To view the accompanying Technical Perspective, visit doi.acm.org/10.1145/3607825

Model Counting Meets Distinct Elements

By A. Pavan, N.V. Vinodchandran, Arnab Bhattacharyya, and Kuldeep S. Meel

Abstract

Constraint satisfaction problems (CSPs) and data stream models are two powerful abstractions to capture a wide variety of problems arising in different domains of computer science. Developments in the two communities have mostly occurred independently and with little interaction between them. In this work, we seek to investigate whether bridging the seeming communication gap between the two communities may pave the way to richer fundamental insights. To this end, we focus on two foundational problems: model counting for CSPs and the computation of the number of distinct elements in a data stream, also known as the zeroth frequency moment (F_0) of a data stream.

Our investigations lead us to observe striking similarity in the core techniques employed in the algorithmic frameworks that have evolved separately for model counting and distinct elements computation. We design a recipe for the translation of algorithms developed for distinct elements estimation to that of model counting, resulting in new algorithms for model counting. We then observe that algorithms in the context of distributed streaming can be transformed into distributed algorithms for model counting. We next turn our attention to viewing streaming from the lens of counting and show that framing distinct elements estimation as a special case of #DNF counting allows us to obtain a general recipe for a rich class of streaming problems, which had been subjected to case-specific analysis in prior works.

1. INTRODUCTION

Constraint Satisfaction Problems (CSP's) and a data stream model are two core themes in computer science with a diverse set of applications in topics including probabilistic reasoning, networks, databases, and verification. Model counting and computation of zeroth frequency moment (F_0) are fundamental problems for CSPs and a data stream model respectively. This paper is motivated by our observation that despite the usage of similar algorithmic techniques for the two problems, the developments in the two communities have, surprisingly, evolved separately, and rarely has a paper from one community been cited by the other.

Given a set of constraints φ over a set of variables in a finite domain \mathcal{D} , the problem of model counting is to estimate the number of solutions of φ . We are often interested when φ is restricted to a special class of representations such as Conjunctive Normal Form (CNF) and Disjunctive Normal Form (DNF). A data stream over a domain [N] is represented by $\mathbf{a} = \langle a_1, a_2, \cdots, a_m \rangle$ where each item a_i is a subset of [N]. The *zeroth frequency moment*, denoted as F_0 , of \mathbf{a} is the number of distinct domain elements appearing in \mathbf{a} ,

that is, $|\cup_i a_i|$ (traditionally, a_i s are singletons; we will also be interested in the case when a_i s are sets). The fundamental nature of model counting and F_0 computation over data streams has led to intense interest from theoreticians and practitioners alike in the respective communities for the past few decades.

The starting point of this work is the confluence of two viewpoints. The first viewpoint contends that some of the algorithms for model counting can conceptually be thought of as operating on the stream of the solutions of the constraints. The second viewpoint contends that a stream can be viewed as a DNF formula, and the problem of F_0 estimation is similar to model counting. These viewpoints make it natural to believe that algorithms developed in the streaming setting can be directly applied to model counting, and vice versa. We explore this connection and indeed design new algorithms for model counting inspired by algorithms for estimating F_0 in data streams. By exploring this connection further, we design new algorithms to estimate F_0 for streaming sets that are succinctly represented by constraints. It is worth noting that the two communities focus on seemingly different efficiency objectives: in streaming algorithms, space complexity is of major concern while in the context of model counting, time (especially NP query complexity in the context of CNF formulas) is of primary concern. Therefore, it is striking to observe that our transformation recipe leads to the design of *efficient* algorithms for F_0 estimation as well as model counting wherein efficient is measured by the concern of the corresponding community. We further investigate this observation and demonstrate that the space complexity of streaming algorithms provides an upper bound on the query complexity of model counting algorithms.

To put our contributions in context, we briefly survey the historical development of algorithmic frameworks in both model counting and F_0 estimation and point out the similarities.

1.1. Model counting

The complexity-theoretic study of model counting was initiated by Valiant who showed that this problem, in general, is #P-complete.³² This motivated researchers to investigate approximate model counting and in particular to design (ε , δ)-approximation schemes. The complexity of approximate

The original version of this paper is entitled "Model Counting Meets F_0 Estimation," and was published in the *Proceedings of the 40th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems, Virtual Event*, China, June 20–25, 2021; https://dl.acm. org/doi/10.1145/3452021.3458311

model counting depends on its representation. When the model φ is represented as a CNF formula φ , designing an efficient (ε , δ)-approximation is NP-hard.³⁰ In contrast, when it is represented as a DNF formula, model counting admits an fully polynomial-time approximation scheme (FPRAS).^{21, 22} We will use #CNF to refer to the case when φ is a CNF formula and #DNF to refer to the case when φ is a DNF formula.

For #CNF, Stockmeyer³⁰ provided a hashing-based randomized procedure that can compute an (ε , δ)-approximation with running time poly($|\varphi|, 1/\varepsilon, 1/\delta$), given access to an NP oracle. Building on Stockmeyer's approach and motivated by the unprecedented breakthroughs in the design of SAT solvers, researchers have proposed a series of algorithmic improvements that have allowed the hashing-based techniques for approximate model counting to scale to formulas involving hundreds of thousands of variables. The practical implementations substitute the NP oracle with an SAT solver. In the context of model counting, we are primarily interested in time complexity and therefore, the number of NP queries is of key importance. The emphasis on the number of NP calls also stems from practice as the practical implementation of model counting algorithms has shown to spend over 99% of their time in the underlying SAT calls.²⁹

Karp and Luby²¹ proposed the first FPRAS scheme for #DNF, which was improved in subsequent works.^{9, 22} Chakraborty et al.⁵ demonstrated that the hashing-based framework can be extended to #DNF, thereby providing a unified framework for both #CNF and #DNF. Meel et al.^{24, 25} subsequently improved the complexity of the hashing-based approach for #DNF and observed that hashing-based techniques achieve better scalability than Monte Carlo techniques.

1.2. Zeroth frequency moment estimation

Estimating (ε , δ)-approximation of the k^{th} frequency moments (F_k) of a stream is a central problem in a data streaming model.¹ In particular, considerable work has been done in designing algorithms for estimating the 0th frequency moment (F_0), the number of distinct elements in the stream. For streaming algorithms, the primary resource concerns are space complexity and processing time per element. In general, for a streaming algorithm to be considered efficient, these should be poly(log N, $1/\epsilon$) where N is the size of the universe (we assume δ to be a small constant and ignore $O(\log(\frac{1}{\delta}))$ factors in this discussion).

