

On Efficient Range-Summability of IID Random Variables in Two or Higher Dimensions

Jingfan Meng ✉

School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Huayi Wang ✉

School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Jun Xu ✉

School of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Mitsunori Oghihara ✉

Department of Computer Science, University of Miami, Coral Gables, MI, USA

Abstract

d -dimensional (for $d > 1$) efficient range-summability (d D-ERS) of random variables (RVs) is a fundamental algorithmic problem that has applications to two important families of database problems, namely, fast approximate wavelet tracking (FAWT) on data streams and approximately answering range-sum queries over a data cube. Whether there are efficient solutions to the d D-ERS problem, or to the latter database problem, have been two long-standing open problems. Both are solved in this work. Specifically, we propose a novel solution framework to d D-ERS on RVs that have Gaussian or Poisson distribution. Our d D-ERS solutions are the first ones that have polylogarithmic time complexities. Furthermore, we develop a novel k -wise independence theory that allows our d D-ERS solutions to have both high computational efficiencies and strong provable independence guarantees. Finally, we show that under a sufficient and likely necessary condition, certain existing solutions for 1D-ERS can be generalized to higher dimensions.

2012 ACM Subject Classification Theory of computation → Streaming, sublinear and near linear time algorithms

Keywords and phrases fast range-summation, multidimensional data streams, Haar wavelet transform

Digital Object Identifier 10.4230/LIPIcs.ICDT.2023.21

Related Version *Previous Version:* <https://arxiv.org/abs/2110.07753>

Funding This material is based upon work supported by the National Science Foundation under Grant No. CNS-1909048, CNS-2007006, CNS-2051800, and by Keysight Technologies under Grant No. BG005054.

1 Introduction

Efficient range-summability (ERS) of random variables (RVs) is a fundamental algorithmic problem that has been studied for nearly two decades [5, 21, 6, 15]. This problem has so far been defined only in one dimension (1D) as follows. Let $X_0, X_1, \dots, X_{\Delta-1}$ be a list of *underlying RVs* each of which has the same *target distribution* X . Here, the (index) universe size Δ is typically a large number (say $\Delta = 2^{64}$). A 1D-ERS problem calls for the following oracle for answering range-sum queries over (realizations of) these underlying RVs. At initialization, the oracle chooses a *random outcome* ω from the *sample space* Ω , which *mathematically determines* the (values of the) realizations $X_0(\omega), X_1(\omega), \dots, X_{\Delta-1}(\omega)$; here the phrase “mathematically determines” emphasizes that (an implementation of) the oracle does not actually realize these RVs (and pay the $O(\Delta)$ time cost) at initialization. Thereafter, given any query range $[l, u) \triangleq \{l, l+1, \dots, u-1\}$ that lies in the universe $[0, \Delta)$, the oracle



© Jingfan Meng, Huayi Wang, Jun Xu, and Mitsunori Oghihara;
licensed under Creative Commons License CC-BY 4.0

26th International Conference on Database Theory (ICDT 2023).

Editors: Floris Geerts and Brecht Vandevoort; Article No. 21; pp. 21:1–21:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

is required to return $S[l, u] \triangleq \sum_{i=l}^{u-1} X_i(\omega)$, the sum of the realizations of all underlying RVs in the range. This requirement is called the *consistency* requirement, which is one of the two essential requirements for the ERS oracle. We will show that such an ERS oracle can be efficiently implemented using hash functions. With such an implementation, the outcome ω corresponds to the seeds of these hash functions.

The other essential requirement is *correct distribution*, which has two aspects. The first aspect is that the underlying RVs $X_0, X_1, \dots, X_{\Delta-1}$ each has the same target (marginal) distribution X . The second aspect is that these RVs should satisfy certain independence guarantees. Ideally, it is desired for these RVs to be mutually independent, but this comes at a high storage cost as we will elaborate shortly. In practice, another type of independence guarantee, namely k -wise independence (in the sense that any subset of k underlying RVs are independent), is good enough for most applications when $k \geq 4$. We will show that our solution for ERS in $d > 1$ dimensions can provide k -wise independence guarantee at a small storage cost of $O(\log^d \Delta)$ for an arbitrarily large k .

A straightforward but naive way to answer a range-sum query, say over $[l, u)$, is simply to sum up the realization of every underlying RV $X_l(\omega), X_{l+1}(\omega), \dots, X_{u-1}(\omega)$ in the query range. This solution, however, has a time complexity of $O(\Delta)$ when $u-l$ is $O(\Delta)$. In contrast, an efficient solution should be able to do so with only $O(\text{polylog}(\Delta))$ time complexity. Indeed, all existing ERS solutions [2, 5, 21, 6, 15] have $O(\log \Delta)$ time complexity.

1.1 Related Work on 1D-ERS

There are in general two families of solutions to the ERS problem in 1D, following two different approaches. The first approach is based on error correction codes (ECC). Solutions taking this approach include BCH3 [21], EH3 [5], and RM7 [2]. This approach has two drawbacks. First, it works only when the target distribution X is Rademacher ($\Pr[X = 1] = \Pr[X = 0] = 0.5$, aka. single-step random walk). Second, although it guarantees 3-wise (in the case of BCH3 and EH3) or 7-wise (in the case of RM7) independence among the underlying RVs, almost all empirical independence beyond that is destroyed. In addition, RM7 is very slow in practice [21].

The second approach is based on a data structure called dyadic simulation tree (DST), which we will describe in Subsection 3.1. The DST-based approach was first briefly mentioned in [6] and later fully developed in [15]. The DST-based approach is better than the ECC-based approach in two aspects. First, it supports a wider range of target distributions including Gaussian, Cauchy, Rademacher [15], and Poisson (see Appendix C of [14]). Second, it provides stronger independence guarantees at a low computational cost. For example, when implemented using the tabulation hashing scheme [24], it guarantees 5-wise independence at a much lower computational cost than RM7 [15]. We will describe a nontrivial generalization of this result to 2D in Section 4.

1.2 ERS in Higher Dimensions

In this work, we formulate the ERS problems in $d > 1$ dimensions (dD), which we denote as dD -ERS, and propose the first-ever solutions to dD -ERS. A dD -ERS problem is similarly defined on a d -dimensional universe $[0, \Delta)^d$ that contains Δ^d integral points. Each dD point $\vec{\mathbf{i}} \in [0, \Delta)^d$ is associated with an RV $X_{\vec{\mathbf{i}}}$, and every such RV has the same target (marginal) distribution X . Here, for ease of presentation, we assume Δ is the same on each dimension and is a power of 2, but our solutions can work without these two assumptions. Let $\vec{\mathbf{l}} = (l_1, l_2, \dots, l_d)^T$ and $\vec{\mathbf{u}} = (u_1, u_2, \dots, u_d)^T$ be two dD points in $[0, \Delta)^d$ such that

$l_j < u_j$ for each dimension $j = 1, 2, \dots, d$. We define $[\vec{\mathbf{l}}, \vec{\mathbf{u}}]$ as the dD rectangular range “cornered” by these two points in the sense $[\vec{\mathbf{l}}, \vec{\mathbf{u}}] \triangleq [l_1, u_1) \times [l_2, u_2) \times \dots \times [l_d, u_d)$, where \times is the Cartesian product.

A dD -ERS problem calls for the following oracle. At initialization, the oracle chooses an outcome ω that *mathematically determines* the realization $X_{\vec{\mathbf{i}}}(\omega)$ for each $\vec{\mathbf{i}} \in [0, \Delta)^d$. Thereafter, given any dD range $[\vec{\mathbf{l}}, \vec{\mathbf{u}}]$, the oracle needs to return in $O(\text{polylog}(\Delta))$ time $S([\vec{\mathbf{l}}, \vec{\mathbf{u}}]) \triangleq \sum_{\vec{\mathbf{i}} \in [\vec{\mathbf{l}}, \vec{\mathbf{u}}]} X_{\vec{\mathbf{i}}}(\omega)$, the sum of the realizations of all underlying RVs in this dD range. Unless otherwise stated, the vectors that appear in the sequel are assumed to be column vectors. We write them in boldface and with a rightward arrow on the top like in “ $\vec{\mathbf{x}}$ ”.

Several 1D-ERS solutions have been proposed as an essential building block for efficient solutions to several database problems. In two such database problems that we will describe in Section 2, their 1D solutions, both proposed in [7], can be readily generalized to dD if their underlying 1D-ERS oracles can be generalized to dD . In fact, in [16], authors stated explicitly that the only missing component for their solutions of the 1D database problems to be generalized to 2D was an efficient 2D-ERS oracle where X is the Rademacher distribution. However, until this paper, *no solution* to any dD -ERS problem for $d > 1$ has been proposed.

1.3 Our dD -ERS Solutions

In this paper, we propose novel solutions to the two dD -ERS problems wherein the target distributions are Gaussian and Poisson respectively. We refer to these two problems as dD Gaussian-ERS and dD Poisson-ERS, respectively. Both solutions generalize the corresponding DST-based 1D-ERS solutions to higher dimensions and have a low time complexity of $O(\log^d \Delta)$ per range-sum query. Our dD Gaussian-ERS solution, in particular, is based on the Haar wavelet transform (HWT), since DST is equivalent to HWT when (and only when) the target distribution X is Gaussian, as will be shown in Subsection 3.2.

Furthermore, we identify a sufficient condition that, if satisfied by the target distribution X , guarantees that the corresponding DST-based 1D-ERS solution can be generalized to a dD -ERS solution. We prove (see Appendix A of [14]) that Gaussian and Poisson are two “nice” distributions that satisfy this sufficient condition. We will also show that, for all such “nice” distributions (including those we might discover in the future), this generalization process (from 1D to dD) follows a universal algorithmic framework that can be characterized as the Cartesian product of d DSTs. We will also provide strong evidence that X “being nice” is likely necessary for this DST generalization (from 1D to dD) to be feasible (see Section 5).

