \leftarrow \lambda \cdot z$  ▷  $\lambda$  is decreasing, compare (22)
       $\eta \leftarrow \eta \cdot (2 - z) + (1 - z)$  ▷  $\eta$  is increasing, compare (30)
       $\varphi \leftarrow \varphi + \beta_u \cdot \lambda, \psi \leftarrow \psi + \beta_u \cdot \lambda \cdot \eta$ 
      if  $u \leq u_{\min}$  then break
       $y \leftarrow y \cdot (y + 2)$  ▷ compare (21)
     $x' \leftarrow (\alpha \cdot 2^{u_{\max}}) \cdot x$ 
    if  $\psi \leq x'$  then break ▷ stop iteration if  $f(x) \geq 0$ , see (18)
     $x_{\text{old}} \leftarrow x$ 
     $x \leftarrow x \cdot (1 + (\varphi - x') / (\psi + x'))$  ▷ compare (29)
    if  $x \leq x_{\text{old}}$  then break ▷ stop if numerically converged
return  $m \cdot 2^{u_{\max}} \cdot \log 1p(x)$  ▷ compare (19)

```

recursively using (20) according to

$$\eta_{j+1} = \eta_j \left(2 - \frac{2}{2+y_j}\right) + \left(1 - \frac{2}{2+y_j}\right). \quad (30)$$

The Newton iteration (29) can be stopped, if $f(x_t) \geq 0$, equivalent to $\varphi(x_t) \leq \alpha 2^{u_{\max}} x_t$, because we expect the sequence x_t to be increasing and converging towards the root. The case $\varphi(x_t) < \alpha 2^{u_{\max}} x_t$ may happen due to numerical errors. It is reasonable to stop the Newton iteration also in this case, as the numerical error limits have been reached. Similarly, we stop the iteration, if $x_{t+1} \leq x_t$. In practice, only a small number of iterations is needed to satisfy any of the two stop conditions. In all our experiments presented in Section 5, the number of iterations never exceeded 10 when calculating the estimate from ELL sketches. On average, we observed between 5 and 7 iterations, dependent on the ELL parameters and the true distinct count. The whole procedure to compute the ML estimate using Newton's method is summarized by Algorithm 8.

B PROOFS

LEMMA B.1. For ρ_{update} and ϕ as defined in (10) and (11) $\sum_{k=u+1}^{(65-p-t)2^t} \rho_{\text{update}}(k) = \frac{2^t(1-t+\phi(u))-u}{2^{\phi(u)}}$ holds.

PROOF. This formula can be proven by induction. For $u = (65-p-t)2^t$ both sides are zero. If the identity holds for u , it can be shown that it also holds for $u-1$. Since $\phi(u) - \phi(u-1) \in \{0, 1\}$, we show the identity for both possible cases. First, if $\phi(u) = \phi(u-1)$:

$$\begin{aligned} \sum_{k=(u-1)+1}^{(65-p-t)2^t} \rho_{\text{update}}(k) &= \rho_{\text{update}}(u) + \sum_{k=u+1}^{(65-p-t)2^t} \rho_{\text{update}}(k) \\ &= \frac{1}{2^{\phi(u)}} + \frac{2^t(1-t+\phi(u))-u}{2^{\phi(u)}} = \frac{2^t(1-t+\phi(u-1))-(u-1)}{2^{\phi(u-1)}}. \end{aligned}$$

And, second, if $\phi(u) = \phi(u-1) + 1$, which implies according to (11) $u-1 \equiv 0 \pmod{2^t}$ and further $\phi(u) = t+1 + \frac{u-1}{2^t} \Leftrightarrow 2^t(\phi(u) - t - 1) - u + 1 = 0$. With the help of this identity we can show again

$$\begin{aligned} \sum_{k=(u-1)+1}^{(65-p-t)2^t} \rho_{\text{update}}(k) &= \rho_{\text{update}}(u) + \sum_{k=u+1}^{(65-p-t)2^t} \rho_{\text{update}}(k) \\ &= \frac{1}{2^{\phi(u)}} + \frac{2^t(1-t+\phi(u))-u}{2^{\phi(u)}} + \frac{2^t(\phi(u)-t-1)-u+1}{2^{\phi(u)}} \\ &= \frac{2^t(-2t+2\phi(u))-2(u-1)}{2^{\phi(u)}} = \frac{2^t(1-t+\phi(u-1))-(u-1)}{2^{\phi(u-1)}}. \end{aligned}$$

□

LEMMA B.2. The function $f(x)$ as defined in (18) is strictly increasing and concave for $x \geq 0$ and $\alpha > 0$.