The first algorithm for computing F_0 with a constant factor approximation was proposed by Flajolet and Martin, who assumed the existence of hash functions with ideal properties resulting in an algorithm with undesirable space complexity.¹⁵ In their seminal work, Alon, Matias, and Szegedy designed an $O(\log N)$ space algorithm for F_0 with a constant approximation ratio that employs 2-universal hash functions.¹ Subsequent investigations into hashingbased schemes by Gibbons and Tirthapura¹⁶ and Bar-Yossef et al.³ provided (ε , δ)-approximation algorithms with space and time complexity $O(\log N) \cdot \text{poly}(\frac{1}{\varepsilon})$. Later, Bar-Yossef et al. proposed *three algorithms* with improved space and time complexity.² While the three algorithms employ hash functions, they differ conceptually in the usage of relevant random variables for the estimation of F_0 . This line of work resulted in the development of an algorithm with optimal space complexity $O(\log N + \frac{1}{\varepsilon^2})$ and $O(\log N)$ update time to estimate the F_0 of a stream.²⁰

The aforementioned works are in the setting where each data item a_i is an element of the universe. Subsequently, there has been a series of results of estimating F_0 in rich scenarios with a particular focus to handle the cases $a_i \subseteq \{1, 2, \dots, N\}$ such as a list or a multidimensional range.^{3, 26, 31}

1.3. The road to a unifying framework

As mentioned above, the algorithmic developments for model counting and F_0 estimation have largely relied on the usage of hashing-based techniques and yet these developments have, surprisingly, been separate, and rarely has a work from one community been cited by the other. In this context, we wonder whether it is possible to bridge this gap and if such an exercise would contribute to new algorithms for model counting as well as for F_0 estimation? The main conceptual contribution of this work is an affirmative answer to the above question. First, we point out that the two well-known algorithms; Stockmeyer's #CNF algorithm³⁰ which is further refined by Chakraborty et al.⁵ and Gibbons and Tirthapura's F_0 estimation algorithm, ¹⁶ are essentially the same.

The core idea of the hashing-based technique of Stockmeyer's and Chakraborty et al's scheme is to use pairwise independent hash functions to partition the solution space (satisfying assignments of a CNF formula) into roughly equal and small cells, wherein a cell is small if the number of solutions is less than a pre-computed threshold, denoted by Thresh. Then a good estimate for the number of solutions is the number of solutions in an arbitrary cell \times number of cells. To determine the appropriate number of cells, the solution space is iteratively partitioned as follows. At the m^{th} iteration, a hash function with range $\{0, 1\}^m$ is considered resulting in cells $h^{-1}(y)$ for each $y \in \{0, 1\}^m$. An NP oracle can be employed to check whether a particular cell (for example h^{-1} (0^m)) is small by enumerating solutions one by one until we have either obtained Thresh+1 number of solutions or we have exhaustively enumerated all the solutions. If the cell $h^{-1}(0^m)$ is small, then the algorithm outputs $t \times 2^m$ as an estimate where t is the number of solutions in the cell $h^{-1}(0^m)$. If the cell $h^{-1}(0^m)$ is not small, then the algorithm moves on to the next iteration where a hash function with range $\{0, 1\}^{m+1}$ is considered.

We now describe Gibbons and Tirthapura's algorithm for F_0 estimation which we call the Bucketing algorithm. Without loss of generality, we assume that N is a power of two and thus identify [N] with $\{0, 1\}^n$. The algorithm maintains a bucket of size Thresh and starts by picking a hash function $h : \{0, 1\}^n \rightarrow \{0, 1\}^n$. It iterates over sampling levels. At level m, when a data item x comes, if h(x) starts with 0^m , then x is added to the bucket. If the bucket overflows, then the sampling level is increased to m + 1 and all elements x in the bucket other than the ones with $h(x) = 0^{m+1}$ are deleted. At the end of the stream, the value $t \times 2^m$ is output as the estimate where t is the number of elements in the bucket and m is the sampling level.

These two algorithms are conceptually the same. In the Bucketing algorithm, at the sampling level m, it looks at only the first m bits of the hashed value; this is equivalent to considering

a hash function with range $\{0, 1\}^m$. Thus the bucket is nothing but all the elements in the stream that belong to the cell $h^{-1}(0^m)$. The final estimate is the number of elements in the bucket times the number of cells, identical to Chakraborty et al.'s algorithm. In both algorithms, to obtain an (ε, δ) approximation, the Thresh value is chosen as $O(\frac{1}{\varepsilon^2})$. To reduce the error probability to $1/\delta$, the median of $O(\log \frac{1}{\delta})$ independent estimations is output.

1.4. Our contributions

Motivated by the conceptual identity between the two algorithms, we further explore the connections between algorithms for model counting and F_0 estimation.

First, we formalize a recipe to transform streaming algorithms for F_o estimation to those for model counting. Such a transformation yields new (ε , δ)-approximate algorithms for model counting, which are different from currently known algorithms. Our transformation recipe from F_0 estimation to model counting allows us to view the problem of the design of distributed #DNF algorithms through the lens of distributed functional monitoring which is well-studied in a data streaming literature.

Building on the connection between model counting and F_0 estimation algorithms, we design new algorithms to estimate F_0 over structured set streams where each element of the stream is a (succinct representation of a) subset of the universe. Thus, the stream is S_1, S_2, \cdots where each $S_i \subseteq [N]$ and the goal is to estimate the F_0 of the stream, that is, the size of \cup_{S_i} . In this scenario, the goal is to design algorithms whose per-item time (time to process each S) is poly-logarithmic in the size of the universe. Structured set streams that are considered in the literature include 1-dimensional (1D) and multidimensional ranges.^{26, 31} Several interesting problems, including the max-dominance norm⁶ and counting triangles in graphs,³ can be reduced to computing F_0 over such ranges.

We observe that several structured sets can be represented as small DNF formulae and thus F_0 counting over these structured set data streams can be viewed as a special case of #DNF. Using the hashing-based techniques for #DNF, we obtain a general recipe for a rich class of structured sets that include DNF sets, affine spaces, and multidimensional ranges. Prior work on structured sets had to rely on involved analysis for each of the specific instances, while our work provides a general recipe for both analysis and implementation.

A natural question that arises from the transformation recipe is the relationship between the space complexity of the streaming algorithms and the query complexity of the obtained model counting algorithms. We establish a relationship between these two quantities by showing that the space complexity is an upper bound on the query complexity.

2. NOTATION

We will assume the universe $[N] = \{0, 1\}^n$.