Unfortunately, so far we have not found any “nice” distribution other than Gaussian and Poisson. Hence dD -ERS for other target distributions remains an open problem, and is likely not solvable by the (generalized) DST approach. We emphasize this is not a shortcoming of the DST approach: That we have obtained computationally efficient solutions in the cases of Gaussian and Poisson is already a pleasant surprise, as the dD -ERS problem has been open for nearly two decades. Furthermore, we will show that our dD Gaussian-ERS solution leads to computationally efficient solutions to both aforementioned database problems (to be described in Section 2), by answering their calls for a dD Gaussian-ERS or equivalent oracle.

Our dD Gaussian-ERS and Poisson-ERS solutions both support two different types of independence guarantees, at different storage costs. The first type is the ideal case in which the Δ^d underlying RVs are mutually independent. As will be shown in Section 3, we can achieve this ideal case by paying $O(T \log^d \Delta)$ storage cost, where T is the total number of range-sum queries to be answered (i.e., $O(\log^d \Delta)$ storage cost per range query). The second type is also quite strong: The Δ^d underlying RVs are k -wise independent, where the constant k can be arbitrarily large. In Section 4, we propose a k -wise independence scheme that can

provide the second type of guarantees by employing $O(\log^d \Delta)$ k -wise independent hash functions. Its storage cost is quite small: only $O(\log^d \Delta)$ for storing the seeds of these hash functions. We emphasize that the issue of how strong this independence guarantee (among the underlying RVs) needs to be affects only the storage cost of our Gaussian-ERS and Poisson-ERS solutions, and is orthogonal to all other issues described in earlier paragraphs such as the $O(\log^d \Delta)$ time complexity of both solutions and the sufficient and likely necessary condition for a DST-based d D-ERS solution to exist.

This k -wise independence scheme makes our d D-ERS solutions very practically useful for two reasons. First, such a k -wise independent hash function in practice requires a very short seed (not longer than a few kilobytes), and each hash operation can be computed in nanoseconds [3, 19]. Second, most applications of ERS only require the underlying RVs to be 4-wise independent [7, 16].

The contributions of this work can be summarized as follows. First, we provide the first set of answers to the long-standing open question whether there is an efficient solution to *any* d D-ERS problem for $d > 1$. Second, our Gaussian-ERS solution solves a long-standing open problem in data streaming that we will describe next. Third, our k -wise independence theory and hashing scheme make our d D ERS solutions very practically useful.

The rest of the paper is organized as follows. In Section 2, we describe two applications of our d D Gaussian-ERS solutions. In Section 3, we first describe our HWT-based Gaussian-ERS scheme in 1D, and then generalize it to 2D and d D. In Section 4, we describe our k -wise independence theory and scheme. In Section 5, we propose a sufficient and likely necessary condition on the target distribution for the DST approach to be generalized to d D. Finally, we conclude the paper in Section 6.

2 Applications of d D Gaussian-ERS

In this section, we introduce two important applications of our d D Gaussian-ERS solution.

2.1 Fast Approximate Wavelet Tracking

The first application is to the problem of fast approximate wavelet tracking (FAWT) on data streams [7, 4]. We first introduce the FAWT problem in 1D [7], or 1D-FAWT for short. In this problem, the precise system state is comprised of a Δ -dimensional vector \vec{s} , each scalar of which is a counter. The precise system state at any moment of time is determined by a data stream, in which each data item is an update to one such counter (called a *point update*) or all counters in a 1D range (called a *range update*). In 1D-FAWT, \vec{s} is considered a Δ -dimensional signal vector that is constantly “on the move” caused by the updates in the data stream. Let \vec{r} be the (Δ -dimensional) vector of HWT coefficients of \vec{s} . Clearly, \vec{r} is also a “moving target”. We denote as \vec{r}_t the snapshot of \vec{r} at a time t . In 1D-FAWT, the goal is to closely track (the precise value of) \vec{r} over time using a *sketch*, in the sense that at moment t , we can recover from the sketch an estimate \vec{r}'_t of \vec{r}_t , such that $\|\vec{r}_t - \vec{r}'_t\|_2$ is small. An acceptable solution should use a sketch whose size (space complexity) is only $O(\text{polylog}(\Delta))$, and be able to maintain the sketch with a computation time cost of $O(\text{polylog}(\Delta))$ per point or range update.

The first solution to 1D-FAWT was proposed in [7]. It requires the efficient computation of an arbitrary scalar in $H\vec{x}$, where H is the $\Delta \times \Delta$ Haar matrix (to be defined in Subsubsection 3.2.1) and \vec{x} is a Δ -dimensional vector of 4-wise independent Rademacher RVs. A key step of this computation is to compute a range-sum of 4-wise independent Rademacher RVs (in 1D), that is used therein as a Tug-of-War (ToW) sketch [1] for “sketching” the L_2

difference (approximation error) between the signal vector and its FAWT approximation. An aforementioned ECC-based ERS solution is used therein to tackle this Rademacher-ERS problem. Authors of [16] stated that if they could find a solution to this Rademacher-ERS problem in dD , then the 1D-FAWT solution in [7] would become a dD -FAWT solution. The first solution to dD -FAWT, proposed in [4], explicitly bypassed this ERS problem.

We note that the 1D-FAWT solution above continues to work, and its time and space complexities remain the same, if we replace the \vec{x} with a Δ -dimensional vector of 4-wise independent standard Gaussian RVs. This is because, with this replacement, the aforementioned ToW sketch becomes a Gaussian Tug-of-War (GToW) sketch (which maps a data item to a Gaussian RV instead of a Rademacher RV) [10], and ToW and GToW are known to have the same (ϵ, δ) accuracy bound [1, 10] for sketching the L_2 norm of a data stream (used here for sketching the aforementioned L_2 difference). Based on this insight, our dD Gaussian-ERS solution can be used to construct a dD -FAWT solution as follows. We simply change, in the *contingent* dD -FAWT solution proposed in [7], the distribution of all Δ^d underlying 4-wise independent RVs from Rademacher to Gaussian. With this replacement, this *contingent* solution will finally work, *provided* we can solve the resulting dD Gaussian-ERS problem. The latter problem is solved by our k -wise (with $k = 4$ here) independence scheme, to be described in Section 4. The resulting dD -FAWT solution has the same time and space complexity of $O(\log^d \Delta)$ as that proposed in [4] for achieving the same accuracy guarantee.

2.2 Range-Sum Queries over Data Cube

Our second application is to the problem of approximately answering range-sum queries over a data cube [8] that is similarly “on the move” propelled by the (point or range) updates that arrive in a stream. This problem can be formulated as follows. The precise system state is comprised of Δ^d counters, namely $\sigma_{\vec{i}}$ for $\vec{i} \in [0, \Delta]^d$, that are “on the move”. Given a range $[\vec{l}, \vec{u}]$ at moment t , the goal is to approximately compute the sum of counter values in this range $C[\vec{l}, \vec{u}] \triangleq \sum_{\vec{i} \in [\vec{l}, \vec{u}]} \sigma_{\vec{i}}(t)$, where $\sigma_{\vec{i}}(t)$ is the value of the counter $\sigma_{\vec{i}}$ at moment t . A desirable solution to this problem in dD should satisfy three requirements (in which multiplicative terms related to the desired (ϵ, δ) accuracy bound are ignored). First, any range-sum query is answered in $O(\text{polylog}(\Delta))$ time. Second, its space complexity is $O(\text{polylog}(\Delta))$. Third, every point or range update to the system state is processed in $O(\text{polylog}(\Delta))$ time. It has been a long-standing open question whether there is a solution to this problem that satisfies all three requirements when $d > 1$. For example, solutions producing exact answers (to the range queries) [9, 22, 11] all require $O(\Delta^d \log \Delta)$ space and hence do not satisfy the second requirement; and Haar+ tree [12] works only on static data, and hence does not satisfy the third requirement.

In 1D, a solution that satisfies all three requirements (with $d = 1$) was proposed in [7, 6]. It involves 1D-ERS computations on 4-wise independent underlying RVs where the target distribution is either Gaussian or Rademacher, which are tackled using a DST-based (in [6]) or a ECC-based (in [7]) 1D-ERS solution, respectively. As shown in [7, 6], this range-sum query solution can be readily generalized to dD if the ERS computations above can be performed in dD . This gap is again filled by our k -wise ($k = 4$) independence scheme for dD Gaussian-ERS, resulting in the first dD solution that satisfies all three requirements, all with $O(\log^d \Delta)$ (time or space) complexity (ignoring ϵ and δ terms).

In the resulting dD solution, we maintain $O(\log(1/\delta)/\epsilon^2)$ (independent instances of) sketches that each “sketches” the content (counter values) of the data cube. Here we describe only one such sketch, which we denote as A , since these sketches are statistically and functionally identical. At any time t , $A(t)$ should track the current system state, namely

$(\sigma_{\vec{i}}(t))$'s, as follows: $A(t) \triangleq \sum_{\vec{i} \in [0, \Delta]^d} \sigma_{\vec{i}}(t) X_{\vec{i}}$. Here $X_{\vec{i}}$ for $\vec{i} \in [0, \Delta]^d$ are (realizations of) a set of Δ^d 4-wise independent standard Gaussian underlying RVs that have one-to-one correspondences with the set of Δ^d counters as follows: Each $X_{\vec{i}}$ is associated with a counter $\sigma_{\vec{i}}$. If we implement these Δ^d RVs using (an instance of) our dD Gaussian-ERS solution, then we can keep the value of $A(t)$ up-to-date, with a time complexity of $O(\log^d \Delta)$ per point or range update (to the system state). Then, given a query range $[\vec{l}, \vec{u}]$ at time t , we estimate the range-sum of counters $C[\vec{l}, \vec{u}]$ from the sketch $A(t)$ using $A(t) \cdot S[\vec{l}, \vec{u}]$ as the estimator. These $O(\log(1/\delta)/\epsilon^2)$ estimators, one obtained from each sketch, are then combined to produce a final estimation that has the following accuracy guarantee (that is the same as in the 1D case). With probability at least $1 - \delta$, the final estimation deviates from the actual value of $C[\vec{l}, \vec{u}]$ by at most $\epsilon \sqrt{V[\vec{l}, \vec{u}] \|\sigma\|_2}$, where $V[\vec{l}, \vec{u}] \triangleq \prod_{j=1}^d (u_j - l_j)$ is the number of counters in the query range, and $\|\sigma(t)\|_2 \triangleq \left(\sum_{\vec{i} \in [0, \Delta]^d} \sigma_{\vec{i}}^2(t) \right)^{1/2}$ is the L_2 norm of the system state. Since each sketch uses an independent dD Gaussian-ERS scheme instance, our dD solution satisfies all three aforementioned requirements, all with $O(\log^d \Delta \log(1/\delta)/\epsilon^2)$ time and space complexity.

