

Supporting Data Uncertainty in Array Databases

Liping Peng · Yanlei Diao

Received: date / Accepted: date

Abstract Uncertain data management has become a key issue in scientific applications. Recently, array databases have gained popularity for scientific data processing due to performance benefits. In this paper, we address uncertain data management in array databases, which may involve both *value uncertainty* within individual tuples and *position uncertainty* regarding where a tuple should belong in an array given uncertain dimension attributes. In our work, we define the formal semantics of array operations on uncertain data involving both types of uncertainty. To address the new challenge raised by position uncertainty, we propose a suite of storage and evaluation strategies for array operators, with a focus on a novel scheme that bounds the overhead of querying by strategically placing a few replicas of the tuples with large variances. Evaluation results show that our best-performing techniques outperform baselines often by a wide margin while incurring only small storage overhead.

1 Introduction

Uncertain data management has been studied intensively in areas such as sensor networks, information extraction, data cleaning, and business intelligence. Recently, it has also started to play a key role in large-scale scientific applications such as severe weather monitoring [18, 31], computational astrophysics [28], and asteroid threat detection [11]. In particular, recent studies [11, 27, 28] show that almost all scientific data are noisy and uncertain. Therefore, capturing uncertainty in

data processing, from data input to query output, has become a key issue in scientific data management.

We next present two real-world applications that motivated our work. In the first application, *massive astrophysical surveys* such as the Sloan Digital Sky Survey (SDSS) [28] aim to generate observations of 10^8 stars and galaxies at nightly data rates of 0.5TB to 20TB. They are expected to enable real-time detection of transient events and anomalies as well as long-term tracking of objects of interest. However, the observations in digital sky surveys are inherently noisy as the objects can be too dim to be recognized in the captured images. SDSS makes repeated observations of faint objects and derive continuous probability distributions for uncertain attributes, e.g., the location and luminosity of objects, in the data cooking process. For example, the Galaxy table in the SDSS astronomy archive has 297 attributes, out of which 151 attributes are uncertain and described by (multivariate) Gaussian distributions.

In the second domain of *severe weather monitoring*, scientists have developed distributed radar sensor networks for detecting hazardous weather events like tornados and severe storms [18]. The produced radar data are noisy due to environmental noise, electronic device noise, quality issues of the antenna, etc. As a result, recent work [12] has developed a data cooking process to produce continuous probability distributions for key meteorological measures, e.g., wind velocity and reflexivity, for each voxel of the air.

For supporting scientific applications, relational technology has proven useful in some applications like SDSS [28]. However, there is a recent realization that most scientific data naturally reside in multi-dimensional arrays rather than relations. This is because most scientific data are produced to characterize physical phenomena that rely heavily on the notions of “adjacency” and “neighborhood” in a multi-dimensional space. Hence, array databases have recently been developed for scientific data processing [6, 11, 26]. Besides convenient

Liping Peng
University of Massachusetts Amherst, MA, USA
E-mail: lppeng@cs.umass.edu

Yanlei Diao
University of Massachusetts Amherst, MA, USA
E-mail: yanlei@cs.umass.edu

expression of array operations, array databases also offer remarkable performance benefits over relational databases [24]. In particular, the new chunk-based storage scheme enables better alignment of logical locality (i.e., objects close in the logical array) and physical locality (i.e., objects close to each other are likely to be stored in the same chunk). Since many array operations exploit logical locality of data, e.g., finding objects close to a location, their associated physical locality can lead to significant I/O savings.

The increasing popularity of array databases has significant implications on uncertain data management: Recent work on array databases [14, 17, 15] has considered the case that a tuple belongs to a specific cell of an array and some of its value attributes are uncertain, which is referred to as the “*value uncertainty*”. On the other hand, a more complicated case arises when the attributes chosen to be the dimensions of an array are uncertain. For example, the x - y positions of an object in SDSS naturally serve as the dimensions of the array, but they are uncertain and characterized by a bivariate Gaussian distribution. As such, the uncertain location of an object can cause its tuple to belong to multiple cells in the array, referred to as the “*position uncertainty*”. SciDB, a leading effort on array databases, has acknowledged this issue in real-world applications but leaves the solution to future work [26].

In this paper, we provide a thorough treatment to uncertain data management in array databases. We focus on continuous uncertain data because they are a natural fit for scientific data and harder to support than discrete uncertain data due to the difficulty in enumerating the possible values. We assume that tuples are loaded into an array database in a batched, append-only fashion, which is common in scientific applications [6, 26], and each tuple has obtained a (joint) distribution for uncertain attributes through a scientific cooking process, as described above. We then address two key questions: (i) What are the intended answers of array operations on uncertain data that may involve both position and value uncertainty? (ii) What are the storage and evaluation methods for efficient array operations on continuous uncertain data?

Given position uncertainty, naive solutions would replicate a tuple in every possible location in the array, or store the tuple once in a default location but to answer a query, search as widely as the entire array to retrieve all the tuples that satisfy the query with a high probability. These solutions incur both a high I/O cost to read numerous tuples, and a high CPU cost to validate the retrieved tuples by computing their probabilities. By addressing such challenges related to position uncertainty, we make the following contributions:

1. We define the formal semantics of array operations on uncertain data involving both position and value uncertainty (§2). We show that *Subarray* and *Structure-Join* are the two most important array operations that involve position uncer-

tainty; many other array operations can be transformed into (one of) these two.

2. For *Subarray*, we provide native support for its operation on arrays with uncertain dimension attributes (§3). We propose a number of storage and evaluation schemes to deal with position uncertainty. In particular, we focus on a novel scheme, called *store-multiple*, that bounds the overhead of querying by strategically placing a few replicas for the tuples with large variances, which would otherwise make the query region grow very large. We also augment *store-multiple* with a detailed cost model and use it to configure storage for best performance under various workloads.

3. For *Structure-Join* on uncertain dimension attributes, we propose two techniques (§4): The first one integrates existing indexes for relational databases [10, 9, 20] with the *store-multiple* scheme for array databases, and minimizes the join cost by solving a set covering problem. We also propose a new subarray-based evaluation strategy for *Structure-Join*, which works without a pre-built index. This strategy employs tight conditions for running repeated subarray queries on the inner array of the join, as well as a detailed cost model for configuring the storage for best performance.

4. We evaluate our techniques using both synthetic workloads and a case study of the Sloan Digital Sky Survey (SDSS) [28] (§5). For *Subarray*, *store-multiple* outperforms other alternatives due to the bounded overhead of querying and optimized storage based on the cost model. For *Structure-Join*, the index-based join incurs high costs due to the index I/O and the set covering problem, while the subarray-based join outperforms it by using tight conditions for probing the inner array and optimized storage based on the cost model. Our case study shows that for real datasets, the storage overhead of *store-multiple* is rather small: **over 79% tuples have only 1 copy** and over 92% tuples have at most 3 copies (considering that 3 is the common number for replication in big data systems). In addition, our best techniques for *Subarray* and *Structure-Join* are shown to outperform the baselines often by a wide margin.

2 Array Model and Algebra

In this section, we provide background on the array model and array algebra proposed recently [6, 25]. Furthermore, we extend the array model to accommodate uncertain data and formally define the semantics of array algebra under the uncertain data model.

2.1 Array Data Model

Background on the Array Model. An array database contains a collection of arrays. Each array is represented as $\mathbb{A}(\mathbf{D}^d; \mathbf{V}^m)$, where \mathbf{D}^d denotes the d dimension attributes

that define the array, and V^m denotes m value attributes. We sometimes also use the shorthand, \mathbb{A}^d , to denote a d -dimensional array. Consider an example in the Digital Sky Survey domain: $\mathbb{A}^2(x_loc, y_loc; luminosity, color, \dots)$ defines a two-dimensional array using the dimension attributes (x_loc, y_loc) . If a dimension attribute is discrete-valued, the model requires a linear ordering of its values. If a dimension attribute is continuous-valued instead, a user-defined function, e.g., $\lfloor x_loc \rfloor$, is assumed to be available for discretizing the domain of the dimension attribute into an ordered set of values. These ordered values are used as the *index values* in a given dimension.

In an array \mathbb{A}^d , a unique combination of the index values of the d dimensions defines a cell, which can contain multiple tuples. Array cells are addressed by the index values of dimensions, e.g., a single cell addressed by $\mathbb{A}[\lfloor x_loc \rfloor=1, \lfloor y_loc \rfloor=2]$, abbreviated as $\mathbb{A}[1, 2]$, or multiple cells by $\mathbb{A}[1 : \infty, 2 : 4]$. By default, tuples in a cell include both dimension attributes and value attributes. If a dimension attribute is discrete-valued, its values in the tuples can be omitted because they are the same as the index values of the cell (which is not true for continuous attributes). Note that this model is an extension of the SciDB array model [6, 25] as it allows continuous attributes to be dimension attributes.

An Array Model for Uncertain Data. We next extend the array model to accommodate uncertain data. When array data are uncertain, the dimension attributes can be uncertain (e.g., the x - y locations of a galaxy follow a bivariate Gaussian distribution); the value attributes can be uncertain (e.g., the luminosity of a galaxy follows a Gaussian); or both groups of attributes can be uncertain.

Uncertainty of value attributes, referred to as *value uncertainty*, is easy to support: we store a (joint) probability distribution of the uncertain value attributes, instead of fixed values, in each tuple. If there is uncertainty regarding the existence of a tuple, called *existence uncertainty*, we store the existence probability as a special value attribute in the tuple. Then we model value and existence uncertainties jointly using a mixed-type distribution [30], which states that the tuple exists with a certain probability, and if the tuple exists, its uncertain value attributes follow a joint distribution.

Uncertainty of dimension attributes is harder to support because a dimension attribute with multiple possible values can cause a tuple to belong to multiple cells in an array, referred to as *position uncertainty*. If we take the tuple's marginal distribution of each uncertain dimension attribute, we can estimate the *possible range* along that dimension where the tuple may belong. Suppose that a marginal distribution has mean μ and standard deviation σ , we can define the possible range to be $[\mu - k\sigma, \mu + k\sigma]$ with a sufficiently large k chosen based on Chebyshev's inequality or Gaussian properties, e.g., when $x \sim N(\mu, \sigma)$, $\Pr(x \in [\mu - 3\sigma, \mu + 3\sigma]) > 0.99$.




Fig. 1 Array \mathbb{A} with dimension attributes, x_loc and y_loc , and the value attribute $luminosity$, all of which can be uncertain.

In this work, we associate each tuple to a default position in the array, which is the cell indexed by the mean values of uncertain dimension attributes. (In implementation we may consider other options, which we discuss later.) Fig. 1 illustrates an array, $\mathbb{A}(x_loc, y_loc; luminosity)$, where continuous uncertain attributes, x_loc and y_loc , are dimension attributes, and the *floor* function discretizes their values as index values. Tuple t_0 has fixed values for x_loc and y_loc and hence belongs to a single cell. Tuple t_1 , however, has a bivariate Gaussian distribution. Therefore, although it is marked in its default cell, $\mathbb{A}[1, 2]$, with a significant probability it can reside in any cell in a possible range, $\mathbb{A}[0 : 5, 0 : 3]$, marked by the red box in the figure. Similarly, t_2 also has a possible range, $\mathbb{A}[2 : 6, 1 : 4]$, due to uncertain x_loc and y_loc .

To draw an analogy with the relational model, we can translate an array to a relation $\mathbb{R}(D_1, \dots, D_d, V_1, \dots, V_m)$. That is, we treat dimension attributes as value attributes in tuples and store all the tuples in a table with no particular order. The top-right corner in Fig. 1 shows the corresponding relation of array \mathbb{A} in the relational model.

2.2 Array Algebra

We next survey operators in array algebra and define their formal semantics under the uncertain data model. These operators were originally proposed in the Array Functional language (AFL) [22] where all attributes have deterministic values and dimension attributes must be discrete valued. In this work, we extend the semantics of these operators to continuous-valued dimension attributes as well as uncertain data in both dimension and value attributes.

Value-based: The operators in the first category operate only on the value attributes of tuples. An example is *Filter*, which applies predicates to the value attributes of tuples stored in the array. Another example is *Apply*, which applies arithmetic operations to the values of tuples. A third example is *Project*, which projects out some value attributes from existing tuples. Since the above operators operate only on the value attributes of tuples, their semantics of uncertain data processing under the array model is the same as the semantics under the relational model; the semantics of the latter is already defined in previous work [30].

Structure-based: The operators in the second category operate on dimension attributes and optionally on value attributes as well. We examine several common operators:

(1) *Subarray* takes an array \mathbb{A} and a condition θ on the dimension attributes, and returns a new array with the tuples that satisfy the condition θ . Revisit our example array.

$Subarray(\mathbb{A}, 1.5 \leq x_loc \leq 3.3 \text{ and } 2.1 \leq y_loc \leq 4.8)$ will first retrieve tuples from the array block $\mathbb{A}[1 : 3, 2 : 4]$, and then filter those tuples based on the precise condition, $1.5 \leq x_loc \leq 3.3$ and $2.1 \leq y_loc \leq 4.8$. The output array always has the same dimensions as the input, but usually fewer cells and tuples. *Subarray* can be translated into selection in relational algebra, i.e., $Subarray(\mathbb{A}, \theta) \equiv \sigma_\theta(R_{\mathbb{A}})$, where $R_{\mathbb{A}}$ is the relational representation of the array.

When the dimension attributes addressed in the condition θ are uncertain, *Subarray* is semantically equivalent to selection on the uncertain dimension attributes in the relational setting. Hence, we have the following definition:

Definition 1 (Probabilistic Subarray) Given an array \mathbb{A}^d , condition θ on uncertain dimension attributes, and a user-specified probability threshold $\lambda \in (0, 1)$, $Subarray(\mathbb{A}, \theta, \lambda)$ returns an array \mathbb{B}^d where the cell $\mathbb{B}[i_1, \dots, i_d]$ contains each tuple t from $\mathbb{A}[i_1, \dots, i_d]$ that satisfies the condition θ with a probability at least λ , i.e., $\int_\theta f_t(\mathbf{x}) dx \geq \lambda$, where $f_t(\mathbf{x})$ is the tuple's probability density function on the the uncertain dimension attributes.

