

New cardinality estimation algorithms for HyperLogLog sketches

Otmar Ertl
otmar.ertl@gmail.com

February 27, 2017

This paper presents new methods to estimate the cardinalities of data sets recorded by HyperLogLog sketches. A theoretically motivated extension to the original estimator is presented that eliminates the bias for small and large cardinalities. Based on the maximum likelihood principle a second unbiased method is derived together with a robust and efficient numerical algorithm to calculate the estimate. The maximum likelihood approach can also be applied to more than a single HyperLogLog sketch. In particular, it is shown that it gives more precise cardinality estimates for union, intersection, or relative complements of two sets that are both represented by HyperLogLog sketches compared to the conventional technique using the inclusion-exclusion principle. All the new methods are demonstrated and verified by extensive simulations.

1. Introduction

Counting the number of distinct elements in a data stream or large datasets is a common problem in big data processing. In principle, finding the number of distinct elements n with a maximum relative error ε in a data stream requires $\mathcal{O}(n)$ space [1]. However, probabilistic algorithms that achieve the requested accuracy only with high probability are able to drastically reduce space requirements. Many different probabilistic algorithms have been developed over the past two decades [2, 3] until a theoretically optimal algorithm was finally found [4]. Although this algorithm achieves the optimal space complexity of $\mathcal{O}(\varepsilon^{-2} + \log n)$ [1, 5], it is not very efficient in practice [3].

More practicable and already widely used in many applications is the HyperLogLog algorithm [6] with a near-optimal space complexity $\mathcal{O}(\varepsilon^{-2} \log \log n + \log n)$. It has the

This paper together with source code for all presented algorithms and simulations is available at <https://github.com/oertl/hyperloglog-sketch-estimation-paper>. A first version of this paper was published there on April 17, 2016.

nice property that partial results can be easily merged, which is a necessity for distributed environments. Unfortunately, the originally proposed estimation method has some problems to guarantee the same accuracy over the full cardinality range. Therefore, a couple of variants have been developed to correct the original estimate by empirical means [7, 8, 9].

A more theoretically profound estimator for HyperLogLog sketches that does not depend on empirical data and significantly improves the estimation error is the historic inverse probability estimator [3, 10]. It trades memory efficiency for mergeability. The estimator needs to be continuously updated while inserting elements and the estimate depends on the element insertion order. Moreover, the estimator cannot be further used after merging two sketches, which limits its application to single data streams. If this restriction is acceptable, the self-learning bitmap [11] could also be used, which provides a similar trade-off and also needs less space than the original HyperLogLog method to achieve the same precision.

Sometimes not only the number of distinct elements but also a sample of them is needed in order to allow later filtering according to some predicate and estimating the cardinalities of corresponding subsets. In this case the k-minimum values algorithm [12] is the method of choice which also allows set manipulations like the construction of intersections, relative complements, or unions [13]. The latter operation is the only one that is natively supported by HyperLogLog sketches. A sketch that represents the set operation result is not always needed. One approach to estimate the corresponding result cardinality directly is based on the inclusion-exclusion principle, which however can become quite inaccurate, especially for small Jaccard indices of the input sets [13]. Alternatively, the HyperLogLog sketches can be combined with some additional data structure that allows estimating the Jaccard index. Together with the HyperLogLog cardinality estimates the error of the intersection size estimate can be reduced [14, 15], however, at the expense of a significantly larger memory footprint due to the additional data structure.

1.1. Outline

This paper presents new algorithms to extract cardinality information from HyperLogLog sketches. In Section 2 we start with an introduction to HyperLogLog sketches and the corresponding update algorithm. We also discuss how the resulting state can be statistically described and approximated by a Poisson model. In Section 3 we describe the original cardinality estimation algorithm, its shortcomings at low and large cardinalities, and previous approaches to overcome them. We present a simple derivation of the original raw estimator and analyze the root cause for its limited operating range. We derive an improved version of the raw estimator, which leads to a new fast cardinality estimation algorithm that works over the full cardinality range as demonstrated by extensive simulations. In Section 4 we derive a second even more precise cardinality estimation algorithm based on the maximum likelihood principle which is again verified by simulations. As presented in Section 5 the same approach can be generalized to two HyperLogLog sketches, which allows result cardinality estimation of set operations like

intersections or complements. The simulation results show that the estimates are significantly better than those of the conventional approach using the inclusion-exclusion principle. Finally, in Section 6 we discuss open problems and ideas for future work including HyperLogLog’s potential application as locality-sensitive hashing algorithm before we conclude in Section 7.

2. HyperLogLog data structure

The HyperLogLog algorithm collects information of incoming elements into a very compact sketching data structure, that finally allows the estimation of the number of distinct elements. The data structure consists of $m = 2^p$ registers whose number is chosen to be a power of two for performance reasons. p is the precision parameter that directly controls the relative estimation error which scales like $1/\sqrt{m}$ [6]. All registers start with zero initial value. Each element insertion potentially increases the value of one of these registers. The maximum value a register can reach is a natural bound given either by the output size of the used hash algorithm or the space that is reserved for a single register. Common implementations allocate up to 8 bits per register.

2.1. Data element insertion

The insertion of a data element into a HyperLogLog data structure requires the calculation of a $(p + q)$ -bit hash value. The leading p bits of the hash value are used to select one of the 2^p registers. Among the next following q bits, the position of the first 1-bit is determined which is a value in the range $[1, q + 1]$. The value $q + 1$ is used, if all q bits are zeros. If the position of the first 1-bit exceeds the current value of the selected register, the register value is replaced. Algorithm 1 shows the update procedure for inserting a data element into the HyperLogLog sketch.

Algorithm 1 Insertion of a data element D into a HyperLogLog data structure that consists of $m = 2^p$ registers. All registers $\mathbf{K} = (K_1, \dots, K_m)$ start from zero. $\langle \dots \rangle_2$ denotes the binary representation of an integer.

```

procedure INSERTELEMENT( $D$ )
   $\langle a_1, \dots, a_p, b_1, \dots, b_q \rangle_2 \leftarrow (p + q)$ -bit hash value of  $D$        $\triangleright a_i, b_i \in \{0, 1\}$ 
   $k \leftarrow \min(\{s \mid b_s = 1\} \cup \{q + 1\})$                                 $\triangleright k \in \{1, 2, \dots, q + 1\}$ 
   $i \leftarrow 1 + \langle a_1, \dots, a_p \rangle_2$                                         $\triangleright i \in \{1, 2, \dots, m\}$ 
  if  $k > K_i$  then
     $K_i \leftarrow k$ 
  end if
end procedure

```

The described element insertion algorithm makes use of what is known as stochastic averaging [16]. Instead of updating each of all m registers using m independent hash values, which would be an $\mathcal{O}(m)$ operation, only one register is selected and updated, which requires only a single hash function and reduces the complexity to $\mathcal{O}(1)$.

A HyperLogLog sketch can be characterized by a parameter pair (p, q) . The precision parameter p controls the relative error while the second parameter defines the possible value range of a register. A register can take all values starting from 0 to $q+1$, inclusively. The sum $p + q$ corresponds to the number of consumed hash value bits and defines the maximum cardinality that can be tracked. Obviously, if the cardinality reaches values in the order of 2^{p+q} , hash collisions will become more apparent and the estimation error will increase drastically.

Algorithm 1 has some properties which are especially useful for distributed data streams. First, the insertion order of elements has no influence on the final HyperLogLog sketch. Furthermore, any two HyperLogLog sketches with same parameters (p, q) representing two different sets can be easily merged. The HyperLogLog sketch that represents the union of both sets can be easily constructed by taking the register-wise maximum values as demonstrated by Algorithm 2.

Algorithm 2 Merge operation for two HyperLogLog sketches with register values $\mathbf{K}_1 = (K_{11}, \dots, K_{1m})$ and $\mathbf{K}_2 = (K_{21}, \dots, K_{2m})$ representing sets S_1 and S_2 , respectively, to obtain the register values $\mathbf{K}' = (K'_1, \dots, K'_m)$ of a HyperLogLog sketch representing $S_1 \cup S_2$.

```

function MERGE( $\mathbf{K}_1, \mathbf{K}_2$ )                                 $\triangleright \mathbf{K}_1, \mathbf{K}_2 \in \{0, 1, \dots, q+1\}^m$ 
    allocate  $\mathbf{K}' = (K'_1, \dots, K'_m)$                      $\triangleright \mathbf{K}' \in \{0, 1, \dots, q+1\}^m$ 
    for  $i \leftarrow 1, m$  do
         $K'_i \leftarrow \max(K_{1i}, K_{2i})$ 
    end for
    return  $\mathbf{K}'$ 
end function

```

At any time a (p, q) -HyperLogLog sketch can be reduced to a (p', q') -HyperLogLog data structure, if $p' \leq p$ and $p' + q' \leq p + q$ is satisfied (see Algorithm 3). This transformation is lossless in a sense that the resulting HyperLogLog sketch is the same as if all elements would have been recorded by a (p', q') -HyperLogLog sketch right from the beginning.

A $(p, 0)$ -HyperLogLog sketch corresponds to a bit array as used by linear counting [17]. Each register value can be stored by a single bit in this case. Hence, linear counting can be regarded as a special case of the HyperLogLog algorithm for which $q = 0$.

2.2. Joint probability distribution of register values

Under the assumption of a uniform hash function, the probability that the register values $\mathbf{K} = (K_1, \dots, K_m)$ of a HyperLogLog sketch with parameters p and q are equal to $\mathbf{k} = (k_1, \dots, k_m)$ is given by the corresponding probability mass function

$$\rho(\mathbf{k}|n) = \sum_{n_1 + \dots + n_m = n} \binom{n}{n_1, \dots, n_m} \frac{1}{m^n} \prod_{i=1}^m \gamma(k_i | n_i) \quad (1)$$

Algorithm 3 Compression of a (p, q) -HyperLogLog sketch with register values $\mathbf{K} = (K_1, \dots, K_m)$ into a (p', q') -HyperLogLog sketch with $p' \leq p$ and $p' + q' \leq p + q$.

```

function COMPRESS( $\mathbf{K}$ )
    allocate  $\mathbf{K}' = (K'_1, \dots, K'_{2^{p'}})$ 
    for  $i \leftarrow 1, 2^{p'}$  do
         $b \leftarrow (i - 1) \cdot 2^{p-p'}$ 
         $j \leftarrow 1$ 
        while  $j \leq 2^{p-p'} \wedge K_{b+j} = 0$  do
             $j \leftarrow j + 1$ 
        end while
        if  $j = 1$  then
             $K'_i \leftarrow \min(K_{b+j} + p - p', q' + 1)$ 
        else if  $j \leq 2^{p-p'}$  then
             $\langle a_1, \dots, a_{p-p'} \rangle_2 \leftarrow j - 1$ 
             $K'_i \leftarrow \min(\{s \mid a_s = 1\} \cup \{q' + 1\})$ 
        else
             $K'_i \leftarrow 0$ 
        end if
    end for
    return  $\mathbf{K}'$ 
end function

```

where n is the cardinality. The n distinct elements are distributed over all m registers according to a multinomial distribution with equal probabilities. $\gamma(k|n)$ is the probability that the value of a register is equal to k , after it was selected n times by the insertion algorithm

$$\gamma(k|n) := \begin{cases} 1 & n = 0 \wedge k = 0 \\ 0 & n = 0 \wedge 1 \leq k \leq q + 1 \\ 0 & n \geq 1 \wedge k = 0 \\ (1 - \frac{1}{2^k})^n - (1 - \frac{1}{2^{k-1}})^n & n \geq 1 \wedge 1 \leq k \leq q \\ 1 - (1 - \frac{1}{2^q})^n & n \geq 1 \wedge k = q + 1. \end{cases} \quad (2)$$

The order of register values K_1, \dots, K_m is not important for the estimation of the cardinality. More formally, the multiset $\{K_1, \dots, K_m\}$ is a sufficient statistic for n . Since the values of the multiset are all in the range $[0, q + 1]$ the multiset can also be represented as $\{K_1, \dots, K_m\} = 0^{C_0} 1^{C_1} \dots q^{C_q} (q + 1)^{C_{q+1}}$ where C_k is the multiplicity of value k . As a consequence, the multiplicity vector $\mathbf{C} := (C_0, \dots, C_{q+1})$ is also a sufficient statistic for the cardinality. In addition, this vector contains all the information about the HyperLogLog sketch that is required for cardinality estimation. The two HyperLogLog parameters can be obtained by $p = \log_2(\sum_{k=0}^{q+1} C_k)$ and $q = \dim \mathbf{C} - 2$, respectively. Algorithm 4 shows the calculation of the multiplicity vector \mathbf{C} for a given HyperLogLog sketch with register values \mathbf{K} .

Algorithm 4 Multiplicity vector extraction from a (p, q) -HyperLogLog sketch with $m = 2^p$ registers having values $\mathbf{K} = (K_1, \dots, K_m)$.

```

function EXTRACTCOUNTS( $\mathbf{K}$ )                                ▷  $\mathbf{K} \in \{0, 1, \dots, q + 1\}^m$ 
  allocate  $\mathbf{C} = (C_0, \dots, C_{q+1})$ 
   $\mathbf{C} \leftarrow (0, 0, \dots, 0)$ 
  for  $i \leftarrow 1, m$  do
     $C_{K_i} \leftarrow C_{K_i} + 1$ 
  end for
  return  $\mathbf{C}$ 
end function

```

2.3. Poisson approximation

Due to the statistical dependence of the register values, the probability mass function (1) makes further analysis difficult. Therefore, a Poisson model can be used [6], which assumes that the cardinality itself is distributed according to a Poisson distribution

$$n \sim \text{Poisson}(\lambda). \quad (3)$$

Under the Poisson model the distribution of the register values is

$$\begin{aligned}
\rho(\mathbf{k}|\lambda) &= \sum_{n=0}^{\infty} \rho(\mathbf{k}|n) e^{-\lambda} \frac{\lambda^n}{n!} & (4) \\
&= \sum_{n_1=0}^{\infty} \cdots \sum_{n_m=0}^{\infty} \prod_{i=1}^m \gamma(k_i|n_i) e^{-\frac{\lambda}{m}} \frac{\lambda^{n_i}}{n_i! m^{n_i}} \\
&= \prod_{i=1}^m \sum_{n=0}^{\infty} \gamma(k_i|n) e^{-\frac{\lambda}{m}} \frac{\lambda^n}{n! m^n} \\
&= \prod_{k=0}^{q+1} \left(\sum_{n=0}^{\infty} \gamma(k|n) e^{-\frac{\lambda}{m}} \frac{\lambda^n}{n! m^n} \right)^{c_k} \\
&= e^{-c_0 \frac{\lambda}{m}} \left(\prod_{k=1}^q \left(e^{-\frac{\lambda}{m 2^k}} \left(1 - e^{-\frac{\lambda}{m 2^k}} \right) \right)^{c_k} \right) \left(1 - e^{-\frac{\lambda}{m 2^q}} \right)^{c_{q+1}}. & (5)
\end{aligned}$$

Here c_k denotes the multiplicity of value k in the multiset $\{k_1, \dots, k_m\}$. The final factorization shows that the register values K_1, \dots, K_m are independent and identically distributed under the Poisson model. The probability that a register has a value less than or equal to k for a given rate λ is given by

$$P(K \leq k|\lambda) = \begin{cases} 0 & k < 0 \\ e^{-\frac{\lambda}{m 2^k}} & 0 \leq k \leq q \\ 1 & k > q. \end{cases} \quad (6)$$

2.4. Depoissonization

Due to the simpler probability mass function (5), it is easier to find an estimator $\hat{\lambda} = \hat{\lambda}(\mathbf{K})$ for the Poisson rate λ than for the cardinality n under the fixed-size model (1). Depoissonization [18] finally allows to translate the estimates back to the fixed-size model. Assume we have found an unbiased estimator for the Poisson rate

$$\mathbb{E}(\hat{\lambda}|\lambda) = \lambda \quad \text{for all } \lambda \geq 0. \quad (7)$$

We know from (4)

$$\mathbb{E}(\hat{\lambda}|\lambda) = \sum_{n=0}^{\infty} \mathbb{E}(\hat{\lambda}|n) e^{-\lambda} \frac{\lambda^n}{n!} \quad (8)$$

and therefore

$$\sum_{n=0}^{\infty} \mathbb{E}(\hat{\lambda}|n) e^{-\lambda} \frac{\lambda^n}{n!} = \lambda \quad (9)$$

holds for all $\lambda \geq 0$. The unique solution of this equation is given by

$$\mathbb{E}(\hat{\lambda}|n) = n. \quad (10)$$

Hence, the unbiased estimator $\hat{\lambda}$ conditioned on n is also an unbiased estimator for n , which motivates us to use $\hat{\lambda}$ directly as estimator for the cardinality $\hat{n} := \hat{\lambda}$. As simulation results will show later, the Poisson approximation works well over the entire cardinality range, even for estimators that are not exactly unbiased.

3. Original cardinality estimation method

The original cardinality estimator [6] is based on the idea that the number of distinct element insertions a register needs to reach the value k is proportional to $m2^k$. Given that, a rough cardinality estimate can be obtained by averaging the values $\{m2^{K_1}, \dots, m2^{K_m}\}$. In the history of the HyperLogLog algorithm different averaging techniques have been proposed. First, there was the LogLog algorithm using the geometric mean and the SuperLogLog algorithm that enhanced the estimate by truncating the largest register values before applying the geometric mean [19]. Finally, the harmonic mean was found to give even better estimates as it is inherently less sensitive to outliers. The result is the so-called raw estimator which is given by

$$\hat{n}_{\text{raw}} = \alpha_m \frac{m}{\frac{1}{m2^{K_1}} + \dots + \frac{1}{m2^{K_m}}} = \frac{\alpha_m m^2}{\sum_{i=1}^m 2^{-K_i}} = \frac{\alpha_m m^2}{\sum_{k=0}^{q+1} C_k 2^{-k}}. \quad (11)$$

Here α_m is a bias correction factor which was derived for a given number of registers m as [6]

$$\alpha_m := \left(m \int_0^{\infty} \left(\log_2 \left(\frac{2+x}{1+x} \right) \right)^m dx \right)^{-1}. \quad (12)$$

Numerical approximations of α_m for various values of m have been listed in [6]. These approximations are used in many HyperLogLog implementations. However, since the published constants have been rounded to 4 significant digits, these approximations even introduce some bias for very high precisions p . For HyperLogLog sketches that are used in practice with 256 or more registers ($p \geq 8$), it is completely sufficient to use

$$\alpha_\infty := \lim_{m \rightarrow \infty} \alpha_m = \frac{1}{2 \log 2} = 0.72134752\dots \quad (13)$$

as approximation for α_m in (11), because the additional bias is negligible compared to the overall estimation error.

Fig. 1 shows the distribution of the relative error for the raw estimator as function of the cardinality. The chart is based on 10 000 randomly generated HyperLogLog sketches. More details of the experimental setup will be explained later in Section 3.5. Obviously, the raw estimator is biased for small and large cardinalities where it fails to return accurate estimates. In order to cover the entire cardinality range, corrections for small and large cardinalities have been proposed.

As mentioned in Section 2.1, a HyperLogLog sketch with parameters (p, q) can be mapped to a $(p, 0)$ -HyperLogLog sketch. Since $q = 0$ corresponds to linear counting and the reduced HyperLogLog sketch corresponds to a bitset with C_0 zeros, the linear counting cardinality estimator [17] can be used

$$\hat{n}_{\text{small}} = m \log(m/C_0). \quad (14)$$

The corresponding relative estimation error as depicted in Fig. 2 shows that this estimator is convenient for small cardinalities. It was proposed to use this estimator for small cardinalities as long as $\hat{n}_{\text{raw}} \leq \frac{5}{2}m$ where the factor $\frac{5}{2}$ was empirically determined [6].