2.3 A Closer Comparison with Related Work

In this section, per referees' requests, we provide an in-depth comparison of this work with prior works on 1D-FAWT [7, 6], on dD -FAWT [4], and on 1D data cube [6].

We start with explaining how the dD -FAWT solution proposed in [4] manages to avoid confronting the dD -ERS problem. The dD -FAWT solution [4] maintains ToW sketches for groups of wavelet coefficients in the wavelet domain. As explained earlier, each ToW sketch "measures" the L_2 norm (and hence the total energy by squaring) of such a group. By the property of HWT, each point or range update to the system state in the time domain translates into $O(\log^d \Delta)$ updates to the sketches in the wavelet domain; we also use this property in our solution to keep its time complexity below $O(\log^d \Delta)$ as shown in Subsection 3.4. To solve the dD -FAWT using these sketches in the wavelet domain, we need only to identify the groups that are (hierarchical) " L_2 heavy hitters" [4]. In [4], a binary search tree built on these sketches is used to search for such " L_2 heavy hitters" in $O(\log \Delta \cdot \log \log \Delta)$ time. Since this dD -FAWT solution [4] does not involve computing the range sums of the Rademacher RVs underlying the ToW sketches, it does not need to formulate or solve any ERS problem.

As we will elaborate in Section 3, our dD Gaussian-ERS solution works in the same way as the dD -FAWT solution proposed in [4], by shifting the (representations of) input streams and the range queries from the time domain to the wavelet domain. Hence, arguably had the dD -FAWT solution proposed in [4] used the Gaussian ToW (GToW) instead of the ToW sketch, this shift would have resulted in a dD -FAWT solution containing the bulk of our dD Gaussian-ERS solution as an embedded module. However, such an embedded module is still "two hops away" from our dD Gaussian-ERS solution as follows. First, since the objective of and the intuition behind this shift in [4] were to avoid rather than to solve the ERS problem, it would not be easy for the authors of [4] to realize that the embedded module can be extended to a standalone dD Gaussian-ERS solution. Second, without our aforementioned k -wise independence theory and construction, the embedded module does not yet guarantee 4-wise independence among underlying Gaussian RVs that is needed for dD -FAWT.

On a related note, should we try to extend the 1D-FAWT solution proposed in [7], which maintains the ToW sketches in the time domain, to dD without the aforementioned Rademacher-by-Gaussian replacement, the underlying Rademacher RVs would have to be efficiently range-summable to keep the time complexity of each point or range update to

the sketches low. However, this appears to be a tall order for now: For $d > 1$, no ECC-based Rademacher-ERS solution has ever been found as explained earlier, and a DST-based Rademacher-ERS solution is unlikely to exist, as we will show in Section 5.

A referee asked whether the 1D data cube solution proposed in [7, 6] can be extended to dD using the same aforementioned ERS avoidance strategy of maintaining the sketches in the wavelet domain as used in [4]. In retrospect, this solution approach would work, but unlikely to be taken since it is counterintuitive and still “two hops away” (from the right solution) as explained above. Indeed, authors of [7, 6] unsurprisingly took the much more intuitive approach of maintaining sketches in the time domain and as a result had to confront the dD Gaussian or Rademacher-ERS problem as explained in Subsection 2.2.


Now we highlight a key difficulty that we believe has prevented authors of [7, 6, 16] from solving the dD -ERS problem and extending their FAWT and data cube solutions from 1D to dD : The Rademacher or Gaussian RVs underlying the sketches need to be both 4-wise independent and efficiently range-summable, and conventional wisdom (until our work) has it that a magic hash function family is needed to achieve both. Authors of [7, 16] tried to extend a magic hash function family, that induces such Rademacher RVs in 1D, to dD . However, as explained earlier, a dD Rademacher-ERS solution is unlikely to exist. Authors of [6] proposed the 1D-DST that laid the foundation of this work and our prior work [15]. A key innovation of [6] is that the 1D-ERS is achieved via a 1D-DST instead of a magic hash function. However, their DST-based 1D Gaussian-ERS solution still relies on a magic hash function, called Nisan’s PRG (Pseudorandom Generator) [18], to provide 4-wise independence among the underlying Gaussian RVs. The use of Nisan’s PRG [18] however restricts the applicability and the extensibility of the 1D-DST approach, since Nisan’s PRG provides independence guarantees only for memory-constrained applications such as data streaming [10]. It is also not clear whether the 1D-DST approach powered by Nisan’s PRG can be extended to dD . In comparison, in our dD -ERS solutions, both dD -ERS and 4-wise independence are provided by the specially engineered dD -DST. As a result, a magic hash function family is no longer needed, since the hash values produced by a hash function are no longer required to be efficiently range summable.

Finally, we state a key difference between this work and [7, 6, 16, 4] with respect to wavelets. In this work, ERS is the end and wavelets is the means, whereas in [7, 6, 16] it is the other way around. In [4], wavelets is the end, but [4] cleverly avoids using ERS as the means as just explained.

3 Our Solution to dD Gaussian-ERS

In this section, we describe our dD Gaussian-ERS solution that answers a range-sum query in $O(\log^d \Delta)$ time. To explain this solution with best clarity, for now we require it to provide the aforementioned ideal guarantee that the Δ^d underlying RVs are *mutually independent*, with the understanding that this requirement affects only the space complexity of our solution. In the next section, this requirement will be relaxed to these RVs being k -wise independent, and as a result, the space complexity of our solution is reduced to $O(\log^d \Delta)$.

Our solution can be summarized as follows. Let \vec{x} denote the Δ^d underlying standard Gaussian RVs, namely $X_{\vec{i}}$ for $\vec{i} \in [0, \Delta)^d$, arranged (in the dictionary order of \vec{i}) into a Δ^d -dimensional vector. Then, after the dD Haar wavelet transform (HWT) is performed on \vec{x} , we obtain another Δ^d -dimensional vector \vec{w} whose scalars are the HWT coefficients of \vec{x} . Our solution builds on the following two observations. The first observation is that scalars in \vec{x} are i.i.d. standard Gaussian RVs if and only if scalars in \vec{w} are (see Theorem 2).



■ **Figure 1** Illustrations of a general DST and a Gaussian-DST with $\Delta = 4$.

The second observation is that the answer to any dD range-sum query can be expressed as a weighted sum of $O(\log^d \Delta)$ scalars (HWT coefficients) in \vec{w} (see Lemma 10). Our algorithm is simply to *generate and remember* only these $O(\log^d \Delta)$ HWT coefficients (that participate in this range-sum query). Our solution satisfies the correct distribution requirement (with mutual independence guarantee) by the first observation. Since the first observation is true only when the target distribution is Gaussian, this HWT-based solution does not work for any other target distribution.

In the following, we first introduce the concept of the dyadic simulation tree (DST) in 1D in Subsection 3.1. Then, we show that 1D DST is equivalent to 1D HWT in the Gaussian case and present our HWT-based Gaussian-ERS algorithm for 1D, in Subsection 3.2. Finally, we describe our HWT-based Gaussian-ERS algorithms for 2D and dD in Subsection 3.3 and Subsection 3.4, respectively.

3.1 A Brief Introduction to DST

In this section, we briefly introduce the concept of the DST, which as mentioned earlier was proposed in [15] as a general solution approach to the one-dimensional (1D) ERS problems for arbitrary target distributions.

We say that $[l, u]$ is a 1D *dyadic range* if there exist integers $j \geq 0$ and $m \geq 0$ such that $l = j \cdot 2^m$ and $u = (j + 1) \cdot 2^m$. We call the sum on a dyadic range a *dyadic range-sum*. Note that any underlying RV X_i is also a dyadic range-sum (on the dyadic range $[i, i + 1)$). Let each underlying RV X_i have standard Gaussian distribution $\mathcal{N}(0, 1)$. In the following, we focus on how to compute a dyadic range-sum, since any (general) 1D range can be “pieced together” using at most $2 \log_2 \Delta$ dyadic ranges [21]. We illustrate the process of computing dyadic range-sums using a “small universe” example (with $\Delta = 4$) shown in Figure 1a. To begin with, the total sum of the universe $S[0, 4]$ sitting at the root of the tree is generated directly from its distribution $\mathcal{N}(0, 4)$. Then, $S[0, 4]$ is split into two children, the half-range-sums $S[0, 2]$ and $S[2, 4]$, such that RVs $S[0, 2]$ and $S[2, 4]$ sum up to $S[0, 4]$, are (mutually) independent, and each has distribution $\mathcal{N}(0, 2)$. This is done by generating the RVs $S[0, 2]$ and $S[2, 4]$ from a conditional (upon $S[0, 4]$) distribution that will be specified shortly. Afterwards, $S[0, 2]$ is split in a similar way into two i.i.d. underlying RVs X_0 and X_1 , and so is $S[2, 4]$ (into X_2 and X_3). As shown in Figure 1a, the four underlying RVs are the leaves of the DST.