PROOF. It is sufficient to show that $g(x) := -\frac{x}{(1+x)^{c-1}}$ with $c \in \mathbb{Z}^+$ is increasing and concave, which is the case if $g'(x) \geq 0$ and $g''(x) \leq 0$. The first derivative is given by

$$\begin{aligned} g'(x) &= \frac{cx(1+x)^{c-1} - ((1+x)^{c-1})}{((1+x)^{c-1})^2} = \frac{cx(1+x)^{c-1} - \sum_{j=0}^{c-1} x(1+x)^j}{((1+x)^{c-1})^2} \\ &= \frac{x \sum_{j=0}^{c-2} (1+x)^{c-1-j}}{((1+x)^{c-1})^2} \geq 0, \end{aligned}$$

which is nonnegative. The second derivative can be expressed as

$$\begin{aligned} g''(x) &= \frac{c(1+x)^{c-2}((2+x)((1+x)^{c-1}) - cx(1+(1+x)^c))}{((1+x)^{c-1})^3} \\ &= \frac{c(1+x)^{c-2}(2((1+x)^{c-1}) - (c+1)x(1+x)^0 - (c-1)x(1+x)^c)}{((1+x)^{c-1})^3} \\ &= -\frac{c(1+x)^{c-2}x((c-1)(1+x)^c + (c-1)(1+x)^0 - 2\sum_{j=1}^{c-1}(1+x)^j)}{((1+x)^{c-1})^3}. \end{aligned}$$

Since $(1+x)^y$ is convex with respect to y ,

$$\begin{aligned} &\underbrace{(1+x)^c + \dots + (1+x)^c}_{(c-1) \text{ terms}} + \underbrace{(1+x)^0 + \dots + (1+x)^0}_{(c-1) \text{ terms}} \\ &\geq \underbrace{(1+x)^{c-1} + (1+x)^{c-1} + \dots + (1+x)^1 + (1+x)^1}_{2(c-1) \text{ terms}} \end{aligned}$$

holds according to Karamata's inequality, which shows that the last factor of the numerator is nonnegative. Since all factors are nonnegative we have $g''(x) \leq 0$. □

LEMMA B.3. The root \hat{x} of f as defined in (18) can be bracketed by $\exp(\ln(1 + \frac{\sigma_1}{\alpha 2^{u_{\max}}}) \frac{\sigma_0}{\sigma_1}) - 1 \leq \hat{x} \leq \frac{\sigma_0}{\alpha 2^{u_{\max}}}$ where $\sigma_0 := \sum_{j=u_{\min}}^{u_{\max}} \beta_j$ and $\sigma_1 := \sum_{j=u_{\min}}^{u_{\max}} \beta_j 2^{u_{\max}-j}$.

PROOF. We rewrite $f(\hat{x}) = 0$ as

$$\alpha 2^{u_{\max}} \hat{x} - \sigma_0 \frac{\sum_{j=0}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} g(\hat{x}, 2^j)}{\sum_{j=0}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j}} = 0 \quad (31)$$

with $g(x, y) := \frac{xy}{(1+x)^{y-1}}$. As g is convex with respect to y , which follows from $\frac{z}{e^z-1}$ being convex with $z = \ln(1+x)y$, we can apply Jensen's inequality

$$\begin{aligned} \alpha 2^{u_{\max}} \hat{x} - \sigma_0 \cdot g\left(\hat{x}, \frac{\sigma_1}{\sigma_0}\right) &= \alpha 2^{u_{\max}} \hat{x} - \sigma_0 \cdot g\left(\hat{x}, \frac{\sum_{j=0}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} 2^j}{\sum_{j=0}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j}}\right) \\ &\geq \alpha 2^{u_{\max}} \hat{x} - \sigma_0 \frac{\sum_{j=0}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j} g(\hat{x}, 2^j)}{\sum_{j=0}^{u_{\max}-u_{\min}} \beta_{u_{\max}-j}} = 0. \end{aligned}$$

Resolving for \hat{x} finally gives the lower bound. The upper bound results from (31) when using $g(x, y) \leq 1$ that is a consequence of Bernoulli's inequality. □

REFERENCES

- [1] [n.d.]. *Apache Data Sketches: A software library of stochastic streaming algorithms*. Retrieved February 15, 2025 from <https://datasketches.apache.org/>
- [2] [n.d.]. *Apache Data Sketches: Features Matrix for Distinct Count Sketches*. Retrieved February 15, 2025 from <https://datasketches.apache.org/docs/DistinctCountFeaturesMatrix.html>
- [3] [n.d.]. *Processor state control for your EC2 instance*. Retrieved January 5, 2025 from https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/processor_state_control.html
- [4] N. Alon, Y. Matias, and M. Szegedy. 1999. The Space Complexity of Approximating the Frequency Moments. *J. Comput. System Sci.* 58, 1 (1999), 137–147. <https://doi.org/10.1006/jcss.1997.1545>
- [5] D. N. Baker and B. Langmead. 2019. Dashing: fast and accurate genomic distances with HyperLogLog. *Genome Biology* 20, 265 (2019). <https://doi.org/10.1186/s13059-019-1875-0>