Revisiting the above example, $Subarray(\mathbb{A}, 1.5 \leq x_loc \leq 3.3 \text{ and } 2.1 \leq y_loc \leq 4.8)$. When x_loc and y_loc are uncertain, we can no longer restrict the search to only the block $\mathbb{A}[1 : 3, 2 : 4]$. It is because tuples that belong to other cells, e.g., $\mathbb{A}[1, 5]$, may satisfy the *Subarray* condition with a probability larger than λ . Based on the formal semantics, the entire array needs to be searched.

(2) *Structure-Join (SJoin)* in the array model takes as input an array \mathbb{A}^d , a second array \mathbb{B}^d of the same dimensionality, and a join condition θ . $SJoin(\mathbb{A}, \mathbb{B}, \theta)$ returns an array \mathbb{C}^{2d} , where the cell $\mathbb{C}[i_1, \dots, i_d, i_{d+1}, \dots, i_{2d}]$ contains the result of θ -join between the tuples in $\mathbb{A}[i_1, \dots, i_d]$ and the tuples in $\mathbb{B}[i_{d+1}, \dots, i_{2d}]$. The equivalent expression in relational algebra is, $R_{\mathbb{A}} \bowtie_\theta R_{\mathbb{B}}$, where $R_{\mathbb{A}}$ and $R_{\mathbb{B}}$ are the relational representations of \mathbb{A} and \mathbb{B} .

The join condition, θ , has a few common forms: (1) If the dimension attributes are discrete-valued, θ usually specifies equality comparison on the dimension attributes, as in the AFL proposal [22].¹ (2) If the dimension attributes are continuous-valued, equi-join is seldom used. Instead, θ takes

¹ In this case, the output array, $\mathbb{C} = SJoin(\mathbb{A}, \mathbb{B}, \theta)$, can be simplified to have the same dimensionality as \mathbb{A} and \mathbb{B} , where each cell $\mathbb{C}[i_1, \dots, i_d]$ contains the result of $\mathbb{A}[i_1, \dots, i_d] \bowtie_\theta \mathbb{B}[i_1, \dots, i_d]$. This definition is consistent with equi-join in relational algebra where only one copy of the common join attributes is retained.

a form of proximity join. A common form is linear proximity join, $|\mathbb{A}.d_i - \mathbb{B}.d_i| < \delta$ for each dimension attribute d_i . The join condition essentially defines a band region for each pair of join attributes. Another common form of proximity join uses Euclidean distance, $\sum_i (\mathbb{A}.d_i - \mathbb{B}.d_i)^2 < \delta^2$. As noted earlier, we focus on continuous uncertain data in this paper and hence proximity join in later technical sections.

Next we consider the case that the continuous dimension attributes of arrays \mathbb{A} and \mathbb{B} are uncertain. While the tuples have default positions in the array based on their mean values, they may belong to multiple cells with non-zero probabilities. In the face of position uncertainty, the join between \mathbb{A} and \mathbb{B} must return all pairs of tuples that satisfy the join condition θ with a significant probability. To do so, we leverage the semantics of cross-product in the above *SJoin* definition, which involves pairing each cell in \mathbb{A} with each cell in \mathbb{B} and then pairing the tuples within those cells. More specifically, we define probabilistic structure-join as follows:

Definition 2 (Probabilistic Structure-Join) Given arrays \mathbb{A}^d and \mathbb{B}^d , a join condition θ , and a probability threshold λ , $SJoin(\mathbb{A}, \mathbb{B}, \theta, \lambda)$ returns an array \mathbb{C}^{2d} where $\mathbb{C}[i_1, \dots, i_d, i_{d+1}, \dots, i_{2d}]$ contains the result of probabilistic θ -join, $\mathbb{A}[i_1, \dots, i_d] \bowtie_{\theta, \lambda} \mathbb{B}[i_{d+1}, \dots, i_{2d}] = \{(t_1, t_2) | t_1 \in \mathbb{A}[i_1, \dots, i_d], t_2 \in \mathbb{B}[i_{d+1}, \dots, i_{2d}], \iint_\theta f_{t_1}(\mathbf{x}) \cdot f_{t_2}(\mathbf{y}) dx dy \geq \lambda\}$, where $f_{t_1}(\mathbf{x})$ and $f_{t_2}(\mathbf{y})$ are the probability density functions for t_1 and t_2 , respectively.

(3) *Regrid-Aggregation* partitions an input array into non-overlapping blocks, and for each block, applies an aggregate function to all the tuples in the block. The output array has one cell for each block which contains the aggregate value computed. *Regrid* can be viewed as repeated application of the *Subarray* operation to extract each block and then to compute the aggregate within each block.

When the dimension attributes are uncertain, one can use the *Probabilistic Subarray* operator to extract the tuples that belong to each block with non-zero probabilities (usually more than those that physically exist in the block). Note that even if a tuple belongs to a block with a small probability, if its aggregate attribute has a large value, it can still contribute a modest value, which is the product of its attribute value and existence probability, to the aggregate. Hence, the probability threshold for tuple existence in *Subarray* should be set to 0 in theory, or a small value in practice.

(4) *GroupBy-Aggregation* takes three arguments including an input array \mathbb{A}^d , a list of grouping dimensions \mathbb{G}^{d_1} , where $d_1 \leq d$, and an aggregate function. Again, it can be viewed as repeated application of *Subarray* to construct array blocks corresponding to the groups and then computing the aggregate within each block.

As shown in the above discussion, *Subarray* and *Structure-Join* are the two most important primitives in array algebra.

Hence, we focus on efficient implementation of them under data uncertainty in subsequent sections.

We finally show two example queries written for the Sloan Digital Sky Survey (SDSS) [28], where the attributes, *rowc* and *colc*, define a two dimensional array called Galaxy. These queries are written by following the convention of the AFL language, but with syntactic differences due to reasons such as the support of continuous dimension attributes. Query Q_1 computes the average brightness in a subarray region. Query Q_2 finds the regions in the sky space with the observation density greater than a threshold τ , where the observation density for a region is defined to be the sum of the number of observations within δ distance in either direction of each point in the region. To do so, the query first performs a self join of the Galaxy array based on the array structure, then groups the join output into 100 by 100 blocks of $G_1.rowc$ and $G_1.colc$, and finally counts the number of observations per block for density filtering.

```

Q1:AVG (SUBARRAY (Galaxy,  $x_1 < rowc < x_2$  and
                 $y_1 < colc < y_2$ ), brightness)

Q2:GROUPBY ( SJOIN ( Galaxy  $G_1$ , Galaxy  $G_2$ ,
                     $|G_1.rowc - G_2.rowc| < \delta$  and
                     $|G_1.colc - G_2.colc| < \delta$ ),
            [ $G_1.rowc/100$ ] and [ $G_1.colc/100$ ], COUNT cnt)

```

3 Native Support for Subarray

In this section, we focus on the structural operation, *Subarray*. Since *Subarray* is equivalent to selection in relational algebra, there are two options for implementation: The first option is to translate *Subarray* to selection in the relational setting. When the dimension attributes are uncertain, to avoid scanning all tuples in the database, existing work has built various indexes based on statistical quantities such as quantiles [7, 8, 29] and moments [20] of tuple distributions. However, these indexes may not be effective when the filtering power is low and can trigger many index I/O's, as we will shown in § 5. The second option is to build native support of *Subarray* in array databases where logical and physical localities are better-aligned. For instance, *Subarray* that exploits logical locality of data, e.g., looking for adjacent array cells from a point, may need to retrieve only a few relevant physical storage units called chunks. This effect of exploiting physical locality is similar to using a clustered primary index on the tuples in a relational database, but without having to build the index.

Hence, in this work we focus on native support of array operations on uncertain data. The task is challenging due to *position uncertainty*. We aim to reduce both the number of chunks loaded (*I/O cost*), often way beyond the query region, and the number of costly integrations to compute tuple probabilities (*CPU cost*), incurred for both relevant and irrelevant tuples in the loaded chunks.

3.1 Storage and Evaluation Schemes

Given a tuple and a probability distribution on its dimension attributes, the array cells that have non-zero probabilities to contain the tuple form the “*possible range*” of the tuple, as defined in Section 2.1. We next propose a number of storage schemes with the property that we can observe the tuple when any cell of its possible range is covered in a query region.

Store-All: One solution is to store a copy of the tuple in each cell of the tuple’s possible range. Fig. 2(a) depicts the storage of two tuples, t_1 and t_2 , where t_1 is replicated in its possible range $\mathbb{A}[0:5, 0:3]$ (including the red and yellow cells), and t_2 is replicated in $\mathbb{A}[2:6, 1:4]$ (the green and yellow cells), with the overlap region marked in yellow. A query region, $\mathbb{A}[2:2, 3:3]$, is marked by a solid blue box in Fig. 2(d). A major advantage of this scheme is that we can execute the query region directly on the array, without any missed results. The disadvantages include possibly excessive storage overheads and high I/O costs in querying because each logical cell may need many physical chunks to store the replicated tuples.

A similar storage scheme is *store-all with pointers*: store a tuple in its default cell defined by the mean values of its dimension attributes, and add a pointer to this tuple in all other cells in the tuple’s possible range. Thus we avoid repeated storage of all attributes of a tuple. However, the numerous pointers can still incur high storage overheads, as well as frequent random I/O’s at query time as a result of chasing pointers from a given query region to fetch relevant tuples stored outside the region.

Store-Mean: To reduce storage overheads, we next consider storing a tuple only once based on the mean values of its dimension attributes. However, directly running *Subarray* on such storage will lead to missed results: tuples whose mean values are outside the query region but whose possible ranges overlap with the region will be missed. To avoid the problem, the query region must be expanded. In this work, we augment each cell with upper and lower bounds for each dimension, indicating the distance to travel along each dimension in order to find all tuples that could belong to that cell—we call these bounds the *upper and lower fences* for expanding the query region from this cell. This way, the storage overhead is limited to two integers per dimension per cell.

Fig. 2(b) shows the storage layout for tuples t_1 and t_2 . Fig. 2(e) shows the cell $\mathbb{A}[2, 3]$ covered in a query region marked by a solid blue box. The fences for the x dimension, $(-1, 3)$, means that at query time, from this cell we need to walk one step to the left and three steps to the right, while the fences for the y dimension, $(-1, 1)$, indicates walking one step up and one step down. After walking on both dimensions, the expanded query region, marked by a dashed blue

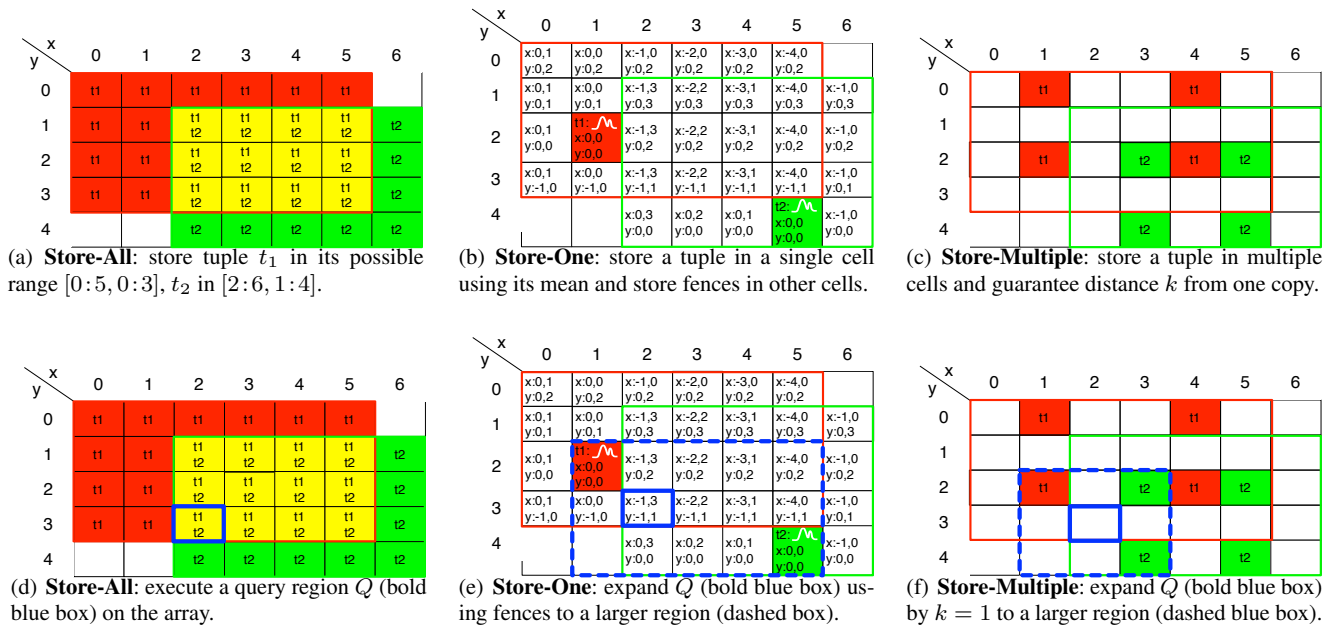


Fig. 2 Alternative storage and evaluation strategies for tuples with uncertain dimension attributes.

box, covers cell $\mathbb{A}[1, 2]$ to retrieve tuple t_1 and cell $\mathbb{A}[5, 4]$ to retrieve t_2 .

To generate fences, whenever a new tuple is inserted into a cell C in the array based on its mean value, we identify every cell in the tuple's possible range, compute its distance from the cell C , then expand its fences if they do not cover the computed distance. At query time, for each cell contained in the query region, we expand it using the upper and lower fences, and take the union of all these expansions to produce a complete *expanded query region*. We perform such expansion using every cell that falls in the query region. Fig. 2(e) depicts the user-specified query region in a solid blue box and the expanded query region using a dashed blue box.