 F_0 Estimation: A data stream **a** over domain [N] can be represented as $\mathbf{a} = a_1, a_2, ..., a_m$ wherein each item $a_i \in [N]$. Let $\mathbf{a}_{u} = \bigcup_{i} \{a_{i}\}$. F_{0} of the stream **a** is $|\mathbf{a}_{u}|$. We are interested in a probably approximately correct scheme that returns an (ε, δ) *estimate c* of F_0 , that is, $\Pr\left|\frac{|\mathbf{a}_{ll}|}{|\mathbf{1}^{l+\varepsilon}x_2|} \le c \le (1+\varepsilon)|\mathbf{a}_{ll}|\right| \ge 1-\delta$. *Model Counting:* Let $X = \{x_1^{r}, x_2^{r}, ..., x_n\}$ be a set of Boolean

variables. For a Boolean formula φ over variables X, let $Sol(\varphi)$ denote the set of all satisfying assignments of φ . The propositional model counting problem is to compute $|Sol(\varphi)|$ for a given formula φ . As in the case of F_{α} , we are interested in a probably approximately correct algorithm that takes as inputs a formula φ , a tolerance $\varepsilon > 0$, and a confidence $\delta \in (0, 1]$, and returns a (ε, δ) -estimate c of $|Sol(\varphi)|$ that is,

$$\Pr\left[\frac{|\operatorname{sol}(\varphi)|}{1+\varepsilon} \le c \le (1+\varepsilon) |\operatorname{Sol}(\varphi)|\right] \ge 1 - \delta.$$

k-wise Independent hash functions: Let $n, m \in \mathbb{N}$ and $\mathcal{H}(n, m) \triangleq$ ${h: \{0, 1\}^n \rightarrow \{0, 1\}^m}$ be a family of hash functions mapping $\{0,1\}^n$ to $\{0,1\}^m$.

DEFINITION 1. A family of hash functions $\mathcal{H}(n, m)$ is k-wise independent, denoted $\mathcal{H}_{k-wise}(n, m)$, if $\forall \alpha_1, \alpha_2, ..., \alpha_k \in \{0, 1\}^m$, for all distinct $x_1, x_2, \dots x_k \in \{0, 1\}^n$,

$$\Pr_{h \in \mathcal{H}(n,m)} \left[\left(h(x_1) = \alpha_1 \right) \land \left(h(x_2) = \alpha_2 \right) \dots \left(h(x_k) = \alpha_k \right) \right] = \frac{1}{2^{km}}$$

Explicit families. An explicit hash family that we use is $\mathcal{H}_{\text{Toeplitz}}$ (n, m), which is known to be 2-wise independent.⁴ The family is defined as follows: $\mathcal{H}_{\text{Toeplitz}}(n,m) \triangleq \{h: \{0,1\}^n \rightarrow \{0,1\}^m\}$ is the family of functions of the form h(x) = Ax + b with $A \in \mathbb{F}_2^{m \times n}$ and $b \in \mathbb{F}_2^{m \times 1}$, and \mathbb{F}_2 is the finite field of size 2. For a hash function $h: \{0, 1\}^n \to \{0, 1\}^m, h_\ell: \{0, 1\}^n \to \{0, 1\}_\ell, \ell \in \{1, ..., m\}, \text{ is the}$ function where $h_{\ell}(y)$ is the first ℓ bits of h(y).

3. FROM STREAMING TO COUNTING

As a first step, we present a unified view of the three hashing-based algorithms proposed in Bar-Yossef et al.² Their first algorithm, the Bucketing algorithm discussed above, is a modification of an F_0 estimation algorithm due to Gibbons and Tirthapura.¹⁶ The second algorithm, which we call Minimum, is based on the idea that if we hash all the items of the stream, then $\mathcal{O}(1/\varepsilon^2)$ -th minimum of the hash values can be used to compute a good estimate of F_0 . The third algorithm, which we call Estimation, chooses a set of k functions, $\{h_1, h_2, ..., h_k\}$, such that each h_i is picked randomly from an $\mathcal{O}(\log(1/\varepsilon))$ -independent hash family. For each hash function h_i , we say that h_i is not *lonely* if there exists $a_i \in \mathbf{a}$ such that $h_i(a_i) = 0$. One can then estimate F_0 of **a** by estimating the number of hash functions that are not lonely.

Algorithm 1, called ComputeF0, presents the overarching architecture of the three proposed algorithms. The architecture of ComputeF0 is fairly simple: it chooses a collection of hash functions using ChooseHashFunctions, calls the subroutine ProcessUpdate for every incoming element of the stream and invokes ComputeEst at the end of the stream to return the F_0 approximation.

ChooseHashFunctions. As shown in Algorithm 2, the hash functions depend on the strategy being implemented. The subroutine PickHashFunctions(\mathcal{H} , t) returns a collection of t independently chosen hash functions from the family \mathcal{H} . We use *H* to denote the collection of hash functions returned, this collection is viewed as either a 1D array or as a 2-dimensional (2D) array. When *H* is a 1D array, *H*[*i*] denotes the *i*th hash function of the collection and when H is a 2D array H[i][j] is the [i, j]th hash function.

ProcessUpdate. For a new item x, the update of S, as shown in Algorithm 3 is as follows:

- Bucketing For a new item x, if $H[i]_{m_i}(x) = 0^{m_i}$, then we add it to S[i] if x is not already present in S[i]. If the size of S[i] is greater than Thresh (which is set to be $O(1/\varepsilon^2)$), then we increment the m_i as in line 8.
- Minimum For a new item x, if H[i](x) is smaller than max S[i], then we replace max S[i] with H[i](x).
- Estimation For a new item *x*, compute z = TrailZero(H[i, j](x)), that is, the number of trailing zeros in H[i, j](x), and replace S[i, j] with *z* if *z* is larger than S[i, j].

ComputeEst. Finally, for each of the algorithms, we estimate F_0 based on the sketch S as described in the sub-routine ComputeEst presented as Algorithm 4. It is crucial to note that the estimation of F_0 is performed solely using the sketch S for the Bucketing and Minimum algorithms. The Estimation algorithm requires an additional parameter r that depends on a loose estimate of F_0 .

Algorithm 1 ComputeF0(n, ε , δ)

1: Thresh $\leftarrow 96/\varepsilon^2$ 2: $t \leftarrow 35 \log(1/\delta)$ 3: $H \leftarrow \text{ChooseHashFunctions}(n, \text{Thresh}, t)$ 4: $S \leftarrow \{\}$ 5: while true do 6: if EndStream then exit; 7: $x \leftarrow input$ () 8: ProcessUpdate(S, H, x, Thresh) 9: $Est \leftarrow \text{ComputeEst}(S, \text{Thresh})$ 10: return Est

Sketch Properties. For each of the three algorithms, their corresponding sketches can be viewed as arrays of size 35 $\log(1/\delta)$. The parameter Thresh is set to $96/\varepsilon^2$.

