



TreeCat: Standalone Catalog Engine for Large Data Systems

Keonwoo Oh

University of Maryland
College Park, Maryland, USA
koh3@umd.edu

Pooja Nilangekar

University of Maryland
College Park, Maryland, USA
poojan@umd.edu

Amol Deshpande

University of Maryland
College Park, Maryland, USA
amol@umd.edu

ABSTRACT

With ever-increasing volume and heterogeneity of data, advent of new specialized compute engines, and demand for complex use cases, large-scale data systems require a performant catalog system that can satisfy diverse needs. We argue that existing solutions, including recent lakehouse storage formats, have fundamental limitations and that there is a strong motivation for a specialized database engine, dedicated to serve as the catalog. We present the design and implementation of *TreeCat*, a database engine that features a hierarchical data model with a path-based query language, a storage format optimized for efficient range queries and versioning, and a correlated scan operation that enables fast query execution. A key performance challenge is supporting concurrent read and write operations from many different clients while providing strict consistency guarantees. To this end, we present a novel MVOCC (multi-versioned optimistic concurrency control) protocol that guarantees serializable isolation. We conduct a comprehensive experimental evaluation comparing our concurrency control scheme with prior techniques, and evaluating our overall system against Hive Metastore, Delta Lake, and Iceberg.

PVLDB Reference Format:

Keonwoo Oh, Pooja Nilangekar, and Amol Deshpande. TreeCat: Standalone Catalog Engine for Large Data Systems. PVLDB, 18(11): 4323 - 4336, 2025. doi:10.14778/3749646.3749696

PVLDB Artifact Availability:

The source code, data, and other artifacts have been made available at <https://github.com/umddb/treecat>.

1 INTRODUCTION

The catalog is an important component of a database management system that is responsible for storing and maintaining metadata such as the properties of logical entities like databases, tables, views, triggers, and schemas; mapping of these entities to the underlying physical storage; and access control information. Metadata operations are in the hot path of most, if not all, database operations. Take, for example, the query execution process of a relational database. The table schema is necessary to first semantically validate the query. Once the query is validated, the query planner uses various metadata, including table schema, storage information, and statistics, to optimize and build a physical query plan that is executed by the execution engine (which usually needs access to the metadata

as well). Without a high-performance catalog that provides fast access to metadata, a database management system cannot operate efficiently.

In a monolithic data system, the catalog is typically tightly coupled with the rest of the system, which works for one or a few data systems. But in a typical enterprise setting, organizations often use many specialized engines, each serving a different function, since no single one-size-fits-all system can meet the requirements of all use cases. Meanwhile, there has been a growing trend toward disaggregation of large-scale data systems into smaller subsystems. This includes the separation of compute and storage layers, the increasing adoption of common table formats (parquet, ORC), and the use of “connectors” to link different systems. These efforts aim to achieve a composable data stack, enabling organizations to store data in a shared storage layer and choose from multiple compute engines. However, as long as each system maintains its own catalog such that the metadata about the “managed” data remain separate, different systems cannot easily share data even if it is stored in the same storage system. *For data management systems to be truly composable, the catalog functionality has to be disaggregated from the rest of the system.* This issue is now widely recognized in the data systems industry, and there are multiple active projects, including lakehouse storage formats [25, 27, 30], and, most recently, catalog services [14, 15, 59, 66], mainly driven by the industry.

We further argue that not only should the catalog functionality exist as a standalone service, but also that it should be unified and not span multiple (semi-)autonomous systems or storage modules. As we show, unified architecture not only leads to significantly reduced latencies, but also reduces duplication of metadata across different data engines, and enables multi-table transactions as well as other complex operations on metadata without sacrificing consistency guarantees. A unified design also makes it possible to build a specialized engine, as we develop in this work, that can meet the specific requirements of a catalog in terms of both functionality and performance.