The advantage of this strategy is small storage overhead in each cell, i.e., only two fences for each dimension, in contrast to *store-all* and *store-all with pointers*. However, the issue is that the expanded query region can grow very large, containing both relevant and irrelevant tuples, which will incur both high I/O cost for fetching all the tuples and high CPU cost for validating them using costly integration based on the precise *Subarray* condition.

Store-Multiple: Finally, we propose a scheme that employs limited replication of tuples and guarantees that from any cell in a tuple's possible range, the query needs to be expanded by at most k cells (steps) along each dimension to find a copy of the tuple. We call k the **step size**, and use it to control both query expansion, by k in each direction on a dimension, and the degree of replication in storage, roughly inversely proportional to k . Most importantly, *store-multiple* overcomes the shortcomings of the previous two schemes:

First, its controlled expansion of the query region, by k cells, is particularly helpful when some tuples have large variances and hence large possible ranges. In other schemes, tuples of large variances will cause them to be replicated in numerous cells (*store-all*) or cause the query region to be expanded based on the largest tuple variance in a wide neighborhood (*store-mean*). Second, *store-multiple* offers the flexibility to configure the parameter k for different workloads to achieve best performance, as we shall show shortly. It is also worth noting that *store-multiple* subsumes both *store-all* and *store-mean*: it becomes *store-all* when $k = 0$, and approximates *store-mean* (without fences) when k is big enough to cover the largest possible range among all tuples.

Fig. 2(c) shows such storage with $k=1$, where tuple t_1 is stored in four cells and t_2 in another four cells. We can verify that for each cell in t_1 's possible region (the red rectangle), we need to walk only one step in both dimensions to find a copy of t_1 . The same guarantee holds for t_2 . Fig. 2(f) shows a query region matching the cell $\mathbb{A}[2, 3]$, marked by the solid blue box, and the expanded region $\mathbb{A}[1 : 3, 2 : 4]$ using $k = 1$, marked by the dashed blue box.

The evaluation of *Subarray* under *store-multiple* includes two steps: (1) I/O step: The original query region is expanded by k cells along both directions on each dimension; all tuples in the expanded query region are read from disk. (2) CPU step: The exact existence probability in the query region is computed for each retrieved tuple based on its distribution and compared with the probability threshold. Since computing the exact existence probability for continuous random variables requires expensive integration, we optimize it by first running fast filters [20] with negligible costs to

prune tuples of low existence probabilities, and building an in-memory hash table to avoid doing integration for different copies of the same tuple.

Two questions remain for *store-multiple*: First, the way to store tuples while guaranteeing the step size k in query expansion is not unique, leading to different degrees of replication of a tuple. How do we find the best layout of tuples under the step size k configuration? Second, given a dataset and typical query workloads, how do we choose the best configuration of k for optimal performance? We address these two issues in §3.2 and §3.3, respectively.

3.2 Tuple Layout under Store-Multiple

Consider the tuple layout in a d -dimensional array \mathbb{A}^d stored using *store-multiple* with a step size configuration $\langle k_1, k_2, \dots, k_d \rangle$. This means that from any cell in the tuple's possible range, walking k_i cells in both directions on the i -th dimension, for $1 \leq i \leq d$, guarantees to find a copy of the tuple. Finding the best way to store tuple copies amounts to a coverage problem, as we define below.

Definition 3 (Covering Cell) Given a d -dimensional array \mathbb{A}^d under *store-multiple* with a step size configuration $\langle k_1, k_2, \dots, k_d \rangle$, the covering range of the walk from a cell $\mathbb{A}[x_1, x_2, \dots, x_d]$ is $\mathbb{A}[x_1 - k_1 : x_1 + k_1, \dots, x_d - k_d : x_d + k_d]$. We also say each cell in $\mathbb{A}[x_1 - k_1 : x_1 + k_1, \dots, x_d - k_d : x_d + k_d]$ is “covered” by the cell $\mathbb{A}[x_1, x_2, \dots, x_d]$.

Definition 4 (Covering Set) A given set of cells \mathcal{S} is covered by a (discrete) set of cells \mathcal{C} if and only if each cell in \mathcal{S} is covered by at least one cell in \mathcal{C} ; \mathcal{C} is called the covering set of \mathcal{S} .

Definition 5 (Problem of Tuple Copy Layout) Given a tuple t , find the minimum covering set \mathcal{C} of its possible range $\mathcal{S} = \mathbb{A}[l_1 : u_1, l_2 : u_2, \dots, l_d : u_d]$ so that placing one copy of the tuple in each cell of \mathcal{C} guarantees the correctness of query expansion using the step size $\langle k_1, k_2, \dots, k_d \rangle$.

We address the problem by first showing the lower bound of the size of a covering set, as shown in the following proposition.

Proposition 1 Given an array \mathbb{A}^d under *store-multiple* with a step size configuration $\langle k_1, k_2, \dots, k_d \rangle$, if a tuple's possible range is $\mathcal{S} = \mathbb{A}[l_1 : u_1, l_2 : u_2, \dots, l_d : u_d]$, at least $\prod_{i=1}^d (\lfloor (u_i - l_i) / (2k_i + 1) \rfloor + 1)$ cells are needed to cover \mathcal{S} .

Proof We can pick a subset of cells from the region $\mathcal{S} = \mathbb{A}[l_1 : u_1, l_2 : u_2, \dots, l_d : u_d]$ as follows: $\mathcal{S}' = \{\mathbb{A}[x_1, x_2, \dots, x_d] \mid \forall i \in \{1, 2, \dots, d\}, x_i = l_i + p_i(2k_i + 1) \text{ and } l_i \leq x_i \leq u_i, \text{ where } p_i \in \{0\} \cup \mathbb{N}\}$. Obviously, the size of the

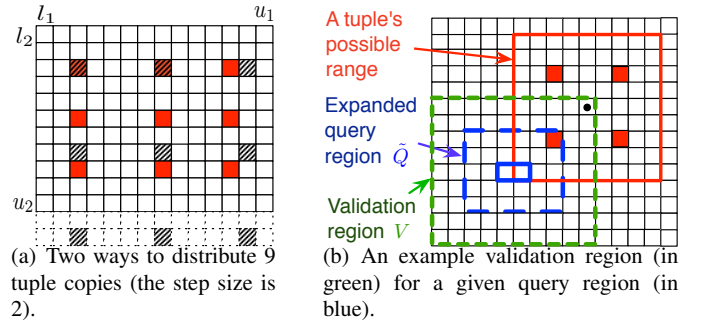


Fig. 3 Tuple copy distribution and validation under *store-multiple*.

set of picked cells $|\mathcal{S}'|$ is $\prod_{i=1}^d (\lfloor (u_i - l_i) / (2k_i + 1) \rfloor + 1)$. Based on Definition 4, if we can prove that at least $|\mathcal{S}'|$ cells are already needed just to cover \mathcal{S}' , then at least $|\mathcal{S}'|$ cells are needed to cover the superset \mathcal{S} .

Let us assume a cell $\mathbb{A}[x_1, x_2, \dots, x_d] \in \mathcal{S}'$ is covered by (the walk from) a cell $\mathbb{A}[y_1, y_2, \dots, y_d]$. This means $y_i - k_i \leq x_i \leq y_i + k_i$ on any dimension i . For any cell $\mathbb{A}[x'_1, x'_2, \dots, x'_d] \in \mathcal{S}' - \{\mathbb{A}[x_1, x_2, \dots, x_d]\}$, there exists a dimension j such that $x_j \neq x'_j$. Without loss of generality, assume $x'_j = x_j + p_j(2k_j + 1)$ where $p_j \in \mathbb{N}$. Then $x'_j \geq y_j - k_j + p_j(2k_j + 1) > y_j + k_j$, which means $\mathbb{A}[x'_1, x'_2, \dots, x'_d]$ does not fall in the covering range of $\mathbb{A}[y_1, y_2, \dots, y_d]$. Therefore, no two cells in \mathcal{S}' can be covered by the same cell. In other words, at least $|\mathcal{S}'|$ cells are needed in order to cover \mathcal{S}' . Then to cover \mathcal{S} , a superset of \mathcal{S}' , at least $|\mathcal{S}'| = \prod_{i=1}^d (\lfloor (u_i - l_i) / (2k_i + 1) \rfloor + 1)$ cells are needed as well.

Given the lower bound on the size, we next consider how to distribute the covering set, i.e., the cells with tuple copies, to achieve this lower bound. To maximize the union of the covering ranges of those tuple copies, we can store them in evenly-spaced cells: on the i -th dimension where the possible range is l_i, u_i , the first copy is stored at $l_i + k_i$ and the other copies are stored $2k_i$ cells away from each other. Fig. 3(a) shows such distribution of tuple copies in a two-dimensional array when $k_1 = k_2 = 2$. The tuple's possible range consists of all the cells within the solid boundary and requires at least 9 copies to be placed. The layout of 9 copies is shown by the shaded cells (ignore the red color for now). However, three copies are stored outside the tuple's possible range, which will increase the chance of reading irrelevant copies when a query region falls outside the tuple's possible range. It is thus desirable to store all copies of a tuple inside its possible range. In our work, when a tuple needs only one copy on the i -th dimension, we store it at the center of its possible range, i.e., $\lfloor (l_i + u_i) / 2 \rfloor$; when it needs more copies, we store the first copy at $l_i + k_i$, the last copy at $u_i - k_i$, and the others (if any) are evenly spaced in between, as shown

symbol	description
T	number of tuples
b	number of bytes per tuple
pr_i	length of a tuple's possible range on the i -th dimension
d	dimensionality of an array
c	chunk size (the I/O unit) in bytes
s_i	length of each cell on the i -th dimension
n_i	number of cells on the i -th dimension
q_i	query region size on the i -th dimension
k_i	step size on the i -th dimension

Table 1 Notation in modeling and analysis.

by the red cells in the figure. Thus we still use the minimum number of copies to cover the tuple's possible range.

3.3 Cost Model of Subarray under Store-Multiple

We next propose a cost model for *Subarray* under the *store-multiple* scheme and use the model to find the optimal step size configuration. The symbols used in the model are summarized in Table 1. Like in SciDB [6], a cell is a logical unit in an array while a chunk is a physical storage unit as well as the I/O unit; tuples in a logical cell can be stored in one or multiple chunks. For *Subarray* evaluation under *store-multiple*, as explained in §3.1, the I/O cost consists of the seek and transfer time of chunks in the expanded query region, while the CPU cost is the product of the number of tuples to be validated and the validation cost per tuple. For simplicity, we assume that the centers of tuples' possible ranges are evenly spread over the whole array. We also begin by assuming that all tuples' possible ranges have the same size, pr_i , on the i -th dimension. Our model can be extended to support possible ranges of variable sizes, as we explain at the end of the section.

I/O Cost: To capture I/O cost, we focus on a key factor, the number of chunks in the expanded query region.

Let us first compute the number of cells with which a tuple's possible range overlaps on the i -th dimension. Obviously this depends on the alignment of the possible range and the cells along this dimension, as shown in Fig. 4. We can chop the possible range into $\lceil pr_i/s_i \rceil$ segments, where the first $\lceil pr_i/s_i \rceil - 1$ segments have length s_i and the last segment has length $r = pr_i - (\lceil pr_i/s_i \rceil - 1)s_i$. Depending on the starting position of the possible range in the first cell, it can overlap with different numbers of cells: when the starting position is in $[0, s_i - r]$, it overlaps with $\lceil pr_i/s_i \rceil$ cells; when the starting position is in $(s_i - r, s_i)$, it overlaps with $\lceil pr_i/s_i \rceil + 1$ cells. Then the expected number of cells the possible range overlaps with is

$$\frac{s_i - r}{s_i} \left\lceil \frac{pr_i}{s_i} \right\rceil + \frac{r}{s_i} \left(\left\lceil \frac{pr_i}{s_i} \right\rceil + 1 \right) = \frac{pr_i}{s_i} + 1 \quad (1)$$




Fig. 4 Illustration of the number of cells that a possible range overlaps with

Calculated in a similar way, the number of cells that overlap with the query region Q on the i -th dimension is $q_i/s_i + 1$, and the number for the expanded query region \tilde{Q} is $q_i/s_i + 1 + 2k_i$.

We next model the number of chunks in the expanded query region \tilde{Q} . It is the product of the number of cells in \tilde{Q} and the average number of chunks per cell. To compute the latter, we first write $u_i - l_i + 1 = pr_i/s_i + 1$ based on Eq.(1), and plug it into Proposition 1 to derive the number of copies per tuple:

$$\prod_{i=1}^d \left(\left\lfloor \frac{pr_i/s_i}{2k_i + 1} \right\rfloor + 1 \right). \quad (2)$$

Then the average number of chunks per cell is the total number of tuple copies in the array divided by the number of cells in the array and then by the number of tuples a chunk can hold, i.e., $\lfloor c/b \rfloor$:

$$T \prod_{i=1}^d \left(\left\lfloor \frac{pr_i/s_i}{2k_i + 1} \right\rfloor + 1 \right) / \prod_{i=1}^d n_i / \lfloor c/b \rfloor. \quad (3)$$

Multiplying this with the number of cells in \tilde{Q} , $\prod_{i=1}^d (q_i/s_i + 1 + 2k_i)$, we get the number of chunks in \tilde{Q} :

$$T \prod_{i=1}^d \left(\left\lfloor \frac{pr_i/s_i}{2k_i + 1} \right\rfloor + 1 \right) / \prod_{i=1}^d n_i / \lfloor c/b \rfloor \cdot \prod_{i=1}^d \left(\frac{q_i}{s_i} + 1 + 2k_i \right) \quad (4)$$

CPU Cost: To capture CPU cost, we model the number of tuples to be validated. Given an expanded query region \tilde{Q} , a tuple is retrieved for validation as long as it has one copy stored in \tilde{Q} . Fig. 3(b) shows an example tuple whose possible range is the red box. The tuple has a copy in \tilde{Q} but its center of the possible range, marked by a black dot, lies outside \tilde{Q} . To handle such tuples, let us define the *validation region*, V , to be the set of cells where the centers of the possible ranges of to-be-validated tuples reside, and model the size of V first. Now consider the i -th dimension of the array: (1) When k_i is large enough that every tuple only needs one copy to cover its possible range, V is simply the expanded query region \tilde{Q} , of size $q_i/s_i + 1 + 2k_i$. (2) When k_i is smaller so that all tuples have more than one copy, tuples that have at least one copy stored in \tilde{Q} need to be validated. It can be derived that in this case V on the i -th dimension is of size $q_i/s_i + 1 + pr_i/s_i$. Summarizing the above two cases

and multiplying the size of V with the number of tuples per cell, we have the number of tuples to be validated as:

$$T / \prod_{i=1}^d n_i \cdot \prod_{i=1}^d \left(\frac{q_i}{s_i} + 1 + z_i \right), \quad (5)$$

where $z_i = 2k_i$ when $pr_i/s_i < 2k_i + 1$ and $z_i = pr_i/s_i$ otherwise.