- [6] R. B. Basat, G. Einziger, S. L. Feibish, J. Moraney, and D. Raz. 2018. Network-wide routing-oblivious heavy hitters. In *Proceedings of the 16th Symposium on Architectures for Networking and Communications Systems (ANCS)*, 66–73. <https://doi.org/10.1145/3230718.3230729>
- [7] P. Boldi, M. Rosa, and S. Vigna. 2011. HyperANF: Approximating the neighbourhood function of very large graphs on a budget. In *Proceedings of the 20th International Conference on World Wide Web (WWW)*, 625–634. <https://doi.org/10.1145/1963405.1963493>
- [8] F. P. Breitwieser, D. N. Baker, and S. L. Salzberg. 2018. KrakenUniq: confident and fast metagenomics classification using unique k-mer counts. *Genome biology* 19, 1 (2018), 1–10. <https://doi.org/10.1186/s13059-018-1568-0>
- [9] Y. Chabchoub, R. Chiky, and B. Dogan. 2014. How can sliding HyperLogLog and EWMA detect port scan attacks in IP traffic? *EURASIP Journal on Information Security* 2014, 5 (2014). <https://doi.org/10.1186/1687-417X-2014-5>
- [10] A. Chen, J. Cao, L. Shepp, and T. Nguyen. 2011. Distinct Counting With a Self-Learning Bitmap. *J. Amer. Statist. Assoc.* 106, 495 (2011), 879–890. <https://doi.org/10.1198/jasa.2011.ap10217>
- [11] V. Clemens, L.-C. Schulz, M. Gartner, and D. Hausheer. 2023. DDoS Detection in P4 Using HYPERLOGLOG and COUNTMIN Sketches. In *Network Operations and Management Symposium (NOMS)*, 1–6. <https://doi.org/10.1109/NOMS56928.2023.10154315>
- [12] E. Cohen. 2015. All-Distances Sketches, Revisited: HIP Estimators for Massive Graphs Analysis. *IEEE Transactions on Knowledge and Data Engineering* 27, 9 (2015), 2320–2334. <https://doi.org/10.1109/TKDE.2015.2411606>
- [13] Y. Collet and M. Kucherawy. 2021. Zstandard Compression and the 'application/zstd' Media Type. RFC 8878. <https://doi.org/10.17487/RFC8878>
- [14] D. R. Cox and E. J. Snell. 1968. A General Definition of Residuals. *Journal of the Royal Statistical Society. Series B (Methodological)* 30, 2 (1968), 248–275. <http://www.jstor.org/stable/2984505>
- [15] Y. Du, H. Huang, Y. Sun, K. Li, B. Zhang, and G. Gao. 2023. A Better Cardinality Estimator with Fewer Bits, Constant Update Time, and Mergeability. In *IEEE Conference on Computer Communications (IEEE INFOCOM)*, 1–10. <https://doi.org/10.1109/INFOCOM53939.2023.10229088>
- [16] M. Durand. 2004. *Combinatoire analytique et algorithmique des ensembles de données*. Ph.D. Dissertation. École Polytechnique, Palaiseau, France. <https://pastel.hal.science/pastel-00000810>
- [17] R. A. L. Elworth, Q. Wang, P. K. Kota, C. J. Barberan, B. Coleman, A. Balaji, G. Gupta, R. G. Baraniuk, A. Shrivastava, and T. J. Treangen. 2020. To Petabytes and beyond: recent advances in probabilistic and signal processing algorithms and their application to metagenomics. *Nucleic Acids Research* 48, 10 (2020), 5217–5234. <https://doi.org/10.1093/nar/gkaa265>
- [18] O. Ertl. 2017. New cardinality estimation algorithms for HyperLogLog sketches. (2017). arXiv:cs.DS/1702.01284
- [19] O. Ertl. 2024. UltraLogLog: A Practical and More Space-Efficient Alternative to HyperLogLog for Approximate Distinct Counting. *Proceedings of the VLDB Endowment* 17, 7 (2024), 1655–1668. <https://doi.org/10.14778/3654621.3654632>
- [20] P. Flajolet, É. Fusy, O. Gandouet, and F. Meunier. 2007. HyperLogLog: the analysis of a near-optimal cardinality estimation algorithm. In *Proceedings of the 13th International Conference on the Analysis of Algorithms (AofA)*, 127–146. <https://doi.org/10.46298/dmcs.3545>
- [21] P. Flajolet and G. N. Martin. 1985. Probabilistic counting algorithms for data base applications. *Journal of computer and system sciences* 31, 2 (1985), 182–209. [https://doi.org/10.1016/0022-0000\(85\)90041-8](https://doi.org/10.1016/0022-0000(85)90041-8)