To this end, the proposed catalog itself can be seen as an independent database (sub-)system specialized for handling metadata operations. We characterize the primary use cases and present our prototype catalog engine, TreeCat, with the following design choices. (1) Given the hierarchical nature of metadata logical modeling and query access patterns, we adopt a *hierarchical data model* along with a *path navigation query language*. (2) For efficient query execution, we employ a storage layout inspired by file systems, implemented using *write-optimized key-value* stores, and leverage batch correlated scan operations. (3) We use versioned storage to support MVCC (multi-version concurrency control) for improved read performance, as well as version control operations such as time travel queries, clone, and snapshot. (4) We present a novel MVOCC mechanism that combines scan range locking and *precision*

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 18, No. 11 ISSN 2150-8097.
doi:10.14778/3749646.3749696

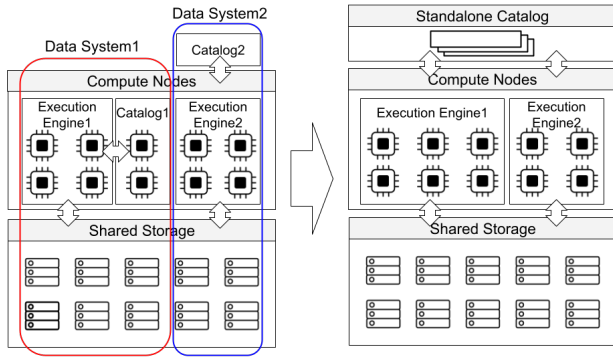


Figure 1: Disaggregating the catalog functionality for higher composability

locking techniques. This approach efficiently manages predicate dependencies while minimizing validation costs. Additionally, we adopt a *commit-time update* technique to avoid aborts caused by conflicts on frequently updated fields, such as statistics. Although we focus on catalogs in this paper, some of our techniques are more generally applicable in any hierarchical database management system. TreeCat is implemented in about 12000 lines of C++ code, and exposes a gPRC interface to serve remote client requests. We present a comprehensive set of experiments illustrating the significant performance benefits of our new MVOCC protocol. We also show that TreeCat outperforms other state-of-the-art systems, including Delta Lake, HMS, and Iceberg.

2 BACKGROUND AND PRIOR WORK

In this section, we categorize and classify existing catalog systems and highlight their shortcomings, which motivate the key requirements and features of TreeCat. The discussion of prior work related to more specific design and implementation details is interspersed in other relevant parts of the paper for better context.

We first distinguish between data catalogs and operational catalogs. **Data catalogs** serve as an organization-level database of metadata about data assets across various systems. They can be viewed as an “information retrieval system” for data assets. The main goal is to make the data assets of an organization more manageable and accessible, providing functions such as keyword search and retrieval of important contextual information. Data catalogs are typically implemented as federated systems that are updated asynchronously through data discovery processes or metadata ingestion pipelines [17, 44, 61, 73]. Because they do not directly serve database engines, they generally lack strict consistency or performance guarantees. There are numerous commercial data catalog software products, including those by Alation, Collibra, Informatica, AWS (Glue), Google (Dataplex), and Microsoft (Azure Data Catalog), as well as open source projects, such as Apache Atlas [23], Metacat [53], and DataHub [16]. **Operational catalogs**, on the other hand, store metadata that are directly accessed by database engines during updates, query planning, and query execution. Although operational catalogs can be a part of, or serve as the main data source for, data catalogs, their primary functions are clearly different. In this paper, we focus exclusively on operational catalogs.

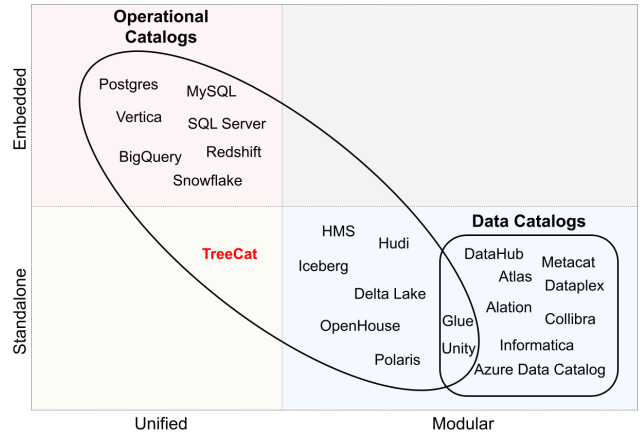


Figure 2: Classification of different catalog systems.