For large cardinalities in the order of 2^{p+q} , for which a lot of registers are already in a saturated state, meaning that they have reached the maximum possible value $q + 1$, the raw estimator underestimates the cardinalities. For the 32-bit hash value case ($p + q = 32$), which was considered in [6], following correction formula was proposed to take these saturated registers into account

$$\hat{n}_{\text{large}} = -2^{32} \log(1 - \hat{n}_{\text{raw}}/2^{32}). \quad (15)$$

The original estimation algorithm as presented in [6] including corrections for small and large cardinalities is summarized by Algorithm 5. The relative estimation error for the original method is shown in Fig. 3. Unfortunately, as can be clearly seen, the ranges where the estimation error is small for \hat{n}_{raw} and \hat{n}_{small} do not overlap. Therefore, the estimation error is much larger near the transition region. To reduce the estimation error for cardinalities close to this region, it was proposed to correct \hat{n}_{raw} for bias. Empirically collected bias correction data is either stored as set of interpolation points [7], as lookup table [8], or as best-fitting polynomial [9]. However, all these empirical approaches treat the symptom and not the cause.

The large range correction formula (15) is not satisfying either as it does not reduce the estimation error. Quite the contrary, it even makes the bias worse. However, instead

Algorithm 5 Original cardinality estimation algorithm for HyperLogLog sketches that use 32-bit hash values ($p + q = 32$) for insertion of data items [6].

```

function ESTIMATECARDINALITY( $\mathbf{K}$ )
     $\hat{n}_{\text{raw}} \leftarrow \alpha_m m^2 \left( \sum_{i=1}^m 2^{-K_i} \right)^{-1}$ 
    if  $\hat{n}_{\text{raw}} \leq \frac{5}{2}m$  then
         $C_0 \leftarrow |\{i | K_i = 0\}|$ 
        if  $C_0 \neq 0$  then
            return  $m \log(m/C_0)$ 
        else
            return  $\hat{n}_{\text{raw}}$ 
        end if
    else if  $\hat{n}_{\text{raw}} \leq \frac{1}{30}2^{32}$  then
        return  $\hat{n}_{\text{raw}}$ 
    else
        return  $-2^{32} \log(1 - \hat{n}_{\text{raw}}/2^{32})$ 
    end if
end function

```

$\triangleright \mathbf{K} \in \{0, 1, \dots, q + 1\}^m, m = 2^p$
 \triangleright raw estimate (11)
 \triangleright small range correction (14)
 \triangleright large range correction (15)




Figure 1: The distribution of the relative estimation error over the cardinality for the raw estimator after evaluation of 10 000 randomly generated HyperLogLog data structures with parameters $p = 12$ and $q = 20$.




Figure 2: The distribution of the relative estimation error for the linear counting estimator after evaluation of 10 000 randomly generated bitmaps of size 2^{12} which correspond to HyperLogLog sketches with parameters $p = 12$ and $q = 0$.




Figure 3: The distribution of the relative estimation error over the cardinality for the original estimation algorithm after evaluation of 10 000 randomly generated HyperLogLog data structures with parameters $p = 12$ and $q = 20$.

of underestimating the cardinalities, they are now overestimated. Another indication for the incorrectness of the proposed large range correction is the fact that it is not even defined for all possible states. For instance, consider a (p, q) -HyperLogLog sketch with $p + q = 32$ for which all registers are equal to the maximum possible value $q + 1$. The raw estimate would be $\hat{n}_{\text{raw}} = \alpha_m 2^{33}$, which is greater than 2^{32} and outside of the domain of the large range correction formula.

A simple approach to avoid the need of any large range correction is to extend the operating range of the raw estimator to larger cardinalities. This can be easily accomplished by increasing $p + q$, which corresponds to using hash values with more bits. Each additional bit doubles the operating range which scales like 2^{p+q} . However, in case $q \geq 31$ the number of possible register values, which are $\{0, 1, \dots, q + 1\}$, is greater than 32 which is no longer representable by 5-bit registers. Therefore, it was proposed to use 6 bits per register in combination with 64-bit hash values [7]. Even larger hash values are needless in practice, because it is unrealistic to encounter cardinalities of order 2^{64} .

3.1. Derivation of the raw estimator

In order to better understand why the raw estimator fails for small and large cardinalities, we start with a brief and simple derivation without the restriction to large cardinalities ($n \rightarrow \infty$) and without using complex analysis as in [6].

Let us assume that the register values have following cumulative distribution function

$$P(K \leq k | \lambda) = e^{-\frac{\lambda}{m2^k}}. \quad (16)$$

For now we ignore that this distribution has infinite support and differs from the register value distribution under the Poisson model (6), whose support is limited to the range $[0, q + 1]$. For a random variable K obeying (16) the expectation of 2^{-K} is given by

$$\begin{aligned} \mathbb{E}(2^{-K}) &= \\ \sum_{k=-\infty}^{\infty} 2^{-k} \left(e^{-\frac{\lambda}{m2^k}} - e^{-\frac{\lambda}{m2^{k-1}}} \right) &= \frac{1}{2} \sum_{k=-\infty}^{\infty} 2^k e^{-\frac{\lambda}{m} 2^k} = \frac{\alpha_{\infty} m \xi(\log_2(\lambda/m))}{\lambda}, \end{aligned} \quad (17)$$

where the function

$$\xi(x) := \log(2) \sum_{k=-\infty}^{\infty} 2^{k+x} e^{-2^{k+x}} \quad (18)$$

is a smooth periodic function with period 1. Numerical evaluations indicate that this function can be bounded by $1 - \varepsilon_{\xi} \leq \xi(x) \leq 1 + \varepsilon_{\xi}$ with $\varepsilon_{\xi} := 9.885 \times 10^{-6}$ (see Fig. 4). This value can also be found using Fourier analysis as shown in Appendix A.

Let K_1, \dots, K_m be a sample distributed according to (16). For large sample sizes $m \rightarrow \infty$ we have asymptotically

$$\mathbb{E} \left(\frac{1}{2^{-K_1} + \dots + 2^{-K_m}} \right) \stackrel{m \rightarrow \infty}{=} \frac{1}{\mathbb{E}(2^{-K_1} + \dots + 2^{-K_m})} = \frac{1}{m \mathbb{E}(2^{-K})}. \quad (19)$$




Figure 4: The deviation of $\xi(x)$ from 1.

Together with (17) we obtain

$$\lambda = \mathbb{E} \left(\frac{\alpha_\infty m^2 \xi(\log_2(\lambda/m))}{2^{-K_1} + \dots + 2^{-K_m}} \right) \quad \text{for } m \rightarrow \infty. \quad (20)$$

Therefore, the asymptotic relative bias of

$$\hat{\lambda} = \frac{\alpha_\infty m^2}{2^{-K_1} + \dots + 2^{-K_m}} \quad (21)$$

is bounded by ε_ξ , which makes this statistic a good estimator for the Poisson parameter. It also corresponds to the raw estimator (11), if the Poisson parameter estimate is used as cardinality estimate (see Section 2.4).

3.2. Limitations of the raw estimator

The raw estimator is based on two prerequisites. First, the number of registers needs to be sufficiently large ($m \rightarrow \infty$). And second, the distribution of register values can be described by (16). However, the latter is not true for small and large cardinalities, which is finally the reason for the bias of the raw estimator.

A random variable K' with cumulative distribution function (16) can be transformed into a random variable K with cumulative distribution function (6) using

$$K = \min(\max(K', 0), q + 1). \quad (22)$$

Therefore, register values K_1, \dots, K_m can be seen as the result after applying this transformation to a sample K'_1, \dots, K'_m of the distribution described by (16). If all registers values fall into the range $[1, q]$, they must be identical to the values K'_1, \dots, K'_m . In other words, the observed register values are also a plausible sample of the assumed distribution described by (16) in this case. Hence, as long as all or at least most register values are in the range $[1, q]$, which is the case if $2^p \ll \lambda \ll 2^{p+q}$, the approximation of

(6) by (16) is valid. This explains why the raw estimator works best for intermediate cardinalities. However, for small and large cardinalities many register values are equal to 0 or $q + 1$, respectively, which cannot be described by (16) and finally leads to the observed bias.

3.3. Corrections to the raw estimator

If we knew the values K'_1, \dots, K'_m for which transformation (22) led to the observed register values K_1, \dots, K_m , we would be able to estimate λ using

$$\hat{\lambda} = \frac{\alpha_\infty m^2}{2^{-K'_1} + \dots + 2^{-K'_m}}. \quad (23)$$

As we have already shown, this estimator is approximately unbiased, because all K'_i follow the assumed distribution. It would be even sufficient, if we knew the multiplicity C'_k of each $k \in \mathbb{Z}$ in $\{K'_1, \dots, K'_m\}$, $C'_k := |\{i | k = K'_i\}|$, because the raw estimator can be also written as

$$\hat{\lambda} = \frac{\alpha_\infty m^2}{\sum_{k=-\infty}^{\infty} C'_k 2^{-k}}. \quad (24)$$

Due to (22), the multiplicities C'_k and the multiplicities C_k for the observed register values have following relationships

$$\begin{aligned} C_0 &= \sum_{k=-\infty}^0 C'_k, \\ C_k &= C'_k, \quad 1 \leq k \leq q, \\ C_{q+1} &= \sum_{k=q+1}^{\infty} C'_k. \end{aligned} \quad (25)$$

The idea is now to find estimates \hat{c}'_k for all $k \in \mathbb{Z}$ and use them as replacements for C'_k in (24). For $k \in [1, q]$ where $C_k = C'_k$ we can use the trivial estimators

$$\hat{c}'_k := C_k, \quad 1 \leq k \leq q. \quad (26)$$

To get estimators for $k \leq 0$ and $k \geq q + 1$, we consider the expectation of C'_k

$$\mathbb{E}(C'_k) = m (P(K' \leq k | \lambda) - P(K' \leq k - 1 | \lambda)) = m e^{-\frac{\lambda}{m 2^k}} \left(1 - e^{-\frac{\lambda}{m 2^k}}\right). \quad (27)$$

From (6) we know that $\mathbb{E}(C_0/m) = e^{-\frac{\lambda}{m}}$ and $\mathbb{E}(1 - C_{q+1}/m) = e^{-\frac{\lambda}{m 2^q}}$, and therefore, we can also write

$$\mathbb{E}(C'_k) = m (\mathbb{E}(C_0/m))^{2^{-k}} \left(1 - (\mathbb{E}(C_0/m))^{2^{-k}}\right) \quad (28)$$

and

$$\mathbb{E}(C'_k) = m (\mathbb{E}(1 - C_{q+1}/m))^{2^{q-k}} \left(1 - (\mathbb{E}(1 - C_{q+1}/m))^{2^{q-k}}\right), \quad (29)$$

which motivates us to use

$$\hat{c}'_k = m (C_0/m)^{2^{-k}} \left(1 - (C_0/m)^{2^{-k}}\right) \quad (30)$$

as estimator for $k \leq 0$ and

$$\hat{c}'_k = m(1 - C_{q+1}/m)^{2^{q-k}} \left(1 - (1 - C_{q+1}/m)^{2^{q-k}}\right) \quad (31)$$

as estimator for $k \geq q + 1$, respectively. This choice of estimators also conserves the mass of zero-valued and saturated registers, because (25) is satisfied, if C'_k is replaced by \hat{c}'_k . Plugging all these estimators into (24) as replacements for C'_k finally gives

$$\hat{\lambda} = \frac{\alpha_\infty m^2}{\sum_{k=-\infty}^{\infty} \hat{c}'_k 2^{-k}} = \frac{\alpha_\infty m^2}{m \sigma(C_0/m) + \sum_{k=1}^q C_k 2^{-k} + m \tau(1 - C_{q+1}/m) 2^{-q}} \quad (32)$$

which we call the improved raw estimator. Here $m \sigma(C_0/m)$ and $2m \tau(1 - C_{q+1}/m)$ are replacements for C_0 and C_{q+1} in the raw estimator (11), respectively. The functions σ and τ are defined as

$$\sigma(x) := x + \sum_{k=1}^{\infty} x^{2^k} 2^{k-1} \quad (33)$$

and

$$\tau(x) := \frac{1}{2} \left(-x + \sum_{k=1}^{\infty} x^{2^{-k}} 2^{-k} \right). \quad (34)$$

We can cross-check the new estimator for the linear counting case with $q = 0$. Using the identity

$$\sigma(x) + \tau(x) = \alpha_\infty \xi(\log_2(\log(1/x))) / \log(1/x), \quad (35)$$

we get

$$\hat{\lambda} = \frac{\alpha_\infty m}{\sigma(C_0/m) + \tau(C_0/m)} = \frac{m \log(m/C_0)}{\xi(\log_2(\log(m/C_0)))} \quad (36)$$

which is as expected almost identical to the linear counting estimator (14), because $\xi(x) \approx 1$ (see Section 3.1).

3.4. Improved raw estimation algorithm

The improved raw estimator (32) can be directly translated into a new cardinality estimation algorithm for HyperLogLog sketches as shown in Algorithm 6. Since the series (33) converges quadratically for all $x \in [0, 1)$ and its terms can be recursively calculated using elementary operations, the function σ can be quickly calculated to machine precision. The case $x = 1$ must be handled separately, because the series diverges and causes an infinite denominator in (32) and therefore a vanishing cardinality estimate. As this case only occurs if all register values are zero ($C_0 = m$), this is exactly what is expected. Among the remaining possible arguments $\{0, \frac{1}{m}, \frac{2}{m}, \dots, \frac{m-1}{m}\}$ we have slowest convergence for $x = \frac{m-1}{m}$. However, even in this case we have a manageable amount of iteration cycles. For example, if double-precision floating-point arithmetic is used, the routine for calculating σ converges after 18 and 26 iteration cycles for $p = 12$ and $p = 20$, respectively.

The calculation of τ is more expensive, because it involves square root evaluations. τ can also be written in a numerically more favorable way as

$$\tau(x) := \frac{1}{3} \left(1 - x - \sum_{k=1}^{\infty} \left(1 - x^{2^{-k}} \right)^2 2^{-k} \right). \quad (37)$$

This representation shows faster convergence for $x > 0$, because $1 - x^{2^{-k}} \leq -\log(x)2^{-k}$. Therefore, the convergence speed is comparable to a geometric series with ratio $1/8$. Disregarding the trivial case $\tau(0) = 0$, $x = \frac{1}{m}$ shows slowest convergence speed among all other possible parameters. For this case and if double-precision floating-point numbers are used, the function τ presented in Algorithm 6 converges after 21 and 22 iteration cycles for considered precisions $p = 12$ and $p = 20$, respectively.

If performance matters σ and τ can also be precalculated for all possible values of C_0 and C_{q+1} . As their value range is $\{0, 1, \dots, m\}$, the function values can be kept in lookup tables of size $m + 1$. In this way a complete branch-free cardinality estimation can be realized. It is also thinkable to calculate τ using the approximation $\tau(x) \approx \alpha_{\infty} / \log(1/x) - \sigma(x)$ which can be obtained from (35). The advantage is that the calculation of σ and the additional logarithm is slightly faster than the calculation of τ . However, for arguments close to 1, where $\alpha_{\infty} / \log(1/x)$ and $\sigma(x)$ are both very large while their difference is very small, special care is needed to avoid numerical cancellation.

The new estimation algorithm is very elegant, because it does neither contain magic numbers nor special cases as the original algorithm. Moreover, the algorithm guarantees monotonicity of the cardinality estimate. The estimate will never become smaller when adding new elements. Although this sounds self-evident, previous approaches that combine estimators for different cardinality ranges or make use of empirical collected data have problems to satisfy this property throughout all cardinalities.

3.5. Experimental setup

To verify the new estimation algorithm, we generated 10 000 different HyperLogLog sketches and filled each of them with 50 billion unique elements. During element insertion, we took snapshots of the sketch state by storing the multiplicity vector at predefined cardinality values, which were roughly chosen according to a geometric series with ratio 1.01. In order to achieve that in reasonable time for all the different HyperLogLog parameter combinations (p, q) we have been interested in, we applied a couple of optimizations.

First, we assumed a uniform hash function. This allows us to simply generate random numbers instead of creating unique elements and calculating their hash values. We used the Mersenne Twister random number generator with a state size of 19 937 bits from the C++ standard library. Second, we used Algorithm 7 for insertions, which does basically the same as Algorithm 1, but additionally keeps track of the multiplicity vector and the minimum register value. In this way the expensive register scan needed to obtain the multiplicity vector as described by Algorithm 4 can be avoided. Furthermore, the knowledge of the minimum register value can be used to abort the insertion procedure

Algorithm 6 Cardinality estimation based on the improved raw estimator. The input is the multiplicity vector $\mathbf{C} = (C_0, \dots, C_{q+1})$ as obtained by Algorithm 4.

```

function ESTIMATECARDINALITY( $\mathbf{C}$ )
   $z \leftarrow m \cdot \tau(1 - C_{q+1}/m)$ 
  for  $k \leftarrow q, 1$  do
     $z \leftarrow 0.5 \cdot (z + C_k)$ 
  end for
   $z \leftarrow z + m \cdot \sigma(C_0/m)$ 
  return  $\alpha_\infty m^2/z$ 
end function

function  $\sigma(x)$ 
  if  $x = 1$  then
    return  $\infty$ 
  end if
   $y \leftarrow 1$ 
   $z \leftarrow x$ 
  repeat
     $x \leftarrow x \cdot x$ 
     $z' \leftarrow z$ 
     $z \leftarrow z + x \cdot y$ 
     $y \leftarrow 2 \cdot y$ 
  until  $z = z'$ 
  return  $z$ 
end function

function  $\tau(x)$ 
  if  $x = 0 \vee x = 1$  then
    return  $0$ 
  end if
   $y \leftarrow 1$ 
   $z \leftarrow 1 - x$ 
  repeat
     $x \leftarrow \sqrt{x}$ 
     $z' \leftarrow z$ 
     $y \leftarrow 0.5 \cdot y$ 
     $z \leftarrow z - (1 - x)^2 \cdot y$ 
  until  $z = z'$ 
  return  $z/3$ 
end function

```

▷ $\sum_{k=0}^{q+1} C_k = m$
 ▷ alternatively, take $m \cdot \tau(1 - C_{q+1}/m)$ from precalculated lookup table

▷ alternatively, take $m \cdot \sigma(C_0/m)$ from precalculated lookup table
 ▷ $\alpha_\infty := 1/(2 \log(2))$

▷ $x \in [0, 1]$

▷ $x \in [0, 1]$

early, if the hash value is not able to increment any register value. Especially at large cardinalities this saves a great number of register accesses and also reduces the number of cache misses. For example, if $K_{\min} = 1$ only half of all insertions will access a register. If $K_{\min} = 2$ it is already just a quarter, and so forth. Finally, instead of repeating the simulation for different HyperLogLog parameters, we did that only once using the quite accurate setting $p = 22$ and $q = 42$. Whenever taking a snapshot, we reduced the sketch state exploiting Algorithm 3 for desired parameters before calculating the multiplicity vector using Algorithm 4.

After all, we got the multiplicity vectors at predefined cardinalities of 10 000 different simulation runs for various HyperLogLog parameter settings. All this data constitutes the basis of the presented results that follow. The generated data allows a quick and simple evaluation of cardinality estimation algorithms by just loading the persisted multiplicity vectors from hard disk and passing them as input to the algorithm.