We now specify the aforementioned conditional distribution used for each split. Suppose the range-sum to split consists of $2n$ underlying RVs, and that its value is equal to z . The lower half-range-sum S_l (the left child in Figure 1a) is generated from the following conditional pdf (or pmf):

$$f(x | z) = \phi_n(x) \phi_n(z - x) / \phi_{2n}(z), \quad (1)$$

$$\begin{matrix} W_0^{-1} \text{-----} \\ W_0^0 \text{-----} \\ W_0^1 \text{-----} \\ W_1^1 \text{-----} \end{matrix} \begin{pmatrix} Y_0 \\ Y_1 \\ Y_2 \\ Y_3 \end{pmatrix} = \begin{pmatrix} 1/2 & 1/2 & 1/2 & 1/2 \\ 1/2 & 1/2 & -1/2 & -1/2 \\ 1/\sqrt{2} & -1/\sqrt{2} & 0 & 0 \\ 0 & 0 & 1/\sqrt{2} & -1/\sqrt{2} \end{pmatrix} \cdot \begin{pmatrix} X_0 \\ X_1 \\ X_2 \\ X_3 \end{pmatrix}$$

■ **Figure 2** An illustration of the HWT formula $\vec{y} = H_4 \vec{x}$.

where $\phi_n(\cdot)$ is the pdf (or pmf) of X^{*n} , the n^{th} convolution power of the target distribution, and $\phi_{2n}(\cdot)$ is the pdf (or pmf) of X^{*2n} . Then, the upper half-range-sum (the right child) is defined as $S_u \triangleq z - S_l$. It was shown in [15] that splitting a (parent) RV using this conditional distribution guarantees that the two resulting RVs S_l and S_u are i.i.d. This guarantee holds regardless of the target distribution. However, *computationally efficient* procedures for generating an RV S_l with distribution $f(x | z)$ are found only when the target distribution is one of the few “nice” distributions: Gaussian, Cauchy, and Rademacher as shown in [15], and Poisson as shown in Appendix C of [14].

Among them, Gaussian distribution has a nice property that an RV S_l with distribution $f(x | z)$ can be generated as a linear combination of z and a “fresh” standard Gaussian RV Y as $S_l \triangleq z/2 + \sqrt{n/2} \cdot Y$, since if we plug Gaussian pdfs $\phi_n(\cdot)$ and $\phi_{2n}(\cdot)$ into (1), $f(x | z)$ is precisely the pdf of $\mathcal{N}(z/2, n/2)$. Here, Y being “fresh” means it is independent of all other RVs.

This linearly decomposable property has a pleasant consequence that *every* dyadic range-sum generated by this 1D Gaussian-DST can be recursively decomposed to a linear combination of some i.i.d. standard Gaussian RVs, as illustrated in Figure 1b. In this example, let Y_0, Y_1, Y_2 and Y_3 be four i.i.d. standard Gaussian RVs. The total sum of the universe $S[0, 4)$ is written as $2Y_0$, because they have the same distribution $\mathcal{N}(0, 4)$. Then, it is split into two half-range-sums $S[0, 2) \triangleq Y_0 + Y_1$ and $S[2, 4) \triangleq Y_0 - Y_1$ using the linear decomposition above with $z = 2Y_0$ and a fresh RV Y_1 . Finally, $S[0, 2)$ and $S[2, 4)$ are similarly split into the four underlying RVs using fresh RVs Y_2 and Y_3 , respectively.

3.2 HWT Representation of 1D Gaussian-DST

In this section, we show that when the target distribution is Gaussian, a DST is mathematically equivalent to a Haar wavelet transform (HWT) in the 1D case. We will also show that this equivalence carries over to higher dimensions. Note that this equivalence does not apply to any target distribution other than Gaussian, and hence the HWT representation cannot replace the role of DST in general. In the following, we describe in Subsubsection 3.2.2 our HWT-based 1D Gaussian-ERS solution that has $O(\log \Delta)$ time complexity, after making some mathematical preparations in Subsubsection 3.2.1.

3.2.1 Mathematical Preliminaries

It is not hard to verify that, if we apply HWT (to be specified soon) to the four underlying RVs shown in Figure 1b, namely $X_0 = Y_0/2 + Y_1/2 + Y_2/\sqrt{2}$, $X_1 = Y_0/2 + Y_1/2 - Y_2/\sqrt{2}$, $X_2 = Y_0/2 - Y_1/2 + Y_3/\sqrt{2}$, and $X_3 = Y_0/2 - Y_1/2 - Y_3/\sqrt{2}$, then the four HWT coefficients we obtain are precisely Y_0, Y_1, Y_2, Y_3 , respectively. In other words, we have $\vec{y} = H_4 \vec{x}$, where $\vec{x} \triangleq (X_0, X_1, X_2, X_3)^T$, $\vec{y} \triangleq (Y_0, Y_1, Y_2, Y_3)^T$, and H_4 is the 4×4 Haar matrix H_4 . This example is illustrated as a matrix-vector multiplication in Figure 2.

The above example in which $\Delta = 4$ can be generalized to an arbitrary universe size Δ (that is a power of 2) as follows. In general, HWT is defined as $\vec{w} = H_\Delta \vec{x}$, where \vec{w} and \vec{x} are both Δ -dimensional vectors, and H_Δ is a $\Delta \times \Delta$ Haar matrix. To simplify notations,

we drop the subscript Δ in the sequel. In wavelet terms, \vec{x} is called a discrete signal vector and \vec{w} is called the *HWT coefficient* vector. Clearly, the i^{th} HWT coefficient is the inner product between \vec{x} and the i^{th} row of H , for $i = 0, 1, \dots, \Delta - 1$. In the wavelet theory, we index each HWT coefficient as W_j^m (instead of W_i) for $m = -1, 0, 1, \dots, \log_2 \Delta - 1$ and $j = 0, 1, \dots, 2^{m^+} - 1$ (where $m^+ \triangleq \max\{0, m\}$) in the dictionary order of (m, j) , and refer to the corresponding row (transposed into a column vector) in H that computes W_j^m as the *HWT vector* $\vec{\psi}_j^m$. Hence we have $W_j^m \triangleq \langle \vec{x}, \vec{\psi}_j^m \rangle$ by definition. In wavelet terms, parameter m is called scale and parameter j is called location. In Figure 2, the 4 HWT coefficients and 4 HWT vectors from top to bottom are on 3 different scales ($-1, 0$, and 1) and are “assigned” 3 different colors accordingly.

We define the *indicator vector* of a 1D range R , denoted as $\mathbb{1}_R$, as a Δ -dimensional 0-1 vector, the i^{th} scalar of which takes value 1 if $i \in R$ and 0 otherwise, for $i = 0, 1, \dots, \Delta - 1$. Throughout this paper, the indicator vectors are the only vectors that are not written in boldface with a rightward arrow on the top. We now specify the HWT vectors. Every HWT vector $\vec{\psi}_j^m$ is normalized such that $\|\vec{\psi}_j^m\|_2 = 1$. The first HWT vector $\vec{\psi}_0^{-1} \triangleq \Delta^{-1/2} \cdot \mathbb{1}_{[0, \Delta]}$ is special: Its corresponding coefficient W_0^{-1} reflects the scaled (by $\Delta^{-1/2}$) range-sum of the entire universe, whereas every other HWT coefficient is the (scaled) difference of two range-sums. Every other HWT vector $\vec{\psi}_j^m$, for $m = 0, 1, \dots, \log_2 \Delta - 1$ and $j = 0, 1, \dots, 2^m - 1$, corresponds to the dyadic range $I_j^m \triangleq [j\Delta/2^m, (j+1)\Delta/2^m)$ in the sense the latter serves as the support of the former: $\vec{\psi}_j^m$ is defined by setting the first half of I_j^m to the value $\sqrt{2^m/\Delta}$, the second half of I_j^m to the value $-\sqrt{2^m/\Delta}$, and the rest of the universe $[0, \Delta) \setminus I_j^m$ to the value 0. Note that $\vec{\psi}_j^m$ has the same number of scalars with value $\sqrt{2^m/\Delta}$ as those with value $-\sqrt{2^m/\Delta}$, so $\langle \vec{\psi}_j^m, \mathbb{1}_{I_j^m} \rangle = 0$. From the definition above, H is known to be orthonormal [17], so the following theorem can be applied to it.

► **Theorem 1** ([13]). *Let M be an $n \times n$ matrix. If M is orthonormal, then it has the following two properties:*

1. $M^T = M^{-1}$, and M^T is also orthonormal.
2. Given any two n -dimensional vectors \vec{x}, \vec{y} , we have $\langle \vec{x}, \vec{y} \rangle = \langle M\vec{x}, M\vec{y} \rangle$.

Let \vec{w} be a Δ -dimensional vector of i.i.d. standard Gaussian RVs. We mathematically define the vector of underlying RVs $\vec{x} = (X_0, X_1, \dots, X_{\Delta-1})^T$ as $\vec{x} \triangleq H^T \vec{w}$. Hence, we have $\vec{w} = H\vec{x}$ by the first property in Theorem 1. The underlying RVs defined this way are i.i.d. standard Gaussian, by the following theorem.

► **Theorem 2** (Proposition 3.3.2 in [25]). *Let $\vec{x} = M\vec{w}$ where M is an orthonormal matrix. Then \vec{x} is a vector of i.i.d. standard Gaussian RVs if and only if \vec{w} is.*

3.2.2 Our HWT-based Algorithm for 1D-ERS

Given any range $[l, u)$, we compute its range-sum $S[l, u)$ as $\langle \vec{w}, H\mathbb{1}_{[l, u)} \rangle$, which is the sum of the HWT coefficients in \vec{w} weighted by the scalars in $H\mathbb{1}_{[l, u)}$. This weighted sum can be computed in $O(\log \Delta)$ time, because, by Theorem 3, the Δ -dimensional vector $H\mathbb{1}_{[l, u)}$ contains only $O(\log \Delta)$ nonzero scalars (weights), and by Remark 5, for each such scalar, its index can be located and its value computed in $O(1)$ time. We refer to the $O(\log \Delta)$ corresponding scalars in \vec{w} whose weights are nonzero as *participating HWT coefficients* in the sequel.