Catalogs can be either embedded or standalone. An **embedded catalog** is integrated within a single data system, and external access to its metadata is restricted to the interfaces provided by the parent system. A **standalone catalog**, on the other hand, exists as an independent service that is accessible by any compatible data system via a standard interface. Catalogs can be further differentiated based on their architectural designs. A catalog with a **modular architecture** is composed of components spread across multiple autonomous services and systems. On the other hand, a catalog with a **unified architecture** consists of tightly integrated components forming a single system. It is important to note that a unified system is not necessarily centralized; it can have a distributed architecture for scalability.

Traditional RDBMSs, including both transactional and analytical systems, have embedded catalogs with unified architectures. In transactional RDBMSs, catalog metadata is stored as a predefined set of relations [34, 48, 58]. Apart from a specialized in-memory cache for low-latency reads, all metadata use the same underlying storage representation and concurrency control mechanism as the data. Analytical RDBMSs and data warehouses have different implementations designed to handle increasingly decentralized scenarios. The catalog of a typical data warehouse functions as a specialized subsystem that decouples metadata from data, since the primary data processing infrastructure is not optimized for transactional workloads. For example, Snowflake, a cloud-native data warehouse, stores its metadata in FoundationDB [78], a distributed transactional key-value store. Many of its features, including automatic data clustering, data pruning during query execution, access control, and zero-copy clone, rely on catalog objects persisted in FoundationDB [47, 67]. Regardless of system performance, tight integration with the rest of the system makes it *difficult to extend the catalog functionality beyond a single system* [6, 47, 74].

Large-scale data processing engines [5, 62, 71] from the big data community use standalone catalogs with modular architectures. These systems originally emerged with a highly disaggregated architecture consisting of an independent compute engine, a distributed storage system, and a metadata service, each designed to be highly scalable. As part of the Hive ecosystem, HMS (Hive Metastore) [71]

was commonly used as the standalone catalog system. Backed by a transactional RDBMS (typically MySQL or PostgreSQL), HMS stores information about databases, tables, access control information, statistics, and partitions (groups of file objects organized by common attribute values). However, HMS does not store file-level metadata in the back-end RDBMS. Instead, it uses the shared storage system (distributed file system or object store) as an index on the files by grouping them in the same partition under the same directory (hence, it is modular). To scan a table, HMS must retrieve the paths of relevant partitions from the back-end RDBMS and call *list()* on each partition directory to obtain the list of files to scan. In an environment where a single dataset contains over a million files, the high latency of *list()* operation of cloud object stores like S3 becomes a serious performance bottleneck [3]. In addition, even retrieving the list of a large number of partitions from HMS is slow due to an inefficient Thrift server implementation and suboptimal back-end database schema [42].

As the limitations of HMS became clear, lakehouse storage formats, namely Delta Lake, Hudi, and Iceberg [3, 25, 27] emerged as alternatives. The specific implementations differ, but the common idea is to store metadata about file objects as separate **metadata files** along with the data files in shared storage. By directly fetching and processing these metadata files, the performance bottlenecks of HMS can be avoided. Furthermore, additional file metadata, such as statistics, can be stored in metadata files and utilized for optimizations like file pruning, further improving system performance. Delta Lake uses a data structure called Delta log, which consists of a sequence of delta JSON files (with incrementing numerical IDs) and periodic checkpoint files [3]. Iceberg organizes metadata into a hierarchy of files: a metadata file contains a list of manifest list files; a manifest list file contains a list of manifest files; and a manifest file contains a list of file objects [29]. Hudi identifies files by the timestamps at which they are first added and stores metadata in HFile format, enabling fast file pruning and listing [24].