Finally we combine the I/O and CPU costs by plugging in unit cost measurements, including the seek and transfer time per chunk and per tuple validation time using integration.

A Generalized Model. We next relax two assumptions made previously in our model: (1) When tuples have different possible range sizes, we can group tuples based on the possible range size. The runtime of a query will be a weighted sum of runtime over each group of tuples, where the number of tuples per group serves as the weight; (2) When tuples are not evenly distributed in the domain, we can feed statistics of tuples' mean positions and the query position into our model to get a more accurate estimate: instead of using $T / \prod_{i=1}^d n_i$, which is the average number of tuples per cell, we can use the number of tuples in each cell of the query. In practice, we can collect above-mentioned statistics when a batch of tuples comes in. For instance, SDSS [28] updates the scanned image of the sky on a nightly basis and can build the statistics as a nightly observation is being produced. If domain knowledge shows that the statistics do not change drastically from day to day, we can also re-use statistics collected in the past.

Implementation. Given the cost model and basic statistics of tuples' possible range sizes and query sizes, at data loading time we estimate the costs of representative queries by running our model for a wide range of step size configurations (which runs fast), and choose the configuration that offers the best performance.

We evaluate the effectiveness of our model for doing so in §5. In addition, our implementation of *store-multiple* supports *store-all* and *store-mean* with only modest changes. In particular, it allows evaluation algorithms that use only the mean copy of each tuple, ignoring other copies.

4 Support for Structure Join

In this section, we focus on the *Structure-Join* operator under position uncertainty. The default evaluation strategy, as stated in Definition 2, creates all pairs of tuples from the two input arrays and evaluates an integral for each pair of tuples, which is prohibitively expensive. To improve performance, we propose to integrate existing indexes for relational databases [10,9,20] with the *store-multiple* scheme for array databases and derive an index-based technique (in §4.1). We further propose a new subarray-based evaluation

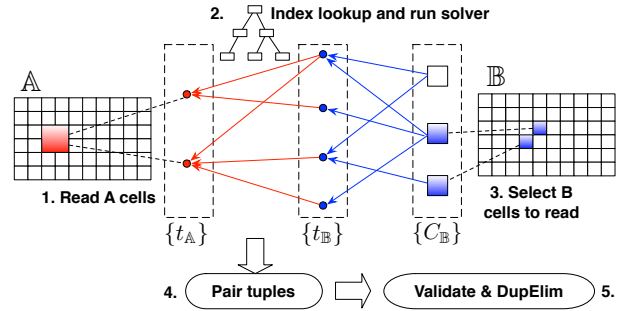


Fig. 5 Illustration of index-based join.

strategy, which does not require an index, and model-based optimization to achieve its best performance (in §4.2).

In our discussion below, we focus on linear proximity join, that is, $SJoin(\mathbb{A}^d, \mathbb{B}^d, \theta, \lambda)$ where $\theta = \bigwedge_{i=1}^d |\mathbb{A}.d_i - \mathbb{B}.d_i| < \delta$, which is the most common in scientific applications. Non-linear proximity join based on Euclidean distance, e.g., $\sum_i (\mathbb{A}.d_i - \mathbb{B}.d_i)^2 < \delta^2$, can be first relaxed to linear proximity join, and then followed by additional filtering using exact integration based on θ . Moreover, we define the selectivity of a probabilistic $SJoin$ to be $\frac{|C|}{|\mathbb{A}| \cdot |\mathbb{B}|}$, where $|\mathbb{A}|$, $|\mathbb{B}|$, $|C|$ are tuple counts of inputs \mathbb{A} and \mathbb{B} , and the output C .

4.1 Index-based Join

Recent work has proposed new indexes on continuous uncertain data in relational databases to improve query performance [10,9,20]. A natural way to use indexes in *Structure-Join* is to perform index nested loops join (INLJ): the outer relation is scanned once; for each tuple in the outer relation, the index on the inner relation is probed to find the candidates. Integrating INLJ with the *store-multiple* scheme, which we proposed for handling uncertain data, raises a new question: Since tuples may have multiple copies stored in the array database, how can we minimize the chance of producing duplicate results from multiple copies of the same tuple and hence the associated I/O and CPU costs?

We propose an algorithm named index-based joins (IBJ) and present it as Algorithm 1. Our main idea is as follows: Consider $SJoin(\mathbb{A}^d, \mathbb{B}^d, \theta, \lambda)$ where both \mathbb{A} (the outer) and \mathbb{B} (the inner) are stored using *store-multiple*. Assume that there is a pre-built index on the continuous uncertain join attribute(s) in \mathbb{B} , and the leaf nodes in the index store the possible range of each tuple. As shown in Fig. 5, the array \mathbb{A} is scanned once, by reading one block at a time (the “for” loop in Line 1) where each block can contain one or multiple cells. Each tuple, t_A , in the current block of \mathbb{A} triggers an index lookup to find its candidate tuples from \mathbb{B} (Line 7), $\{t_B\}$, as depicted by the left mapping structure in Fig. 5. Each candidate tuple t_B may be stored in multiple cells, $\{C_B\}$, as de-

Algorithm 1 Index-based Joins (IBJ)

```

1: for each read block  $R_A$  in  $\mathbb{A}$  do
2:    $tmap.clear()$ ;
3:    $subsets.clear()$ ;
4:   for each cell  $C_A$  in  $R_A$  do
5:      $loadToMemory(C_A)$ ;
6:     for each tuple  $t_A$  in  $C_A$  do
7:        $candidates \leftarrow findCandidates(t_A)$ ;
8:       for each tuple  $t_B$  in  $candidates$  do
9:          $tmap.get(t_B).add(t_A)$ ;
10:         $cells \leftarrow findCells(t_B)$ ;
11:        for each cell  $C$  in  $cells$  do
12:           $subsets.get(C).add(t_B)$ ;
13:         $toRead \leftarrow setCover(subsets)$ ;
14:        for each cell  $C_B$  in  $toRead$  do
15:           $loadToMemory(C_B)$ ;
16:          for each tuple  $t_B$  in  $C_B$  do
17:            for each tuple  $t_A$  in  $tmap.get(t_B)$  do
18:               $validate(t_A, t_B)$ ;
19:             $tmap.remove(t_B)$ ;
20:  $removeDuplicates()$ ;

```

pictured by the right mapping structure; these C_B cells can be computed from t_B 's possible range stored in the index (Line 10). After processing all tuples in the current block of \mathbb{A} , two mapping structures are complete (maintained in Line 9 and Line 12). We want to read a subset of the \mathbb{B} cells to minimize I/O while guaranteeing *coverage*: every candidate \mathbb{B} tuple resides in at least one of those cells. Then the \mathbb{A} tuples and their candidates \mathbb{B} tuples are paired and validated (Line 18) if the index returns a superset of true matches [10, 9] (otherwise not needed if the index returns only the true matches [20]). Duplicate matches due to existence of tuple copies can be removed using an in-memory hash table (Line 20).

Given the two mapping structures, deciding which \mathbb{B} cells to read to minimize I/O while guaranteeing coverage amounts to a set covering problem (SCP). Denote the union of candidates for all probing tuples in the current block as U . Each cell C_B from \mathbb{B} stores a subset of U . Using the number of chunks as the weight of C_B , denoted by $\omega(C_B)$, our problem is a variant of the classic SCP:

$$\begin{aligned}
& \text{minimize} && \sum_{C_B \in \mathbb{B}} \omega(C_B) \cdot I_{C_B} \\
& \text{subject to} && \sum_{C_B: t_B \in C_B} I_{C_B} \geq 1 && \text{for all } t_B \in U \\
& && I_{C_B} \in \{0, 1\} && \text{for all } C_B \in \mathbb{B}
\end{aligned}$$

This is an NP-hard problem but can be solved efficiently by many techniques, e.g., the integer linear program and greedy algorithm (Line 13).

The above algorithm for *Structure-Join*, called index-based join (IBJ), has several configuration choices. (1) *Storage schemes*: Although IBJ generally works for input arrays with any step size configuration, its performance is better

shown when the outer array uses a large step size so that each tuple has fewer copies and thus the chance of duplicate index lookups for the same tuple is reduced. It turns out IBJ also prefers a large step size for the inner to minimize I/O cost. We provide a detailed explanation in evaluation in §5. (2) *Memory allocation*: In this algorithm, the memory is shared among the read block of \mathbb{A} , the two mapping structures, the index, and the cache of \mathbb{B} cells (or chunks). Since the sizes of both mappings increase with the size of the \mathbb{A} block, these three data structures share a memory quota and then the block size can be automatically determined by keeping adding new \mathbb{A} cells until reaching the memory quota. The index and the cache of \mathbb{B} have their own memory quotas. These three parts compete for memory: if more memory is given to the \mathbb{A} block and the mappings, fewer blocks are needed and SCP is more effective in reducing the I/O of reading \mathbb{B} ; the more memory is given to the index, the lower I/O cost in index lookups; and the more is given to cache \mathbb{B} cells, the lower I/O cost in reading \mathbb{B} cells. We evaluate different memory allocation schemes in §5.

4.2 Subarray-based Join

The index-based join requires pre-built indexes, which may not always be available, and can consume excessive memory due to the use of the tuple-level mapping. We next present a new evaluation strategy of *Structure-Join*, called subarray-based join (SBJ).

Similar to block nested loops joins, *Structure-Join* can be transformed into iterative *Subarray* operations on the inner array, for each block of the outer array. Assume that the smaller array, \mathbb{A} , is the outer. For each cell C_A , we do the following: (1) Load it into memory, form a subarray condition θ_{C_A} on the inner array \mathbb{B} , and run the *Subarray* query on \mathbb{B} . (2) Pair tuples in C_A with those tuples retrieved by *Subarray* on \mathbb{B} . (3) The final validation phase computes the exact probability for each tuple pair (t_A, t_B) to satisfy the join condition and compares it with the threshold λ . We describe the subarray condition θ_{C_A} and the full algorithm in §4.2.1 and present a cost model for optimization in §4.2.2.

4.2.1 Subarray Conditions and the SBJ Algorithm

The subarray condition θ_{C_A} for each outer cell C_A must produce all join results while being as tight as possible to maximize performance. Below we propose several necessary conditions for linear proximity join that guarantee the retrieval of all join results.

Given a tuple t_A , let $(l_{t_A.d_i}, u_{t_A.d_i})$ denote the lower and upper bounds of its possible range on the i -th dimension. Similarly, we have $(l_{t_B.d_i}, u_{t_B.d_i})$ for tuple t_B . Then we have:




Fig. 6 Illustration of $\tilde{p} = \tilde{p}_{1,1} + \tilde{p}_{1,2} + \tilde{p}_2$.

Proposition 2 For any tuple pair $(t_{\mathbb{A}}, t_{\mathbb{B}})$ returned by $SJoin(\mathbb{A}^d, \mathbb{B}^d, \bigwedge_{i=1}^d |\mathbb{A}.d_i - \mathbb{B}.d_i| < \delta, \lambda)$, the intervals $(l_{t_{\mathbb{A}}, d_i} - \delta, u_{t_{\mathbb{A}}, d_i} + \delta)$ and $(l_{t_{\mathbb{B}}, d_i}, u_{t_{\mathbb{B}}, d_i})$ overlap on any dimension i ($i = 1, \dots, d$).