Algorithm 7 Extension of Algorithm 1 that keeps track of the multiplicity vector $\mathbf{C} = (C_0, \dots, C_{q+1})$ and the minimum register value K_{\min} during insertion. Initially, $K_{\min} = 0$ and $\mathbf{C} = (m, 0, 0, \dots, 0)$.

```

procedure INSERTELEMENT( $D$ )
   $\langle a_1, \dots, a_p, b_1, \dots, b_q \rangle_2 \leftarrow (p+q)$ -bit hash value of  $D$        $\triangleright a_i, b_i \in \{0, 1\}$ 
   $k \leftarrow \min(\{s \mid b_s = 1\} \cup \{q+1\})$                                  $\triangleright k \in \{1, 2, \dots, q+1\}$ 
  if  $k > K_{\min}$  then
     $i \leftarrow 1 + \langle a_1, \dots, a_p \rangle_2$                                    $\triangleright i \in \{1, 2, \dots, m\}$ 
    if  $k > K_i$  then
       $C_{K_i} \leftarrow C_{K_i} - 1$ 
       $C_k \leftarrow C_k + 1$ 
      if  $K_i = K_{\min}$  then
        while  $C_{K_{\min}} = 0$  do
           $K_{\min} \leftarrow K_{\min} + 1$ 
        end while
      end if
       $K_i \leftarrow k$ 
    end if
  end if
end procedure

```

3.6. Estimation error

Fig. 5 shows the distribution of the relative error of the estimated cardinality using Algorithm 6 compared to the true cardinality for $p = 12$ and $q = 20$. As the mean shows, the error is unbiased over the entire cardinality range. The new approach is able to accurately estimate cardinalities up to 4 billions ($\approx 2^{p+q}$) which is about an order of magnitude larger than the operating range upper bound of the raw estimator (Fig. 1).




Figure 5: Relative error of the improved raw estimator as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 12$ and $q = 20$.

The improved raw estimator beats the precision of methods that apply bias correction on the raw estimator [7, 8, 9]. Based on the simulated data we have empirically determined the bias correction function w_{corr} for the raw estimator (11) that satisfies $n = \mathbb{E}(w_{\text{corr}}(\hat{n}_{\text{raw}})|n)$ for all cardinalities. By definition, the estimator $\hat{n}'_{\text{raw}} := w_{\text{corr}}(\hat{n}_{\text{raw}})$ is unbiased and a function of the raw estimator. Its standard deviation can be compared with that of the improved raw estimator in Fig. 6. For cardinalities smaller than 10 000 the empirical bias correction approach is not very precise. This is the reason why all previous approaches had to switch over to the linear counting estimator at some point. The standard deviation of the linear counting estimator is also shown in Fig. 6. Obviously, the previous approaches cannot do better than given by the minimum of both curves for linear counting and raw estimator. In practice, the standard deviation is even larger, because the choice between both estimators must be made based on an estimate and not on the true cardinality, for which the intersection point of both curves represents the ideal transition point. In contrast, the improved raw estimator performs well over the entire cardinality range.

The new estimation algorithm also works well for other HyperLogLog configurations. First we considered configurations using a 32-bit hash function ($p + q = 32$). The relative estimation error for precisions $p = 8$, $p = 16$, and $p = 22$ are shown in Figs. 7 to 9, respectively. As expected, since $p + q = 32$ is kept constant, the operating range remains more or less the same, while the relative error decreases with increasing precision parameter p . Again, the new algorithm gives essentially unbiased estimates. Only for very high precisions, an oscillating bias becomes apparent (compare Fig. 9), that is caused by approximating the periodic function ξ by a constant (see Section 3.1).

As proposed in [7], the operating range can be extended by replacing the 32-bit hash




Figure 6: Standard deviations of the relative error of different cardinality estimators over the true cardinality for a HyperLogLog sketch with parameters $p = 12$ and $q = 20$.




Figure 7: Relative error of the improved raw estimator as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 8$ and $q = 24$.

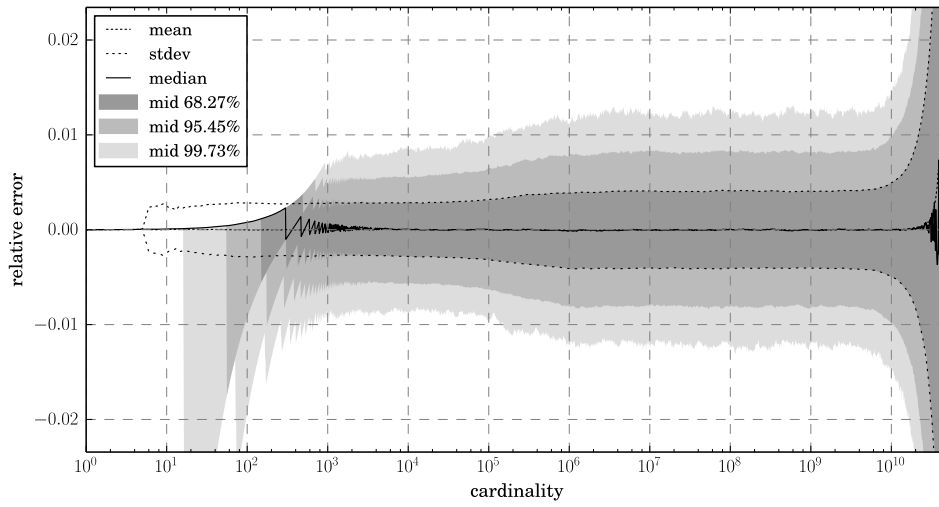


Figure 8: Relative error of the improved raw estimator as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 16$ and $q = 16$.




Figure 9: Relative error of the improved raw estimator as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 22$ and $q = 10$.




Figure 10: Relative error of the improved raw estimator as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 12$ and $q = 52$.

function by a 64-bit hash function. Fig. 10 shows the relative error for such a HyperLogLog configuration with parameters $p = 12$ and $q = 52$. The doubled hash value size shifts the maximum trackable cardinality value towards 2^{64} . As Fig. 10 shows, when compared to the 32-bit hash value case given in Fig. 5, the estimation error remains constant over the entire simulated cardinality range up to 50 billions.

We also evaluated the case $p = 12$ and $q = 14$, which is interesting, because the register values are limited to the range $[0, 15]$. As a consequence, 4 bits are sufficient for representing a single register value. This allows two registers to share a single byte, which is beneficial from a performance perspective. Nevertheless, this configuration still allows the estimation of cardinalities up to 100 millions as shown in Fig. 11, which could be enough for many applications.

3.7. Performance

To evaluate the performance of the improved raw estimation algorithm, we investigated the average computation time to obtain the cardinality from a given multiplicity vector. For different cardinality values we loaded the precalculated multiplicity vectors (see Section 3.5) of 1000 randomly generated HyperLogLog sketches into main memory. The average computation time was determined by cycling over these multiplicity vectors and passing them as input to the algorithm. For each evaluated cardinality value the average execution time was calculated after 100 cycles which corresponds to 100 000 algorithm executions for each cardinality value. The results for HyperLogLog configurations $p = 12, q = 20$ and $p = 12, q = 52$ are shown in Fig. 12. Two variants of Algorithm 6 have been evaluated for which the functions σ and τ have been either calculated on demand




Figure 11: Relative error of the improved raw estimator as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 12$ and $q = 14$.

or taken from a lookup table. All these benchmarks were carried out on an Intel Core i5-2500K clocking at 3.3 GHz.

The results show that the execution times are nearly constant for $q = 52$. Using a lookup table makes not much difference, because the on-demand calculation for σ is very fast and the calculation of τ is rarely needed due to the small probability of saturated registers, $C_{53} = 0$ usually holds for realistic cardinalities. If lookup tables are used, the computation time is as expected independent of the cardinality also for the case $q = 20$. The faster computation times for $q = 20$ compared to the $q = 52$ case can be explained by the much smaller dimension of the multiplicity vector which is equal to $q + 2$. In contrast, if functions σ and τ are calculated on demand, executions take significantly more time for larger cardinalities. The reason is that beginning at cardinality values in the order of 100 000 the probability of saturated registers increases. This makes the calculation of τ necessary, which is much more expensive than that of σ , because it requires more iterations and involves square root evaluations. However, the calculation of τ could be avoided at all, if the HyperLogLog parameters are appropriately chosen. If 2^{p+q} is much larger than the maximum expected cardinality value, the number of saturated registers will be negligible, $C_{q+1} \ll m$. The calculation of τ could be omitted in this case, because $\tau(1 - C_{q+1}/m) \approx 0$.

The numbers presented in Fig. 12 do not include the processing time to extract the multiplicity vector out of the HyperLogLog sketch, which requires a complete scan over all registers and counting the different register values into an array as demonstrated by Algorithm 4. A theoretical lower bound for this processing time can be derived using the maximum memory bandwidth of the CPU, which is 21 GB/s for an Intel Core i5-2500K. If we consider a HyperLogLog sketch with precision $p = 12$ which uses 5 bits per register,

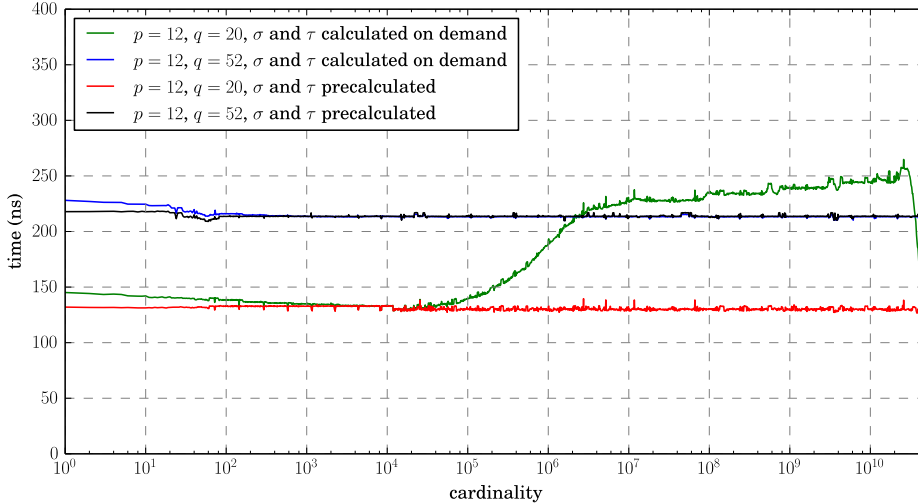


Figure 12: Average computation time as a function of the true cardinality with an Intel Core i5-2500K clocking at 3.3 GHz when estimating the cardinality from HyperLogLog sketches with parameters $p = 12$, $q = 20$ and $p = 12$, $q = 52$, respectively. Both cases, σ and τ precalculated and calculated on demand have been considered.

the total data size of the HyperLogLog sketch is 2.5 kB minimum. Consequently, the transfer time from main memory to CPU will be at least 120 ns. Having this value in mind, the presented numbers for estimating the cardinality from the multiplicity vector are quite satisfying.

4. Maximum likelihood estimation

We know from Section 2.4 that any unbiased estimator for the Poisson parameter is also an unbiased estimator for the cardinality. Moreover, we know that under suitable regularity conditions of the probability mass function the maximum likelihood estimator is asymptotically efficient [20]. This means, if the number of registers m is large enough, the maximum likelihood method should give us an unbiased estimator for the cardinality.

For HyperLogLog sketches that have been obtained without stochastic averaging (see Section 2.1) the maximum likelihood method was already previously considered [21]. The register values are statistically independent by nature in this case, which allows factorization of the joint probability mass function and a straightforward calculation of the maximum likelihood estimate. The practically more relevant case which uses stochastic averaging and which is considered in this paper, would lead to a more complicated likelihood function (compare (1)). However, the Poisson approximation makes the maximum likelihood method feasible again and we are finally able to derive another new robust and efficient cardinality estimation algorithm. Furthermore, in the course of the

derivation we will demonstrate that consequent application of the maximum likelihood method reveals that the cardinality estimate needs to be roughly proportional to the harmonic mean for intermediate cardinality values. The history of the HyperLogLog algorithm shows that the raw estimator (11) was first found after several attempts using the geometric mean [6, 19].

4.1. Log-likelihood function

Using the probability mass function of the Poisson model (5) the log-likelihood and its derivative are given by

$$\log \mathcal{L}(\lambda|\mathbf{K}) = -\frac{\lambda}{m} \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k \log\left(1 - e^{-\frac{\lambda}{m2^k}}\right) + C_{q+1} \log\left(1 - e^{-\frac{\lambda}{m2^q}}\right) \quad (38)$$

and

$$\frac{d}{d\lambda} \log \mathcal{L}(\lambda|\mathbf{K}) = -\frac{1}{\lambda} \left(\frac{\lambda}{m} \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k \frac{\frac{\lambda}{m2^k}}{1 - e^{-\frac{\lambda}{m2^k}}} + C_{q+1} \frac{\frac{\lambda}{m2^q}}{1 - e^{-\frac{\lambda}{m2^q}}} \right). \quad (39)$$

As a consequence, the maximum likelihood estimate for the Poisson parameter is given by

$$\hat{\lambda} = m\hat{x}, \quad (40)$$

if \hat{x} denotes the root of the function

$$f(x) := x \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k \frac{\frac{x}{2^k}}{1 - e^{\frac{x}{2^k}}} + C_{q+1} \frac{\frac{x}{2^q}}{1 - e^{\frac{x}{2^q}}}. \quad (41)$$

This function can also be written as

$$f(x) := x \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k h\left(\frac{x}{2^k}\right) + C_{q+1} h\left(\frac{x}{2^q}\right) - (m - C_0), \quad (42)$$

where the function $h(x)$ is defined as

$$h(x) := 1 - \frac{x}{e^x - 1}. \quad (43)$$

$h(x)$ is strictly increasing and concave as can be seen in Fig. 13. For nonnegative values this function ranges from $h(0) = 0$ to $h(x \rightarrow \infty) = 1$. Since the function $f(x)$ is also strictly increasing, it is obvious that there exists a unique root \hat{x} for which $f(\hat{x}) = 0$. The function is nonpositive at 0 since $f(0) = C_0 - m \leq 0$ and, in case $C_{q+1} < m$ which implies $\sum_{k=0}^q \frac{C_k}{2^k} > 0$, the function is at least linearly increasing. $C_{q+1} = m$ corresponds to the case with all registers equal to the maximum value $q + 1$, for which the maximum likelihood estimate would be positive infinite.

It is easy to see that the estimate $\hat{\lambda}$ remains equal or becomes larger, when inserting an element into the HyperLogLog sketch following Algorithm 1. An update potentially




Figure 13: The function $h(x)$.

changes the multiplicity vector (C_0, \dots, C_{q+1}) to $(C_0, \dots, C_i - 1, \dots, C_j + 1, \dots, C_{q+1})$ where $i < j$. Writing (42) as

$$f(x) := C_0 x + C_1 \left(h\left(\frac{x}{2^1}\right) + \frac{x}{2^1} - 1 \right) + C_2 \left(h\left(\frac{x}{2^2}\right) + \frac{x}{2^2} - 1 \right) + \dots \\ \dots + C_q \left(h\left(\frac{x}{2^q}\right) + \frac{x}{2^q} - 1 \right) + C_{q+1} \left(h\left(\frac{x}{2^q}\right) - 1 \right). \quad (44)$$

shows that the coefficient of C_i is larger than the coefficient of C_j in case $i < j$. Keeping x fixed during an update decreases $f(x)$. As a consequence, since $f(x)$ is increasing, the new root and hence the estimate must be larger than prior the update.

For the special case $q = 0$, which corresponds to the already mentioned linear counting algorithm, (41) can be solved analytically. In this case, the maximum likelihood method under the Poisson model leads directly to the linear counting estimator (14). Due to this fact we could expect that maximum likelihood estimation under the Poisson model also works well for the more general HyperLogLog case.

4.2. Inequalities for the maximum likelihood estimate

In the following we derive lower and upper bounds for \hat{x} . Applying Jensen's inequality on h in (42) gives an upper bound for $f(x)$:

$$f(x) \leq x \sum_{k=0}^q \frac{C_k}{2^k} + (m - C_0) \cdot h\left(x \cdot \frac{\sum_{k=1}^q \frac{C_k}{2^k} + \frac{C_{q+1}}{2^q}}{m - C_0}\right) - (m - C_0). \quad (45)$$

The left-hand side is zero, if \hat{x} is inserted. Resolution for \hat{x} finally gives the lower bound

$$\hat{x} \geq \frac{m - C_0}{\sum_{k=1}^q \frac{C_k}{2^k} + \frac{C_{q+1}}{2^q}} \log\left(1 + \frac{\sum_{k=1}^q \frac{C_k}{2^k} + \frac{C_{q+1}}{2^q}}{\sum_{k=0}^q \frac{C_k}{2^k}}\right). \quad (46)$$

This bound can be weakened using $\log(1+x) \geq \frac{2x}{x+2}$ for $x \geq 0$ which results in

$$\hat{x} \geq \frac{m - C_0}{C_0 + \frac{3}{2} \sum_{k=1}^q \frac{C_k}{2^k} + \frac{C_{q+1}}{2^{q+1}}}. \quad (47)$$

Using the monotonicity of h , the lower bound

$$f(x) \geq x \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k h\left(\frac{x}{2^{K'_{\max}}}\right) + C_{q+1} h\left(\frac{x}{2^{K'_{\max}}}\right) - (m - C_0) \quad (48)$$

can be found, where $K'_{\max} := \min(K_{\max}, q)$ and $K_{\max} := \max(\{k | C_k > 0\})$. Again, inserting \hat{x} and transformation gives

$$\hat{x} \leq 2^{K'_{\max}} \log\left(1 + \frac{m - C_0}{2^{K'_{\max}} \sum_{k=0}^q \frac{C_k}{2^k}}\right) \quad (49)$$

as upper bound which can be weakened using $\log(1+x) \leq x$ for $x \geq 0$

$$\hat{x} \leq \frac{m - C_0}{\sum_{k=0}^q \frac{C_k}{2^k}}. \quad (50)$$

If the HyperLogLog sketch is in the intermediate range, where $C_0 = C_{q+1} = 0$ the bounds (47) and (50) differ only by a constant factor and both are proportional to the harmonic mean of $2^{K_1}, \dots, 2^{K_m}$. Hence, consequent application of the maximum likelihood method would have directly suggested to use a cardinality estimator that is proportional to the harmonic mean without knowing the raw estimator (11) in advance.

4.3. Computation of the maximum likelihood estimate

Since f is concave and increasing, both, Newton-Raphson iteration and the secant method, will converge to the root, provided that the function is negative for the chosen starting points. We will use the secant method to derive the new cardinality estimation algorithm. Even though the secant method has the disadvantage of slower convergence, a single iteration is simpler to calculate as it does not require the evaluation of the first derivative. An iteration step of the secant method can be written as

$$x_i = x_{i-1} - (x_{i-1} - x_{i-2}) \frac{f(x_{i-1})}{f(x_{i-1}) - f(x_{i-2})}. \quad (51)$$

If we set x_0 equal to 0, for which $f(x_0) = -(m - C_0)$, and x_1 equal to one of the derived lower bounds (46) or (47), the sequence (x_0, x_1, x_2, \dots) is monotone increasing. Using the definitions

$$\Delta x_i := x_i - x_{i-1} \quad (52)$$

and

$$g(x) := f(x) + (m - C_0) = x \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k h\left(\frac{x}{2^k}\right) + C_{q+1} h\left(\frac{x}{2^q}\right) \quad (53)$$

the iteration scheme can also be written as

$$\Delta x_i = \Delta x_{i-1} \frac{(m - C_0) - g(x_{i-1})}{g(x_{i-1}) - g(x_{i-2})}, \quad (54)$$

$$x_i = x_{i-1} + \Delta x_i. \quad (55)$$

The iteration can be stopped, if $\Delta x_i \leq \delta \cdot x_i$. Since the expected statistical error for the HyperLogLog data structure scales according to $\frac{1}{\sqrt{m}}$ [6], it makes sense to choose $\delta = \frac{\varepsilon}{\sqrt{m}}$ with some constant ε . For all results presented later in Section 4.5 we used $\varepsilon = 10^{-2}$.