Storing the metadata alongside the data has benefits such as a simplified architecture, good scalability due to the shared storage system, and speedup from parallel metadata processing. However, the notion that the adoption of the lakehouse storage format results in a simpler architecture than HMS is only partially true. A separate system must still be deployed to manage the metadata about high-level catalog objects, such as databases and tables; lakehouse storage formats simply offload most of the catalog functions to the shared storage system. More recently, several efforts have focused on providing these services through Hive-compatible APIs or REST APIs [14, 15, 59, 66]. *Due to the modular architecture, it is still quite difficult to maintain consistency of the metadata.* As discussed in Section 5.4, no lakehouse storage format currently supports multi-table transactions. Lastly, every metadata operation involves reading and writing metadata files in the shared storage system, which has non-trivial overhead. The high latency of these operations can limit the write throughput under high contention [38].

3 HIGH-LEVEL DESIGN

In this section, we present the high-level data model and the basic API of TreeCat. We also characterize the primary use cases with concrete examples.

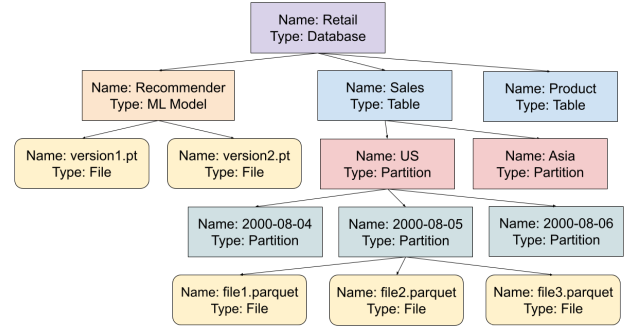


Figure 3: An example retail database, which contains heterogeneous assets, including relational data and ML models.

3.1 Data Model

A catalog needs to manage metadata about numerous logical objects and their relationships, which are almost always organized into nested logical groups, naturally forming a multi-level hierarchy. It also needs to maintain arbitrary statistics and clustering information (for cost-based optimizations) and must incorporate explicit versioning semantics to support operations such as time-travel and cloning. Further, it should support a sufficiently powerful query language and/or API so that the clients don't end up replicating the functionality, while being flexible and extensible enough to support diverse data sets and database engines. This naturally leads us to adopt a *hierarchical data model* for TreeCat. We use an example Retail database shown in Figure 3 to illustrate this point.

Example 1: The retail database primarily contains relational data, but also includes other types of assets, such as ML models. Furthermore, different types of assets have different nested substructures. While an ML model consists of different versions, each identified by a tag, data in a relational table are partitioned by a set of attribute values. Leaf objects correspond to physical data objects in the storage system.

We further note that: (1) Each group has common contextual information that is important for handling the constituent entities. (2) A predicate is often evaluated against such information, allowing entire groups of irrelevant objects to be pruned out during the data retrieval process. (3) There is data locality where entities in the same group are likely to be accessed together. These points are illustrated by the example use case that we present later. We now proceed with formal definitions of TreeCat's data model.

- (1) Every node in the hierarchy corresponds to an object of a type with a set of properties.
- (2) There are 2 types of objects:
 - (a) A non-leaf object has a parent and a set of children objects. It is uniquely identified by its path, $/oid_1/oid_2/.../oid_n$, where oid_n is its object id and the object with the path of $/oid_1/oid_2/.../oid_{n-1}$ is its parent.
 - (b) A leaf object has one or more parents, but **does not** have any child object. A leaf object with multiple parents has multiple paths. Given one of its paths, $/oid_1/oid_2/.../oid_n$, oid_n is one of its object ids and the object with the path of $/oid_1/oid_2/.../oid_{n-1}$ is one of its parents. Unlike non-leaf

objects, a leaf object is **immutable** and can only be added or removed.

- (3) Every change (object addition, removal, or update) is associated with a *vid* (version id) that determines its globally consistent chronological order.