- [22] M. J. Freitag and T. Neumann. 2019. Every Row Counts: Combining Sketches and Sampling for Accurate Group-By Result Estimates. In *Proceedings of the 9th Conference on Innovative Data Systems Research (CIDR)*.
- [23] S. Heule, M. Nunkesser, and A. Hall. 2013. HyperLogLog in Practice: Algorithmic Engineering of a State of the Art Cardinality Estimation Algorithm. In *Proceedings of the 16th International Conference on Extending Database Technology (EDBT)*, 683–692. <https://doi.org/10.1145/2452376.2452456>
- [24] S. Janson, J. Lumbroso, and R. Sedgewick. 2024. Bit-Array-Based Alternatives to HyperLogLog. In *Proceedings of the 35th International Conference on Probabilistic, Combinatorial and Asymptotic Methods for the Analysis of Algorithms (AofA)*, 5:1–5:19. <https://doi.org/10.4230/LIPIcs.AofA.2024.5>
- [25] M. Karppa and R. Pagh. 2022. HyperLogLogLog: Cardinality Estimation With One Log More. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, 753–761. <https://doi.org/10.1145/3534678.3539246>
- [26] K. J. Lang. 2017. Back to the Future: an Even More Nearly Optimal Cardinality Estimation Algorithm. (2017). arXiv:cs.DS/1708.06839
- [27] J. Lu, H. Chen, J. Zhang, T. Hu, P. Sun, and Z. Zhang. 2023. Virtual self-adaptive bitmap for online cardinality estimation. *Information Systems* 114, 102160 (2023). <https://doi.org/10.1016/j.is.2022.102160>
- [28] G. Marçais, B. Solomon, R. Patro, and C. Kingsford. 2019. Sketching and Sublinear Data Structures in Genomics. *Annual Review of Biomedical Data Science* 2, 1 (2019), 93–118. <https://doi.org/10.1146/annurev-biodatasci-072018-021156>
- [29] R. Morris. 1978. Counting large numbers of events in small registers. *Commun. ACM* 21, 10 (1978), 840–842. <https://doi.org/10.1145/359619.359627>
- [30] T. Ohayon. 2021. ExtendedHyperLogLog: Analysis of a new Cardinality Estimator. (2021). arXiv:cs.DS/2106.06525
- [31] C. Pavlopoulou, M. J. Carey, and V. J. Tsotras. 2022. Revisiting Runtime Dynamic Optimization for Join Queries in Big Data Management Systems. In *Proceedings of the 25th International Conference on Extending Database Technology (EDBT)*. <https://doi.org/10.5441/002/edbt.2022.01>
- [32] O. Peters. [n.d.]. *PolymurHash*. Retrieved February 15, 2025 from <https://github.com/orlp/polymur-hash>
- [33] S. Pettie and D. Wang. 2021. Information Theoretic Limits of Cardinality Estimation: Fisher Meets Shannon. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, 556–569. <https://doi.org/10.1145/3406325.3451032>
- [34] S. Pettie, D. Wang, and L. Yin. 2021. Non-Mergeable Sketching for Cardinality Estimation. In *48th International Colloquium on Automata, Languages, and Programming (ICALP)*, Vol. 198, 104:1–104:20. <https://doi.org/10.4230/LIPIcs.ICALP.2021.104>
- [35] B. W. Priest, R. Pearce, and G. Sanders. 2018. Estimating Edge-Local Triangle Count Heavy Hitters in Edge-Linear Time and Almost-Vertex-Linear Space. In *Proceedings of the IEEE High Performance Extreme Computing Conference (HPEC)*. <https://doi.org/10.1109/HPEC.2018.8547721>
- [36] W. R. Punter, O. Papapetrou, and M. Garofalakis. 2023. OmniSketch: Efficient Multi-Dimensional High-Velocity Stream Analytics with Arbitrary Predicates. *Proceedings of the VLDB Endowment* 17, 3 (2023), 319–331. <https://doi.org/10.14778/3632093.3632098>
- [37] J. Qin, D. Kim, and Y. Tung. 2016. LogLog-Beta and More: A New Algorithm for Cardinality Estimation Based on LogLog Counting. (2016). arXiv:cs.DS/1612.02284
- [38] B. Scheuermann and M. Mauve. 2007. Near-optimal compression of probabilistic counting sketches for networking applications. In *Proceedings of the 4th ACM International Workshop on Foundations of Mobile Computing (FOMC)*.