4.4. Maximum likelihood estimation algorithm

To get a fast cardinality estimation algorithm, it is crucial to minimize evaluation costs for (53). A couple of optimizations allow significant reduction of computational effort:

- Only a fraction of all count values C_k is nonzero. If we denote $K_{\min} := \min(\{k | C_k > 0\})$ and $K_{\max} := \max(\{k | C_k > 0\})$, it is sufficient to loop over all indices in the range $[K_{\min}, K_{\max}]$.
- The sum $\sum_{k=0}^q \frac{C_k}{2^k}$ in (53) can be precalculated and reused for all function evaluations.
- Many programming languages allow the efficient multiplication and division by any integral power of two using special functions, such as `ldexp` in C/C++ or `scalb` in Java.
- The function $h(x)$ only needs to be evaluated at points $\left\{ \frac{x}{2^{K'_{\max}}}, \frac{x}{2^{K'_{\max}-1}}, \dots, \frac{x}{2^{K'_{\min}}} \right\}$ where $K'_{\max} := \min(K_{\max}, q)$. This series corresponds to a geometric series with ratio two. A straightforward calculation using (43) is very expensive because of the exponential function. However, if we know $h\left(\frac{x}{2^{K'_{\max}}}\right)$ all other required function values can be easily obtained using the identity

$$h(4x) = \frac{x + h(2x)(1 - h(2x))}{x + (1 - h(2x))}. \quad (56)$$

This recursive formula is stable in a sense that the relative error of $h(4x)$ is smaller than that of $h(2x)$ as shown in Appendix B.

- If x is smaller than 0.5, the function $h(x)$ can be well approximated by a Taylor series around $x = 0$

$$h(x) = \frac{x}{2} - \frac{x^2}{12} + \frac{x^4}{720} - \frac{x^6}{30240} + \mathcal{O}(x^8), \quad (57)$$

which can be optimized for numerical evaluation using Estrin's scheme and $x' := \frac{x}{2}$ and $x'' := x'x'$

$$h(x) = x' - x''/3 + (x''x'') (1/45 - x''/472.5) + \mathcal{O}(x^8). \quad (58)$$

The smallest argument for which h needs to be evaluated is $\frac{x}{2^{K'_{\max}}}$. If $C_{q+1} = 0$, we can find an upper bound for the smallest argument using (49)

$$\frac{x}{2^{K'_{\max}}} \leq \frac{\hat{x}}{2^{K'_{\max}}} \leq \log \left(1 + \frac{\sum_{k=0}^{K'_{\max}} C_k}{2^{K'_{\max}} \sum_{k=0}^{K'_{\max}} \frac{C_k}{2^k}} \right) \leq \log 2 \approx 0.693. \quad (59)$$

In practice, $\frac{x}{2^{K'_{\max}}} \leq 0.5$ is satisfied most of the time as long as only a few registers are saturated, that is $C_{q+1} \ll m$. In case $\frac{x}{2^{K'_{\max}}} > 0.5$, $h\left(\frac{x}{2^\kappa}\right)$ is calculated instead with $\kappa = 2 + \lfloor \log_2(x) \rfloor$. By definition, $\frac{x}{2^\kappa} \leq 0.5$ which allows using the Taylor series approximation. $h\left(\frac{x}{2^{K'_{\max}}}\right)$ is finally obtained after $\kappa - K'_{\max}$ iterations using (56). As shown in Appendix C, a small approximation error of h does not have much impact on the error of the maximum likelihood estimate as long as most registers are not saturated.

Putting all these optimizations together finally gives the new cardinality estimation algorithm presented as Algorithm 8. The algorithm requires mainly only elementary operations. For very large cardinalities it makes sense to use the strong (46) instead of the weak lower bound (47) as second starting point for the secant method. The stronger bound is a much better approximation especially for large cardinalities, where the extra logarithm evaluation is amortized by savings in the number of iteration cycles. Therefore, the presented algorithm switches over to the stronger bound, if

$$\frac{\sum_{k=1}^q \frac{C_k}{2^k} + \frac{C_{q+1}}{2^q}}{\sum_{k=0}^q \frac{C_k}{2^k}} > 1.5 \quad (60)$$

is satisfied. The threshold value of 1.5 was found to be a reasonable choice in order to reduce the computation time for large cardinalities significantly.

4.5. Estimation error

We have investigated the estimation error of the maximum likelihood estimation algorithm for the same HyperLogLog configurations as for the improved raw estimation algorithm in Section 3.6. Figs. 14 to 19 show very similar results for same HyperLogLog parameters. What is different for the maximum likelihood estimation approach is a somewhat smaller median bias with less oscillations around zero for small cardinalities. The standard deviation of the relative error is also slightly better for the maximum likelihood estimator than for the improved raw estimator as shown in Fig. 6. Furthermore, contrary to the improved raw estimator which reveals a small oscillating bias for the mean (see Fig. 9), the maximum likelihood estimator seems to be completely unbiased (see Fig. 17).

4.6. Performance

We also measured the performance of Algorithm 8 using the same test setup as described in Section 3.7. The results for HyperLogLog configurations $p = 12, q = 20$ and $p =$

Algorithm 8 Cardinality estimation based on the maximum likelihood principle. The input is the multiplicity vector $\mathbf{C} = (C_0, \dots, C_{q+1})$ as obtained by Algorithm 4.

```

function ESTIMATECARDINALITY( $\mathbf{C}$ )
     $\triangleright \sum_{k=0}^{q+1} C_k = m$ 
    if  $C_{q+1} = m$  then
        return  $\infty$ 
    end if
     $K_{\min} \leftarrow \min(\{k | C_k > 0\})$ 
     $K'_{\min} \leftarrow \max(K_{\min}, 1)$ 
     $K_{\max} \leftarrow \max(\{k | C_k > 0\})$ 
     $K'_{\max} \leftarrow \min(K_{\max}, q)$ 
     $z \leftarrow 0$ 
    for  $k \leftarrow K'_{\max}, K'_{\min}$  do
         $z \leftarrow 0.5 \cdot z + C_k$ 
    end for
     $z \leftarrow z \cdot 2^{-K'_{\min}}$ 
     $c \leftarrow C_{q+1}$ 
    if  $q \geq 1$  then
         $c \leftarrow c + C_{K'_{\max}}$ 
    end if
     $g_{\text{prev}} \leftarrow 0$ 
     $a \leftarrow z + C_0$ 
     $b \leftarrow z + C_{q+1} \cdot 2^{-q}$ 
     $m' \leftarrow m - C_0$ 
    if  $b \leq 1.5 \cdot a$  then
         $x \leftarrow m' / (0.5 \cdot b + a)$ 
         $\triangleright$  weak lower bound (47)
    else
         $x \leftarrow m' / b \cdot \log(1 + b/a)$ 
         $\triangleright$  strong lower bound (46)
    end if

```

Algorithm 8 (continued)

```
 $\Delta x \leftarrow x$   
while  $\Delta x > x \cdot \delta$  do ▷ secant method iteration  
     $\delta = \varepsilon / \sqrt{m}, \varepsilon = 10^{-2}$   
     $\kappa \leftarrow 2 + \lfloor \log_2(x) \rfloor$   
     $x' \leftarrow x \cdot 2^{-\max(K'_{\max}, \kappa) - 1}$  ▷  $x' \in [0, 0.25]$   
     $x'' \leftarrow x' \cdot x'$   
     $h \leftarrow x' - x''/3 + (x'' \cdot x'') \cdot (1/45 - x''/472.5)$  ▷ Taylor approximation (58)  
    for  $k \leftarrow (\kappa - 1), K'_{\max}$  do  
         $h \leftarrow \frac{x' + h \cdot (1-h)}{x' + (1-h)}$  ▷ calculate  $h(\frac{x}{2^k})$ , see (56),  
        at this point  $x' = \frac{x}{2^{k+2}}$   
         $x' \leftarrow 2x'$   
    end for  
     $g \leftarrow c \cdot h$  ▷ compare (53)  
    for  $k \leftarrow (K'_{\max} - 1), K'_{\min}$  do  
         $h \leftarrow \frac{x' + h \cdot (1-h)}{x' + (1-h)}$  ▷ calculate  $h(\frac{x}{2^k})$ , see (56),  
        at this point  $x' = \frac{x}{2^{k+2}}$   
         $g \leftarrow g + C_k \cdot h$   
         $x' \leftarrow 2x'$   
    end for  
     $g \leftarrow g + x \cdot a$   
    if  $g > g_{\text{prev}} \wedge m' \geq g$  then  
         $\Delta x \leftarrow \Delta x \cdot \frac{m' - g}{g - g_{\text{prev}}}$  ▷ see (54)  
    else  
         $\Delta x \leftarrow 0$   
    end if  
     $x \leftarrow x + \Delta x$   
     $g_{\text{prev}} \leftarrow g$   
end while  
return  $m \cdot x$   
end function
```




Figure 14: Relative error of the maximum likelihood estimates as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 12$ and $q = 20$.

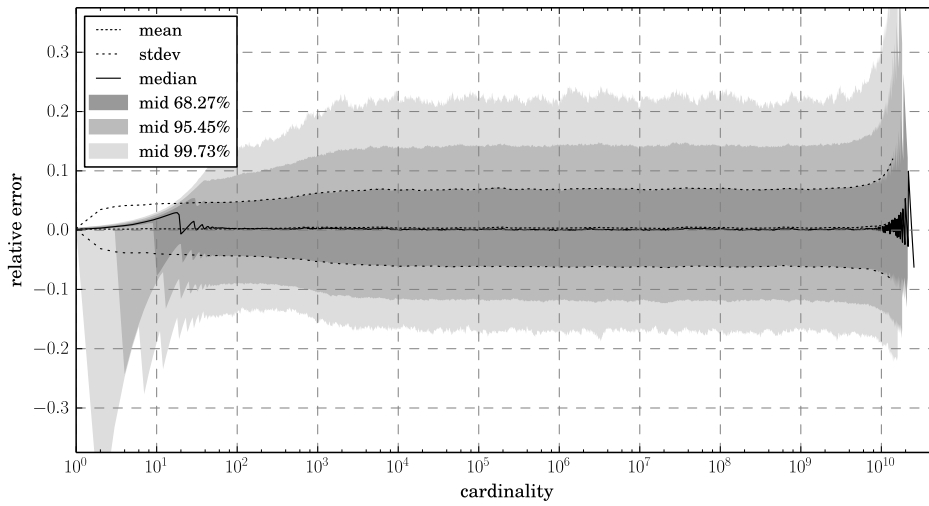


Figure 15: Relative error of the maximum likelihood estimates as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 8$ and $q = 24$.




Figure 16: Relative error of the maximum likelihood estimates as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 16$ and $q = 16$.




Figure 17: Relative error of the maximum likelihood estimates as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 22$ and $q = 10$.




Figure 18: Relative error of the maximum likelihood estimates as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 12$ and $q = 52$.




Figure 19: Relative error of the maximum likelihood estimates as a function of the true cardinality for a HyperLogLog sketch with parameters $p = 12$ and $q = 14$.

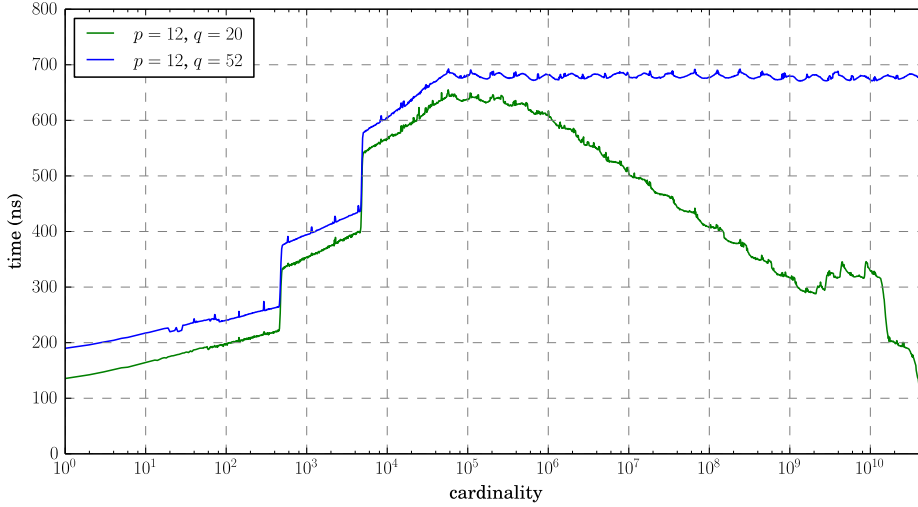


Figure 20: Average execution time of the maximum likelihood estimation algorithm as a function of the true cardinality with an Intel Core i5-2500K clocking at 3.3 GHz for HyperLogLog sketches with parameters $p = 12$, $q = 20$ and $p = 12$, $q = 52$, respectively.

$12, q = 52$ are shown in Fig. 20. The average computation time for the maximum likelihood algorithm shows a different behavior than for the improved raw estimation algorithm (compare Fig. 12). The average execution time is larger for most cardinalities, but nevertheless, it never exceeds 700 ns which is still fast enough for many applications. The steps in the chart can be explained by different numbers of iteration cycles until the secant method is stopped. For example, at a cardinality value around 5000 the average number of required cycles until the stop criterion is satisfied increases abruptly from two to three. More than three iteration cycles have never been observed for any cardinality estimate in this performance test.

5. Cardinality estimation of set intersections and complements

While the union of two sets that are represented by HyperLogLog sketches can be straightforwardly computed using Algorithm 2, the computation of cardinalities of other set operations like intersections and complements is more challenging. The conventional approach uses the inclusion-exclusion principle

$$\begin{aligned}
 |S_1 \setminus S_2| &= |S_1 \cup S_2| - |S_2|, \\
 |S_2 \setminus S_1| &= |S_1 \cup S_2| - |S_1|, \\
 |S_1 \cap S_2| &= |S_1| + |S_2| - |S_1 \cup S_2|,
 \end{aligned} \tag{61}$$

and the fact that HyperLogLog sketches can be easily merged using Algorithm 2. Unfortunately, the estimation error does not scale well for this approach. Especially for small

Jaccard indices, the relative estimation error can become very large [13]. In the worst case, the estimate could be negative without artificial restriction to nonnegative values. Therefore, it was proposed to combine HyperLogLog sketches with minwise hashing [14, 15], which improves the estimation error, even though at the expense of significant more space consumption.

It was recently pointed out without special focus on HyperLogLog sketches, that the application of the maximum likelihood method to the joint likelihood function of two probabilistic data structures, which represent the intersection operands, gives better intersection size estimates [22]. For HyperLogLog sketches recorded without stochastic averaging (compare Section 2.1) this was shown in [15], where also an algorithm based on the maximum likelihood principle was outlined. However, in practice, HyperLogLog sketches are recorded using stochastic averaging because of the much cheaper element insertions. Motivated by the good results we have obtained for a single HyperLogLog sketch using the maximum likelihood method in combination with the Poisson approximation, we are tempted to apply this approach also for the estimation of set operation result sizes.

Assume two given HyperLogLog sketches with register values \mathbf{K}_1 and \mathbf{K}_2 representing the sets S_1 and S_2 , respectively. The goal is to find estimates for the cardinalities of the pairwise disjoint sets $X = S_1 \cap S_2$, $A = S_1 \setminus S_2$, and $B = S_2 \setminus S_1$. The Poisson approximation allows us to assume that pairwise distinct elements are inserted into the HyperLogLog sketches representing S_1 and S_2 at rates λ_a and λ_b , respectively. Furthermore, we assume that further unique elements are inserted into both HyperLogLog sketches simultaneously at rate λ_x . We expect that good estimates $\hat{\lambda}_a$, $\hat{\lambda}_b$, and $\hat{\lambda}_x$ for the rates are also good estimates for the cardinalities of A , B , and X .

5.1. Joint log-likelihood function

In order to get maximum likelihood estimators for $\hat{\lambda}_a$, $\hat{\lambda}_b$, and $\hat{\lambda}_x$ we need to derive the joint probability distribution of two HyperLogLog sketches. Under the Poisson model the register values are independent and identically distributed. Therefore, we first derive the joint probability distribution for a single register that has value K_1 in the first HyperLogLog sketch representing S_1 and value K_2 in the second HyperLogLog sketch representing S_2 .

The HyperLogLog sketch that represents S_1 can be thought to be constructed from two HyperLogLog sketches representing A and X and merging both using Algorithm 2. Analogously, the HyperLogLog sketch for S_2 could have been obtained from sketches for B and X . Let K_a , K_b , and K_x be the value of the considered register in the HyperLogLog sketch representing A , B , and X , respectively. The corresponding values in sketches for S_1 and S_2 are given by

$$K_1 = \max(K_a, K_x), \quad K_2 = \max(K_b, K_x). \quad (62)$$

Their joint cumulative probability function is given as

$$\begin{aligned}
P(K_1 \leq k_1 \wedge K_2 \leq k_2) &= P(\max(K_a, K_x) \leq k_1 \wedge \max(K_b, K_x) \leq k_2) \\
&= P(K_a \leq k_1 \wedge K_b \leq k_2 \wedge K_x \leq \min(k_1, k_2)) \\
&= P(K_a \leq k_1) P(K_b \leq k_2) P(K_x \leq \min(k_1, k_2)). \tag{63}
\end{aligned}$$

Here the last transformation used the independence of K_a , K_b , and K_x , because by definition, the sets A , B , and X are pairwise disjoint. Furthermore, under the Poisson model K_a , K_b , and K_x can be described by (6). If we take into account that pairwise distinct elements are added to A , B , and X at rates λ_a , λ_b , and λ_x , respectively, the probability that a certain register has a value less than or equal to k_1 in the first HyperLogLog sketch and simultaneously a value less than or equal to k_2 in the second one can be written as

$$P(K_1 \leq k_1 \wedge K_2 \leq k_2) = \begin{cases} 0 & k_1 < 0 \vee k_2 < 0 \\ e^{-\frac{\lambda_a}{m2^{k_1}} - \frac{\lambda_b}{m2^{k_2}} - \frac{\lambda_x}{m2^{\min(k_1, k_2)}}} & 0 \leq k_1 \leq q \wedge 0 \leq k_2 \leq q \\ e^{-\frac{\lambda_b + \lambda_x}{m2^{k_2}}} & 0 \leq k_2 \leq q < k_1 \\ e^{-\frac{\lambda_a + \lambda_x}{m2^{k_1}}} & 0 \leq k_1 \leq q < k_2 \\ 1 & q < k_1 \wedge q < k_2. \end{cases} \tag{64}$$