Proof We prove by contradiction. Consider a tuple pair $(t_{\mathbb{A}}, t_{\mathbb{B}})$ returned by $SJoin$. Assume that there exists a dimension d_i where $(l_{t_{\mathbb{A}}, d_i} - \delta, u_{t_{\mathbb{A}}, d_i} + \delta)$ and $(l_{t_{\mathbb{B}}, d_i}, u_{t_{\mathbb{B}}, d_i})$ do not overlap, i.e., $l_{t_{\mathbb{A}}, d_i} - \delta > u_{t_{\mathbb{B}}, d_i}$ or $u_{t_{\mathbb{A}}, d_i} + \delta < l_{t_{\mathbb{B}}, d_i}$. Without loss of generality, let us assume $l_{t_{\mathbb{A}}, d_i} - \delta > u_{t_{\mathbb{B}}, d_i}$. Below we focus on computing probability $p = \iint_{\theta} f_{t_{\mathbb{A}}}(\mathbf{x}) f_{t_{\mathbb{B}}}(\mathbf{y}) dx dy$ where the integration domain θ is $\{(\mathbf{x}, \mathbf{y}) \mid \bigwedge_{i=1}^d |x_i - y_i| < \delta\}$. We start with finding an upper bound of p . Relaxing the join condition by only considering dimension d_i , we have

$$p < \iint_{|x_i - y_i| < \delta} f_{t_{\mathbb{A}}}(\mathbf{x}) f_{t_{\mathbb{B}}}(\mathbf{y}) dx dy = \iint_{|x_i - y_i| < \delta} f_{t_{\mathbb{A}}, d_i}(x_i) f_{t_{\mathbb{B}}, d_i}(y_i) dx_i dy_i.$$

It means the probability for $(t_{\mathbb{A}}, t_{\mathbb{B}})$ to satisfy the join predicate is upper bounded by the probability for their values on dimension d_i to satisfy the join predicate on dimension d_i , denoted as \tilde{p} . The integration domain is colored in Fig. 6 and partitioned into three parts. Denote the probability mass of each partition as $\tilde{p}_{1,1}$, $\tilde{p}_{1,2}$ and \tilde{p}_2 respectively. Below we derive the upper bound for each of them by applying the assumption.

$$\begin{aligned} \tilde{p}_{1,1} &= \int_{-\infty}^{l_{t_{\mathbb{A}}, d_i} - \delta} f_{t_{\mathbb{B}}, d_i}(y_i) \left(\int_{y_i - \delta}^{y_i + \delta} f_{t_{\mathbb{A}}, d_i}(x_i) dx_i \right) dy_i \\ &< \int_{-\infty}^{l_{t_{\mathbb{A}}, d_i} - \delta} f_{t_{\mathbb{B}}, d_i}(y_i) \left(\int_{-\infty}^{l_{t_{\mathbb{A}}, d_i}} f_{t_{\mathbb{A}}, d_i}(x_i) dx_i \right) dy_i \\ &= \frac{\epsilon}{2} \int_{-\infty}^{l_{t_{\mathbb{A}}, d_i} - \delta} f_{t_{\mathbb{B}}, d_i}(y_i) dy_i < \frac{\epsilon}{2} \end{aligned}$$

$$\begin{aligned} \tilde{p}_{1,2} &= \int_{l_{t_{\mathbb{A}}, d_i} - \delta}^{l_{t_{\mathbb{A}}, d_i} + \delta} f_{t_{\mathbb{B}}, d_i}(y_i) \left(\int_{y_i - \delta}^{l_{t_{\mathbb{A}}, d_i}} f_{t_{\mathbb{A}}, d_i}(x_i) dx_i \right) dy_i \\ &< \int_{u_{t_{\mathbb{B}}, d_i}}^{+\infty} f_{t_{\mathbb{B}}, d_i}(y_i) \left(\int_{-\infty}^{l_{t_{\mathbb{A}}, d_i}} f_{t_{\mathbb{A}}, d_i}(x_i) dx_i \right) dy_i \\ &= \frac{\epsilon}{2} \int_{u_{t_{\mathbb{B}}, d_i}}^{+\infty} f_{t_{\mathbb{B}}, d_i}(y_i) dy_i = \frac{\epsilon}{2} \cdot \frac{\epsilon}{2} = \frac{\epsilon^2}{4} \\ \tilde{p}_2 &= \int_{l_{t_{\mathbb{A}}, d_i}}^{+\infty} f_{t_{\mathbb{A}}, d_i}(x_i) \left(\int_{x_i - \delta}^{x_i + \delta} f_{t_{\mathbb{B}}, d_i}(y_i) dy_i \right) dx_i \\ &< \int_{l_{t_{\mathbb{A}}, d_i}}^{+\infty} f_{t_{\mathbb{A}}, d_i}(x_i) \left(\int_{u_{t_{\mathbb{B}}, d_i}}^{+\infty} f_{t_{\mathbb{B}}, d_i}(y_i) dy_i \right) dx_i \\ &= \frac{\epsilon}{2} \int_{l_{t_{\mathbb{A}}, d_i}}^{+\infty} f_{t_{\mathbb{A}}, d_i}(x_i) dx_i < \frac{\epsilon}{2} \end{aligned}$$

Finally we have $p < \tilde{p} = \tilde{p}_{1,1} + \tilde{p}_{1,2} + \tilde{p}_2 < \epsilon + \epsilon^2/4 \ll \lambda$, which means $(t_{\mathbb{A}}, t_{\mathbb{B}})$ can never be in the join result. Then we reach a contradiction and thus the assumption is wrong.

The proposition states a way to find a superset of the join answers: for each tuple $t_{\mathbb{A}}$ from \mathbb{A} , expand its possible range by δ on each dimension, denoted by $I_{t_{\mathbb{A}}}$, then pair $t_{\mathbb{A}}$ with all tuples $t_{\mathbb{B}}$ from \mathbb{B} whose possible ranges overlap with $I_{t_{\mathbb{A}}}$.

When \mathbb{A} is stored using *store-mean*, we use the above result to form a subarray condition on \mathbb{B} , for each cell $C_{\mathbb{A}} \in \mathbb{A}$. The next proposition shows how to do so, i.e., by relaxing the condition using the minimum lower bound and maximum upper bound of possible ranges of all tuples in $C_{\mathbb{A}}$.

Proposition 3 (Subarray for Store-mean) Consider $SJoin(\mathbb{A}^d, \mathbb{B}^d, \bigwedge_{i=1}^d |\mathbb{A}.d_i - \mathbb{B}.d_i| < \delta, \lambda)$ when \mathbb{A} is under *store-mean*. For a cell $C_{\mathbb{A}}$, a subarray condition $\theta_{C_{\mathbb{A}}}$ that returns all join results is:

$$\bigwedge_{i=1}^d \min_{t_{\mathbb{A}} \in C_{\mathbb{A}}} l_{t_{\mathbb{A}}, d_i} - \delta < \mathbb{B}.d_i < \max_{t_{\mathbb{A}} \in C_{\mathbb{A}}} u_{t_{\mathbb{A}}, d_i} + \delta.$$

When \mathbb{A} is stored using *store-multiple*, we do not need to relax the join condition as aggressively, e.g., to accommodate the large possible ranges of some tuples. Instead, we can bound the relaxation using the step size of \mathbb{A} and δ . Given the step size $\langle k_1, k_2, \dots, k_d \rangle$ of array \mathbb{A} , we define some notation:

- Denote the value range of cell $C_{\mathbb{A}}$ on dimension d_i as $(l_{C_{\mathbb{A}}, d_i}, u_{C_{\mathbb{A}}, d_i})$.
- For any cell $C_{\mathbb{A}} = \mathbb{A}[x_1, \dots, x_d]$, two cells bound the expansion from $C_{\mathbb{A}}$ by the step size of \mathbb{A} , denoted as $C_{\mathbb{A}}^- = \mathbb{A}[x_1 - k_1, \dots, x_d - k_d]$ and $C_{\mathbb{A}}^+ = \mathbb{A}[x_1 + k_1, \dots, x_d + k_d]$.

Proposition 4 (Subarray for Store-multiple) Consider $SJoin(\mathbb{A}^d, \mathbb{B}^d, \bigwedge_{i=1}^d |\mathbb{A}.d_i - \mathbb{B}.d_i| < \delta, \lambda)$ when \mathbb{A} is under

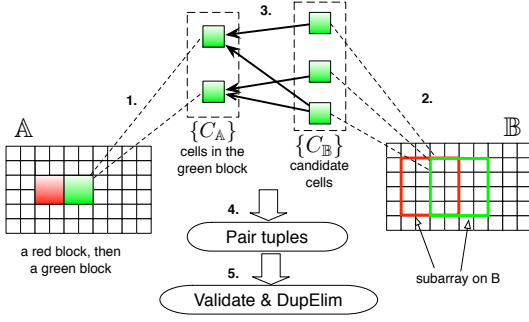


Fig. 7 Illustration of subarray-based joins

store-multiple. For cell C_A , a subarray condition θ_{C_A} that returns all join results is:

$$\bigwedge_{i=1}^d l_{C_A^-.d_i} - \delta < \mathbb{B}.d_i < u_{C_A^+.d_i} + \delta.$$

Proof Let S_{t_A} denote the set of cells that store a copy of t_A , i.e., $S_{t_A} = \{C_A | t_A \in C_A\}$. Below we first prove that $(l_{t_A.d_i}, u_{t_A.d_i}) \subseteq \bigcup_{C_A \in S_{t_A}} (l_{C_A^-.d_i}, u_{C_A^+.d_i})$: When t_A only needs one copy to cover its possible range on dimension d_i , assume the copy is stored at C_A , then $(l_{t_A.d_i}, u_{t_A.d_i}) \subseteq (l_{C_A^-.d_i}, u_{C_A^+.d_i})$ because otherwise it needs at least two copies; When t_A has more than one copies on dimension d_i , according to §3.2, the first copy and the last copy are stored k_i cells away from the lower and upper bounds of t_A 's possible range respectively, depicted by Fig. 3(b). So $l_{t_A.d_i} = \min_{C_A \in S_{t_A}} l_{C_A^-.d_i}$ and $u_{t_A.d_i} = \max_{C_A \in S_{t_A}} u_{C_A^+.d_i}$, which means $(l_{t_A.d_i}, u_{t_A.d_i}) = \bigcup_{C_A \in S_{t_A}} (l_{C_A^-.d_i}, u_{C_A^+.d_i})$. Combining the two cases, we have $(l_{t_A.d_i}, u_{t_A.d_i}) \subseteq \bigcup_{C_A \in S_{t_A}} (l_{C_A^-.d_i}, u_{C_A^+.d_i})$. Then for any tuple t_B , if its possible range $(l_{t_B.d_i}, u_{t_B.d_i})$ overlaps with $(l_{t_A.d_i} - \delta, u_{t_A.d_i} + \delta)$, which is a necessary condition for t_B being a true match of t_A according to Proposition 2, it must also overlap with $\bigcup_{C_A \in S_{t_A}} (l_{C_A^-.d_i} - \delta, u_{C_A^+.d_i} + \delta)$. This means that t_B will be returned by at least one of the subarray queries formed for all cells in S_{t_A} , say $Subarray(\mathbb{B}, \theta_{C_{A_0}}, \lambda)$. In this way, we guarantee that no result is missing.

This proposition states that for each cell C_A , the subarray condition on the inner array \mathbb{B} can be formed by expanding C_A first by the step size of \mathbb{A} and then by δ .

We present subarray-based join (SBJ) as Algorithm 2 and illustrate it with Fig. 7. It processes one block of the outer at a time (Line 1 in Algorithm 2; marked as Step 1 in Fig. 7, with a red block followed by a green block of \mathbb{A}).

Algorithm 2 Subarray-based Join (SBJ)

```

1: for each read block  $R_A$  in  $\mathbb{A}$  do
2:    $toRead.clear(); map.clear();$ 
3:   for each cell  $C_A$  in  $R_A$  do
4:      $loadToMemory(C_A);$ 
5:      $Q \leftarrow formQueryRegion(C_A);$ 
6:      $S \leftarrow Subarray(\mathbb{B}, Q);$ 
7:     for each cell  $C_B$  in  $S$  do
8:        $toRead.add(C_B);$ 
9:        $map.get(C_B).add(C_A);$ 
10:    for each cell  $C_B$  in  $toRead$  do
11:       $loadToMemory(C_B);$ 
12:      for each cell  $C_A$  in  $map.get(C_B)$  do
13:        for each tuple  $t_A$  in  $C_A$  do
14:          for each tuple  $t_B$  in  $C_B$  do
15:             $validate(t_A, t_B);$ 
16:   $removeDuplicates();$ 

```

For each cell C_A in the current block, the algorithm forms a *Subarray* query and runs it on the inner array \mathbb{B} (Line 6; Step 2). We call the \mathbb{B} cells returned by the *Subarray* query for each C_A the **candidate cells** of C_A . Since the candidate cells of different outer cells may overlap, as an optimization to save I/O, the algorithm maintains the union of the candidate cells of all outer cells in the current block (Line 8), in $\{C_B\}$ in Fig. 7. To avoid nonviable pairs of tuples, the algorithm maintains a hash map that maps a cell C_B to only those \mathbb{A} cells whose candidate cells include C_B , i.e., the mapping structure in Fig. 7 (Line 9; Step 3). Then the algorithm reads relevant cells of \mathbb{B} and pairs tuples accordingly (Line 10-14; Step 4). It is optional to apply quick filters to the paired tuples to reduce later CPU cost. It finally does validation using the join condition and removes duplicates (Line 15-16; Step 5).

As shown in Fig. 7, the memory is shared by (1) the read block of \mathbb{A} , (2) the cell-level mapping structure, which is much more compact than the tuple-level mapping in the index-based join, and (3) the cache of \mathbb{B} cells (or chunks). Since first two items grow together, they share a memory quota. Then the block size is determined by loading more cells from \mathbb{A} until the memory quota is used up. The rest of the memory is given to the cache of \mathbb{B} . Intuitively, the more memory is given to data structures (1) and (2), the fewer batches and fewer duplicate reads of the same \mathbb{B} cell due to the overlapping subarray regions on \mathbb{B} .

4.2.2 A Cost Model for Optimization

Next we build a cost model for SBJ under the *store-multiple* scheme which can be used to find the optimal step size during data loading given basic data statistics. We use the symbols in Table 1 with subscripts to distinguish inner and outer arrays.

I/O cost: We model the numbers of \mathbb{A} and \mathbb{B} chunks read in I/O and later translate them to seek and transfer times.

First consider the outer array \mathbb{A} . Its number of chunks, denoted by $\|\mathbb{A}\|$, is the total number of tuple copies, denoted by $|\mathbb{A}|$, divided by the number of tuple copies per chunk. Based on Eq. (2) in §3.3, we have:

$$|\mathbb{A}| = T_{\mathbb{A}} \prod_{i=1}^d \left(\left\lfloor \frac{pr_{\mathbb{A},i}/s_{\mathbb{A},i}}{2k_{\mathbb{A},i} + 1} \right\rfloor + 1 \right), \quad \|\mathbb{A}\| = |\mathbb{A}| / \lfloor c/b_{\mathbb{A}} \rfloor.$$

Now consider the inner array \mathbb{B} . Each cell in \mathbb{B} may be read multiple times as it can exist in the results of *Subarray* queries formed from different \mathbb{A} blocks. Hence, the I/O cost for reading \mathbb{B} is the product of (1) the number of \mathbb{A} blocks, $\alpha_{R_{\mathbb{A}}}$, (2) the number of \mathbb{B} cells to read per \mathbb{A} block, denoted by $\beta_{R_{\mathbb{A}}}$, and (3) the number of chunks per \mathbb{B} cell, $\|\mathbb{C}_{\mathbb{B}}\|$. Below we model each of them in order.