- [39] R. Sedgewick. 2022. HyperBit: A Memory-Efficient Alternative to HyperLogLog. (2022). <https://www.birs.ca/workshops/2022/22w5004/files/BobSedgewick/HyperBit.pdf> Analytic and Probabilistic Combinatorics Workshop at the Banff International Research Station (BIRS) for Mathematical Innovation and Discovery.
- [40] R. Stanojevic. 2007. Small Active Counters. In *IEEE Conference on Computer Communications (IEEE INFOCOM)*, 2153–2161. <https://doi.org/10.1109/INFCOM.2007.249>
- [41] D. Ting. 2014. Streamed Approximate Counting of Distinct Elements: Beating Optimal Batch Methods. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 442–451. <https://doi.org/10.1145/2623330.2623669>
- [42] D. Ting. 2019. Approximate distinct counts for billions of datasets. In *Proceedings of the International Conference on Management of Data (SIGMOD)*, 69–86. <https://doi.org/10.1145/3299869.3319897>
- [43] R. Urban. [n.d.]. *SMhasher: Hash function quality and speed tests*. Retrieved February 15, 2025 from <https://github.com/rurban/smhasher>
- [44] A. Vaneev. [n.d.]. *Komihash*. Retrieved February 15, 2025 from <https://github.com/avaneev/komihash/tree/b27fd681308f92a1fae617b4ecd0981cc69d31a0>
- [45] D. Wang and S. Pettie. 2023. Better Cardinality Estimators for HyperLogLog, PCSA, and Beyond. In *Proceedings of the 42nd ACM Symposium on Principles of Database Systems (PODS)*, 317–327. <https://doi.org/10.1145/3584372.3588680>
- [46] J. Wires, S. Ingram, Z. Drudi, N. J. A. Harvey, and A. Warfield. 2014. Characterizing storage workloads with counter stacks. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 335–349. <https://www.usenix.org/conference/osdi14/technical-sessions/presentation/wires>
- [47] Q. Xiao, S. Chen, M. Chen, and Y. Ling. 2015. Hyper-Compact Virtual Estimators for Big Network Data Based on Register Sharing. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 417–428. <https://doi.org/10.1145/2745844.2745870>
- [48] Q. Xiao, S. Chen, Y. Zhou, and J. Luo. 2020. Estimating Cardinality for Arbitrarily Large Data Stream With Improved Memory Efficiency. *IEEE/ACM Transactions on Networking* 28, 2 (2020), 433–446. <https://doi.org/10.1109/TNET.2020.2970860>
- [49] Y. Zhao, S. Guo, and Y. Yang. 2016. Hermes: An Optimization of HyperLogLog Counting in real-time data processing. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 1890–1895. <https://doi.org/10.1109/IJCNN.2016.7727430>
- [50] W. Yi. [n.d.]. *Wyhash*. Retrieved February 15, 2025 from <https://github.com/wangyi-fudan/wyhash>
- [51] Y. Yu and G. M. Weber. 2022. HyperMinHash: MinHash in LogLog Space. *IEEE Transactions on Knowledge & Data Engineering* 34, 01 (2022), 328–339. <https://doi.org/10.1109/TKDE.2020.2981311>