To provide the aforementioned ideal guarantee of *mutual independence* (among the Δ underlying RVs), for each such participating HWT coefficient (which is a standard Gaussian RV), we generate the RV and remember its realization (in memory) if it has never been

generated before (say for answering an earlier range-sum query), or retrieve its realization from memory otherwise. The space complexity of this algorithm is $O(\min\{T \log \Delta, \Delta\})$, since each of the T range-sum queries involves $O(\log \Delta)$ participating HWT coefficients. This algorithm satisfies the aforementioned consistency requirement, because $\langle \vec{\mathbf{w}}, H\mathbb{1}_{[l,u]} \rangle = \langle H\vec{\mathbf{x}}, H\mathbb{1}_{[l,u]} \rangle = \langle \vec{\mathbf{x}}, \mathbb{1}_{[l,u]} \rangle = X_l + X_{l+1} + \cdots + X_{u-1}$. The second equation above is by the second property in Theorem 1.

► **Theorem 3.** *Given any range $[l, u] \subseteq [0, \Delta)$, $H\mathbb{1}_{[l,u]}$ contains at most $2 \log_2 \Delta + 2$ nonzero scalars.*

Theorem 3 is a straightforward corollary of Lemma 4, since H has only $\log_2 \Delta + 1$ scales.

► **Lemma 4.** *Given any range $[l, u] \subseteq [0, \Delta)$, $H\mathbb{1}_{[l,u]}$ contains at most 2 nonzero scalars on each scale.*

Proof. On scale $m = -1$, there is only one HWT coefficient anyway, so the claim trivially holds. We next prove the claim for any fixed $m \geq 0$. For each HWT vector $\vec{\psi}_j^m$, $j = 0, 1, \dots, 2^m - 1$, we denote the corresponding HWT coefficient as $r_j^m \triangleq \langle \vec{\psi}_j^m, \mathbb{1}_{[l,u]} \rangle$. It is not hard to verify that the relationship between the range $[l, u)$ and the dyadic range I_j^m must be one of the following three cases.

1. I_j^m and $[l, u)$ are disjoint. In this case, $r_j^m = 0$.
2. $I_j^m \subseteq [l, u)$. In this case, $r_j^m = \langle \vec{\psi}_j^m, \mathbb{1}_{I_j^m} \rangle = 0$ as explained in the second last sentence above Theorem 1.
3. Otherwise, I_j^m partially intersects $[l, u)$. This case may happen only to at most two (I_j^m)'s: the one that covers l and the one that covers $u - 1$. In this case, r_j^m can be nonzero. ◀

► **Remark 5.** Each scalar r_j^m (in $H\mathbb{1}_{[l,u]}$) that may be nonzero can be identified and computed in $O(1)$ time as follows. Note r_j^m may be nonzero only in the case (3) above, in which j is equal to either $\lfloor l2^m/\Delta \rfloor$ or $\lfloor (u-1)2^m/\Delta \rfloor$. As a result, if $r_j^m \neq 0$, its value can be computed in two steps [22]. First, intersect $[l, u)$ with the first half and the second half of I_j^m , respectively. Second, scale the size of the first intersection minus the size of the second by $\sqrt{2^m/\Delta}$, as was explained by the third last sentence above Theorem 1.

The following lemma is a special case of Lemma 4 where $l = u - 1$. This lemma holds, because in case (3) above, for $m = -1, 0, 1, \dots, \log_2 \Delta - 1$, there exists a unique dyadic interval I_j^m that covers l (namely, the one with $j = \lfloor l2^m/\Delta \rfloor$).

► **Lemma 6.** *Given any $l \in [0, \Delta)$, $H\mathbb{1}_{\{l\}}$ has exactly one nonzero scalar on each scale.*

3.3 Range-Summable Gaussian RVs in 2D

In the following, we describe in Subsubsection 3.3.2 our 2D Gaussian-ERS solution that has $O(\log^2 \Delta)$ time complexity, after making some mathematical preparations in Subsubsection 3.3.1.

3.3.1 Mathematical Preliminaries

Like in the 1D case, our 2D Gaussian-ERS solution builds on the 2D-HWT $\vec{\mathbf{w}} = H^{\otimes 2}\vec{\mathbf{x}}$. Here the vector $\vec{\mathbf{x}}$ is comprised of the Δ^2 underlying RVs $X_{\vec{\mathbf{i}}}$ for $\vec{\mathbf{i}} \in [0, \Delta)^2$, listed in the dictionary order; and the vector $\vec{\mathbf{w}}$ is comprised of the resulting Δ^2 2D-HWT coefficients. The $\Delta^2 \times \Delta^2$ 2D-HWT matrix $H^{\otimes 2}$ is the self *Kronecker product* (defined next) of the $\Delta \times \Delta$ 1D-HWT matrix H .

21:12 On Efficient Range-Summability of IID RVs in Multiple Dimensions

► **Definition 7.** Let A be a $p \times q$ matrix and B be a $t \times v$ matrix. Then their Kronecker product $A \otimes B$ is the following $pt \times qv$ matrix.

$$A \otimes B \triangleq \begin{pmatrix} a_{11}B & \cdots & a_{1q}B \\ \vdots & \ddots & \vdots \\ a_{p1}B & \cdots & a_{pq}B \end{pmatrix}.$$

We now state two theorems concerning the Kronecker product.

► **Theorem 8** (Theorem 13.3 in [13]). Let P, Q, R, T be four matrices such that the matrix products $P \cdot R$ and $Q \cdot T$ are well-defined. Then $(P \otimes Q) \cdot (R \otimes T) = (P \cdot R) \otimes (Q \cdot T)$.

► **Theorem 9** (Corollary 13.8 in [13]). The Kronecker product of two orthonormal matrices is also orthonormal.

Now we describe the Δ^2 2D HWT coefficients and the order in which they are listed in $\vec{\mathbf{w}}$. Recall that in the 1D case, each HWT coefficient takes the form W_j^m , where m is the scale, and the j is the location. In the 2D case, each dimension has its own pair of scale and location parameters that is independent of the other dimension. For convenience of presentation, we refer to these two dimensions as vertical (the first) and horizontal (the second), respectively. We denote the vertical scale and location pair as m_1 and j_1 , and the horizontal pair as m_2 and j_2 . Each HWT coefficient takes the form $W_{j_1, j_2}^{m_1, m_2}$. In the 2D case, there are $(\log_2 \Delta + 1)^2$ scales, namely (m_1, m_2) for $m_1, m_2 = -1, 0, 1, \dots, \log_2 \Delta - 1$. At each scale (m_1, m_2) , there are $n_{m_1, m_2} \triangleq 2^{m_1+1} + 2^{m_2+1}$ locations, namely (j_1, j_2) for $j_1 = 0, 1, \dots, 2^{m_1+1} - 1$ and $j_2 = 0, 1, \dots, 2^{m_2+1} - 1$.

We now give a 2D example in which $\Delta = 4$. In this 2D example, there are $\Delta^2 = 16$ HWT coefficients. To facilitate the “color coding” of different scales, we arrange the 16 HWT coefficients into a 4×4 matrix \mathbf{W} shown in Figure 3. \mathbf{W} is the only matrix that we write in boldface in order to better distinguish it from its scalars ($W_{j_1, j_2}^{m_1, m_2}$)’s. Figure 3 contains three differently colored rows of heights 1, 1, and 2 respectively, that correspond to vertical scales $m_1 = -1, 0, 1$ respectively, and contains three differently colored columns that correspond to the three horizontal scales. Their “Cartesian product” contains 9 “color cells” that correspond to the 9 different scales (values of (m_1, m_2)). For example, the cell colored in pink corresponds to scale $(1, 1)$ and contains 4 HWT coefficients $W_{0,0}^{1,1}, W_{0,1}^{1,1}, W_{1,0}^{1,1}, W_{1,1}^{1,1}$. The vector $\vec{\mathbf{w}}$ is defined from \mathbf{W} by flattening its 16 scalars in the row-major order, as shown at the bottom of Figure 3.

$$\begin{pmatrix} W_{0,0}^{-1,-1} & W_{0,0}^{-1,0} & W_{0,0}^{-1,1} & W_{0,1}^{-1,1} \\ W_{0,0}^{0,-1} & W_{0,0}^{0,0} & W_{0,0}^{0,1} & W_{0,1}^{0,1} \\ W_{0,0}^{1,-1} & W_{0,0}^{1,0} & W_{0,0}^{1,1} & W_{0,1}^{1,1} \\ W_{1,0}^{1,-1} & W_{1,0}^{1,0} & W_{1,0}^{1,1} & W_{1,1}^{1,1} \end{pmatrix}$$

$$\left(W_{0,0}^{-1,-1} \ W_{0,0}^{-1,0} \ W_{0,0}^{-1,1} \ W_{0,1}^{-1,1} \ W_{0,0}^{0,-1} \ W_{0,0}^{0,0} \ W_{0,0}^{0,1} \ W_{0,1}^{0,1} \ W_{0,0}^{1,-1} \ W_{0,0}^{1,0} \ W_{0,0}^{1,1} \ W_{0,1}^{1,1} \ W_{1,0}^{1,-1} \ W_{1,0}^{1,0} \ W_{1,0}^{1,1} \ W_{1,1}^{1,1} \right)$$

■ **Figure 3** The 2D-HWT coefficients, arranged both as a matrix \mathbf{W} and as a flattened vector $\vec{\mathbf{w}}^T$.

Like in the 1D case, let $\vec{\mathbf{w}}$ be a vector of Δ^2 i.i.d. standard Gaussian RVs. As explained earlier, the vector $\vec{\mathbf{x}}$ of Δ^2 underlying RVs are *mathematically* defined as $\vec{\mathbf{x}} \triangleq (H^{\otimes 2})^T \vec{\mathbf{w}}$. The RVs in $\vec{\mathbf{x}}$ are i.i.d. standard Gaussian by Theorem 2, because $H^{\otimes 2}$ is an orthonormal matrix by Theorem 9.