- Bucketing The element S[i] is a tuple $\langle \ell_i, m_i \rangle$ where ℓ_i is a list of size at most Thresh, where $\ell_i = \{x \in \mathbf{a} | H[i]_{m_i}(x) = 0^{m_i}\}$. We use S[i](0) to denote ℓ_i and S[i](1) to denote m_i .
- Minimum Each S[i] holds the lexicographically Thresh many smallest elements of $\{H[i](x) \mid x \in \mathbf{a}\}$.
- Estimation Each S[i] holds a tuple of size Thresh. The *j*'th entry of this tuple is the largest number of trailing zeros in any element of $H[i, j](\mathbf{a})$.

Algorithm 2 ChooseHashFunctions(*n*, Thresh, *t*)

```
1: switch AlgorithmType do
        case AlgorithmType==Bucketing
2:
3:
             H \leftarrow \mathsf{PickHashFunctions}(\mathcal{H}_{\mathsf{Toeplitz}}(n, n), t)
4:
         case AlgorithmType==Minimum
             H \leftarrow \mathsf{PickHashFunctions}(\mathcal{H}_{\mathsf{Toeplitz}}(n, 3n), t)
5:
         case AlgorithmType==Estimation
6:
7:
              s \leftarrow 10 \log(1/\varepsilon)
                   \leftarrow PickHashFunctions(\mathcal{H}_{s-wise} (n, n), t \times
8:
             H
    Thresh)
    return H
```

Algorithm 3 ProcessUpdate(S, H, x, Thresh)

```
1: for i \in [1, |H|] do
 2:
          switch AlgorithmType do
 3:
               case Bucketing
 4:
                    m_i = \mathcal{S}[i](0)
                    if H[i]_{m}(x) = 0^{m_{i}} then
 5:
 6:
                        \mathcal{S}[i](0) \leftarrow \mathcal{S}[i](0) \cup \{x\}
 7:
                        if size(S[i](0)) > Thresh then
 8:
                             S[i](1) \leftarrow S[i](1) + 1
 9:
                            for y \in S do
                                 if H[i]_{m_{i}+1}(y) \neq 0^{m_{i}+1} then
10:
11:
                                        Remove(\mathcal{S}[i](0), y)
12:
               case Minimum
13:
                    if size(S[i]) < Thresh then
14:
                         S[i].Append(H[i](x))
15:
                    else
16:
                         j \leftarrow \arg \max(S[i])
17:
                         if \mathcal{S}[i](j) > H[i](x) then
18:
                              \mathcal{S}[i](j) \leftarrow H[i](x)
19:
                case Estimation
20:
                    for i \in [1, \text{Thresh}] do
                        S[i, j] \leftarrow \max(S[i, j], \operatorname{TrailZero}(H[i, j](x)))
21:
22: return S
```

3.1. A recipe for transformation

Observe that for each of the algorithms, the final computation of F_0 estimation depends on the sketch S. Therefore, as long as for two streams **a** and **â**, if their corresponding sketches (and the hash functions chosen) match, then the three schemes presented above would return the same estimates. The recipe for a transformation of streaming algorithms to model counting algorithms is based on the following insight:

- 1. Capture the relationship $\mathcal{P}(\mathcal{S}, H, \mathbf{a}_u)$ between the sketch \mathcal{S} , set of hash functions H, and set \mathbf{a}_u at the end of stream.
- 2. View the formula φ as a symbolic representation of the unique set \mathbf{a}_u represented by the stream \mathbf{a} such that $Sol(\varphi) = \mathbf{a}_u$.
- Given a formula φ and set of hash functions *H*, design an algorithm to construct sketch S such that the property P(S, H, Sol(φ)) holds. Using the sketch S, |Sol(φ)| can be estimated.

Algorithm 4 ComputeEst(S, Thresh)



By applying the above recipe to the three F_0 estimation algorithms, we can derive corresponding model counting algorithms. In particular, applying the above recipe to the Bucketing algorithm leads us to the state-of-the-art hashingbased model counting algorithm, ApproxMC, proposed by Chakraborty et al.⁵ Applying the above recipe to Minimum and Estimation allows us to obtain different model counting schemes. In this extended abstract, we illustrate this transformation for the Minimum-based algorithm.

3.2. Example application of recipe: Minimum-based algorithm

We showcase the application of the recipe in the context of a minimum-based algorithm. For a given multiset **a** (e.g., a data stream or solutions to a model), we now specify the property $\mathcal{P}(S, H, \mathbf{a}_u)$ as follows: The sketch *S* is an array of sets indexed by members of *H* that holds lexicographically *p* minimum elements of *H* [*i*](\mathbf{a}_u) where *p* is min $(\frac{96}{\varepsilon^2}, |\mathbf{a}_u|)$. \mathcal{P} is the property that specifies this relationship.

The following lemma due by Bar-Yossef et al.² establishes the relationship between the property \mathcal{P} and the number of distinct elements of a multiset. Let max(S_i) denote the largest element of the set S_i .

LEMMA 1. Let $\mathbf{a} \subseteq \{0, 1\}^n$ be a multiset. Let $H \subseteq \mathcal{H}_{\text{Toeplitz}}(n, 3n)$ where each H[i] is independently drawn from $\mathcal{H}_{\text{Toeplitz}}(n, 3n)$, and $|H| = O(\log 1/\delta)$. Let S be such that $\mathcal{P}(S, H, a_u)$ holds. Let $c = Median\left\{\frac{p \cdot 2^m}{\max(S[i])}\right\}_i$. Then

$$\Pr\left[\frac{|\mathbf{a}_u|}{(1+\varepsilon)} \le c \le (1+\varepsilon)|\mathbf{a}_u|\right] \ge 1-\delta.$$

Therefore, we can transform the minimum algorithm for F_0 estimation to that of model counting given access to a subroutine that can compute S such that $\mathcal{P}(S, H, Sol(\varphi))$ holds. The following proposition establishes the existence and complexity of such a subroutine, called FindMin.

PROPOSITION 1. There is an algorithm FindMin that, given φ over n variables, $h \in \mathcal{H}_{\text{Toeplitz}}(n, m)$, and p as input, returns a set, $\mathcal{B} \subseteq h(\text{Sol}(\varphi))$ so that if $|h(\text{Sol}(\varphi))| \leq p$, then $\mathcal{B} = h(\text{Sol}(\varphi))$, otherwise \mathcal{B} is the p lexicographically minimum elements of $h(\text{Sol}(\varphi))$. Moreover, if φ is a CNF formula, then FindMin makes $\mathcal{O}(p \cdot m)$ calls to an NP oracle, and if φ is a DNF formula with kterms, then FindMin takes $\mathcal{O}(m^3 \cdot n \cdot k \cdot p)$ time.