The joint probability mass function for both register values can be calculated using

$$\begin{aligned}
\rho(k_1, k_2) &= P(K_1 \leq k_1 \wedge K_2 \leq k_2) - P(K_1 \leq k_1 - 1 \wedge K_2 \leq k_2) \\
&\quad - P(K_1 \leq k_1 \wedge K_2 \leq k_2 - 1) + P(K_1 \leq k_1 - 1 \wedge K_2 \leq k_2 - 1), \tag{65}
\end{aligned}$$

which finally gives

$$\rho(k_1, k_2) = \begin{cases} e^{-\frac{\lambda_a + \lambda_x}{m} - \frac{\lambda_b}{m2^{k_2}}} \left(1 - e^{-\frac{\lambda_b}{m2^{k_2}}}\right) & 0 = k_1 < k_2 \leq q \\ e^{-\frac{\lambda_a + \lambda_x}{m}} \left(1 - e^{-\frac{\lambda_b}{m2^q}}\right) & 0 = k_1 < k_2 = q + 1 \\ e^{-\frac{\lambda_a + \lambda_x}{m2^{k_1}} - \frac{\lambda_b}{m2^{k_2}}} \left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^{k_1}}}\right) \left(1 - e^{-\frac{\lambda_b}{m2^{k_2}}}\right) & 1 \leq k_1 < k_2 \leq q \\ e^{-\frac{\lambda_a + \lambda_x}{m2^{k_1}} - \frac{\lambda_a + \lambda_x}{m2^{k_1}}} \left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^{k_1}}}\right) \left(1 - e^{-\frac{\lambda_b}{m2^q}}\right) & 1 \leq k_1 < k_2 = q + 1 \\ e^{-\frac{\lambda_b + \lambda_x}{m} - \frac{\lambda_a}{m2^{k_1}}} \left(1 - e^{-\frac{\lambda_a}{m2^{k_1}}}\right) & 0 = k_2 < k_1 \leq q \\ e^{-\frac{\lambda_b + \lambda_x}{m}} \left(1 - e^{-\frac{\lambda_a}{m2^q}}\right) & 0 = k_2 < k_1 = q + 1 \\ e^{-\frac{\lambda_b + \lambda_x}{m2^{k_2}} - \frac{\lambda_a}{m2^{k_1}}} \left(1 - e^{-\frac{\lambda_b + \lambda_x}{m2^{k_2}}}\right) \left(1 - e^{-\frac{\lambda_a}{m2^{k_1}}}\right) & 1 \leq k_2 < k_1 \leq q \\ e^{-\frac{\lambda_b + \lambda_x}{m2^{k_2}} - \frac{\lambda_b + \lambda_x}{m2^{k_2}}} \left(1 - e^{-\frac{\lambda_b + \lambda_x}{m2^{k_2}}}\right) \left(1 - e^{-\frac{\lambda_a}{m2^q}}\right) & 1 \leq k_2 < k_1 = q + 1 \\ e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m}} & 0 = k_1 = k_2 \\ e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m2^k}} \left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^k}} - e^{-\frac{\lambda_b + \lambda_x}{m2^k}} + e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m2^k}}\right) & 1 \leq k_1 = k_2 = k \leq q \\ 1 - e^{-\frac{\lambda_a + \lambda_x}{m2^q}} - e^{-\frac{\lambda_b + \lambda_x}{m2^q}} + e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m2^q}} & k_1 = k_2 = q + 1. \end{cases} \quad (66)$$

The logarithm of the joint probability mass function can be written using Iverson bracket notation ($[\text{true}] := 1$, $[\text{false}] := 0$) as

$$\begin{aligned} \log(\rho(k_1, k_2)) = & \log\left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^{k_1}}}\right) [1 \leq k_1 < k_2] + \log\left(1 - e^{-\frac{\lambda_a}{m2^{\min(k_1, q)}}}\right) [k_2 < k_1] \\ & + \log\left(1 - e^{-\frac{\lambda_b + \lambda_x}{m2^{k_2}}}\right) [1 \leq k_2 < k_1] + \log\left(1 - e^{-\frac{\lambda_b}{m2^{\min(k_2, q)}}}\right) [k_1 < k_2] \\ & + \log\left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^{\min(k_1, q)}}} - e^{-\frac{\lambda_b + \lambda_x}{m2^{\min(k_1, q)}}} + e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m2^{\min(k_1, q)}}}\right) [1 \leq k_1 = k_2] \\ & - \frac{\lambda_a}{m2^{k_1}} [k_1 \leq q] - \frac{\lambda_b}{m2^{k_2}} [k_2 \leq q] - \frac{\lambda_x}{m2^{\min(k_1, k_2)}} [k_1 \leq q \vee k_2 \leq q]. \end{aligned} \quad (67)$$

Since the values for different registers are independent under the Poisson model, we are now able to write the joint probability mass function for all registers in both Hyper-LogLog sketches

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \prod_{i=1}^m \rho(k_{1i}, k_{2i}). \quad (68)$$

The maximum likelihood estimates $\hat{\lambda}_a$, $\hat{\lambda}_b$, and $\hat{\lambda}_x$ are obtained by maximization of

the log-likelihood function given by

$$\log \mathcal{L}(\lambda_a, \lambda_b, \lambda_x | \mathbf{K}_1, \mathbf{K}_2) = \sum_{i=1}^m \log(\rho(K_{1i}, K_{2i})). \quad (69)$$

Insertion of (67) results in

$$\begin{aligned} \log \mathcal{L}(\lambda_a, \lambda_b, \lambda_x | \mathbf{K}_1, \mathbf{K}_2) = & \\ & \sum_{k=1}^q \log\left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^k}}\right) C_{1k}^< + \log\left(1 - e^{-\frac{\lambda_b + \lambda_x}{m2^k}}\right) C_{2k}^< \\ & + \sum_{k=1}^{q+1} \log\left(1 - e^{-\frac{\lambda_a}{m2^{\min(k,q)}}}\right) C_{1k}^> + \log\left(1 - e^{-\frac{\lambda_b}{m2^{\min(k,q)}}}\right) C_{2k}^> \\ & + \sum_{k=1}^{q+1} \log\left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^{\min(k,q)}}} - e^{-\frac{\lambda_b + \lambda_x}{m2^{\min(k,q)}}} + e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m2^{\min(k,q)}}}\right) C_k^= \\ & - \frac{\lambda_a}{m} \sum_{k=0}^q \frac{C_{1k}^< + C_k^= + C_{1k}^>}{2^k} - \frac{\lambda_b}{m} \sum_{k=0}^q \frac{C_{2k}^< + C_k^= + C_{2k}^>}{2^k} - \frac{\lambda_x}{m} \sum_{k=0}^q \frac{C_{1k}^< + C_k^= + C_{2k}^<}{2^k}, \end{aligned} \quad (70)$$

where $C_{1k}^<$, $C_{1k}^>$, $C_{2k}^<$, $C_{2k}^>$, and $C_k^=$ are defined as

$$\begin{aligned} C_{1k}^< &:= |\{i | k = K_{1i} < K_{2i}\}|, \\ C_{1k}^> &:= |\{i | k = K_{1i} > K_{2i}\}|, \\ C_{2k}^< &:= |\{i | k = K_{2i} < K_{1i}\}|, \\ C_{2k}^> &:= |\{i | k = K_{2i} > K_{1i}\}|, \\ C_k^= &:= |\{i | k = K_{1i} = K_{2i}\}|. \end{aligned} \quad (71)$$

These $5(q+2)$ values represent a sufficient statistic for estimating λ_a , λ_b , and λ_x . Actually, the number of values could be further reduced, because of the invariants $C_{10}^> = C_{20}^> = C_{1,q+1}^< = C_{2,q+1}^< = 0$, $\sum_{k=0}^{q+1} C_{1k}^< = \sum_{k=0}^{q+1} C_{2k}^>$, $\sum_{k=0}^{q+1} C_{2k}^< = \sum_{k=0}^{q+1} C_{1k}^>$, and $\sum_{k=0}^{q+1} C_{1k}^< + C_k^= + C_{2k}^< = m$.

Since (70) is a generalization of (38) for two HyperLogLog sketches, the log-likelihood function for a single HyperLogLog sketch can be obtained by considering special cases. For example, assume $\lambda_x = 0$, which means that both sketches represent disjoint sets, (70) can be split into the sum of two unary functions with parameters λ_a and λ_b each of which correspond to (38) as expected. Or, consider two sketches representing identical sets. In this case all registers are equal, which means that $C_{1k}^> = C_{2k}^> = C_{1k}^< = C_{2k}^< = 0$ for all k , and therefore the maximum likelihood method yields $\hat{\lambda}_a = \hat{\lambda}_b = 0$ and the value for $\hat{\lambda}_x$ will be equal to the single HyperLogLog maximum likelihood estimate (38).

The log-likelihood function (70) does not always have a strict global maximum point. For example, if all register values of the first HyperLogLog sketch are larger than the corresponding values in the second HyperLogLog sketch, that is $C_{1k}^< = C_k^= = C_{2k}^> = 0$

for all k , the function can be rewritten as sum of two functions, one dependent on λ_a and the other dependent on $\lambda_b + \lambda_x$. The maximum is obtained, if $\lambda_a = \hat{\lambda}_1$ and $\lambda_b + \lambda_x = \hat{\lambda}_2$. Here $\hat{\lambda}_1$ and $\hat{\lambda}_2$ are the maximum likelihood cardinality estimates for the first and second HyperLogLog sketch, respectively. This means that the maximum likelihood method makes no clear statement about the intersection size in this case. The estimate for λ_x could be anything between 0 and $\hat{\lambda}_2$. For comparison, the inclusion-exclusion approach would give $\hat{\lambda}_2$ as estimate for λ_x . This result is questionable, because there is no evidence that sets S_1 and S_2 really have common elements in this case.

The inclusion-exclusion method does not use all the available information given by the sufficient statistic (71), because the estimator is a function of the three vectors $(\mathbf{C}_1^< + \mathbf{C}^= + \mathbf{C}_1^>)$, $(\mathbf{C}_2^< + \mathbf{C}^= + \mathbf{C}_2^>)$, and $(\mathbf{C}_1^> + \mathbf{C}^= + \mathbf{C}_2^>)$. In contrast, the maximum likelihood method uses more information as it incorporates all the individual values of the sufficient statistic.

5.2. Computation of the maximum likelihood estimates

The maximum likelihood estimates can be obtained by maximizing (70). Since the three parameters are all nonnegative, this is a constrained optimization problem. In order to get rid of these constraints, we use the transformation $\lambda = e^\varphi$. This mapping has also the nice property that relative accuracy limits are translated into absolute ones, because $\Delta\varphi = \Delta\lambda/\lambda$. Many optimization algorithm implementations allow the definition of absolute limits rather than relative ones.

The transformed log-likelihood function can be written as

$$\begin{aligned}
f(\varphi_a, \varphi_b, \varphi_x) &:= \log \mathcal{L}(e^{\varphi_a}, e^{\varphi_b}, e^{\varphi_x} | \mathbf{K}_1, \mathbf{K}_2) = \\
&+ \sum_{k=1}^q C_{1k}^< \log(z_{xk} + y_{xk}z_{ak}) + C_{2k}^< \log(z_{xk} + y_{xk}z_{bk}) \\
&+ \sum_{k=1}^q C_{1k}^> \log(z_{ak}) + C_{2k}^> \log(z_{bk}) + C_k^= \log(z_{xk} + y_{xk}z_{ak}z_{bk}) \\
&+ C_{1,q+1}^> \log(z_{aq}) + C_{2,q+1}^> \log(z_{bq}) + C_{q+1}^= \log(z_{xq} + y_{xq}z_{aq}z_{bq}) \\
&- \sum_{k=0}^q (C_{1k}^< + C_k^= + C_{1k}^>) x_{ak} + (C_{2k}^< + C_k^= + C_{2k}^>) x_{bk} + (C_{1k}^< + C_k^= + C_{2k}^<) x_{xk}.
\end{aligned} \tag{72}$$

Here we introduced the following expressions for simplification:

$$x_{*k} := \frac{e^{\varphi_*}}{m2^k}, \quad y_{*k} := e^{-x_{*k}}, \quad z_{*k} := 1 - y_{*k}. \tag{73}$$

Quasi-Newton methods are commonly used to find the maximum of multi-dimensional functions. They require the calculation of the gradient $\nabla f = \left(\frac{\partial f}{\partial \varphi_a}, \frac{\partial f}{\partial \varphi_b}, \frac{\partial f}{\partial \varphi_x} \right)$ which is, using

$$\frac{\partial x_{*k}}{\partial \varphi_*} = x_{*k}, \quad \frac{\partial y_{*k}}{\partial \varphi_*} = -x_{*k}y_{*k}, \quad \frac{\partial z_{*k}}{\partial \varphi_*} = x_{*k}y_{*k}, \tag{74}$$

given by

$$\begin{aligned}
\frac{\partial f}{\partial \varphi_a} &= \sum_{k=1}^q C_{1k}^{<} \frac{y_{xk} x_{ak} y_{ak}}{z_{xk} + y_{xk} z_{ak}} + C_k^{=} \frac{y_{xk} x_{ak} y_{ak} z_{bk}}{z_{xk} + y_{xk} z_{ak} z_{bk}} + C_{1k}^{>} \frac{x_{ak} y_{ak}}{z_{ak}} \\
&\quad + C_{q+1}^{=} \frac{y_{xq} x_{aq} y_{aq} z_{bq}}{z_{xq} + y_{xq} z_{aq} z_{bq}} + C_{1,q+1}^{>} \frac{x_{aq} y_{aq}}{z_{aq}} - \sum_{k=0}^q (C_{1k}^{<} + C_k^{=} + C_{1k}^{>}) x_{ak}, \\
\frac{\partial f}{\partial \varphi_b} &= \sum_{k=1}^q C_{2k}^{<} \frac{y_{xk} x_{bk} y_{bk}}{z_{xk} + y_{xk} z_{bk}} + C_k^{=} \frac{y_{xk} z_{ak} x_{bk} y_{bk}}{z_{xk} + y_{xk} z_{ak} z_{bk}} + C_{2k}^{>} \frac{x_{bk} y_{bk}}{z_{bk}} \\
&\quad + C_{q+1}^{=} \frac{y_{xq} z_{aq} x_{bq} y_{bq}}{z_{xq} + y_{xq} z_{aq} z_{bq}} + C_{2,q+1}^{>} \frac{x_{bq} y_{bq}}{z_{bq}} - \sum_{k=0}^q (C_{2k}^{<} + C_k^{=} + C_{2k}^{>}) x_{bk}, \\
\frac{\partial f}{\partial \varphi_x} &= \sum_{k=1}^q C_{1k}^{<} \frac{x_{xk} y_{xk} y_{ak}}{z_{xk} + y_{xk} z_{ak}} + C_k^{=} \frac{x_{xk} y_{xk} (y_{ak} + z_{ak} y_{bk})}{z_{xk} + y_{xk} z_{bk} z_{ak}} + C_{2k}^{<} \frac{x_{xk} y_{xk} y_{bk}}{z_{xk} + y_{xk} z_{bk}} \\
&\quad + C_{q+1}^{=} \frac{x_{xq} y_{xq} (y_{aq} + z_{aq} y_{bq})}{z_{xq} + y_{xq} z_{aq} z_{bq}} - \sum_{k=0}^q (C_{1k}^{<} + C_k^{=} + C_{2k}^{<}) x_{xk}.
\end{aligned} \tag{75}$$

The calculation of (72) and its derivatives requires some care when calculating y_{*k} and z_{*k} . Since x_{*k} is nonnegative, we have $y_{*k}, z_{*k} \in [0, 1]$. In order to reduce the numerical error of z_{*k} for small x_{*k} , it is essential to use the function $\text{expm1}(x) := e^x - 1$ that is available in most programming languages. If x_{*k} is smaller than $\log(2)$, we calculate y_{*k} and z_{*k} using $z_{*k} = -\text{expm1}(-x_{*k})$ and $y_{*k} = 1 - z_{*k}$, respectively, and not as defined in (73). In this way the numerical error of both, y_{*k} and z_{*k} , is minimized and still only a single exponential function needs to be evaluated. Apart from that, the numerical evaluation of (72) is straightforward. Apparently, the arguments of all logarithms are in range $[0, 1]$. The case that some argument vanishes, which would cause the logarithm to be negative infinite, does not occur in practice. Consider for example the logarithm evaluation associated with $C_{1k}^{<}$ which is only relevant if $C_{1k}^{<} > 0$. In this case however, we can be certain that the cardinality of $A \cup X$ is at least 1. Therefore, it is expected that at least one of the two maximum likelihood estimates $\hat{\lambda}_a$ and $\hat{\lambda}_x$ is at least in the order of 1. If the optimization algorithm starts with appropriate initial values, function evaluations for which $\max(\lambda_a, \lambda_x) \ll 1$ are not expected. Furthermore, the argument of the logarithm can be bounded by $z_{xk} + y_{xk} z_{ak} \geq \max(z_{xk}, z_{ak}) = 1 - \exp\left(-\frac{\max(\lambda_a, \lambda_x)}{m2^k}\right) \geq \min\left(\frac{1}{2}, \frac{\max(\lambda_a, \lambda_x)}{m2^{k+1}}\right)$. For the last inequality we used $1 - e^{-x} \geq \frac{1}{2} \min(1, x)$. The derived lower bound shows that the argument of the logarithm is large enough to be accurately represented by double-precision floating-point numbers, provided that λ_a and λ_x are not both substantially smaller than 1. Similar argumentation holds for all other logarithmic terms and also all divisions that appear in (75).

Algorithm 9 demonstrates the calculation of the estimates given the register values \mathbf{K}_1 and \mathbf{K}_2 of both HyperLogLog sketches. First, the sufficient statistic consisting of $\mathbf{C}_1^{>}$, $\mathbf{C}_1^{<}$, $\mathbf{C}_2^{>}$, $\mathbf{C}_2^{<}$, and $\mathbf{C}^{=}$ is extracted. Next, a case is distinguished, where all registers have a value equal to zero in at least one of both HyperLogLog sketches, that is

$C_{10}^< + C_0^= + C_{20}^< = m$. In this case it is certain that the HyperLogLog sketches represent disjoint sets and their corresponding cardinality estimates can be used for $\hat{\lambda}_a$ and $\hat{\lambda}_b$, respectively.

For the general case, the three-dimensional function $f(\varphi_a, \varphi_b, \varphi_x)$ needs to be optimized numerically. A very popular algorithm for such nonlinear optimization problems is the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [23]. In particular, we used the implementation provided by the Dlib C++ library [24]. Good initial guess values are important to ensure fast convergence for the optimization algorithm. An obvious choice are the cardinality estimates obtained by application of the inclusion-exclusion principle (61). However, in order to ensure that their logarithms are all defined, we require that the initial values are not smaller than 1. Within the optimization loop the function f and its gradient ∇f need to be evaluated using (72) and (75) as needed by the optimization algorithm. The optimization loop is continued as long as the absolute change of at least one parameter is larger than a predefined threshold δ . The threshold is again set proportional to $1/\sqrt{m}$ (compare Section 4.3). For the results presented below we have used $\delta = \varepsilon/\sqrt{m}$ with $\varepsilon = 10^{-2}$.

5.3. Results

To evaluate the estimation error of the new joint cardinality estimation approach we applied Algorithm 9 to HyperLogLog sketch pairs with parameters $p = 20$ and $q = 44$ and with known cardinalities $|A|$, $|B|$, and $|X|$. HyperLogLog sketches with predefined cardinalities can be quickly generated by taking a precalculated multiplicity vector with known cardinality, setting the register values accordingly, and shuffling them. Using this approach we generated three different HyperLogLog sketches with known cardinalities $|A|$, $|B|$, and $|X|$. Then we merged the HyperLogLog sketches representing the disjoint sets A and X using Algorithm 2 in order to get a sketch for $S_1 = A \cup X$. A second sketch for $S_2 = A \cup X$ was built analogously. Since three independent multiplicity vectors are needed to construct a single example and we had simulated 10 000 different HyperLogLog sketches previously as described in Section 3.5, we were able to construct 3333 different independent HyperLogLog pairs for a given set of known true cardinalities $|A|$, $|B|$, and $|X|$.

Table 1 lists all the different cardinality configurations for which we have evaluated the joint cardinality estimation algorithm. Among the considered cases there are also cardinalities that are small compared to the number of registers in order to prove that the new approach also covers the small cardinality range where many register values are equal to zero. The last column of Table 1 shows the average number of iterations of the BFGS algorithm until the stop criterion was satisfied. Each iteration step involved a function (72) and a gradient (75) evaluation.