3.3.2 Our HWT-Based Algorithm for 2D-ERS

Our 2D-ERS algorithm (that guarantees mutual independence among the underlying RVs) is similar to the 1D-ERS algorithm described earlier. Given any 2D range $[\vec{l}, \vec{u}] \triangleq [l_1, u_1] \times [l_2, u_2]$, where $\vec{l} = (l_1, l_2)^T$ and $\vec{u} = (u_1, u_2)^T$, we compute its range-sum $S[\vec{l}, \vec{u}]$ as $\langle \vec{w}, H^{\otimes 2} \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle$. Here the 2D indicator vector $\mathbb{1}_{[\vec{l}, \vec{u}]}$ is defined as the result of flattening the following $\Delta \times \Delta$ matrix in row-major order: For $\vec{i} \in [0, \Delta]^2$, the \vec{i}^{th} scalar in the matrix takes value 1 if $\vec{i} \in [\vec{l}, \vec{u}]$ and takes value 0 otherwise. This return value $\langle \vec{w}, H^{\otimes 2} \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle$ can be computed in $O(\log^2 \Delta)$ time, since it involves generating, and computing the weighted sum of, $O(\log^2 \Delta)$ participating HWT coefficients according to Lemma 10. The space complexity is $O(\min\{T \log^2 \Delta, \Delta^2\})$ for remembering the realizations of the $O(\log^2 \Delta)$ participating HWT coefficients (per query) like that explained earlier in the 1D case. Our 2D-ERS algorithm meets the consistency requirement, because $\langle \vec{w}, H^{\otimes 2} \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle = \langle H^{\otimes 2} \vec{x}, H^{\otimes 2} \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle = \langle \vec{x}, \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle = \sum_{(i_1, i_2) \in [\vec{l}, \vec{u}]} X_{i_1, i_2}$.

► **Lemma 10.** *For any 2D range $[\vec{l}, \vec{u}] \subseteq [0, \Delta]^2$, $H^{\otimes 2} \mathbb{1}_{[\vec{l}, \vec{u}]}$ has $O(\log^2 \Delta)$ nonzero scalars.*

Proof. It is not hard to verify $\mathbb{1}_{[\vec{l}, \vec{u}]} = \mathbb{1}_{[l_1, u_1]} \otimes \mathbb{1}_{[l_2, u_2]}$. By Theorem 8, $H^{\otimes 2} \mathbb{1}_{[\vec{l}, \vec{u}]} = (H \otimes H) \cdot (\mathbb{1}_{[l_1, u_1]} \otimes \mathbb{1}_{[l_2, u_2]}) = (H \mathbb{1}_{[l_1, u_1]}) \otimes (H \mathbb{1}_{[l_2, u_2]})$. By Theorem 3, both $H \mathbb{1}_{[l_1, u_1]}$ and $H \mathbb{1}_{[l_2, u_2]}$ have $O(\log \Delta)$ nonzero scalars, so their Kronecker product has $O(\log^2 \Delta)$ nonzero scalars. ◀

3.4 Generalization to Higher Dimensions

Our HWT-based Gaussian-ERS solution, just like HWT itself, can be naturally generalized to higher dimensions as follows. In dimension $d > 2$, we continue to have the inverse HWT formula $\vec{x} \triangleq M^T \vec{w}$, where \vec{x} is the vector of Δ^d underlying RVs (arranged in dictionary order of \vec{i}), \vec{w} is the vector of their HWT coefficients (that are i.i.d. standard Gaussian RVs), and M is the $\Delta^d \times \Delta^d$ HWT matrix in dD . Here $M \triangleq \underbrace{H \otimes \cdots \otimes H}_d$, where H is the 1D

Haar matrix described above. Since M is orthonormal by Theorem 9, the RVs in \vec{x} are i.i.d. standard Gaussian by Theorem 2.

In our dD -ERS algorithm (that guarantees mutual independence among the underlying RVs), given a dD range $[\vec{l}, \vec{u}] \triangleq [l_1, u_1] \times [l_2, u_2] \times \cdots \times [l_d, u_d]$, its range-sum $S[\vec{l}, \vec{u}]$ can be computed as $\langle \vec{w}, M \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle$, because $\langle \vec{w}, M \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle = \langle M \vec{x}, M \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle = \langle \vec{x}, \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle = \sum_{\vec{i} \in [\vec{l}, \vec{u}]} X_{\vec{i}}$. The weighted sum $\langle \vec{w}, M \mathbb{1}_{[\vec{l}, \vec{u}]} \rangle$ can be computed in $O(\log^d \Delta)$ time, because the weight vector $M \mathbb{1}_{[\vec{l}, \vec{u}]} = M \cdot (\mathbb{1}_{[l_1, u_1]} \otimes \mathbb{1}_{[l_2, u_2]} \otimes \cdots \otimes \mathbb{1}_{[l_d, u_d]}) = (H \mathbb{1}_{[l_1, u_1]}) \otimes (H \mathbb{1}_{[l_2, u_2]}) \otimes \cdots \otimes (H \mathbb{1}_{[l_d, u_d]})$ has only $O(\log^d \Delta)$ nonzero scalars (weights) by Theorem 3 and the property of Kronecker product. Hence, we need to generate and remember only $O(\log^d \Delta)$ corresponding participating HWT coefficients. As a result, our dD -ERS algorithm has $O(\min\{T \log^d \Delta, \Delta^d\})$ space complexity.

4 k -wise Independence Theory

In this section, in all subsequent paragraphs, we assume $d = 2$ (2D) for notational simplicity. All our statements and proofs can be readily generalized to higher dimensions. Recall that, for guaranteeing mutual independence among the Δ^d underlying RVs, our HWT-based dD Gaussian-ERS needs to *remember* (the realization of) every participating HWT coefficient that was generated for answering a past range-sum query, which can lead to high storage overhead when the number of queries T is large. In this section we propose a k -wise

independence theory and scheme that guarantees that the Δ^d underlying Gaussian RVs are k -wise independent. It does so by using $O(\log^d \Delta)$ k -wise independent hash functions (described next) instead. This scheme has the same time complexity of $O(\log^d \Delta)$ as the idealized Gaussian-ERS solution, and a much smaller space complexity of $O(\log^d \Delta)$, for storing the seeds of $O(\log^d \Delta)$ k -wise independent hash functions. This scheme significantly extends its 1D version proposed in [15]. Finally, we note this scheme works also for our Poisson-ERS solution. We however will not explain how it works in this paper, since doing so would involve drilling down to the messy and lengthy detail of the Cartesian product of $d > 1$ DSTs (since we cannot use the relatively clean and simple d D HWT in the Poisson case).

A k -wise independent hash function $h(\cdot)$ has the following property: Given an arbitrary set of k distinct keys i_1, i_2, \dots, i_k , their hash values $h(i_1), h(i_2), \dots, h(i_k)$ are independent. Such hash functions are very computationally efficient when k is a small number such as $k = 2$ (roughly 2 nanoseconds per hash) and $k = 4$ (several nanoseconds per hash) [3, 23, 19]. Typically, the hash values are (uniform random) integers. We can map them to Gaussian RVs using a deterministic function $g(\cdot)$ such as the Box-Muller transform [20].

Recall (from Figure 3) that the Δ^2 HWT coefficients in the vector $\vec{\mathbf{w}}$ are on $(\log_2 \Delta + 1)^2$ different scale pairs, namely (m_1, m_2) for $m_1, m_2 = -1, 0, 1, \dots, \log_2 \Delta - 1$. Our scheme uses $(\log_2 \Delta + 1)^2$ independent k -wise independent hash functions that we denote as $h_{m_1, m_2}(\cdot)$, for $m_1, m_2 = -1, 0, 1, \dots, \log_2 \Delta - 1$. During the initialization phase, we uniformly randomly seed these $(\log_2 \Delta + 1)^2$ hash functions; once seeded, they are fixed thereafter as usual. As mentioned earlier, these seeds correspond to the outcome ω that fixes (mathematically defines) the HWT coefficient vector $\vec{\mathbf{w}}$.

Our scheme can be stated *literally in one sentence*: Each such (seeded and fixed) $h_{m_1, m_2}(\cdot)$ is solely responsible for hash-generating any HWT coefficient on scale (m_1, m_2) that is *participating* (as defined earlier) in answering a range-sum query. In other words, for any scale $m_1, m_2 = -1, 0, 1, \dots, \log_2 \Delta - 1$, and location $j_1 = 0, 1, \dots, 2^{m_1} - 1$, $j_2 = 0, 1, \dots, 2^{m_2} - 1$, the value of the HWT coefficient $W_{j_1, j_2}^{m_1, m_2}$ is *mathematically defined* as $g(h_{m_1, m_2}(j_1, j_2))$, where $g(\cdot)$ is the aforementioned deterministic function (that maps an integer to a Gaussian RV). Hence our scheme has a much lower space complexity of $O(\log^2 \Delta)$, for remembering the seeds of the $O(\log^2 \Delta)$ hash functions.

The following theorem states that our scheme achieves its intended objective of ensuring that the Δ^2 underlying RVs *mathematically defined* by it are k -wise independent. In this theorem and proof, we denote the vector of Δ^2 HWT coefficients and the vector of Δ^2 underlying RVs both mathematically defined by our k -wise scheme as $\vec{\mathbf{v}}$ and $\vec{\mathbf{z}}$, respectively. We do so to distinguish this vector pair from the original vector pair $\vec{\mathbf{w}}$ and $\vec{\mathbf{x}}$ that are mathematically defined by the idealized scheme (that guarantees mutual independence). Recall that $\vec{\mathbf{z}} = M^T \vec{\mathbf{v}}$ and $\vec{\mathbf{x}} = M^T \vec{\mathbf{w}}$, where $M = H^{\otimes 2}$ is the 2D HWT matrix, and that $\vec{\mathbf{x}}$ is comprised of i.i.d. standard Gaussian RVs.