Equipped with Proposition 1, we are now ready to present the algorithm, called ApproxModelCountMin, for model counting. Since the complexity of FindMin is PTIME when φ is in DNF, we have ApproxModelCountMin as a FPRAS for DNF formulas.

Algorithm 5 ApproxModelCountMin(φ , ε , δ)

1: $t \leftarrow 35 \log(1/\delta)$ 2: $H \leftarrow \text{PickHashFunctions}(\mathcal{H}_{\text{Toeplitz}}(n, 3n), t)$ 3: $S \leftarrow \{\}$ 4: Thresh $\leftarrow \frac{96}{\varepsilon^2}$ 5: **for** $i \in [1, \overline{t}]$ **do** 6: $S[\overline{i}] \leftarrow \text{FindMin}(\varphi, H[\overline{i}], \text{Thresh})$ 7: $Est \leftarrow \text{Median}\left(\left\{\frac{\text{Thresh} \times 2^{3n}}{\max\{s[\overline{i}]\}}\right\}_{i}\right)$ THEOREM 1. Given φ , ε , δ , ApproxModelCountMin returns Est such that

$$\Pr\left(\frac{\left|\mathsf{Sol}(\varphi)\right|}{1+\varepsilon} \le Est \le (1+\varepsilon)\left|\mathsf{Sol}(\varphi)\right|\right) \ge 1-\delta.$$

If φ is a CNF formula, then ApproxModelCountMin is a polynomial-time algorithm that makes $\mathcal{O}\left(\frac{1}{\varepsilon^2} n \log\left(\frac{1}{\delta}\right)\right)$ calls to an NP oracle. If φ is a DNF formula, then ApproxModelCountMin is an FPRAS.

We now give proof of Proposition 1 by describing the subroutine FindMin.

PROOF. We first present the algorithm when the formula φ is a DNF formula. Adapting the algorithm for the case of CNF can be done by using similar ideas. Let $\phi = T_1 \lor T_2 \lor \cdots \lor T_k$ be a DNF formula over n variables where T_i is a term. Let $h: \{0, 1\}^n \to \{0, 1\}^m$ be a linear hash function in $\mathcal{H}_{\text{Toeplitz}}(n, m)$ defined by a $m \times n$ binary matrix A. Let C be the set of hashed values of the satisfying assignments for $\varphi: C = \{h(x) \mid x \models \varphi\} \subseteq \{0, 1\}^m$. Let C_p be the first p elements of C in the lexicographic order. Our goal is to compute C_p .

We illustrate an algorithm with running time $O(m^3np)$ to compute C_p when the formula is just a term T. This algorithm can easily be generalized to formulas with k-terms. Let T be a term with width w (number of literals) and $C = \{Ax \mid x \models T\}$. By fixing the variables in T we get a vector b_T and an $N \times (n - w)$ matrix A_T so that $C = \{A_T x + b_T \mid x \in \{0, 1\}^{(n-w)}\}$. Both A_T and b_T can be computed from A and T in linear time. Let $h_T(x)$ be the transformation $A_T x + b_T$.

We will compute C_p iteratively as follows: assuming we have computed the $(q-1)^{th}$ minimum of C, we will compute the q^{th} minimum using a prefix-search strategy. We will use a subroutine to solve the following basic prefix-search primitive: Given any l bit string $y_1 \dots y_l$, is there an $x \in \{0, 1\}^{n-w}$ so that $y_1 \dots y_l$ is a prefix for some string in $\{h_T(x)\}$? This task can be performed using Gaussian elimination over an $(l+1) \times (n-w)$ binary matrix and can be implemented in time $O(l^2(n-w))$.

Let $y = y_1 \dots y_m$ be the $(q-1)^{th}$ minimum in C. Let r_1 be the rightmost 0 of y. Then using the above-mentioned procedure we can find the lexicographically smallest string in the range of h_T that extends $y_1 \dots y_{(r-1)}$ 1 if it exists. If no such string exists in C, find the index of the next 0 in y and repeat the procedure. In this manner the q^{th} minimum can be computed using O(m) calls to the prefix-searching primitive resulting in an $O(m^3n)$ time algorithm. Invoking the above procedure p times results in an algorithm to compute C_n in $O(m^3np)$ time.

We now discuss distributed DNF counting problem.

3.3. Distributed DNF counting

Consider the problem of *distributed DNF counting*. In this setting, there are *k* sites that can each communicate with a central coordinator. The input DNF formula φ is partitioned into *k* DNF subformulas $\varphi_1, ..., \varphi_k$, where each φ_i is a subset of the terms of the original φ , with the *j*'th site receiving only φ_j . The goal is for the coordinator to obtain an (ϵ, δ) -approximation of the number of solutions to φ , while minimizing the total number of bits communicated between the sites

and the coordinator. Distributed algorithms for sampling and counting solutions to CSP's have been studied recently in other models of distributed computation.¹¹⁻¹⁴ From a practical perspective, given the centrality of #DNF in the context of probabilistic databases,^{27, 28} a distributed DNF counting algorithm would entail applications in distributed probabilistic databases.

From our perspective, distributed DNF counting falls within the *distributed functional monitoring* framework formalized by Cormode et al.⁷ Here, the input is a stream **a** which is partitioned arbitrarily into sub-streams $\mathbf{a}_1, ..., \mathbf{a}_k$ that arrive at each of *k* sites. Each site can communicate with the central coordinator, and the goal is for the coordinator to compute a function of the joint stream **a** while minimizing the total communication. This general framework has several direct applications and has been studied extensively (see Cormode et al.,⁸ Huang et al.,¹⁸ Woodruff and Zhang³³ and the references therein). In distributed DNF counting problem, each sub-stream \mathbf{a}_i corresponds to the set of satisfying assignments to each subformula φ_i , while the function to be computed is F_0 .

The algorithms discussed in Section 3 can be extended to the distributed setting. We briefly describe the distributed implementation of the minimum-based algorithm.

Distributed implementation of the minimum-based algorithm. The coordinator chooses hash functions H[1], ..., H[t] from $\mathcal{H}_{\mathsf{Toeplitz}}(n, 3n)$ and sends them to the k sites. Each site runs the FindMin algorithm for each hash function and sends the outputs to the coordinator. So, the coordinator receives sets S[i, j], consisting of the Thresh lexicographically smallest hash values of the solutions to φ_j . The coordinator then extracts S[i], the Thresh lexicographically smallest elements of $S[i, 1] \cup \cdots \cup S[i, k]$, and proceeds with the rest of the algorithm ApproxModelCountMin. The communication cost is $O(kn/\varepsilon^2 \cdot \log(1/\delta))$ to account for the k sites sending the outputs of their FindMin invocations. The time complexity for each site is polynomial in n, ε^{-1} , and $\log(\delta^{-1})$.