Tables 2 to 4 compare the relative estimation error for cardinalities $|A|$, $|B|$, and $|X|$ to that of the conventional approach using single sketch cardinality estimation together with (61). The mean, the standard deviation, and the root-mean-square error have been calculated from the cardinality estimates of 3333 examples. Furthermore, we calculated an improvement factor which represents the root-mean-square error ratio between both

Algorithm 9 Joint cardinality estimation.

```

function ESTIMATECARDINALITIES(  $\mathbf{K}_1, \mathbf{K}_2$  )
     $\mathbf{C}_1^> \leftarrow (0, \dots, 0)$ 
     $\mathbf{C}_1^< \leftarrow (0, \dots, 0)$ 
     $\mathbf{C}_2^> \leftarrow (0, \dots, 0)$ 
     $\mathbf{C}_2^< \leftarrow (0, \dots, 0)$ 
     $\mathbf{C}^= \leftarrow (0, \dots, 0)$ 
    for  $i \leftarrow 1, m$  do
        if  $K_{1i} < K_{2i}$  then
             $C_{1K_{1i}}^< \leftarrow C_{1K_{1i}}^< + 1$ 
             $C_{2K_{2i}}^> \leftarrow C_{2K_{2i}}^> + 1$ 
        else if  $K_{1i} > K_{2i}$  then
             $C_{1K_{1i}}^> \leftarrow C_{1K_{1i}}^> + 1$ 
             $C_{2K_{2i}}^< \leftarrow C_{2K_{2i}}^< + 1$ 
        else
             $C_{K_{1i}}^= \leftarrow C_{K_{1i}}^= + 1$ 
        end if
    end for
     $\hat{\lambda}_{ax} \leftarrow \text{ESTIMATECARDINALITY}(\mathbf{C}_1^< + \mathbf{C}^= + \mathbf{C}_1^>)$ 
     $\hat{\lambda}_{bx} \leftarrow \text{ESTIMATECARDINALITY}(\mathbf{C}_2^< + \mathbf{C}^= + \mathbf{C}_2^>)$ 
    if  $C_{10}^< + C_0^= + C_{20}^< = m$  then
         $\hat{\lambda}_a \leftarrow \hat{\lambda}_{ax}, \hat{\lambda}_b \leftarrow \hat{\lambda}_{bx}, \hat{\lambda}_x \leftarrow 0$ 
        return  $(\hat{\lambda}_a, \hat{\lambda}_b, \hat{\lambda}_x)$ 
    end if
     $\hat{\lambda}_{abx} \leftarrow \text{ESTIMATECARDINALITY}(\mathbf{C}_1^> + \mathbf{C}^= + \mathbf{C}_2^>)$ 
     $\varphi_a \leftarrow \log(\max(1, \hat{\lambda}_{abx} - \hat{\lambda}_{bx}))$ 
     $\varphi_b \leftarrow \log(\max(1, \hat{\lambda}_{abx} - \hat{\lambda}_{ax}))$ 
     $\varphi_x \leftarrow \log(\max(1, \hat{\lambda}_{ax} + \hat{\lambda}_{bx} - \hat{\lambda}_{abx}))$ 
    repeat
         $\varphi'_a \leftarrow \varphi_a, \varphi'_b \leftarrow \varphi_b, \varphi'_x \leftarrow \varphi_x$ 
         $\varphi_a, \varphi_b, \varphi_x \leftarrow \text{OPTIMIZATIONSTEP}(\varphi_a, \varphi_b, \varphi_x)$ 
    until  $\max(|\varphi_a - \varphi'_a|, |\varphi_b - \varphi'_b|, |\varphi_x - \varphi'_x|) \leq \delta$ 
     $\hat{\lambda}_a \leftarrow e^{\varphi_a}, \hat{\lambda}_b \leftarrow e^{\varphi_b}, \hat{\lambda}_x \leftarrow e^{\varphi_x}$ 
    return  $(\hat{\lambda}_a, \hat{\lambda}_b, \hat{\lambda}_x)$ 
end function

```

 $\triangleright \mathbf{K}_1, \mathbf{K}_2 \in \{0, 1, \dots, q+1\}^m$
 $\triangleright \mathbf{C}_1^> = (C_{10}^>, \dots, C_{1,q+1}^>)$
 $\triangleright \mathbf{C}_1^< = (C_{10}^<, \dots, C_{1,q+1}^<)$
 $\triangleright \mathbf{C}_2^> = (C_{20}^>, \dots, C_{2,q+1}^>)$
 $\triangleright \mathbf{C}_2^< = (C_{20}^<, \dots, C_{2,q+1}^<)$
 $\triangleright \mathbf{C}^= = (C_0^=, \dots, C_{q+1}^=)$
 \triangleright use Algorithm 8

 $\triangleright \Leftrightarrow \min(K_{1i}, K_{2i}) = 0 \forall i$
 \triangleright start optimization

 \triangleright calculate f and ∇f using (72) and (75) as needed by optimizer

 \triangleright stop criterion

 $\delta = \varepsilon / \sqrt{m}, \varepsilon = 10^{-2}$

approaches. Since we only observed values greater than one, the new maximum likelihood estimation approach gives better estimates for all investigated cases. For some cases the improvement factor is even clearly greater than two. Due to the square root scaling of the error, this means that we would need four times more registers to get the same error when using the conventional approach. As the results suggest the joint estimation algorithm works well over the entire cardinality range without the need of special handling of small cardinalities.

We also investigated, if the new approach could give better estimates for the union operation. The conventional approach merges both HyperLogLog sketches using Algorithm 2 and estimates the union size using single sketch cardinality estimation. The joint cardinality estimation algorithm provides another opportunity. The union can also be estimated by simply summing up the three estimates for sets $|A|$, $|B|$, and $|X|$. The corresponding results are shown in Table 5. As can be clearly seen, the joint estimation algorithm is also able to improve the cardinality estimation of unions by a significant amount.

6. Future work

As described in Section 2.4 an unbiased estimator for the rate in the Poisson model, is also unbiased estimator for the cardinality. We have shown that this approach works well for the improved raw estimator as well as for the maximum likelihood estimator even though both are only approximately unbiased. Therefore, it would be interesting what conditions on an approximately unbiased Poisson rate estimator are sufficient to guarantee approximate unbiasedness, if used as cardinality estimator.

Using the maximum likelihood method we have been able to improve the cardinality estimates for the results of set operations between two HyperLogLog sketches. Unfortunately, joint cardinality estimation is much more expensive than for a single HyperLogLog sketch, because it requires maximization of a multi-dimensional function. Since we have found the improved raw estimator which is almost as precise as the maximum likelihood estimator for the single HyperLogLog case, we could imagine that there also exists a faster algorithm for the two HyperLogLog case. It is expected that such a new algorithm makes use of all the information given by the sufficient statistic (71).

The presented maximum likelihood method could also be used to estimate the cardinality of set operations between more than two HyperLogLog sketches. However, further research is necessary to determine, if this is feasible from a practical point of view. As for the inclusion-exclusion principle, the effort would scale at least exponentially with the number of involved HyperLogLog sketches.

The maximum likelihood method can also be used to estimate distance measures such as the Jaccard distance of two sets that are represented as HyperLogLog sketches. This directly leads to the question whether the HyperLogLog algorithm could be used for locality-sensitive hashing [25, 26]. Various locality-sensitive hashing algorithms have been proposed in the past. Among the most popular ones are the SimHash [27] and the minwise hashing [28] algorithms whose hash collision probabilities are a function of the

Table 1: List of true cardinalities of the pairwise disjoint sets A , B , and X , for which maximum likelihood estimation from two HyperLogLog sketches with parameters $p = 20$ and $q = 44$ representing sets $S_1 = A \cup X$ and $S_2 = B \cup X$ was investigated. The last column shows the average number of iteration cycles until the stop criterion of the BFGS algorithm was satisfied.

#	true cardinalities			Jaccard index	cardinality ratio	avg number iterations
	$ A $	$ B $	$ X $	$ S_1 \cap S_2 / S_1 \cup S_2 $	$ S_1 / S_2 $	
1	1 598 728 349	251 188 153	2 742 179	1.480E-3	6.365	36.207
2	13 049 219 720	5 945 604 476	103 608 860	5.425E-3	2.195	38.736
3	144 199	1608	3457	2.316E-2	89.676	25.777
4	36 085 381	8 525 744	3 516 662	7.307E-2	4.233	42.431
5	652 906 563	42 736 075	23 996 891	3.335E-2	15.278	42.134
6	74 772	2617	235	3.027E-3	28.572	22.507
7	148 568	12 597	7147	4.246E-2	11.794	30.143
8	3 837 563 445	277 467 992	104 644 949	2.480E-2	13.831	42.335
9	84 256	1624	402	4.659E-3	51.882	22.287
10	122 704 431	36 085 381	32 994 301	1.720E-1	3.400	42.099
11	22 160 727	6 987 227	586 521	1.973E-2	3.172	41.717
12	48 156 078	4 333 904	1 313 150	2.441E-2	11.112	42.234
13	184 928	117 006	27 370	8.311E-2	1.581	34.904
14	17 149	2540	1530	7.211E-2	6.752	21.734
15	19 279 012	2 939 988	873 249	3.782E-2	6.558	42.000
16	1 602 292	23 812	2277	1.398E-3	67.289	30.128
17	99 786	1923	2805	2.684E-2	51.891	25.060
18	61 756 895	744 726	974 256	1.535E-2	82.926	41.038
19	10 277 123 441	352 310 724	1 783 650 614	1.437E-1	29.171	40.825
20	134 200 029	11 958 008	15 335 333	9.496E-2	11.223	41.642
21	110 226	1205	691	6.163E-3	91.474	22.834
22	773 240 139	13 474 584	20 465 059	2.535E-2	57.385	41.597
23	307 179	5046	8549	2.665E-2	60.876	29.473
24	74 032	16 317	6343	6.560E-2	4.537	29.328
25	50 221	4898	958	1.708E-2	10.253	23.960
26	680 932	24 290	109 135	1.340E-1	28.033	35.867
27	1 097 816	452 822	48 262	3.018E-2	2.424	39.268
28	84 256	2919	2210	2.472E-2	28.865	25.247
29	13 049 219 720	721 215 037	145 319 400	1.044E-2	18.093	41.151
30	59 347 163	1 287 276	286 512	4.703E-3	46.103	39.462
31	119 095 712	5 235 826	586 521	4.695E-3	22.746	39.894
32	1 521 135 167	30 469 684	325 353 043	1.733E-1	49.923	39.975
33	1 783 650 614	21 296 023	389 170 221	1.774E-1	83.755	39.633
34	32 092	4054	272	7.469E-3	7.916	22.065
35	47 102 261 762	2 264 762 985	1 174 389 309	2.324E-2	20.798	42.336
36	15 837	13 915	1441	4.620E-2	1.138	24.242
37	3 083 083 913	1 174 389 309	162 128 255	3.668E-2	2.625	42.291
38	464	305	14	1.788E-2	1.521	13.086
39	43 163 436	1 054 980	8 441 331	1.603E-1	40.914	39.841
40	3857	3224	87	1.214E-2	1.196	17.622
41	5 329 186 306	4 114 387 435	429 886 036	4.354E-2	1.295	41.013
42	2 138 265	34 751	6343	2.910E-3	61.531	32.249
43	9 703 090	3 733 007	18 023	1.340E-3	2.599	35.445
44	721 215 037	220 709 637	6 325 446	6.671E-3	3.268	39.260
45	26 507 497	8 275 003	1 479 690	4.081E-2	3.203	42.088
46	35 805	1707	146	3.877E-3	20.975	20.765
47	750 499 260	447 341 132	3 089 958	2.573E-3	1.678	37.744
48	8 441 331	1 634 498	53 310	5.263E-3	5.164	37.918
49	5 235 826	485 486	64 405	1.113E-2	10.785	38.540
50	93 072	52 782	358	2.449E-3	1.763	26.491
51	5 172 455 725	3 208 269 483	72 414 817	8.567E-3	1.612	39.867
52	1 697 082 357	131 555 758	4 509 877	2.460E-3	12.900	37.680

Table 2: The mean, the standard deviation, and the root-mean-square error of the relative estimation error for all cases given in Table 1 when estimating $|A|$ using the conventional approach and the maximum likelihood method, respectively.

#	conventional approach			maximum likelihood method			improvement rmse ratio
	mean	stdev	rmse	mean	stdev	rmse	
1	1.289E-5	1.164E-3	1.164E-3	9.746E-6	1.036E-3	1.036E-3	1.123
2	-2.285E-5	1.417E-3	1.417E-3	-2.523E-6	1.087E-3	1.087E-3	1.304
3	-7.719E-6	7.118E-4	7.118E-4	-8.021E-6	7.060E-4	7.061E-4	1.008
4	-6.395E-5	1.319E-3	1.320E-3	-5.023E-5	1.107E-3	1.108E-3	1.191
5	1.925E-5	1.101E-3	1.101E-3	2.878E-5	1.044E-3	1.044E-3	1.055
6	-1.374E-5	7.271E-4	7.272E-4	-1.275E-5	7.170E-4	7.171E-4	1.014
7	-1.073E-5	7.741E-4	7.742E-4	-8.752E-6	7.340E-4	7.340E-4	1.055
8	3.529E-5	1.106E-3	1.107E-3	3.141E-5	1.043E-3	1.043E-3	1.061
9	1.297E-5	7.218E-4	7.220E-4	1.095E-5	7.153E-4	7.154E-4	1.009
10	-3.688E-6	1.461E-3	1.461E-3	-4.449E-6	1.245E-3	1.245E-3	1.173
11	-5.030E-5	1.272E-3	1.273E-3	-4.314E-5	1.033E-3	1.034E-3	1.231
12	5.651E-6	1.127E-3	1.127E-3	6.821E-6	1.028E-3	1.028E-3	1.096
13	-1.255E-5	1.169E-3	1.169E-3	-2.606E-6	9.041E-4	9.041E-4	1.293
14	-8.608E-7	8.447E-4	8.447E-4	-1.319E-5	7.684E-4	7.685E-4	1.099
15	-3.134E-5	1.153E-3	1.153E-3	-2.042E-5	1.013E-3	1.013E-3	1.138
16	-1.300E-5	7.731E-4	7.732E-4	-1.091E-5	7.646E-4	7.646E-4	1.011
17	-1.344E-5	7.384E-4	7.385E-4	-1.337E-5	7.229E-4	7.230E-4	1.021
18	7.711E-6	1.022E-3	1.022E-3	1.155E-5	1.007E-3	1.007E-3	1.014
19	-2.510E-6	1.173E-3	1.173E-3	-7.793E-6	1.087E-3	1.087E-3	1.079
20	1.899E-5	1.241E-3	1.242E-3	2.087E-5	1.141E-3	1.141E-3	1.088
21	-4.536E-6	7.130E-4	7.130E-4	-3.447E-6	7.058E-4	7.058E-4	1.010
22	5.285E-6	1.058E-3	1.058E-3	1.037E-5	1.036E-3	1.036E-3	1.022
23	7.841E-6	7.298E-4	7.299E-4	8.400E-6	7.127E-4	7.128E-4	1.024
24	-1.129E-5	8.831E-4	8.832E-4	-4.758E-6	7.851E-4	7.851E-4	1.125
25	-1.347E-5	7.880E-4	7.881E-4	-8.379E-6	7.453E-4	7.453E-4	1.057
26	-1.130E-5	8.663E-4	8.664E-4	-1.471E-5	8.195E-4	8.197E-4	1.057
27	1.349E-5	1.079E-3	1.079E-3	1.924E-5	8.349E-4	8.351E-4	1.292
28	-1.781E-5	7.610E-4	7.613E-4	-1.347E-5	7.360E-4	7.362E-4	1.034
29	7.786E-6	1.085E-3	1.085E-3	2.323E-6	1.054E-3	1.054E-3	1.030
30	6.713E-6	1.017E-3	1.017E-3	1.244E-5	9.969E-4	9.970E-4	1.020
31	1.063E-5	1.090E-3	1.091E-3	1.496E-5	1.053E-3	1.053E-3	1.036
32	-2.000E-5	1.200E-3	1.200E-3	-1.625E-5	1.114E-3	1.114E-3	1.077
33	4.975E-6	1.219E-3	1.219E-3	2.154E-6	1.141E-3	1.141E-3	1.068
34	1.519E-5	7.680E-4	7.681E-4	1.355E-5	7.205E-4	7.206E-4	1.066
35	-1.724E-5	1.082E-3	1.082E-3	-2.074E-5	1.038E-3	1.038E-3	1.043
36	-2.077E-5	1.193E-3	1.193E-3	-1.578E-5	8.946E-4	8.947E-4	1.333
37	-9.366E-7	1.385E-3	1.385E-3	1.554E-5	1.112E-3	1.112E-3	1.246
38	-1.520E-5	1.078E-3	1.078E-3	-2.170E-5	8.660E-4	8.663E-4	1.244
39	6.763E-6	1.210E-3	1.210E-3	2.194E-5	1.128E-3	1.128E-3	1.072
40	1.625E-5	1.130E-3	1.131E-3	6.289E-6	8.675E-4	8.676E-4	1.303
41	-3.313E-6	1.665E-3	1.665E-3	1.743E-5	1.181E-3	1.181E-3	1.409
42	-1.286E-5	8.168E-4	8.170E-4	-1.212E-5	8.034E-4	8.035E-4	1.017
43	-1.348E-5	1.300E-3	1.300E-3	2.815E-5	9.965E-4	9.969E-4	1.304
44	5.351E-5	1.283E-3	1.284E-3	2.553E-5	1.048E-3	1.048E-3	1.225
45	-7.546E-5	1.332E-3	1.335E-3	-5.078E-5	1.074E-3	1.075E-3	1.241
46	-4.855E-6	7.332E-4	7.332E-4	-4.521E-6	7.145E-4	7.145E-4	1.026
47	7.932E-6	1.475E-3	1.475E-3	2.192E-5	1.062E-3	1.063E-3	1.389
48	-1.825E-5	1.144E-3	1.144E-3	1.667E-6	9.695E-4	9.695E-4	1.180
49	-2.281E-6	9.927E-4	9.927E-4	5.282E-6	9.255E-4	9.255E-4	1.073
50	1.395E-5	1.034E-3	1.034E-3	-1.046E-5	7.980E-4	7.981E-4	1.296
51	-9.423E-6	1.529E-3	1.529E-3	7.494E-6	1.098E-3	1.098E-3	1.392
52	1.672E-5	1.094E-3	1.094E-3	8.713E-6	1.016E-3	1.016E-3	1.077

Table 3: The mean, the standard deviation, and the root-mean-square error of the relative estimation error for all cases given in Table 1 when estimating $|B|$ using the conventional approach and the maximum likelihood method, respectively.