We first model $\alpha_{R_{\mathbb{A}}}$. Assume that a memory quota of K chunks is given to the \mathbb{A} block and its mapping. Then the number of cells in each \mathbb{A} block, $n_{R_{\mathbb{A}}}$, is $K / (\|\mathbb{C}_{\mathbb{A}}\| + \|\mathcal{M}_{C_{\mathbb{A}}}\|)$, where $\|\mathbb{C}_{\mathbb{A}}\|$ is the number of chunks per \mathbb{A} cell and $\|\mathcal{M}_{C_{\mathbb{A}}}\|$ is the number of chunks for the mapping entries per \mathbb{A} cell. It is easy to see that

$$\|\mathbb{C}_{\mathbb{A}}\| = \|\mathbb{A}\| / \prod_{i=1}^d n_{\mathbb{A},i}.$$

According to Proposition 4, the subarray condition formed for cell $C_{\mathbb{A}}$ expands $C_{\mathbb{A}}$ by \mathbb{A} 's step size and then by δ , so the length of the *Subarray* query on dimension d_i is $(1 + 2k_{\mathbb{A},i})s_{\mathbb{A},i} + 2\delta$. It amounts to $((1 + 2k_{\mathbb{A},i})s_{\mathbb{A},i} + 2\delta) / s_{\mathbb{B},i} + 1$ cells in the \mathbb{B} array according to Eq. (1). When running this query on \mathbb{B} , the number of candidate cells of $C_{\mathbb{A}}$ is:

$$\beta_{C_{\mathbb{A}}} = \prod_{i=1}^d \left(\frac{(1 + 2k_{\mathbb{A},i})s_{\mathbb{A},i} + 2\delta}{s_{\mathbb{B},i}} + 1 + 2k_{\mathbb{B},i} \right) \quad (6)$$

Assuming that each mapping entry has b_{map} bytes, we have:

$$\|\mathcal{M}_{C_{\mathbb{A}}}\| = \beta_{C_{\mathbb{A}}} \cdot \frac{b_{map}}{c}.$$

We then get the number of \mathbb{A} blocks as the total number of cells divided by the number of cells in each $R_{\mathbb{A}}$ block:

$$\alpha_{R_{\mathbb{A}}} = \frac{\prod_{i=1}^d n_{\mathbb{A},i}}{n_{R_{\mathbb{A}}}} = \frac{(\|\mathbb{C}_{\mathbb{A}}\| + \|\mathcal{M}_{C_{\mathbb{A}}}\|) \prod_{i=1}^d n_{\mathbb{A},i}}{K}$$

We next model the second factor, $\beta_{R_{\mathbb{A}}}$. For the current read block $R_{\mathbb{A}}$, we take the union of \mathbb{B} cells returned by the *Subarray* query formed for each \mathbb{A} cell. This union is equivalent to the set of \mathbb{B} cells returned by a single *Subarray* query formed for the entire read block $R_{\mathbb{A}}$. Hence, similar to Eq. (6), we can get $\beta_{R_{\mathbb{A}}}$ as follows:

$$\beta_{R_{\mathbb{A}}} = \prod_{i=1}^d \left(\frac{(n_{R_{\mathbb{A}}}^{\frac{1}{d}} + 2k_{\mathbb{A},i})s_{\mathbb{A},i} + 2\delta}{s_{\mathbb{B},i}} + 1 + 2k_{\mathbb{B},i} \right).$$

We can get the last factor $\|\mathbb{C}_{\mathbb{B}}\|$ in the same way as $\|\mathbb{C}_{\mathbb{A}}\|$.

CPU cost: The CPU cost is the product of the number of tuple pairs to be validated, which we will model below, and the validation cost per tuple pair. According to our algorithm, tuples in each cell $C_{\mathbb{A}}$ are paired with the tuples in $C_{\mathbb{A}}$'s candidate cells and all such tuple pairs need to be validated. Therefore, the number of tuple pairs is the product of (1) the number of tuple copies in \mathbb{A} , (2) the number of candidate cells per \mathbb{A} cell, and (3) the number of tuple copies per \mathbb{B} cell. Using Eq. (6), we compute the product as:

$$|\mathbb{A}| \cdot \prod_{i=1}^d \left(\frac{(1 + 2k_{\mathbb{A},i})s_{\mathbb{A},i} + 2\delta}{s_{\mathbb{B},i}} + 1 + 2k_{\mathbb{B},i} \right) \cdot \frac{|\mathbb{B}|}{\prod_{i=1}^d n_{\mathbb{B},i}}$$

Finally, the combined IO and CPU model allows us to find the optimal step sizes for inner and outer arrays if *SJoin* is the key workload. Statistics needed are the distribution of tuples' possible ranges and common distance values in joins. Collecting such statistics is a common task of the query optimizer and we can leverage a large body of work on relational DBMSs in our work.

5 Experiments

We evaluate our techniques for *Subarray* and *Structure-Join* using both a wide range of synthetic workloads with controlled properties and real datasets from the Sloan Digital Sky Survey (SDSS) [28].

5.1 Evaluation of Subarray

Experimental Setup on Synthetic Datasets. We generate synthetic workloads based on the statistics of dimension attributes (*rowc*, *colc*) in SDSS. The value of each dimension attribute in SDSS can be described by a Gaussian distribution, (μ, σ) . Here, μ is specified by the value of attribute *rowc* (or *colc*) and determines where the center of a tuple's possible range is located along dimension *rowc* (or *colc*); σ is specified by the value of attribute *rowcErr* (or *colcErr*) and determines how wide a tuple's possible range is along dimension *rowc* (or *colc*); without loss of generality, we consider a tuple's possible range per dimension to be $\mu \pm 3\sigma$. Putting the SDSS dataset into an array with each cell addressed by $(\lfloor rowc \rfloor, \lfloor colc \rfloor)$, i.e., each cell is a square of side 1, a tuple's possible range takes $2.5 \times 2.5 (=6.25)$ cells on average.

Our synthetic datasets are in both 2D and 3D spaces, each of which contains 2,000,000 tuples and is loaded into a 2D array with $250 \times 250 (=62500)$ cells or a 3D array with $40 \times 40 \times 40 (=64000)$ cells. The CPU cost per integration in the 3D space is much higher than that in the 2D space. We

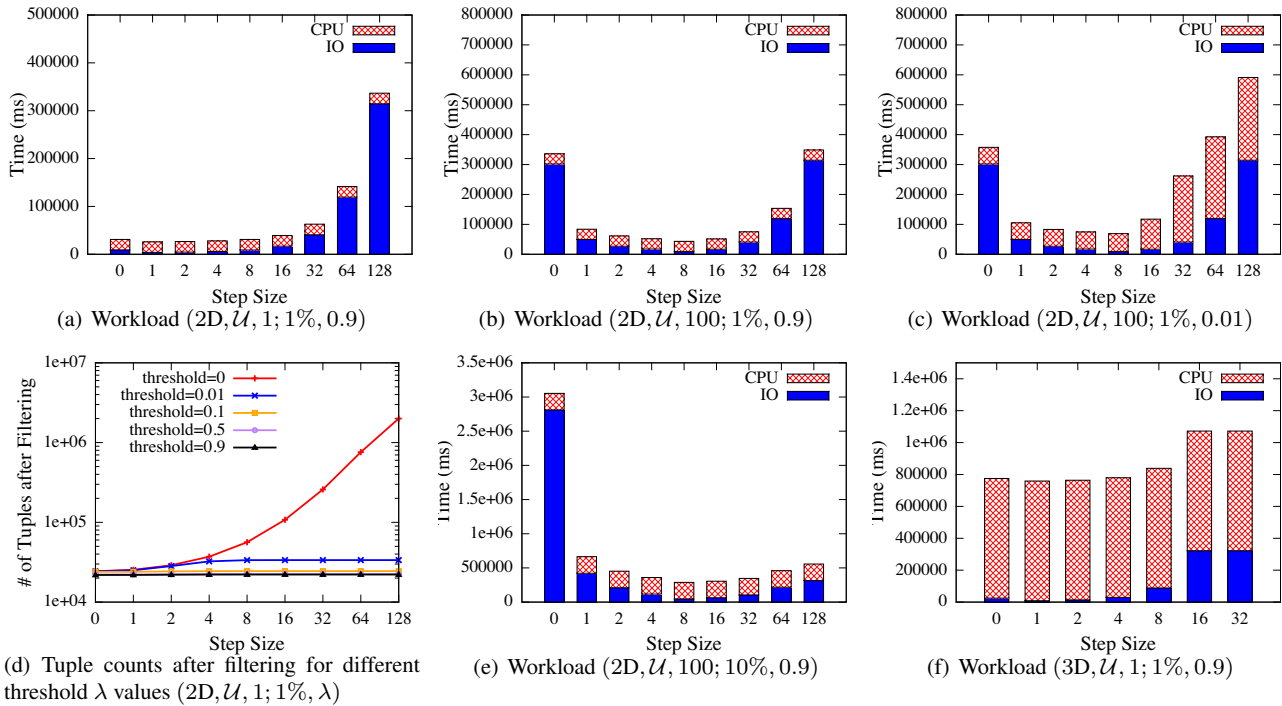


Fig. 8 Cost breakdown of *Subarray* with varied step sizes for various workloads.

generate μ values using the distribution, D_μ , which is set to a uniform distribution over the domain or a Gaussian distribution with more tuples clustered at the center; we collect the top 10 frequent σ values from the SDSS dataset and rescale the σ values by a factor S_σ to obtain various average possible range sizes: 6.25 (i.e., the average possible range size in SDSS), 100 and 625 cells (to simulate the scenarios when cells have finer granularities), where $S_\sigma = 1, 16$ and 100 correspondingly. We generate a dataset for each combination of D_μ and S_σ configurations. To generate queries, we vary the query size, q , between 0.01% and 10% of the domain. The threshold, λ , prunes tuples based on the existence probability in the query region. Usually the user wants only the tuples with high existence probabilities; we use $\lambda=0.9$ to represent this workload. We also tested $\lambda=0.01$ (e.g., needed if we have an aggregate after subarray). Table 2 shows a summary of all parameter settings. Recall that the evaluation of *Subarray* includes an I/O step and a CPU step. We optimize the CPU step by first running fast filters [20] with negligible costs before computing the expensive integral for the exact existence probability of each retrieved tuple.

Expt 1: Cost Breakdown. Our *store-multiple* scheme has a parameter, *step size* k , which determines both the degree of replication and query expansion. We start by showing how the *Subarray* processing cost changes as k varies. Fig. 8(a) shows results for the default workload, (2D, $D_\mu=\mathcal{U}$, $S_\sigma=1$; $q=1\%$, $\lambda=0.9$), whose average possible range size is $2.5 \times 2.5 = 6.25$ cells, while Fig. 8(b) shows results for $S_\sigma=100$,

	Parameter	Default Value	Other Values
Data	dimensionality	2	3
	D_μ , distribution of μ	uniform (\mathcal{U})	normal (\mathcal{N})
	S_σ , scale factor of σ	1	16, 100
Query	q , query range / domain	1%	0.01%, 0.1%, 10%
	λ , probability threshold	0.9	0.01

Table 2 Parameters in *Subarray* experiments using synthetic datasets.

with average possible range size being $25 \times 25 = 625$ cells and magnified trends. To further study the effect of λ and q , we tune their values for Fig. 8(b) one at a time: we change λ from 0.9 to 0.01 and show the cost breakdown in Fig. 8(c); we also vary q from 1% to 10% and show the cost breakdown in Fig. 8(e). Finally Fig. 8(f) shows the cost breakdown for a 3D workload. The overall trends are:

(1) *The I/O cost first decreases and then increases with the step size.* I/O is determined by both the number of cells in the expanded query region and the number of chunks per cell. When k is small, which means more aggressive replication of tuples, the expanded query region is small, but the number of chunks per cell is large and has a stronger impact on I/O. As k grows larger, fewer tuples are replicated, so each cell is smaller. But the expanded query region becomes very wide and affects I/O cost more. So the optimal I/O cost appears in the middle of the spectrum of k .