A straightforward implementation of the Bucketing algorithm leads to a distributed DNF counting algorithm whose communication cost is $\tilde{O}(k(n + 1/\varepsilon^2) \cdot \log(1/\delta))$ and time complexity per site is polynomial in n, ε^{-1} , and $\log(\delta^{-1})$. Similarly, the estimation-based algorithm leads to a distributed algorithm with $\tilde{O}(k(n + 1/\varepsilon^2) \log(1/\delta))$ communication cost. However, we do not know a polynomial time algorithm to implement the last algorithm on DNF terms.

3.4. Lower bound

The communication cost for the Bucketing and Estimationbased algorithms is nearly optimal in their dependence on k and ε . Woodruff and Zhang³³ showed that the randomized communication complexity of estimating F_0 up to a $1 + \varepsilon$ factor in the distributed functional monitoring setting is $\Omega(k/\varepsilon^2)$. We can reduce the F_0 estimation problem to distributed DNF counting. Namely, if for the F_0 estimation problem, the *j*'th site receives items $a_1, ..., a_m \in [N]$, then for the distributed DNF counting problem, φ_j is a DNF formula on $\lceil \log_2 N \rceil$ variables whose solutions are exactly $a_1, ..., a_m$ in their binary encoding. Thus, we immediately get an $\Omega(k/\varepsilon^2)$ lower bound for the distributed DNF counting

ective, distributed DNF counting falls ad functional monitoring framework forle et al.⁷ Here, the input is a stream **a** d arbitrarily into sub-streams **a**, ..., **a** is to design efficient algorithms (both in terms of memory and processing time per item) for computing $|\cup_i S_i|$ —the number of distinct elements in the union of all the sets in the stream. We call this problem *F* computation over struc-

the stream. We call this problem F_0 computation over structured set streams. We discuss two types of structured sets *DNF Sets* and *Affine Spaces*. As we mentioned in the introduction, other structured sets studied in the literature are single and multi-dimensional ranges. Our techniques also give algorithms for estimating F_0 of such structured set streams, which we omit in this extended abstract.

problem. Finding the optimal dependence on *N* for k > 1

In this section, we consider the structured set streaming

model where each item S_i of the stream is a succinct rep-

resentation of a set over the universe $U = \{0, 1\}^n$. Our goal

remains an interesting open question.

4. FROM COUNTING TO STREAMING

4.1. DNF sets

A particular representation we are interested in is where each set is presented as the set of satisfying assignments to a DNF formula. Let φ be a DNF formula over *n* variables. Then the DNF set corresponding to φ be the set of satisfying assignments of φ . The *size* of this representation is the number of terms in the formula φ . A stream over DNF sets is a stream of DNF formulas $\varphi_1, \varphi_2, \ldots$ Given such a DNF stream, the goal is to estimate $|\bigcup_i S_i|$ where S_i the DNF set represented by φ_i . This quantity is the same as the number of satisfying assignments of the formula $\forall_i \varphi_i$. We show that the algorithms described in the previous section carry over to obtain (ϵ , δ) estimation algorithms for this problem with space and per-item time poly($1/\epsilon$, *n*, *k*, log($1/\delta$)) where *k* is the size of the formula.

THEOREM 2. There is a streaming algorithm to compute an (ϵ, δ) approximation of F_0 over DNF sets. This algorithm takes space $O\left(\frac{n}{\varepsilon^2} \cdot \log \frac{1}{\delta}\right)$ and processing time $O\left(n^4 \cdot k \cdot \frac{1}{\varepsilon^2} \cdot \log \frac{1}{\delta}\right)$ per item where k is the size (number of terms) of the corresponding DNF formula.

PROOF. We show how to adapt the Minimum-value based algorithm from Section 3.2 to this setting. The algorithm picks a hash function $h \in \mathcal{H}_{\text{Toenlitz}}(n, 3n)$ and maintains the set \mathcal{B} consisting of t lexicographically minimum elements of the set $\{h (Sol(\varphi_1 \lor ... \lor \varphi_{i-1}))\}$ after processing *i*-1 items. When φ_i arrives, it computes the set \mathcal{B}' consisting of the t lexicographically minimum values of the set $\{h (Sol(\varphi_i))\}\$ and subsequently updates \mathcal{B} by computing the *t* lexicographically smallest elements from $\mathcal{B} \cup \mathcal{B}'$. By Proposition 1, the computation of \mathcal{B}' can be done in time *O* $(n^4 \cdot k \cdot t)$ where k is the number of terms in φ_{i} . Updating \mathcal{B} can be done in $O(t \cdot n)$ time. Thus, the update time for item φ_i is $O(n^4 \cdot k \cdot t)$. For obtaining an (ε, δ) -approximation, we set $t = O\left(\frac{1}{\lambda}\right)$ and repeat the procedure $O\left(\log\frac{1}{\lambda}\right)$ times and take the median value. Thus the update time for the item φ is $O\left(n^4 \cdot k \cdot \frac{1}{c^2} \cdot \log \frac{1}{c}\right)$. For analyzing space, each hash function uses $O(\tilde{n})$ bits and the algorithm stores $O(\frac{1}{\epsilon^2})$ minimums, resulting in overall space usage of $O\left(\frac{n}{\varepsilon^2} \cdot \log \frac{1}{\delta}\right)$. The proof of correctness follows from Lemma 1. Instead of the Minimum-value based algorithm, we could also adapt the Bucketing-based algorithm to obtain an algorithm with similar space and time complexities.

4.2. Affine spaces

Another example of a structured stream is where each item of the stream is an affine space represented by Ax = B where A is a Boolean matrix and B is a zero-one vector. Without loss of generality, we may assume that A is a $n \times n$ matrix. Thus an affine stream consists of $\langle A_1, B_1 \rangle$, $\langle A_2, B_2 \rangle$, ..., where each $\langle A_i, B_i \rangle$ succinctly represents a set $\{x \in \{0, 1\}^n | A_i x = B_i\}$. Here operations are over the finite field of size 2. For an $n \times n$ Boolean matrix A and a zero-one vector B, let Sol $\langle A, B \rangle$) denote the set of all x that satisfy Ax = B.