#	conventional approach			maximum likelihood method			improvement rmse ratio
	mean	stdev	rmse	mean	stdev	rmse	
1	3.035E-5	3.698E-3	3.698E-3	1.082E-5	1.576E-3	1.576E-3	2.346
2	-4.702E-5	2.302E-3	2.302E-3	-2.960E-6	1.216E-3	1.216E-3	1.894
3	-4.868E-5	9.643E-3	9.643E-3	-9.106E-6	7.152E-3	7.152E-3	1.348
4	-9.060E-5	3.298E-3	3.299E-3	-4.553E-5	2.022E-3	2.023E-3	1.631
5	-1.553E-4	5.737E-3	5.739E-3	-2.657E-5	3.482E-3	3.482E-3	1.648
6	-5.567E-5	5.240E-3	5.240E-3	-2.250E-5	3.184E-3	3.184E-3	1.646
7	-3.752E-5	3.541E-3	3.541E-3	-1.491E-5	2.453E-3	2.453E-3	1.444
8	8.603E-5	5.481E-3	5.482E-3	4.731E-5	3.133E-3	3.133E-3	1.749
9	1.340E-4	7.143E-3	7.144E-3	4.159E-5	4.489E-3	4.489E-3	1.591
10	3.432E-5	3.125E-3	3.125E-3	3.788E-5	2.216E-3	2.217E-3	1.410
11	-5.585E-5	2.734E-3	2.734E-3	-3.566E-5	1.453E-3	1.454E-3	1.881
12	2.653E-5	5.045E-3	5.045E-3	5.265E-6	2.734E-3	2.734E-3	1.845
13	2.204E-5	1.546E-3	1.546E-3	3.773E-5	1.110E-3	1.111E-3	1.392
14	-1.820E-5	2.754E-3	2.754E-3	-7.144E-5	1.920E-3	1.921E-3	1.434
15	-3.790E-5	3.875E-3	3.875E-3	2.103E-5	2.186E-3	2.186E-3	1.773
16	-2.034E-4	9.797E-3	9.799E-3	-6.704E-5	4.728E-3	4.728E-3	2.073
17	-1.441E-4	7.286E-3	7.287E-3	-1.370E-4	5.392E-3	5.394E-3	1.351
18	1.020E-4	1.335E-2	1.335E-2	9.586E-5	8.986E-3	8.987E-3	1.485
19	7.563E-5	8.080E-3	8.081E-3	-8.047E-5	6.350E-3	6.351E-3	1.272
20	-5.421E-5	5.135E-3	5.135E-3	-1.619E-5	3.587E-3	3.587E-3	1.432
21	2.219E-5	9.444E-3	9.444E-3	8.959E-5	6.521E-3	6.522E-3	1.448
22	9.566E-5	1.069E-2	1.069E-2	1.510E-4	7.528E-3	7.529E-3	1.420
23	-1.068E-4	8.116E-3	8.117E-3	-7.135E-5	5.904E-3	5.905E-3	1.375
24	-2.241E-5	2.316E-3	2.316E-3	1.616E-6	1.617E-3	1.617E-3	1.433
25	-3.027E-5	3.229E-3	3.229E-3	2.059E-5	2.058E-3	2.058E-3	1.569
26	2.436E-4	6.135E-3	6.140E-3	1.380E-4	4.705E-3	4.707E-3	1.304
27	4.772E-5	1.982E-3	1.982E-3	6.298E-5	1.200E-3	1.202E-3	1.649
28	-2.276E-4	5.571E-3	5.575E-3	-1.191E-4	3.904E-3	3.906E-3	1.427
29	3.200E-5	6.155E-3	6.155E-3	-4.422E-5	3.202E-3	3.203E-3	1.922
30	-2.787E-4	9.641E-3	9.645E-3	-6.863E-5	4.882E-3	4.882E-3	1.976
31	-1.264E-4	7.032E-3	7.033E-3	-3.750E-5	3.127E-3	3.127E-3	2.249
32	-9.557E-6	1.114E-2	1.114E-2	3.899E-5	8.776E-3	8.776E-3	1.270
33	3.484E-4	1.465E-2	1.465E-2	-5.508E-5	1.162E-2	1.162E-2	1.261
34	-3.365E-5	2.864E-3	2.864E-3	-4.537E-5	1.809E-3	1.810E-3	1.583
35	1.726E-5	6.839E-3	6.839E-3	-2.484E-5	3.976E-3	3.976E-3	1.720
36	-3.148E-6	1.305E-3	1.305E-3	4.279E-6	9.825E-4	9.825E-4	1.328
37	-4.329E-5	2.538E-3	2.538E-3	-3.505E-6	1.502E-3	1.502E-3	1.690
38	-5.696E-6	1.388E-3	1.388E-3	-1.530E-5	9.998E-4	9.999E-4	1.388
39	-2.851E-4	1.011E-2	1.012E-2	-2.377E-4	7.885E-3	7.889E-3	1.282
40	1.558E-5	1.273E-3	1.273E-3	3.272E-6	9.479E-4	9.479E-4	1.343
41	2.107E-5	1.959E-3	1.959E-3	4.807E-5	1.281E-3	1.281E-3	1.529
42	-4.961E-5	9.654E-3	9.654E-3	-1.783E-5	5.001E-3	5.001E-3	1.930
43	-3.169E-5	2.480E-3	2.480E-3	7.632E-5	1.215E-3	1.218E-3	2.037
44	1.128E-4	2.823E-3	2.825E-3	2.291E-5	1.365E-3	1.365E-3	2.069
45	-9.187E-5	2.873E-3	2.875E-3	-2.030E-5	1.628E-3	1.628E-3	1.766
46	2.855E-6	4.558E-3	4.558E-3	1.449E-5	2.769E-3	2.769E-3	1.646
47	-3.915E-5	2.128E-3	2.128E-3	-1.559E-5	1.168E-3	1.168E-3	1.822
48	-6.558E-5	3.329E-3	3.330E-3	3.537E-5	1.408E-3	1.408E-3	2.364
49	-6.749E-5	4.581E-3	4.581E-3	9.343E-6	2.176E-3	2.176E-3	2.105
50	1.981E-5	1.507E-3	1.507E-3	-2.312E-5	9.520E-4	9.522E-4	1.582
51	-1.831E-5	2.086E-3	2.086E-3	9.137E-6	1.204E-3	1.204E-3	1.733
52	9.727E-5	5.202E-3	5.203E-3	-3.350E-6	2.075E-3	2.075E-3	2.507

Table 4: The mean, the standard deviation, and the root-mean-square error of the relative estimation error for all cases given in Table 1 when estimating $|X|$ using the conventional approach and the maximum likelihood method, respectively.

#	conventional approach			maximum likelihood method			improvement rmse ratio
	mean	stdev	rmse	mean	stdev	rmse	
1	-2.437E-3	3.264E-1	3.264E-1	-7.470E-4	1.116E-1	1.116E-1	2.926
2	3.662E-3	1.197E-1	1.198E-1	1.131E-3	4.016E-2	4.017E-2	2.982
3	4.081E-5	4.504E-3	4.504E-3	2.310E-5	3.369E-3	3.370E-3	1.337
4	1.653E-4	7.648E-3	7.649E-3	5.631E-5	4.399E-3	4.399E-3	1.739
5	2.478E-4	1.009E-2	1.009E-2	1.911E-5	6.060E-3	6.060E-3	1.666
6	3.974E-4	5.767E-2	5.767E-2	3.048E-5	3.455E-2	3.455E-2	1.669
7	7.918E-5	5.995E-3	5.996E-3	4.075E-5	4.110E-3	4.110E-3	1.459
8	-1.996E-4	1.427E-2	1.428E-2	-9.701E-5	7.874E-3	7.874E-3	1.813
9	-5.706E-4	2.869E-2	2.869E-2	-1.980E-4	1.788E-2	1.788E-2	1.604
10	-1.700E-6	3.211E-3	3.211E-3	-4.518E-6	2.292E-3	2.292E-3	1.401
11	3.700E-4	3.070E-2	3.070E-2	1.282E-4	1.357E-2	1.357E-2	2.262
12	-1.204E-4	1.631E-2	1.631E-2	-5.151E-5	8.539E-3	8.539E-3	1.910
13	-5.165E-5	5.661E-3	5.662E-3	-1.138E-4	3.608E-3	3.610E-3	1.568
14	4.061E-6	4.318E-3	4.318E-3	9.206E-5	2.998E-3	2.999E-3	1.440
15	1.723E-5	1.267E-2	1.267E-2	-1.810E-4	6.876E-3	6.878E-3	1.841
16	2.057E-3	1.020E-1	1.020E-1	6.307E-4	4.869E-2	4.869E-2	2.096
17	1.115E-4	4.986E-3	4.987E-3	1.067E-4	3.727E-3	3.728E-3	1.338
18	-9.800E-5	1.022E-2	1.022E-2	-9.339E-5	6.878E-3	6.878E-3	1.486
19	1.839E-5	1.854E-3	1.854E-3	4.840E-5	1.567E-3	1.568E-3	1.182
20	3.215E-5	4.010E-3	4.010E-3	3.232E-6	2.844E-3	2.844E-3	1.410
21	-7.916E-5	1.641E-2	1.641E-2	-1.967E-4	1.130E-2	1.130E-2	1.452
22	-1.385E-5	7.100E-3	7.100E-3	-5.038E-5	5.032E-3	5.032E-3	1.411
23	7.833E-5	4.726E-3	4.726E-3	5.787E-5	3.479E-3	3.480E-3	1.358
24	4.137E-5	5.476E-3	5.476E-3	-1.895E-5	3.676E-3	3.676E-3	1.490
25	1.634E-4	1.615E-2	1.615E-2	-9.279E-5	9.954E-3	9.954E-3	1.622
26	-7.280E-5	1.460E-3	1.462E-3	-5.170E-5	1.225E-3	1.226E-3	1.193
27	-5.498E-4	1.718E-2	1.719E-2	-6.873E-4	8.998E-3	9.024E-3	1.905
28	2.729E-4	7.261E-3	7.266E-3	1.323E-4	5.069E-3	5.071E-3	1.433
29	-1.751E-4	3.005E-2	3.005E-2	2.033E-4	1.501E-2	1.501E-2	2.002
30	1.183E-3	4.327E-2	4.329E-2	2.398E-4	2.168E-2	2.168E-2	1.996
31	1.164E-3	6.230E-2	6.231E-2	3.699E-4	2.689E-2	2.689E-2	2.317
32	-2.450E-5	1.430E-3	1.431E-3	-2.895E-5	1.295E-3	1.295E-3	1.104
33	-4.084E-5	1.275E-3	1.275E-3	-1.952E-5	1.185E-3	1.185E-3	1.076
34	1.550E-4	4.123E-2	4.123E-2	3.301E-4	2.466E-2	2.466E-2	1.672
35	-3.329E-5	1.308E-2	1.308E-2	4.784E-5	7.446E-3	7.446E-3	1.756
36	-3.854E-6	1.020E-2	1.020E-2	-6.474E-5	6.265E-3	6.266E-3	1.629
37	3.896E-4	1.700E-2	1.701E-2	1.020E-4	8.108E-3	8.109E-3	2.097
38	7.621E-5	2.619E-2	2.619E-2	2.817E-4	1.552E-2	1.553E-2	1.687
39	2.555E-5	1.557E-3	1.557E-3	1.941E-5	1.358E-3	1.358E-3	1.147
40	-7.173E-4	3.911E-2	3.912E-2	-2.736E-4	2.339E-2	2.339E-2	1.672
41	1.411E-4	1.590E-2	1.590E-2	-1.127E-4	7.549E-3	7.550E-3	2.106
42	3.007E-4	5.277E-2	5.277E-2	1.263E-4	2.709E-2	2.709E-2	1.948
43	1.080E-2	4.686E-1	4.688E-1	-1.491E-2	1.752E-1	1.758E-1	2.666
44	-4.057E-3	9.101E-2	9.110E-2	-9.202E-4	3.193E-2	3.194E-2	2.852
45	3.479E-4	1.503E-2	1.503E-2	-5.124E-5	7.429E-3	7.429E-3	2.024
46	-2.737E-5	5.257E-2	5.257E-2	-1.604E-4	3.123E-2	3.124E-2	1.683
47	5.752E-3	2.671E-1	2.672E-1	2.353E-3	8.080E-2	8.083E-2	3.306
48	2.003E-3	9.923E-2	9.925E-2	-1.091E-3	3.601E-2	3.603E-2	2.755
49	3.401E-4	3.392E-2	3.393E-2	-2.383E-4	1.544E-2	1.544E-2	2.197
50	-3.934E-3	1.954E-1	1.955E-1	2.387E-3	9.488E-2	9.491E-2	2.060
51	1.753E-3	8.153E-2	8.155E-2	5.423E-4	2.943E-2	2.944E-2	2.770
52	-2.605E-3	1.488E-1	1.489E-1	3.309E-4	5.340E-2	5.340E-2	2.788

Table 5: The mean, the standard deviation, and the root-mean-square error of the relative estimation error for all cases given in Table 1 when estimating $|A \cup B \cup X|$ using the conventional approach and the maximum likelihood method, respectively.

#	conventional approach			maximum likelihood method			improvement rmse ratio
	mean	stdev	rmse	mean	stdev	rmse	
1	1.163E-5	1.012E-3	1.012E-3	8.772E-6	9.034E-4	9.035E-4	1.121
2	-1.038E-5	1.016E-3	1.016E-3	3.493E-6	8.070E-4	8.070E-4	1.259
3	-7.037E-6	6.876E-4	6.876E-4	-7.312E-6	6.821E-4	6.821E-4	1.008
4	-5.192E-5	1.019E-3	1.020E-3	-4.161E-5	8.631E-4	8.641E-4	1.180
5	1.650E-5	1.004E-3	1.004E-3	2.517E-5	9.518E-4	9.521E-4	1.055
6	-1.391E-5	7.010E-4	7.012E-4	-1.295E-5	6.912E-4	6.913E-4	1.014
7	-8.919E-6	6.862E-4	6.862E-4	-7.111E-6	6.494E-4	6.494E-4	1.057
8	3.280E-5	1.010E-3	1.011E-3	2.927E-5	9.527E-4	9.531E-4	1.061
9	1.253E-5	7.048E-4	7.049E-4	1.055E-5	6.984E-4	6.984E-4	1.009
10	3.805E-6	1.010E-3	1.010E-3	3.503E-6	8.690E-4	8.690E-4	1.162
11	-4.332E-5	9.735E-4	9.745E-4	-3.800E-5	8.005E-4	8.014E-4	1.216
12	4.256E-6	1.013E-3	1.013E-3	5.272E-6	9.235E-4	9.235E-4	1.097
13	-3.509E-6	7.194E-4	7.194E-4	2.487E-6	5.749E-4	5.750E-4	1.251
14	-2.582E-6	6.996E-4	6.997E-4	-1.257E-5	6.349E-4	6.350E-4	1.102
15	-3.034E-5	9.707E-4	9.712E-4	-2.122E-5	8.524E-4	8.527E-4	1.139
16	-1.289E-5	7.608E-4	7.609E-4	-1.083E-5	7.524E-4	7.525E-4	1.011
17	-1.249E-5	7.067E-4	7.068E-4	-1.242E-5	6.917E-4	6.918E-4	1.022
18	7.195E-6	9.943E-4	9.943E-4	1.093E-5	9.802E-4	9.803E-4	1.014
19	2.711E-6	9.912E-4	9.912E-4	-1.781E-6	9.221E-4	9.221E-4	1.075
20	1.482E-5	1.045E-3	1.045E-3	1.645E-5	9.611E-4	9.612E-4	1.088
21	-4.709E-6	7.008E-4	7.008E-4	-3.638E-6	6.937E-4	6.937E-4	1.010
22	6.309E-6	1.014E-3	1.014E-3	1.118E-5	9.930E-4	9.930E-4	1.022
23	7.916E-6	7.000E-4	7.000E-4	8.464E-6	6.833E-4	6.833E-4	1.024
24	-9.715E-6	6.942E-4	6.943E-4	-4.613E-6	6.154E-4	6.154E-4	1.128
25	-1.191E-5	7.084E-4	7.085E-4	-7.291E-6	6.704E-4	6.704E-4	1.057
26	-1.194E-5	7.321E-4	7.322E-4	-1.511E-5	6.912E-4	6.914E-4	1.059
27	6.183E-6	7.710E-4	7.710E-4	1.030E-5	6.082E-4	6.082E-4	1.268
28	-1.747E-5	7.198E-4	7.200E-4	-1.331E-5	6.960E-4	6.962E-4	1.034
29	7.131E-6	1.020E-3	1.020E-3	2.010E-6	9.907E-4	9.907E-4	1.030
30	6.217E-6	9.911E-4	9.911E-4	1.180E-5	9.716E-4	9.717E-4	1.020
31	1.030E-5	1.041E-3	1.041E-3	1.443E-5	1.005E-3	1.005E-3	1.035
32	-2.061E-5	9.899E-4	9.901E-4	-1.756E-5	9.216E-4	9.217E-4	1.074
33	1.806E-7	1.007E-3	1.007E-3	-2.245E-6	9.463E-4	9.463E-4	1.064
34	1.080E-5	6.820E-4	6.821E-4	9.360E-6	6.404E-4	6.405E-4	1.065
35	-1.607E-5	1.012E-3	1.012E-3	-1.933E-5	9.709E-4	9.711E-4	1.042
36	-1.213E-5	7.013E-4	7.014E-4	-9.095E-6	5.683E-4	5.684E-4	1.234
37	2.138E-6	1.008E-3	1.008E-3	1.365E-5	8.303E-4	8.304E-4	1.214
38	-9.866E-6	6.895E-4	6.896E-4	-1.378E-5	5.766E-4	5.767E-4	1.196
39	3.927E-6	1.007E-3	1.007E-3	1.633E-5	9.401E-4	9.402E-4	1.071
40	7.043E-6	6.909E-4	6.909E-4	1.536E-6	5.692E-4	5.693E-4	1.214
41	1.313E-5	1.005E-3	1.005E-3	2.453E-5	7.734E-4	7.737E-4	1.298
42	-1.253E-5	8.012E-4	8.013E-4	-1.181E-5	7.880E-4	7.881E-4	1.017
43	-4.050E-6	9.741E-4	9.741E-4	2.150E-5	7.501E-4	7.504E-4	1.298
44	3.990E-5	1.011E-3	1.012E-3	1.861E-5	8.356E-4	8.358E-4	1.211
45	-6.193E-5	1.006E-3	1.008E-3	-4.384E-5	8.200E-4	8.212E-4	1.228
46	-4.593E-6	6.975E-4	6.975E-4	-4.264E-6	6.797E-4	6.797E-4	1.026
47	5.173E-6	1.002E-3	1.002E-3	1.395E-5	7.649E-4	7.651E-4	1.310
48	-1.525E-5	9.629E-4	9.630E-4	1.355E-6	8.188E-4	8.188E-4	1.176
49	-3.941E-6	9.006E-4	9.006E-4	2.911E-6	8.391E-4	8.391E-4	1.073
50	6.400E-6	7.031E-4	7.032E-4	-9.157E-6	5.623E-4	5.624E-4	1.250
51	2.300E-6	1.005E-3	1.005E-3	1.270E-5	7.704E-4	7.705E-4	1.305
52	1.605E-5	1.015E-3	1.015E-3	8.640E-6	9.421E-4	9.422E-4	1.077

angular and Jaccard distances, respectively. A generalization of the latter method is b -bit minwise hashing that improves memory efficiency by only storing the lowest b bits of the minimum hash value [29]. The probability that two different sets are mapped to equal hash values K_1 and K_2 is roughly

$$P(K_1 = K_2) \approx 1 - (1 - \frac{1}{2^b})D. \quad (76)$$

The HyperLogLog algorithm itself can be regarded as hashing algorithm as it maps sets to register values. For sufficiently large cardinalities we can use the Poisson approximation and assume that the number of zero-valued HyperLogLog registers can be ignored. Furthermore, if the HyperLogLog parameter q is chosen large enough, the number of saturated registers can be ignored as well. As a consequence, we can assume that the distribution of register values follows (16) and the probability that a register has the same value for two different sets is (compare (66))

$$P(K_1 = K_2) = \sum_{k=-\infty}^{\infty} e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m2^k}} \left(1 - e^{-\frac{\lambda_a + \lambda_x}{m2^k}} - e^{-\frac{\lambda_b + \lambda_x}{m2^k}} + e^{-\frac{\lambda_a + \lambda_b + \lambda_x}{m2^k}} \right). \quad (77)$$

Using the approximation $\sum_{k=-\infty}^{\infty} e^{-\frac{x}{2^k}} - e^{-\frac{y}{2^k}} \approx 2\alpha_\infty (\log(y) - \log(x))$ (compare (90) in Appendix A) we get

$$P(K_1 = K_2) \approx 1 + 2\alpha_\infty \log\left(1 - \frac{1}{2}D + \frac{1}{4}D^2 \frac{\lambda_a \lambda_b}{(\lambda_a + \lambda_b)^2}\right) \quad (78)$$

where $D = \frac{\lambda_a + \lambda_b}{\lambda_a + \lambda_b + \lambda_x}$ is the Jaccard distance. Since $\frac{\lambda_a \lambda_b}{(\lambda_a + \lambda_b)^2}$ is always in the range $[0, \frac{1}{4}]$, the probability for equal register values can be bounded by

$$1 + 2\alpha_\infty \log(1 - \frac{1}{2}D) \lesssim P(K_1 = K_2) \lesssim 1 + 2\alpha_\infty \log(1 - \frac{1}{2}D + \frac{1}{16}D^2). \quad (79)$$

As shown in Fig. 21 the bounds are very close, especially for small Jaccard distances, where the probability can be well approximated by

$$P(K_1 = K_2) \approx 1 - \alpha_\infty D. \quad (80)$$

This dependency on the Jaccard distance is very similar to that of minwise hashing (76), which makes the HyperLogLog algorithm an interesting candidate for locality-sensitive hashing with respect to the Jaccard similarity. The memory efficiency of different hashing approaches can be measured by a storage factor that is the variance of the distance estimator multiplied by the number of bits used for storing the hash signature [29]. For a hash algorithm that maps two different sets to the same b -bit hash value with probability $1 - aD$ with some constant $a \in (0, 1]$, the storage factor is given by $bD(\frac{1}{a} - D)$. Particularly, for the case $D = 0.3$, this gives 6.72 for conventional minwise hashing with 32-bit hash values ($b = 32, a \approx 1$), 0.51 for 1-bit minwise hashing ($b = 1, a = \frac{1}{2}$), and 1.96 for the HyperLogLog algorithm with 6-bit registers ($b = 6, a = \alpha_\infty$). This means that the HyperLogLog algorithm is not as memory-efficient as 1-bit minwise hashing, but significantly better than 32-bit minwise hashing, when estimating the Jaccard distance by just counting equal registers. However, in contrast to 1-bit minwise hashing




Figure 21: The approximate probability range of equal register values as a function of the Jaccard distance.

the HyperLogLog sketch contains additional information which allows estimating the cardinality or merging two HyperLogLog hash signatures. Furthermore, the method described in Section 5 would allow a more accurate estimation of the Jaccard distance by using the estimates for intersection and union sizes. This could be used for additional more precise filtering when searching for similar items.