(2) *The CPU cost does not change with the step size when the probability threshold λ is high.* The CPU cost de-

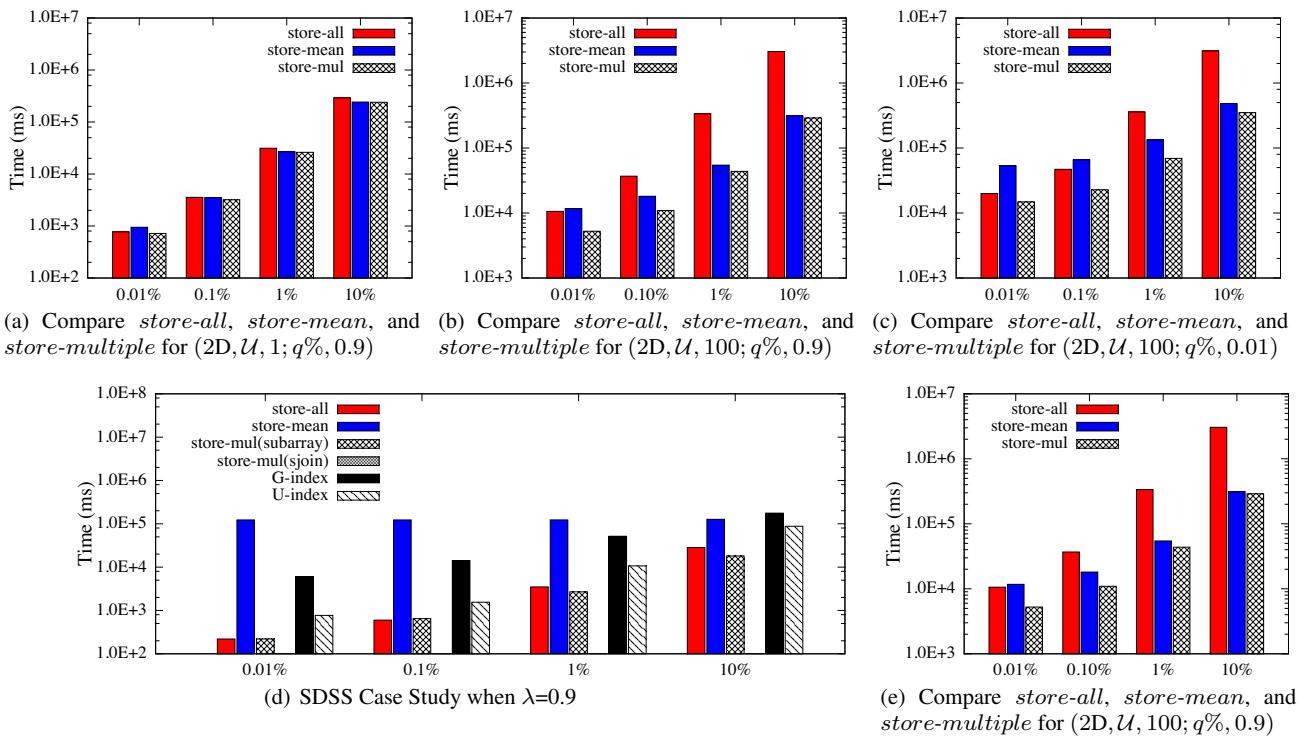


Fig. 9 Evaluation results of *Subarray* on synthetic and SDSS datasets.

depends on the number of tuples that passed the quick filter and need to be validated using expensive integration. Fig. 8(d) shows the number of tuples that pass the filter. When λ is sufficiently high, ≥ 0.1 in this figure, the filter can drop most irrelevant tuples, so the number of tuples after filtering does not change with the step size, or the number of tuples retrieved from storage. We examined the filter’s effect using multiple datasets and our observation is consistent.

It can be seen from the above cost breakdown plots that, between CPU cost and I/O cost, which one is dominating depends on many factors including λ (by comparing Fig. 8(b) and Fig. 8(c)), the system constants like the per integration cost (by comparing Fig. 8(a) and Fig. 8(f)), and the step size configuration (by comparing bars within each plot). To show the challenge of finding the optimal step size, we compare Fig. 8(a) and 8(b) and note that *the optimal step size shifts right when the average possible range increases*. We also examined other workloads and observed that *the optimal step size shifts left when λ is very small and increases with the query region size*.

Expt 2: Model Accuracy. We now use the cost model in §3.3 to determine the step size to be used when loading data into an array. We assume that the user can provide basic statistics including the σ distribution in the data and common *Subarray* sizes. We denote the optimal step size \mathbf{k}^* , and the step size returned by our model $\tilde{\mathbf{k}}$. We measure the per-

formance loss of our model, $\frac{Cost(\tilde{\mathbf{k}}) - Cost(\mathbf{k}^*)}{Cost(\mathbf{k}^*)}$. When tuples’ mean values, μ , are normally distributed around the center of the array, the center of the query region matters as the data density varies. For such datasets, we pick around 10 query centers (3×3 for 2D datasets and $2 \times 2 \times 2$ for 3D datasets) that evenly scattered over the array, and report on the average.

Table 2 shows 96 combinations of parameters. Our model returns the optimal step size (i.e., has no performance loss) in 43 out of 48 workloads when the tuples’ μ values are uniformly distributed and in 40 out of 48 workloads when the tuples’ μ values are normally distributed. In the cases when our model selects a suboptimal step size, the average performance loss is 2.72% which shows that our model is very effective in configuring the *store-mul* scheme.

Expt 3: Comparing Schemes. We now use the step size returned by the model to configure *store-multiple* and compare it to *store-all* and *store-mean* with fences for *Subarray* evaluation. The results are shown in a log scale in Fig. 9(a)-9(c) for different workloads. Each plot shows four queries with different query region sizes.

In all cases, *store-multiple* works the best. In comparison, when all tuples have small possible ranges, the three storage schemes do not differ much because *store-all* incurs only a small storage overhead and the expanded query region for *store-mean* is also very constrained, as shown for

the dataset in Fig. 9(a). However, for datasets when $S_\sigma = 100$, *store-all* often incurs tremendous storage overheads and I/O costs in querying, as shown in Fig. 9(e) and 9(c). Moreover, *store-multiple* outperforms *store-mean* considerably when the query region q is small, e.g., $q < 1\%$, which is the common case, due to a more constrained expanded query region. When q grows larger, e.g., $q=10\%$, their difference is reduced because the optimal step size of *store-multiple* tends to be larger. This means that infrequent replication of tuples works fine if q is large, and most tuples have only one copy under *store-multiple*, similar to *store-mean*.

5.1.1 A Case Study using SDSS datasets.

We finally evaluate *Subarray* using three SDSS [28] datasets on two dimension attributes *rowc* and *colc*, each containing 1.89, 10.3 and 30.2 million tuples.

5.2 Evaluation of Structure-Join

Experimental Setup. We generate synthetic workloads for *Structure-Join* as before with a few changes: For efficiency, we first consider 1D arrays of 1000 cells and datasets of 100,000 tuples. We will show results on two-dimensional arrays and larger datasets in the case study. We consider $SJoin(\mathbb{A}_1, \mathbb{A}_2, |\mathbb{A}_1.x - \mathbb{A}_2.x| < \delta, \lambda)$, where \mathbb{A}_1 and \mathbb{A}_2 are loaded from the same dataset, and we fix δ to 1% of the domain for proximity join. The memory size is about 176 chunks, i.e., 30% of the data size. We use a state-of-the-art index on continuous uncertain data [20] whenever possible, e.g., in the Index-based Join and as an in-memory filter in Subarray-based Join (detailed later). This index returns only true matches on 1D arrays, so validation is not needed.

Expt 4: Index-Based Join (IBJ). We first study the Index-based Join, which is sensitive to: (1) the selectivity, which we control using the probability threshold λ , (2) the storage scheme, and (3) the memory allocation scheme. We first use $\lambda=0.9$ or 0.01 to represent high or low selectivity. IBJ failed on all datasets when $\lambda=0.01$ regardless of memory allocation and storage schemes, because a single tuple can have so many matches that its mapping entries do not fit in memory. Hence, we consider only $\lambda=0.9$ below.

We next consider the effect of the storage scheme. As stated in §3.3, when array \mathbb{A} is under *store-multiple*, it can still be processed as *store-mean*. Our question is whether IBJ works better by considering multiple tuple copies or only the copy at mean (ignoring other copies). We observe that for the outer array, it is always better to use only the copy at mean because it avoids duplicate index lookups for the same outer tuple. For the inner array, Fig. 10(a) shows the performance of *store-multiple* with different k values and that of *store-mean* on the $S_\sigma = 10$ dataset, where each line denotes a

different memory allocation scheme (discussed shortly). As is seen, IBJ has the best performance when the inner array uses *store-mean* under all memory allocation schemes. The reason is that after the index lookups return all candidate tuples, fetching all of them from the inner array requires more I/O when nonviable tuples have a copy stored in the to-read cells (chunks), which is less likely to happen with a lower degree of replication. The observation also holds for other datasets. Hence, for better performance, IBJ should consider only the tuple copy at the mean in the underlying storage.

Next we examine memory allocation among: (i) the outer block and associated mappings; (ii) the cache of index nodes; and (iii) the cache of inner cells (or chunks). Like caching inner tuples in an index-based join in relational databases, the third factor is least effective, so we fix it to 1 chunk to focus on others. Fig. 5 shows the result with one line per memory allocation scheme in the form of cache size of factor (i, ii, iii). When the size of factor (i) is small (e.g., 50), a small inner step size (e.g., 1) may not be able to run because each candidate tuple is mapped to more cells and the mapping entries may not fit in memory. Overall, the more memory is given to factor (i), the more tuples per read block and thus IBJ is more likely to benefit from solving the set covering problem.

Expt 5: Subarray-Based Join (SBJ). The performance of SBJ is affected by the I/O cost for running repeated *Subarray* queries on the inner array, and the CPU cost for validating paired tuples. To reduce CPU cost, as we pair tuples from an outer cell and its candidate cell, we implement a filter [20] to prune nonviable pairs quickly. SBJ enjoys a memory allocation scheme of giving most memory to the outer block and its mapping, which is used below.

We first demonstrate that SBJ's performance is sensitive to the storage scheme. Fig. 10(b) shows various combinations of the outer step size, k_{out} , and inner step size, k_{in} , with $\lambda = 0.9$. Each line represents a fixed value of k_{out} , and the x -axis varies values of k_{in} , with the optimal inner step size circled. There are two main trends: (1) For a fixed k_{out} , the optimal inner step size k_{in}^* is in the middle of its spectrum. As explained in Expt 1, the inner I/O first decreases and then increases with its step size. (2) Once k_{in} is fixed, the optimal k_{out}^* also occurs in the middle (e.g., $k_{out}^*=16$), because it achieves the best tradeoff between (a) pairing and filtering costs for the same outer tuple, which decreases with larger k_{out} , and (b) the number of candidate cells to consider, which increases with k_{out} due to the enlarged expanded subarray region.

Next we show that the cost model in §4.2.2 can predicate the performance of SBJ so that given basic statistics, we can use it to choose the optimal step size configuration during data loading (if SJoin is known to the key workload). We again use the performance loss to evaluate the model accuracy. The results are shown in Table 3, where $\langle k_{out} \rangle; \langle k_{in} \rangle$

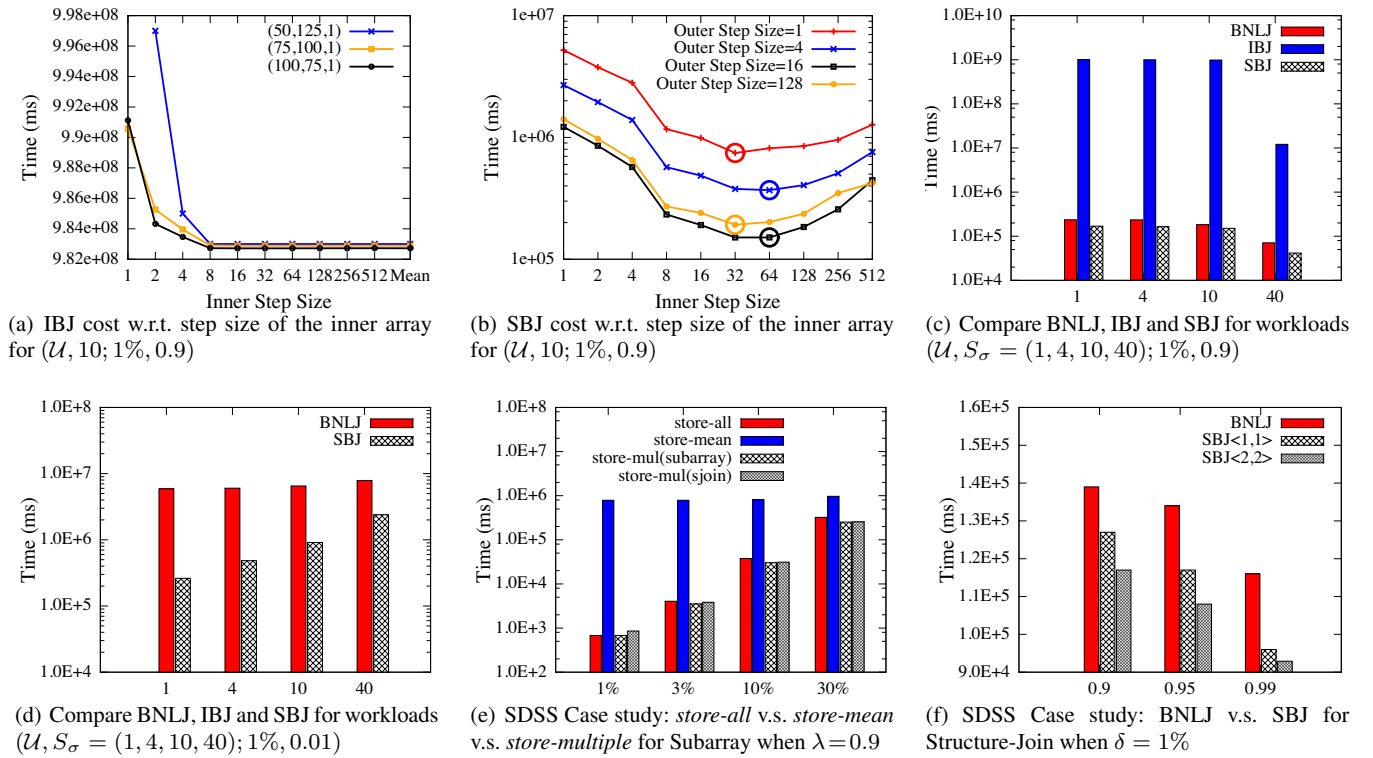


Fig. 10 Performance results of the *Structure-Join* evaluation and the Case Study using SDSS.

denotes the outer and inner step sizes. The model returns the optimal step sizes in most cases and the overall performance loss is within 6% (if any).