► **Theorem 11.** *The vector $\vec{\mathbf{z}}$ is comprised of k -wise independent standard Gaussian RVs.*

Proof. It suffices to prove that *any* k distinct scalars in $\vec{\mathbf{z}}$ – say the $(i_1)^{th}, (i_2)^{th}, \dots, (i_k)^{th}$ scalars – are i.i.d. standard Gaussian. Let $\vec{\mathbf{z}}'$ be the k -dimensional vector comprised of these k scalars. Let $(M^T)'$ be the $k \times \Delta^2$ matrix formed by the $(i_1)^{th}, (i_2)^{th}, \dots, (i_k)^{th}$ rows in M^T . Then, we have $\vec{\mathbf{z}}' = (M^T)' \vec{\mathbf{v}}$. Now let the random vector $\vec{\mathbf{x}}'$ be defined as $(M^T)' \vec{\mathbf{w}}$. Then $\vec{\mathbf{x}}'$ is comprised of k i.i.d. standard Gaussian RVs, as its scalars are a subset of those of $\vec{\mathbf{x}}$. Hence, to prove that the scalars in $\vec{\mathbf{z}}'$ are i.i.d. standard Gaussian RVs, it suffices to prove the claim that $\vec{\mathbf{z}}'$ has the same distribution as $\vec{\mathbf{x}}'$.

We prove this claim using Proposition 12. To this end, we first write $\vec{\mathbf{z}}'$ and $\vec{\mathbf{x}}'$ each as the sum of $N = (\log_2 \Delta + 1)^2$ independent random vectors. Recall that in Subsection 3.3, we have classified the HWT coefficients in $\vec{\mathbf{w}}$ and $\vec{\mathbf{v}}$, and the columns of M^T (called HWT vectors

there) into N different (m_1, m_2) scales (colors in Figure 3). Recall that n_{m_1, m_2} scalars in \vec{w} and \vec{v} , and accordingly n_{m_1, m_2} columns of M^T , have scale (m_1, m_2) . Let \vec{w}_{m_1, m_2} and \vec{v}_{m_1, m_2} be the n_{m_1, m_2} -dimensional vectors comprised of the coefficients classified to scale (m_1, m_2) in \vec{w} and \vec{v} , respectively. Let $(M^T)'_{m_1, m_2}$ be the $k \times n_{m_1, m_2}$ matrix comprised of the columns of $(M^T)'$ classified to scale (m_1, m_2) . Then, we have $\vec{z}' = \sum_{(m_1, m_2)} (M^T)'_{m_1, m_2} \vec{v}_{m_1, m_2}$ and $\vec{x}' = \sum_{(m_1, m_2)} (M^T)'_{m_1, m_2} \vec{w}_{m_1, m_2}$, where both summations are over all N scales. The N summands in the RHS of the first equation are independent random vectors, because for each scale $(m_1, m_2) \in [-1, \log_2 \Delta]^2$, all scalars in \vec{v}_{m_1, m_2} are generated by the same per-scale hash function $h_{m_1, m_2}(\cdot)$, which is independent of all $N - 1$ other per-scale hash functions. The same can be said about the N summands in the RHS of the second equation, since \vec{w} is comprised of i.i.d. RVs by design. To prove this claim using Proposition 12, it remains to prove the fact that for each scale $(m_1, m_2) \in [-1, \log_2 \Delta]^2$, the pair of random vectors $(M^T)'_{m_1, m_2} \vec{v}_{m_1, m_2}$ and $(M^T)'_{m_1, m_2} \vec{w}_{m_1, m_2}$ have the same distribution.

This fact can be proved as follows. Note that for each scale $(m_1, m_2) \in [-1, \log_2 \Delta]^2$, each row in $(M^T)'_{m_1, m_2}$ has exactly one nonzero scalar, since the corresponding row in M^T , or equivalently the corresponding column in M , has exactly one nonzero scalar at each scale (m_1, m_2) , due to Lemma 13. Therefore, although the number of columns in $(M^T)'_{m_1, m_2}$ can be as many as $O(\Delta^2)$, at most k of them (one for each row), say the $(\alpha_1)^{th}, (\alpha_2)^{th}, \dots, (\alpha_k)^{th}$ columns, contain nonzero scalars. Then, $(M^T)'_{m_1, m_2} \vec{v}_{m_1, m_2}$ is a function of only the $(\alpha_1)^{th}, (\alpha_2)^{th}, \dots, (\alpha_k)^{th}$ scalars in \vec{v}_{m_1, m_2} , and these k scalars are i.i.d. Gaussian RVs since they are all generated by the same k -wise independent hash function $h_{m_1, m_2}(\cdot)$. Note that $(M^T)'_{m_1, m_2} \vec{w}_{m_1, m_2}$ is the same function of the $(\alpha_1)^{th}, (\alpha_2)^{th}, \dots, (\alpha_k)^{th}$ scalars in \vec{w}_{m_1, m_2} , which are i.i.d. Gaussian RVs by design. Hence, $(M^T)'_{m_1, m_2} \vec{v}_{m_1, m_2}$ has the same distribution as $(M^T)'_{m_1, m_2} \vec{w}_{m_1, m_2}$. ◀

► **Proposition 12.** *Suppose random vectors \vec{x} and \vec{z} each is the sum of N independent random vectors as follows: $\vec{x} = \vec{x}_1 + \vec{x}_2 + \dots + \vec{x}_N$ and $\vec{z} = \vec{z}_1 + \vec{z}_2 + \dots + \vec{z}_N$. Then, \vec{x} and \vec{z} have the same distribution if each pair of components \vec{x}_i and \vec{z}_i have the same distribution, for $i = 1, 2, \dots, N$.*

► **Lemma 13.** *Any column of $M = H^{\otimes 2}$, which is equal to $H^{\otimes 2} \mathbb{1}_{\{\vec{i}\}}$ for some $\vec{i} = (i_1, i_2)^T$, has exactly one nonzero scalar on each 2D scale (m_1, m_2) .*

Proof. The 2D indicator vector can be decomposed to the Kronecker product of two 1D indicator vectors as $\mathbb{1}_{\{\vec{i}\}} = \mathbb{1}_{\{i_1\}} \otimes \mathbb{1}_{\{i_2\}}$, so $H^{\otimes 2} \mathbb{1}_{\{\vec{i}\}} = (H \mathbb{1}_{\{i_1\}}) \otimes (H \mathbb{1}_{\{i_2\}})$ by Theorem 8. The claim above follows from Lemma 6, which implies that $H \mathbb{1}_{\{i_1\}}$ and $H \mathbb{1}_{\{i_2\}}$ each has exactly one nonzero scalar on each 1D scale. ◀

5 Multidimensional Dyadic Simulation

As explained in Subsection 3.1, in one dimension (1D), any dyadic range-sum $S[l, u]$, *no matter what* the target distribution is, can be computed by performing $O(\log \Delta)$ *binary splits* along the path from the root $S[0, \Delta)$ to the node $S[l, u)$ along the dyadic simulation tree (DST). Since we have just *computationally efficiently* generalized the Gaussian-DST approach (equivalent to the HWT-based approach in the 1D Gaussian case) to any dimension $d \geq 2$, we wonder whether we can do the same for all target distributions. By “computationally efficiently”, we mean that a generalized solution should be able to compute any d D range-sum in $O(\log^d \Delta)$ time like in the Gaussian case.

Unfortunately, it appears hard, if not impossible, to generalize the DST approach to d D for arbitrary target distributions. We have identified a sufficient condition on the target distribution for such an efficient generalization to exist. We prove the sufficiency by proposing

a DST-based universal algorithmic framework (described in Appendix C of [14]) that solves the d D-ERS problem for any target distribution satisfying this condition. Unfortunately, so far only two distributions, namely Gaussian and Poisson, are known to satisfy this condition, as is elaborated in Appendix A of [14]. We also describe in Appendix B of [14] two example distributions that do not satisfy this sufficient condition, namely Cauchy and Rademacher. In the following, we specify this condition and explain why it is “almost necessary”.

For ease of presentation, in the following, we fix the number of dimensions d at 2. We assume all underlying RVs, X_{i_1, i_2} for (i_1, i_2) in the 2D universe $[0, \Delta]^2$, are i.i.d. with a certain target distribution X . This assumption is appropriate for our reasoning below about the time complexity of a 2D ERS solution, since as shown earlier this time complexity is not affected by the strength of the independence guarantee provided, in the cases of Gaussian and Poisson. In a 2D universe, any 2D range can be considered the Cartesian product of its horizontal and vertical 1D ranges. We say a 2D range is *dyadic* if and only if its horizontal and vertical 1D ranges are both dyadic. Since any general (not necessarily dyadic) 1D range can be “pieced together” using $O(\log \Delta)$ 1D dyadic ranges [21], it is not hard to show, using the Cartesian product argument, that any general 2D range can be “pieced together” using $O(\log^2 \Delta)$ 2D dyadic ranges. Hence in the following, we focus on the generation of only 2D dyadic range-sums. We assume all underlying RVs, X_{i_1, i_2} for (i_1, i_2) in the 2D universe $[0, \Delta]^2$, are i.i.d. with a certain target distribution X .

We need to introduce some additional notations. We define each horizontal strip-sum $S_i^H \triangleq X_{i,0} + X_{i,1} + \dots + X_{i,\Delta-1}$ for $i \in [0, \Delta)$ as the sum of range $[i, i+1) \times [0, \Delta)$, and each vertical strip-sum $S_i^V \triangleq X_{0,i} + X_{1,i} + \dots + X_{\Delta-1,i}$ for $i \in [0, \Delta)$ as the sum of range $[0, \Delta) \times [i, i+1)$. We denote as S the total sum of all underlying RVs in the universe, *i.e.*, $S \triangleq \sum_{i_1=0}^{\Delta-1} \sum_{i_2=0}^{\Delta-1} X_{i_1, i_2} = \sum_{i=0}^{\Delta-1} S_i^H = \sum_{i=0}^{\Delta-1} S_i^V$.