PROPOSITION 2. Given (A, B), $h \in \mathcal{H}_{\text{Toeplitz}}(n, 3n)$, and t as input, there is an algorithm, AffineFindMin, that returns a set, $\mathcal{B} \subseteq h(\text{Sol}(\langle A, B \rangle))$ so that if $|h(\text{Sol}(\langle A, B \rangle))| \leq t$, then $\mathcal{B} = h$ $(\text{Sol}(\langle A, B \rangle))$, otherwise \mathcal{B} is the t lexicographically minimum elements of $h(\text{Sol}(\langle A, B \rangle))$. The time taken by this algorithm is $O(n^4t)$ and the space taken by the algorithm is O(tn).

The above proposition together with the minimum-based algorithm gives the following theorem.

THEOREM 3. There is a streaming algorithm that computes a (ϵ, δ) -approximation of F_0 over affine spaces. This algorithm takes space $O\left(\frac{n}{\epsilon^2} \cdot \log(1/\delta)\right)$ and processing time of $O\left(n^4 \frac{1}{\epsilon^2} \log(1/\delta)\right)$ per item.

5. RELATING SKETCH SPACE COMPLEXITY AND NP QUERY COMPLEXITY

Our investigations reveal surprising connections between algorithms for F_0 estimation and model counting that are of interest to two different communities. It is noteworthy that the two communities often have different concerns: in the context of model counting, one is focused on the NP-query complexity while in the context of streaming, the focus is on the space complexity. This begs the question of whether the connections are a matter of happenstance or there is an inherent relationship between the space complexity in the context of streaming and the query complexity for model counting. We detail our investigations on the existence of such a relationship.

In the following, we will fold the hash function *h* also in the sketch *S*. With this simplification, instead of writing $P(S, h, Sol(\varphi))$ we write $P(S, Sol(\varphi))$.

We first introduce some complexity-theoretic notation. For a complexity class C, a language L belongs to the complexity class $\exists \cdot C$ if there is a polynomial $q(\cdot)$ and a language $L' \in C$ such that for every $x, x \in L \Leftrightarrow \exists y, |y| \le q(|x|), \langle x, y \rangle \in L'$.

Consider a streaming algorithm for F_0 that constructs a sketch such that $P(S, a_u)$ holds for some property P using which we can estimate $|a_u|$, where the size of S is polylogarithmic in the size of the universe and polynomial in $1/\varepsilon$. Now consider the following *Sketch-Language*

$$L_{sketch} = \{ \langle \varphi, S \rangle | P(S, Sol(\varphi)) \text{ holds} \}.$$

THEOREM 4. If L_{sketch} belongs to the complexity class C, then there exists a FP^{3.C} model counting algorithm that estimates the

number of satisfying assignments of a given formula φ . The number of queries made by the algorithm is bounded by the sketch size.

Let us apply the above theorem to the minimum-based algorithm. The sketch language consists of tuples of the form $\langle \varphi, \langle h, v_1, \dots, v_i \rangle$ where $\{v_1, \dots, v_i\}$ is the set of *t* lexicographically smallest elements of the set $h(\text{Sol}(\varphi))$. It can be seen that this language is in coNP. Since $\exists \cdot \text{coNP}$ is the same as the class Σ_2^P , we obtain a $\text{FP}^{\Sigma_2^P}$ algorithm. Since $t = O(1/\varepsilon^2)$ and *h* maps from *n*-bit strings to 3n-bit strings, it follows that the size of the sketch is $O(n/\varepsilon^2)$. Thus the number of queries made by the algorithm is $O(n/\varepsilon^2)$.

Interestingly, all the model counting algorithms that were obtained following our recipe are probabilistic polynomialtime algorithms that make queries to languages in NP. The above generic transformation gives a deterministic polynomialtime algorithm that makes queries to a $\Sigma_2^{\rm p}$ language. Precisely characterizing the properties of the sketch that lead to probabilistic algorithms making only NP queries is an interesting direction to explore.

6. CONCLUSION AND FUTURE OUTLOOK

Our investigation led to a diverse set of results that unify over two decades of work in model counting and F_0 estimation. The viewpoint presented in this work has the potential to spur several new interesting research directions.

Higher Moments. There has been a long line of work on the estimation of higher moments, that is, F_k over data streams. A natural direction of future research is to adapt the notion of F_k in the context of the model of counting and explore its applications. We expect extensions of the framework and recipe presented in this work to derive algorithms for higher frequency moments in the context of model counting.

Sparse XORs. In the context of model counting, the performance of underlying SAT solvers strongly depends on the size of XORs. The standard constructions lead to XORs of size $\Theta(n)$ and an interesting line of research has focused on the design of sparse XOR-based hash functions^{10,17,19} culminating in showing that one can use hash functions of the form h(x) = Ax + b wherein each entry of the *m*-th row of *A* is 1 with probability $O(\frac{\log m}{m})^{23}$ Such XORs were shown to improve runtime efficiency. In this context, a natural direction would be to explore the usage of sparse XORs in the context of F_0 estimation.

Acknowledgments

We thank the anonymous reviewers of PODS 21 for their valuable comments. We are grateful to Phokion Kolaitis for suggesting exploration beyond the transformation recipe that led to results in Section 5. We thank Wim Martens for providing valuable suggestions on an earlier version of the manuscript. Bhattacharyya was supported in part by the NRF Fellowship Programme [NRF-NRFFAI1-2019-0002] and an Amazon Research Award. Meel was supported in part by the NRF Fellowship Programme [NRF-NRFFAI1-2019-0004] and the AI Singapore Programme [AISG-RP-2018-005], and NUS ODPRT Grant [R-252-000-685-13]. Vinod was supported in part by NSF CCF-2130608, NSF CCF-184908, and NSF HDR:TRIPODS-1934884 awards. Pavan was supported in part by NSF CCF-2130536, NSF CCF-1849053, and NSF HDR:TRIPODS-1934884 awards.

References

- Alon, N., Matias, Y., Szegedy, M. The space complexity of approximating the frequency moments. *J. Comput. Syst. Sci.* 58, 1 (1999), 137–147.
- Bar-Yossef, Z., Jayram, T.S., Kumar, R., Sivakumar, D., Trevisan, L. Counting distinct elements in a data stream. In *Volume 2483 of Proceedings of RANDOM* (2002). Soringer Cambridge USA 1–10
- (2002), Springer, Cambridge, USA, 1–10.
 Bar-Yossef, Z., Kumar, R., Sivakumar, D. Reductions in streaming algorithms, with an application to counting triangles in graphs. In *Proceedings of SODA* (2002), ACM/SIAM, NY, 623–632.
- Carter, J.L., Wegman, M.N. Universal classes of hash functions. In Proceedings of the ^{gm} Annual ACM Symposium on Theory of Computing (1977), ACM, NY, 106–112.
- Chakraborty, S., Meel, K.S., Vardi, M.Y. Algorithmic improvements in approximate counting for probabilistic inference: From linear to logarithmic SAT calls. In *Proceedings of IJCAI* (2016), IJCAI/AAAI Press, New York, USA.
- Cormode, G., Muthukrishnan, S. Estimating dominance norms of multiple data streams. In *Proceedings* of ESA, Volume 2832 of Lecture Notes in Computer Science. G.D. Battista and U. Zwick, eds. Springer, Budapest, Hungary, 2003, 148–160.
- Cormode, G., Muthukrishnan, S., Yi, K. Algorithms for distributed functional monitoring. ACM Trans. Algorithms (TALG) 7, 2 (2011), 1–20.
- Cormode, G., Muthukrishnan, S., Yi, K., Zhang, Q. Continuous sampling from distributed streams. J. ACM (JACM) 59, 2 (2012), 1–25.
- 9. Dagum, P., Karp, R., Luby, M., Ross, S. An optimal algorithm for monte carlo

estimation. SIAM J. Comput. 29, 5 (2000), 1484–1496.