Since HyperLogLog sketches can be efficiently constructed, because only a single hash function evaluation is needed for each item, the preprocessing step would be very fast. In contrast, preprocessing is very costly for minwise hashing, because of the many required permutations [29]. To overcome this problem, one permutation hashing was proposed [30] which keeps for a predefined number of bins the minimum of all values that are mapped to this bin. Actually, this is very similar to the HyperLogLog algorithm. The bins correspond to HyperLogLog registers which keep the first one-bit position of the minimum value instead of the minimum value itself. This close connection was also pointed out in [10], where both methods have been considered as variants of k -partition minwise hashing. The only difference is that the HyperLogLog algorithm uses base-2 ranks while one permutation hashing uses full ranks.

7. Conclusion

We have presented new algorithms for the estimation of cardinalities from HyperLogLog sketches based on the Poisson approximation. For the estimation from a single sketch, we have developed two different algorithms that are inherently unbiased over the full cardinality range without dependence on empirically determined bias correction data. The first uses the original estimator extended by theoretically motivated correction terms. Due to its simplicity we believe that it has the potential to become the standard textbook cardinality estimation algorithm for HyperLogLog sketches. The second is based on the maximum likelihood method and solves the corresponding optimization problem

using the secant method. The maximum likelihood method was also successfully applied to the estimation of set operation result sizes where the operands are represented by HyperLogLog sketches. The new approach improves the cardinality estimates of set intersections, relative complements, as well as unions significantly when compared to the conventional approach using the inclusion-exclusion principle.

A. Analysis of $\xi(x)$

The Fourier series of the periodic function

$$\xi(x) := \log(2) \sum_{k=-\infty}^{\infty} 2^{k+x} e^{-2^{k+x}}, \quad (81)$$

which has a period equal to 1, is

$$\xi(x) = \frac{a_0}{2} + \operatorname{Re} \left(\sum_{l=1}^{\infty} a_l e^{2\pi i l x} \right) \quad (82)$$

with coefficients

$$\begin{aligned} a_l &= 2 \int_0^1 \xi(x) e^{-2\pi i l x} dx = 2 \log(2) \int_0^1 \sum_{k=-\infty}^{\infty} 2^{k+x} e^{-2^{k+x}} e^{-2\pi i l x} dx = \\ &= 2 \log(2) \sum_{k=-\infty}^{\infty} \int_k^{k+1} 2^x e^{-2^x} e^{-2\pi i l x} dx = 2 \log(2) \int_{-\infty}^{\infty} 2^x e^{-2^x} e^{-2\pi i l x} dx. \end{aligned} \quad (83)$$

The variable transformation $y = 2^x$ yields

$$a_l = 2 \int_0^{\infty} e^{-y} y^{-\frac{2\pi i l}{\log(2)}} dy = 2\Gamma \left(1 - \frac{2\pi i l}{\log(2)} \right) \quad (84)$$

where Γ denotes the gamma function. Obviously, the constant term in the Fourier series is equal to 1, because $a_0 = 2$.

To investigate the maximum deviation of $\xi(x)$ from this value we consider all further coefficients a_l with $l \geq 1$. Using the identity $|\Gamma(1 + ix)| = \sqrt{\frac{\pi x}{\sinh(\pi x)}}$ we are able to write for the absolute values of the coefficients

$$|a_l| = 2 \sqrt{\frac{bl}{\sinh(bl)}} \quad \text{with } b := \frac{2\pi^2}{\log 2}. \quad (85)$$

In particular, the amplitude of the first harmonic is $|a_1| \approx 9.884 \times 10^{-6}$. Clearly, the maximum deviation of $\xi(x) - 1$ from the first harmonic must be smaller than $\sum_{l=2}^{\infty} |a_l|$. The ratio of subsequent coefficients is given by

$$\frac{|a_{l+1}|}{|a_l|} = \sqrt{\frac{l+1}{l}} \sqrt{\frac{\sinh(bl)}{\sinh(b(l+1))}} = \sqrt{\frac{l+1}{l}} \sqrt{\frac{1}{\cosh(b) + \frac{\sinh(b)}{\tanh(bl)}}}. \quad (86)$$

For $l \geq 2$ we have $\sqrt{\frac{l+1}{l}} \leq \sqrt{\frac{3}{2}}$. Together with $\tanh(x) \leq 1$ we obtain

$$\frac{|a_{l+1}|}{|a_l|} \leq \sqrt{\frac{3}{2}} \sqrt{\frac{1}{\cosh(b) + \sinh(b)}} = \sqrt{\frac{3}{2e^b}} \quad (87)$$

which leads to $|a_l| \leq |a_2| \left(\sqrt{\frac{3}{2e^b}}\right)^{l-2}$ for $l \geq 2$ and further to

$$\sum_{l=2}^{\infty} |a_l| \leq |a_2| \sum_{l=0}^{\infty} \left(\sqrt{\frac{3}{2e^b}}\right)^l = 2 \sqrt{\frac{2b}{\sinh(2b)}} \frac{1}{1 - \sqrt{\frac{3}{2e^b}}} \approx 9.154 \times 10^{-12}. \quad (88)$$

As a consequence, the maximum deviation of $\xi(x)$ from 1 is bounded by

$$9.884 \times 10^{-6} \leq |a_1| - \sum_{l=2}^{\infty} |a_l| \leq \max_x (|\xi(x) - 1|) \leq |a_1| + \sum_{l=2}^{\infty} |a_l| \leq 9.885 \times 10^{-6}. \quad (89)$$

An interesting approximation formula can be derived from $\xi(x) \approx 1$ by integrating on both sides:

$$\sum_{k=-\infty}^{\infty} e^{-2^{k+y}} - e^{-2^{k+x}} \approx x - y. \quad (90)$$

B. Numerical stability of recursion formula for $h(x)$

In order to investigate the error propagation of a single recursion step using (56) we define $h_1 := h(2x)$ and $h_2 := h(4x)$. The recursion formula simplifies to

$$h_2 = \frac{x + h_1(1 - h_1)}{x + (1 - h_1)}. \quad (91)$$

If h_1 is approximated by $\tilde{h}_1 = h_1(1 + \varepsilon_1)$ with relative error ε_1 , the recursion formula will give an approximation for h_2

$$\tilde{h}_2 = \frac{x + \tilde{h}_1(1 - \tilde{h}_1)}{x + (1 - \tilde{h}_1)}. \quad (92)$$

The corresponding relative error ε_2 is given by

$$\varepsilon_2 = \frac{\tilde{h}_2}{h_2} - 1. \quad (93)$$

Combination of (91), (92), and (93) yields for its absolute value

$$|\varepsilon_2| = |\varepsilon_1| \frac{\left| \frac{h_1(1-2h_1)}{x+h_1(1-h_1)} + \frac{h_1}{x+1-h_1} - \varepsilon_1 \frac{h_1^2}{x+h_1(1-h_1)} \right|}{\left| 1 - \varepsilon_1 \frac{h_1}{x+1-h_1} \right|} \quad (94)$$

and the triangle inequality leads to

$$|\varepsilon_2| \leq |\varepsilon_1| \frac{\left| \frac{h_1(1-2h_1)}{x+h_1(1-h_1)} + \frac{h_1}{x+1-h_1} \right| + |\varepsilon_1| \frac{h_1^2}{x+h_1(1-h_1)}}{\left| 1 - \varepsilon_1 \frac{h_1}{x+1-h_1} \right|}. \quad (95)$$

By numerical means it is easy to show that the inequalities $\left| \frac{h_1(1-2h_1)}{x+h_1(1-h_1)} + \frac{h_1}{x+1-h_1} \right| \leq 0.517$, $\frac{h_1^2}{x+h_1(1-h_1)} \leq 0.436$, and $\frac{h_1}{x+1-h_1} \leq 0.529$ hold for all $x \geq 0$. Therefore, if we additionally assume, for example, $|\varepsilon_1| \leq 0.1$, we get

$$|\varepsilon_2| \leq |\varepsilon_1| \frac{0.517 + 0.1 \cdot 0.436}{1 - 0.1 \cdot 0.529} \leq |\varepsilon_1| \cdot 0.592, \quad (96)$$

which means that the relative error is decreasing in each recursion step and the recursive calculation of h is numerically stable.

C. Error caused by approximation of $h(x)$

According to (42) the exact estimate \hat{x} fulfills

$$\hat{x} \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k h\left(\frac{\hat{x}}{2^k}\right) + C_{q+1} h\left(\frac{\hat{x}}{2^q}\right) = m - C_0. \quad (97)$$

If h is not calculated exactly but approximated by \tilde{h} with maximum relative error $\varepsilon_h \ll 1$

$$\left| \tilde{h}(x) - h(x) \right| \leq \varepsilon_h h(x) \quad (98)$$

the solution of the equation will be off by some relative error ε_x :

$$\hat{x}(1 + \varepsilon_x) \sum_{k=0}^q \frac{C_k}{2^k} + \sum_{k=1}^q C_k \tilde{h}\left(\frac{\hat{x}(1 + \varepsilon_x)}{2^k}\right) + C_{q+1} \tilde{h}\left(\frac{\hat{x}(1 + \varepsilon_x)}{2^q}\right) = m - C_0. \quad (99)$$

Due to (98) there exists some $\alpha \in [-\varepsilon_h, \varepsilon_h]$ for which

$$\begin{aligned} \hat{x}(1 + \varepsilon_x) \sum_{k=0}^q \frac{C_k}{2^k} + (1 + \alpha) \sum_{k=1}^q C_k h\left(\frac{\hat{x}(1 + \varepsilon_x)}{2^k}\right) + \\ + (1 + \alpha) C_{q+1} h\left(\frac{\hat{x}(1 + \varepsilon_x)}{2^q}\right) = m - C_0. \end{aligned} \quad (100)$$

Since $h'(x) \in [0, 0.5]$ for $x \geq 0$, there exists a $\beta \in [0, 0.5]$ for which

$$\begin{aligned} \hat{x}(1 + \varepsilon_x) \sum_{k=0}^q \frac{C_k}{2^k} + (1 + \alpha) \left(\sum_{k=1}^q C_k h\left(\frac{\hat{x}}{2^k}\right) + \frac{C_k}{2^k} \hat{x} \varepsilon_x \beta \right) + \\ + (1 + \alpha) \left(C_{q+1} h\left(\frac{\hat{x}}{2^q}\right) + \frac{C_{q+1}}{2^q} \hat{x} \varepsilon_x \beta \right) = m - C_0. \end{aligned} \quad (101)$$

Subtracting (97) multiplied by $(1 + \alpha)$ from (101) and resolving ε_x gives

$$\varepsilon_x = \alpha \frac{\hat{x} \sum_{k=0}^q \frac{C_k}{2^k} - (m - C_0)}{\hat{x} \left(\sum_{k=0}^q \frac{C_k}{2^k} + (1 + \alpha) \beta \left(\sum_{k=1}^q \frac{C_k}{2^k} + \frac{C_{q+1}}{2^q} \right) \right)}. \quad (102)$$

Using $|\alpha| \leq \varepsilon_h$, $\beta \geq 0$, and (50) the absolute value of the relative error can be bounded by

$$|\varepsilon_x| \leq |\varepsilon_h| \frac{(m - C_0) - \hat{x} \sum_{k=0}^q \frac{C_k}{2^k}}{\hat{x} \sum_{k=0}^q \frac{C_k}{2^k}}. \quad (103)$$

Furthermore, using (47) we finally get

$$|\varepsilon_x| \leq \frac{|\varepsilon_h|}{2} \frac{\sum_{k=1}^q \frac{C_k}{2^k} + \frac{C_{q+1}}{2^q}}{\sum_{k=0}^q \frac{C_k}{2^k}} \leq \frac{|\varepsilon_h|}{2} \left(1 + \frac{\frac{C_{q+1}}{2^q}}{\sum_{k=0}^q \frac{C_k}{2^k}} \right) \leq \frac{|\varepsilon_h|}{2} \left(1 + \frac{C_{q+1}}{m - C_{q+1}} \right). \quad (104)$$

Hence, as long as most registers are not saturated ($C_{q+1} \ll m$), the relative error ε_x of the calculated estimate using the approximation $\tilde{h}(x)$ for $h(x)$ has the same order of magnitude as ε_h .

References

- [1] Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. *Journal of Computer and System Sciences*, 58(1):137–147, 1999.
- [2] Ahmed Metwally, Divyakant Agrawal, and Amr El Abbadi. Why go logarithmic if we can go linear?: Towards effective distinct counting of search traffic. In *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology*, pages 618–629, Nantes, France, March 2008.
- [3] Daniel Ting. 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*, pages 442–451, New York, NY, USA, August 2014.
- [4] Daniel M. Kane, Jelani Nelson, and David P. Woodruff. An optimal algorithm for the distinct elements problem. In *Proceedings of the 29th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 41–52, Indianapolis, IN, USA, June 2010.
- [5] Piotr Indyk and David Woodruff. Tight lower bounds for the distinct elements problem. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 283–288, Cambridge, MA, USA, October 2003.

- [6] Philippe Flajolet, Éric Fusy, Olivier Gandouet, and Frédéric Meunier. Hyperloglog: The analysis of a near-optimal cardinality estimation algorithm. In *Proceedings of the 13th Conference on Analysis of Algorithms*, pages 127–146, Juan des Pins, France, June 2007.
- [7] Stefan Heule, Marc Nunkesser, and Alexander Hall. 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*, pages 683–692, Genoa, Italy, March 2013.
- [8] Lee Rhodes. System and method for enhanced accuracy cardinality estimation, September 24 2015. US Patent 20,150,269,178.
- [9] Salvatore Sanfilippo. Redis new data structure: The HyperLogLog. <http://antirez.com/news/75>, 2014.
- [10] Edith Cohen. All-distances sketches, revisited: HIP estimators for massive graphs analysis. In *Proceedings of the 33rd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pages 88–99, New York, NY, USA, 2014.
- [11] Aiyou Chen, Jin Cao, Larry Shepp, and Tuan Nguyen. Distinct counting with a self-learning bitmap. *Journal of the American Statistical Association*, 106(495):879–890, 2011.
- [12] Kevin Beyer, Peter J. Haas, Berthold Reinwald, Yannis Sismanis, and Rainer Gemulla. On synopses for distinct-value estimation under multiset operations. In *Proceedings of the 26th ACM SIGMOD International Conference on Management of Data*, pages 199–210, Beijing, China, June 2007.
- [13] Anirban Dasgupta, Kevin Lang, Lee Rhodes, and Justin Thaler. A framework for estimating stream expression cardinalities. *arXiv preprint arXiv:1510.01455*, 2015.
- [14] Andrew Pascoe. Hyperloglog and MinHash - A union for intersections. <http://tech.adroll.com/media/hllminhash.pdf>, 2013.
- [15] Reuven Cohen, Liran Katzir, and Aviv Yehezkel. A minimal variance estimator for the cardinality of big data set intersection. *arXiv preprint arXiv:1606.00996*, 2016.
- [16] Philippe Flajolet and G. Nigel Martin. Probabilistic counting algorithms for database applications. *Journal of Computer and System Sciences*, 31(2):182 – 209, 1985.
- [17] Kyu-Young Whang, Brad T. Vander-Zanden, and Howard M. Taylor. A linear-time probabilistic counting algorithm for database applications. *ACM Transactions on Database Systems*, 15(2):208–229, 1990.
- [18] Philippe Jacquet and Wojciech Szpankowski. Analytical depoissonization and its applications. *Theoretical Computer Science*, 201(1):1–62, 1998.

- [19] Marianne Durand and Philippe Flajolet. Loglog counting of large cardinalities. In *Proceedings of the 11th Annual European Symposium on Algorithms*, pages 605–617, Budapest, Hungary, September 2003.
- [20] George Casella and Roger L. Berger. *Statistical inference*. Duxbury, Pacific Grove, CA, USA, 2nd edition, 2002.
- [21] Peter Clifford and Ioana A. Cosma. A statistical analysis of probabilistic counting algorithms. *Scandinavian Journal of Statistics*, 39(1):1–14, 2012.
- [22] Daniel Ting. Towards optimal cardinality estimation of unions and intersections with sketches. In *Proceedings of the 22nd ACM SIGKDD Conference on Knowledge and Data Mining*, pages 1195–1204, San Francisco, CA, USA, August 2016.
- [23] William H. Press. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge University Press, 2007.
- [24] Davis E. King. Dlib-ml: A machine learning toolkit. *Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [25] Jure Leskovec, Anand Rajaraman, and Jeffrey David Ullman. *Mining of massive datasets*. Cambridge University Press, 2014.
- [26] Jingdong Wang, Heng Tao Shen, Jingkuan Song, and Jianqiu Ji. Hashing for similarity search: A survey. *arXiv preprint arXiv:1408.2927*, 2014.
- [27] Moses S. Charikar. Similarity estimation techniques from rounding algorithms. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing*, pages 380–388, Montreal, Canada, May 2002.
- [28] Andrei Z. Broder. On the resemblance and containment of documents. In *Proceedings of the Compression and Complexity of Sequences*, pages 21–29, Positano, Italy, June 1997.
- [29] Ping Li and Arnd Christian König. Theory and applications of b-bit minwise hashing. *Communications of the ACM*, 54(8):101–109, 2011.
- [30] Ping Li, Art Owen, and Cun hui Zhang. One permutation hashing. *Advances in Neural Information Processing Systems*, 25:3113–3121, 2012.