Expt 6: Comparison of Join Algorithms. We now use the step size returned by the model to configure SBJ, and compare it to IBJ and block nested loops joins (BNLJ) where both inner and outer arrays are stored using *store-mean*. The results are shown in Fig. 10(c) and Fig. 10(d) with one group of bars per dataset. As is mentioned in Expt 4, IBJ failed when $\lambda = 0.01$ due to the huge tuple-level mapping used. When $\lambda = 0.9$, IBJ still works poorly due to the tremendous index I/Os. For all datasets tested, SBJ outperforms BNLJ, e.g., 25% better when $\lambda=0.9$ in average and 85.8% better when $\lambda=0.01$. This is because SBJ does not incur much storage overhead and can effectively limit the number of inner cells to be loaded and the number of tuples to be filtered to save the CPU cost.

5.3 A Case Study using SDSS

We finally evaluate our techniques using the SDSS [28] dataset, which uses two dimension attributes *rowc* and *colc* and includes 1,893,685 tuples. Most tuples have small possible ranges; the average possible range size is 2.5 for both *rowc* and *colc*. The subarray model in §3.3 suggests $\langle 0, 0 \rangle$ as the step size configuration when $q = 1\%$ of the domain size per

D_μ	S_σ	λ	optimal step size	model step size	perf. loss
$\mathcal{U}(0, 1000)$	1	0.9	$\langle 2 \rangle; \langle 4 \rangle$	$\langle 4 \rangle; \langle 4 \rangle$	5.3%
		0.01	$\langle 2 \rangle; \langle 2 \rangle$	$\langle 2 \rangle; \langle 2 \rangle$	0%
	4	0.9	$\langle 8 \rangle; \langle 8 \rangle$	$\langle 8 \rangle; \langle 16 \rangle$	3.6%
		0.01	$\langle 8 \rangle; \langle 8 \rangle$	$\langle 8 \rangle; \langle 8 \rangle$	0%
	10	0.9	$\langle 16 \rangle; \langle 32 \rangle$	$\langle 16 \rangle; \langle 32 \rangle$	0%
		0.01	$\langle 8 \rangle; \langle 8 \rangle$	$\langle 16 \rangle; \langle 16 \rangle$	0%
	40	0.9	$\langle 32 \rangle; \langle 32 \rangle$	$\langle 32 \rangle; \langle 32 \rangle$	0%
		0.01	$\langle 32 \rangle; \langle 32 \rangle$	$\langle 32 \rangle; \langle 32 \rangle$	0%

Table 3 SBJ Model Accuracy when $\delta = 1\%$

dimension and $\langle 1, 1 \rangle$ for other query sizes, while the SBJ model in §4.2.2 suggests $\langle 2, 2 \rangle$ when $\delta = 1\%$ of the domain size is taken as a common workload for proximity join. When the two models do not return the same step size, the user should choose the right model by considering the importance (e.g., frequency and cost) of each type of query.

Regarding storage overheads, for the step size configuration $\langle 1, 1 \rangle$, 79.28% of tuples have only one copy and 92.36% of tuples have at most three copies, and the numbers for $\langle 2, 2 \rangle$ are 90.84% and 98.82%. As such, *store-multiple* incurs only a very modest storage overhead when most tuples have concentrated distributions.

The *Subarray* performance is shown in Fig. 10(e) which is similar to Fig. 9(a) except that the performance of *store-mean* with fences is orders of magnitude slower than other

schemes when $q=1\%$, 3% or 10% , and loses to *store-multiple* by 73% when $q=30\%$. This is because its expanded query region almost covers the entire array due to the existence of very large variance tuples (e.g., 2039.78^2), and such tuples are not in the top 10 frequent values we used to generate our synthetic datasets. Although our subarray and join models return different step sizes, when $q=10\%$ and 30% , they both improve *store-all* by over 17% and 21% , respectively.

The next query finds neighbors (distance within 1% of domain size) for 0.01% of the *Galaxy* tuples using *Structure-Join*. We compare SBJ only to the baseline BNLJ, since IBJ is shown to work poorly in §5.2. For multi-dimensional join, the filter [20] returns a superset of true matches. Thus, unlike the experiments in §5.2, the validation involves an expensive integral per tuple pair and contributes to the CPU cost. The results are shown in Fig. 10(f). When the input array is configured by the SBJ model, SBJ achieves 15.8% , 19.4% and 20.2% improvement over BNLJ when $\lambda=0.9$, 0.95 and 0.99 . When the array is configured by the subarray model, the performance gains are 8.6% , 12.7% and 17.2% , respectively. Overall, SBJ's performance increases fast with higher threshold λ .

6 Related Work

Most relevant techniques have been discussed in earlier sections. Below, we survey several broader areas.

Probabilistic processing under the array model. Recent work [15] observes that correlations in array data are mostly restricted to local areas and proposes a unified model for modeling both correlated data and physical storage. Monte Carlo processing has also been studied for join and sampling for uncertain array data [14]. As stated earlier, this line of work focuses on only value uncertainty in array data but not position uncertainty, i.e., it does not consider the fact that uncertain attributes can be used as dimension attributes.

Indexing and storing multi-dimensional uncertain data. Recent work has addressed indexing uncertain data. Some existing indexes in relational probabilistic databases [10, 9, 7, 8, 29, 13, 20] can be used in our index-based join or serve as filters to reduce validation (CPU) cost in the evaluation of subarray and subarray-based join. These indexes may not be effective when the filtering power is low (e.g., when there is an aggregate after the subarray or structure-join, the probability threshold needs to be a small number like 0.01 because a tuple with low existence probability may contribute a modest value to the aggregate) and can trigger many random I/O's since they are often secondary indexes. In contrast, we aim to provide native support in the array model, where logical and physical localities are better-aligned and the effect of exploiting physical locality is similar to using a clustered primary index on the tuples in a relational database, but without having to build the index.

Other indexes [5, 1, 21, 3, 2] are designed for similarity and nearest-neighbor queries, not directly applicable to our work. Probabilistic databases use secondary storage to record query lineage and efficiently compute tuple existence probabilities [19], but their focus is on discrete random variables in the relational model, not on continuous random variables in the array model.

Redundant storage for efficient query processing. Also related is the work on using redundant storage for answering point enclosure and range queries in an I/O efficient way [23, 16, 4]. To map to that work, we can translate our subarray query in two steps: First, find all tuples whose possible ranges (bounding boxes of tuples' distributions) intersect the query rectangle, which however cannot be simply solved by point enclosure and range queries [4]. Second, compute the existence probabilities of candidate tuples and validate them against a probability threshold, which is CPU-intensive and not considered in prior work (while our work does).

7 Conclusions

In this paper, we addressed uncertain data management in array databases, which may involve both *value uncertainty* and *position uncertainty*. To support array operations under position uncertainty, we proposed a number of storage and evaluation schemes for *Subarray*, in particular, the *store-multiple* scheme, and building on that, the index-based join and subarray-based join for *Structure-Join*. Evaluation results show that for *Subarray*, *store-multiple* outperforms other alternatives by using a cost model to configure the storage and bounding the overhead of querying. For *Structure-Join*, the subarray-based join outperforms the index-based join by configuring the storage for the workload and avoiding many overheads in processing. Our case study using SDSS shows that for real datasets, the storage overhead of the *store-multiple* scheme is small and our best techniques for *Subarray* and *Structure-Join* outperform the baselines often by a wide margin.

We plan to further extend our work by (1) considering statistics-driven storage and approximate query processing to further improve performance, and (2) integrating our techniques, which are fundamentally based on tuple-level replication, in big data systems that inherently maintain replicas for fault tolerance reasons.

References

1. Uncertain spatial data handling: Modeling, indexing and query. *Computers & Geosciences*, 33(1):42–61, 2007.
2. P. K. Agarwal, B. Aronov, S. Har-Peled, J. M. Phillips, K. Yi, and W. Zhang. Nearest-neighbor searching under uncertainty ii. In *PODS*, 2013.

3. P. K. Agarwal, A. Efrat, S. Sankararaman, and W. Zhang. Nearest-neighbor searching under uncertainty. In *PODS*, 2012.
4. L. Arge, V. Samoladas, and K. Yi. Optimal external memory planar point enclosure. *Algorithmica*, 54(3):337–352, May 2009.
5. C. Bohm, A. Pryakhin, and M. Schubert. The gauss-tree: Efficient object identification in databases of probabilistic feature vectors. In *Data Engineering, 2006. ICDE '06. Proceedings of the 22nd International Conference on*, pages 9–9, April 2006.
6. P. G. Brown. Overview of SciDB: large scale array storage, processing and analysis. In *SIGMOD Conference*, pages 963–968, 2010.
7. R. Cheng, D. V. Kalashnikov, and S. Prabhakar. Evaluating probabilistic queries over imprecise data. In *SIGMOD Conference*, pages 551–562, 2003.
8. R. Cheng, S. Singh, and S. Prabhakar. U-dbms: a database system for managing constantly-evolving data. In *VLDB '05: Proceedings of the 31st international conference on Very large data bases*, pages 1271–1274. VLDB Endowment, 2005.
9. R. Cheng, S. Singh, S. Prabhakar, R. Shah, J. S. Vitter, and Y. Xia. Efficient join processing over uncertain data. In *CIKM*, pages 738–747, 2006.
10. R. Cheng, Y. Xia, S. Prabhakar, R. Shah, and J. S. Vitter. Efficient indexing methods for probabilistic threshold queries over uncertain data. In *VLDB*, pages 876–887, 2004.
11. P. Cudré-Mauroux, H. Kimura, K.-T. Lim, J. Rogers, R. Simakov, E. Soroush, P. Velikhov, D. L. Wang, M. Balazinska, J. Becla, D. J. DeWitt, B. Heath, D. Maier, S. Madden, J. M. Patel, M. Stonebraker, and S. B. Zdonik. A demonstration of scidb: A science-oriented dbms. *PVLDB*, 2(2):1534–1537, 2009.
12. Y. Diao, B. Li, A. Liu, L. Peng, C. Sutton, T. Tran, and M. Zink. Capturing data uncertainty in high-volume stream processing. In *CIDR*, 2009.
13. T. Ge. Join queries on uncertain data: Semantics and efficient processing. In *ICDE*, 2011.
14. T. Ge, D. Grabiner, and S. B. Zdonik. Monte carlo query processing of uncertain multidimensional array data. In *ICDE*, pages 936–947, 2011.
15. T. Ge and S. B. Zdonik. A*-tree: A structure for storage and modeling of uncertain multidimensional arrays. *PVLDB*, 3(1):964–974, 2010.
16. J. M. Hellerstein, E. Koutsoupias, D. P. Miranker, C. H. Papadimitriou, and V. Samoladas. On a model of indexability and its bounds for range queries. *J. ACM*, 49(1):35–55, Jan. 2002.
17. H. Kimura, S. Madden, and S. B. Zdonik. Upi: A primary index for uncertain databases. *PVLDB*, 3(1):630–637, 2010.
18. J. F. Kurose, E. Lyons, D. McLaughlin, D. Pepyne, B. Philips, D. Westbrook, and M. Zink. An end-user-responsive sensor network architecture for hazardous weather detection, prediction and response. In *AINTEC*, pages 1–15, 2006.
19. D. Olteanu and J. Huang. Secondary-storage confidence computation for conjunctive queries with inequalities. In *Proceedings of the 2009 ACM SIGMOD International Conference on Management of Data*, SIGMOD '09, pages 389–402, New York, NY, USA, 2009. ACM.
20. L. Peng, Y. Diao, and A. Liu. Optimizing probabilistic query processing on continuous uncertain data. *PVLDB*, 4, 2011.
21. B. E. Rutenberg and A. K. Singh. Indexing the earth movers distance using normal distributions. In *Proceedings of the VLDB Endowment*, Vol. 5, No. 3, pages 205–216, 2012.
22. SciDB. Scidb array functional language 11.06. [http://trac.scidb.org/wiki/Docs/Release_11.06/ArrayFunctionalLanguage_11.06/](http://trac.scidb.org/wiki/Docs/Release_11.06/ArrayFunctionalLanguage_11.06)
23. T. K. Sellis, N. Roussopoulos, and C. Faloutsos. The r+-tree: A dynamic index for multi-dimensional objects. In *Proceedings of the 13th International Conference on Very Large Data Bases*, VLDB '87, pages 507–518, San Francisco, CA, USA, 1987. Morgan Kaufmann Publishers Inc.
24. M. Stonebraker, C. Bear, U. Çetintemel, M. Chemiack, T. Ge, N. Hachem, S. Harizopoulos, J. Lifter, J. Rogers, and S. B. Zdonik. One size fits all? part 2: Benchmarking studies. In *CIDR*, pages 173–184, 2007.
25. M. Stonebraker, J. Becla, D. J. DeWitt, K.-T. Lim, D. Maier, O. Ratzesberger, and S. B. Zdonik. Requirements for science data bases and scidb. In *CIDR*, 2009.
26. M. Stonebraker, P. Brown, A. Poliakov, and S. Raman. The architecture of scidb. In *SSDBM*, pages 1–16, 2011.
27. D. Suciuc, A. Connolly, and B. Howe. Embracing uncertainty in large-scale computational astrophysics. In *Proceedings of the 3rd International Workshop on Management of Uncertain Data (MUD)*, 2009.
28. A. S. Szalay, P. Z. Kunszt, A. Thakar, J. Gray, D. R. Slutz, and R. J. Brunner. Designing and mining multi-terabyte astronomy archives: The sloan digital sky survey. In *SIGMOD Conference*, pages 451–462, 2000.
29. Y. Tao, X. Xiao, and R. Cheng. Range search on multidimensional uncertain data. *ACM Trans. Database Syst.*, 32(3), Aug. 2007.
30. T. T. L. Tran, L. Peng, Y. Diao, A. McGregor, and A. Liu. Claro: modeling and processing uncertain data streams. *VLDB J.*, 21(5):651–676, 2012.
31. T. T. L. Tran, L. Peng, B. Li, Y. Diao, and A. Liu. Pods: a new model and processing algorithms for uncertain data streams. In *SIGMOD Conference*, pages 159–170, 2010.