- Érmon, S., Gomes, C.P., Sabharwal, A., Selman, B. Low-density parity constraints for hashing-based discrete integration. In *Proceedings of ICML* (2014), JMLR, Beijing, China, 271–279.
- Feng, W., Hayes, T.P., Yin, Y. Distributed symmetry breaking in sampling (optimal distributed randomly coloring with fewer colors). arXiv preprint arXiv:1802.06953 (2018).
- Feng, W., Sun, Y., Yin, Y. What can be sampled locally? *Distrib. Comput. 33* (2018), 1–27.
- Feng, W., Yin, Y. On local distributed sampling and counting. In Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing (2018), ACM, NY, 189–198.
- Fischer, M., Ghaffari, M. A simple parallel and distributed sampling technique: Local glauber dynamics. In 3^{2rd} International Symposium on Distributed Computing (2018) Schloss Dagstuhl - Leibniz-Zentrum für Informatik. New Orleans. USA.
- Flajolet, P., Martin, G.N. Probabilistic counting algorithms for data base applications. *J. Comput. Syst. Sci. 31*, 2 (1985), 182–209.
- Gibbons, P.B., Tirthapura, S. Estimating simple functions on the union of data streams. In *Proceedings of SPAA*. A. L. Rosenberg, ed. ACM, NY, 2001, 281–291.
- Gomes, C.P., Hoffmann, J., Sabharwal, A., Selman, B. From sampling to model counting. In *Proceedings of IJCAI* (2007), IJCAI/AAAI Press,
- Hyderabad, India, 2293–2299. 18. Huang, Z., Yi, K., Zhang, Q. Randomized algorithms for tracking distributed

count, frequencies, and ranks. In *Proceedings of PODS* (2012), ACM, Scottsdale, USA 295–306.

- Ivrii, A., Malik, S., Meel, K.S., Vardi, M.Y. On computing minimal independent support and its applications to sampling and counting. *Constraints An Int. J. 21*, 1 (2016), 41–58.
- Kane, D.M., Nelson, J., Woodruff, D.P. An optimal algorithm for the distinct elements problem. In *Proceedings of PODS* (2010), ACM, NY, 41–52.
- Karp, R., Luby, M. Monte-carlo algorithms for enumeration and reliability problems. In *Proceedings* of *FOCS* (1983), IEEE Computer Society, Arizona, USA.
- Karp, Ř.M., Luby, M., Madras, N. Monte-carlo approximation algorithms for enumeration problems. J. Algorithms 10, 3 (1989), 429–448.
- Meel, K.S., Akshay, S. Sparse hashing for scalable approximate model counting: Theory and practice. In *Proceedings of LICS* (2020) ACM, Saarbrücken, Germany.
- Meel, K.S., Shrotri, A.Á., Vardi, M.Y. On hashing-based approaches to approximate dnf-counting. In *Proceedings of FSTTCS* (2017) Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Kanpur, India.
- 25. Meel, K.S., Shrotri, A.A., Vardi, M.Y. Not all fprass are equal: Demystifying fprass for dnf-counting (extended

A. Pavan, Iowa State University, Ames, Iowa, USA.

N.V. Vinodchandran, University of Nebraska-Lincoln, Lincoln, Nebraska, USA. abstract). In *Volume 8 of Proceedings* of *IJCAI* (2019), IJCAI, Macau, China. Pavan, A., Tirthapura, S. Rangeefficient counting of distinct elements

26.

- in a massive data stream. *SIAM J. Comput.* 37, 2 (2007), 359–379. 27. Ré, C., Suciu, D. Approximate lineage for probabilistic databases. *Proc. VLDB*
- Endowment 1, 1 (2008), 797–808.
 Senellart, P. Provenance and probabilities in relational databases.
- ACM SIGMOD Rec. 46, 4 (2018), 5–15.
 Soos, M., Meel, K.S. Bird: Engineering an efficient cnf-xor sat solver and its applications to approximate model counting. In Proceedings of AAAI Conference on Artificial Intelligence (AAAI) (2019) AAAI Press,
- Honolulu, USA.
 Stockmeyer, L. The complexity of approximate counting. In *Proceedings* of STOC (1983), ACM, Boston, 118–126.
- Tirthapura, S., Woodruff, D.P. Rectangle-efficient aggregation in spatial data streams. In *Proceedings* of *PODS* (2012), ACM, NY, 283–294.
- Valiant, L. The complexity of enumeration and reliability problems. SIAM J. Comput. 8, 3 (1979), 410–421.
- Woodruff, D.P., Zhang, Q. Tight bounds for distributed functional monitoring. In Proceedings of the 44th Annual ACM Symposium on Theory of Computing (2012), ACM, New York, USA 941–960.

Arnab Bhattacharyya and Kuldeep S. Meel, National University of Singapore, Singapore.

This work is licensed under a Creative Commons Attribution International 4.0 License. https://creativecommons.org/licenses/by/4.0/

ACM Student Research Competition



Attention: Undergraduate and Graduate Computing Students

Association for Computing Machinery Advancing Computing as a Science & Profession

The ACM Student Research Competition (SRC) offers a unique forum for undergraduate and graduate students to present their original research before a panel of judges and attendees at well-known ACM-sponsored and co-sponsored conferences. The SRC is an internationally recognized venue enabling undergraduate and graduate students to earn many tangible and intangible rewards from participating:

- Awards: cash prizes, medals, and ACM student memberships
- **Prestige:** Grand Finalists receive a monetary award and a Grand Finalist certificate that can be framed and displayed
- Visibility: opportunities to meet with researchers in their field of interest and make important connections
- Experience: opportunities to sharpen communication, visual, organizational, and presentation skills in preparation for the SRC experience

Learn more about ACM Student Research Competitions: https://src.acm.org