Now we are ready to state this sufficient condition. For ease of presentation, we break it down into two parts. The first part, stated in the following formula, states that the vector of vertical strip-sums and the vector of horizontal strip-sums in $[0, \Delta]^2$ are conditionally independent given the total sum S .

$$(S_0^V, S_1^V, \dots, S_{\Delta-1}^V) \perp\!\!\!\perp (S_0^H, S_1^H, \dots, S_{\Delta-1}^H) \mid S. \quad (2)$$

The second part is that this conditional independence relation holds for the two corresponding vectors in any 2D dyadic range (that is not necessarily a square). Intuitively, this condition says that how a 2D dyadic range-sum is split horizontally is conditionally independent (upon this 2D range-sum) of how it is split vertically. Roughly speaking, this condition implies that the 1D-DST governing the horizontal splits is conditionally independent of the other 1D-DST governing the vertical splits. Hence, our DST-based universal algorithmic framework for 2D can be viewed as the Cartesian product of the two 1D-DSTs, as will be elaborated in Appendix C of [14].

In the following, we offer some intuitive evidence why this condition is likely necessary. Without loss of generality, we consider the generation of an arbitrary horizontal strip-sum $S_{i_1}^H$ conditional on the vector of vertical strip-sums $(S_0^V, S_1^V, \dots, S_{\Delta-1}^V)$. Suppose (2) does not hold, which means $(S_0^H, S_1^H, \dots, S_{\Delta-1}^H)$ is not conditionally independent of $(S_0^V, S_1^V, \dots, S_{\Delta-1}^V)$ given S . Then the distribution of $S_{i_1}^H$ is arguably parameterized by the values (realizations) of all Δ vertical strip-sums $S_0^V, S_1^V, \dots, S_{\Delta-1}^V$, since $S_{i_1}^H$ and any vertical strip-sum $S_{i_2}^V$ for $i_2 \in [0, \Delta)$ are in general dependent RVs by Theorem 14 (See Appendix D of [14] for its nontrivial proof). Hence, unless some magic happens (which we cannot rule out rigorously), to generate (realize) the RV $S_{i_1}^H$, conceivably we need to first realize all Δ RVs $(S_0^V, S_1^V, \dots, S_{\Delta-1}^V)$, the time complexity of which is $\Omega(\Delta)$.

► **Theorem 14.** *For any (i_1, i_2) in $[0, \Delta]^2$, $S_{i_1}^H$ and $S_{i_2}^V$ are dependent RVs unless the target distribution X satisfies $\Pr[X = c] = 1$ for some constant c .*

6 Conclusion

In this work, we propose novel solutions to dD -ERS for RVs that have Gaussian or Poisson distribution. Our solutions are the first ones that compute any multi-dimensional range-sum in polylogarithmic time. Our dD Gaussian-ERS scheme solves the long-standing open problem of efficiently answering approximate range-sum queries over a multidimensional data cube. We develop a novel k -wise independence theory that provides both high computational efficiencies and strong provable independence guarantees. Finally, we show that when the underlying distribution satisfies a sufficient and likely necessary condition, its DST-based 1D-ERS solution can be generalized to higher dimensions.

References

- 1 Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96, pages 20–29, New York, NY, USA, 1996. Association for Computing Machinery. doi:10.1145/237814.237823.
- 2 A. R. Calderbank, A. Gilbert, K. Levchenko, S. Muthukrishnan, and M. Strauss. Improved range-summable random variable construction algorithms. In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '05, pages 840–849, USA, 2005. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070550>.
- 3 J. Lawrence Carter and Mark N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143–154, 1979. doi:10.1016/0022-0000(79)90044-8.
- 4 Graham Cormode, Minos Garofalakis, and Dimitris Sacharidis. Fast approximate wavelet tracking on streams. In Yannis Ioannidis, Marc H. Scholl, Joachim W. Schmidt, Florian Matthes, Mike Hatzopoulos, Klemens Boehm, Alfons Kemper, Torsten Grust, and Christian Boehm, editors, *Advances in Database Technology – EDBT 2006*, pages 4–22, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg. doi:10.1007/11687238_4.
- 5 Joan Feigenbaum, Sampath Kannan, Martin J. Strauss, and Mahesh Viswanathan. An approximate L1-difference algorithm for massive data streams. *SIAM Journal on Computing*, 32(1):131–151, 2002. doi:10.1137/S0097539799361701.
- 6 Anna C. Gilbert, Sudipto Guha, Piotr Indyk, Yannis Kotidis, S. Muthukrishnan, and Martin J. Strauss. Fast, small-space algorithms for approximate histogram maintenance. In *Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing*, STOC '02, pages 389–398, New York, NY, USA, 2002. Association for Computing Machinery. doi:10.1145/509907.509966.
- 7 Anna C. Gilbert, Yannis Kotidis, S. Muthukrishnan, and Martin J. Strauss. One-pass wavelet decompositions of data streams. *IEEE Trans. on Knowl. and Data Eng.*, 15(3):541–554, March 2003. doi:10.1109/TKDE.2003.1198389.
- 8 J. Gray, A. Bosworth, A. Lyaman, and H. Pirahesh. Data cube: a relational aggregation operator generalizing GROUP-BY, CROSS-TAB, and SUB-TOTALS. In *Proceedings of the Twelfth International Conference on Data Engineering*, pages 152–159, 1996. doi:10.1109/ICDE.1996.492099.
- 9 Nabil Ibtehad, M. Kaykobad, and M. Sohel Rahman. Multidimensional segment trees can do range updates in poly-logarithmic time. *Theoretical Computer Science*, 854:30–43, 2021. doi:10.1016/j.tcs.2020.11.034.
- 10 Piotr Indyk. Stable distributions, pseudorandom generators, embeddings, and data stream computation. *J. ACM*, 53(3):307–323, May 2006. doi:10.1145/1147954.1147955.
- 11 Mehrdad Jahangiri, Dimitris Sacharidis, and Cyrus Shahabi. SHIFT-SPLIT: I/O efficient maintenance of wavelet-transformed multidimensional data. In *Proceedings of the 2005 ACM SIGMOD International Conference on Management of Data*, SIGMOD '05, pages 275–286, New York, NY, USA, 2005. Association for Computing Machinery. doi:10.1145/1066157.1066189.

- 12 Panagiotis Karras and Nikos Mamoulis. The Haar+ tree: A refined synopsis data structure. In *2007 IEEE 23rd International Conference on Data Engineering*, pages 436–445, 2007. doi:10.1109/ICDE.2007.367889.
- 13 Alan J. Laub. *Matrix analysis – for scientists and engineers*. SIAM, 2005. URL: <http://bookstore.siam.org/ot91/>.
- 14 Jingfan Meng, Huayi Wang, Jun Xu, and Mitsunori Ogihara. On efficient range-summability of IID random variables in two or higher dimensions (extended version). *CoRR*, abs/2110.07753, 2021. arXiv:2110.07753.
- 15 Jingfan Meng, Huayi Wang, Jun Xu, and Mitsunori Ogihara. A Dyadic Simulation Approach to Efficient Range-Summability. In Dan Olteanu and Nils Vortmeier, editors, *25th International Conference on Database Theory (ICDT 2022)*, volume 220 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:18, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ICDT.2022.17.
- 16 S. Muthukrishnan and Martin Strauss. Maintenance of multidimensional histograms. In Paritosh K. Pandya and Jaikumar Radhakrishnan, editors, *FST TCS 2003: Foundations of Software Technology and Theoretical Computer Science*, pages 352–362, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. doi:10.1007/978-3-540-24597-1_30.
- 17 Yves Nievergelt. *Multidimensional Wavelets and Applications*, pages 36–72. Birkhäuser Boston, Boston, MA, 1999. doi:10.1007/978-1-4612-0573-9_2.
- 18 Noam Nisan. Pseudorandom generators for space-bounded computation. *Comb.*, 12(4):449–461, 1992. doi:10.1007/BF01305237.
- 19 Mihai Pundefinedtraşcu and Mikkel Thorup. The power of simple tabulation hashing. *J. ACM*, 59(3), June 2012. doi:10.1145/2220357.2220361.
- 20 Christian P. Robert and George Casella. *Monte Carlo Statistical Methods*, page 43. Springer New York, 2004. doi:10.1007/978-1-4757-4145-2_2.
- 21 Florin Rusu and Alin Dobra. Pseudo-random number generation for sketch-based estimations. *ACM Trans. Database Syst.*, 32(2):11–es, June 2007. doi:10.1145/1242524.1242528.
- 22 Rolfé R. Schmidt and Cyrus Shahabi. Propolyne: A fast wavelet-based algorithm for progressive evaluation of polynomial range-sum queries. In Christian S. Jensen, Simonas Šaltenis, Keith G. Jeffery, Jaroslav Pokorný, Elisa Bertino, Klemens Böhn, and Matthias Jarke, editors, *Advances in Database Technology – EDBT 2002*, pages 664–681, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg. doi:10.1007/3-540-45876-X_41.
- 23 Mikkel Thorup and Yin Zhang. Tabulation based 4-universal hashing with applications to second moment estimation. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '04, pages 615–624, USA, 2004. Society for Industrial and Applied Mathematics. URL: <http://dl.acm.org/citation.cfm?id=982792.982884>.
- 24 Mikkel Thorup and Yin Zhang. Tabulation based 5-universal hashing and linear probing. In *Proceedings of the Meeting on Algorithm Engineering and Experiments*, ALENEX '10, pages 62–76, USA, 2010. Society for Industrial and Applied Mathematics. doi:10.1137/1.9781611972900.7.
- 25 Roman Vershynin. *Random Vectors in High Dimensions*, pages 38–69. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018. doi:10.1017/9781108231596.006.