We highlight a few important design decisions that deviate from standard practices. (1) **Immutable leaf objects**: Leaf objects, which usually correspond to physical data objects (e.g., Parquet files), are treated as immutable to enable seamless versioning and sharing — updates can be handled as delete-plus-insert if needed. (2) **Shared parents**: Allowing leaf objects to have multiple parents lets datasets be shared, cloned or branched, without duplicating physical files or metadata. (3) **Global chronological versions**: A single, consistent timeline underlies our concurrency control and powers version-control features such as time-travel queries, snapshots, and clones.

Note that, while Figure 3 shows one possible instantiation of metadata schema, the engine imposes no fixed metadata model—users choose whatever schema best fits their use case.

3.2 Query Language

TreeCat’s query language is based on top-down path navigation, where a query takes the form of a path expression, a sequence of n predicates, each of which is evaluated against the children of context objects that satisfy the preceding path subexpression. A predicate is either a wildcard character or an expression (enclosed with []) that may be evaluated against the object id or object properties. We extend the example in Figure 3 to demonstrate how client systems can utilize TreeCat. Consider how a relational database system would execute the following simple query.

```
SELECT count(*)
FROM retail.sales
WHERE region = 'US' AND date > '2025-01-01'
AND price > 5;
```

After parsing the query, the system has to semantically validate it. First, it retrieves the metadata of the Sales table from TreeCat, using the following metadata query.

```
/[obj_id='retail']/[obj_id='sales'
and obj_type='table']
```

A nonempty query result also validates that the table with the specified database name and table name indeed exists. After the table metadata is retrieved, the system uses it to further check whether column names and types are valid, etc. During the query planning phase, the paths and statistics of the data files that must be scanned during the query execution have to be retrieved. The following query is submitted to retrieve the file-level metadata that satisfy the appropriate predicate filters.

```
/[obj_id='retail']/[obj_id='sales' and
obj_type='table']/[part_val='US']
/[part_val > '2025-01-01']/[stats.price.min > 5]
```

The physical plan can finally be executed, once it is constructed.

We briefly discuss related work, including the APIs of existing standalone catalogs and query languages of similar, more general-purpose database systems. Standalone catalogs, such as HMS and lakehouse storage formats, provide basic RPC routines, such as

getTable() and *getPartitions()*, via Thrift or the REST API for retrieving high-level catalog objects. Retrieving file-level metadata of a given table involves iterating the list of file paths via an iterator interface with an optional predicate filter, rather than a unified query language like TreeCat. Meanwhile, the query languages of general graph databases, including Cypher [50], Gremlin [28], and GSQL [72], and those of XML databases, including XPath [9] and XQuery [76], are more powerful than TreeCat’s query language and can express a wider range of path queries. However, advanced semantics makes efficient implementation of a strong transaction isolation level very difficult. To our knowledge, there are currently no efficient graph database or XML database that fully support serializable isolation level. In our design, we sacrifice the query expressivity, so we can guarantee serializability using techniques that are discussed in Section 5.

3.3 System API

Basic API: Next, we list a basic API that TreeCat supports.

- **startTransaction(mode)**: Starts a new transaction in the specified mode. If the mode is read-only, it returns the latest *vid* at which a snapshot of the catalog can be read. If the mode is read-write, it creates a new transaction session and returns *transaction_id* and *read_vid*.
- **executeQuery(query, optional transaction_id, optional vid)**: Executes a read query on the catalog using the query language defined above.
- **commit(transaction_id, write_set)**: Validates and commits the transaction identified by *transaction_id*, applying modifications specified by the given *write_set*.
- **snapshot(snapshot_name, vid)**: Creates an immutable snapshot of the catalog data at *vid* and names it *snapshot_name*.
- **clone(src_path, dest_path, optional vid)**: Clones/copies the object at *src_path* and all of its descendants to the *dest_path*. Note that this is a metadata-level operation, and does not involve cloning the physical objects.

Updates: TreeCat can be updated via *commit()* function, which takes the modifications in the form of a write set. A write set is a collection of (path, value, type) tuples, where path specifies the object to modify, value contains the updated value, and type specifies the operation type, as listed below:

- (1) **Add**: a new object of the given path and initialize it with the given value. The operation has a precondition that the parent object exists, but an object with the same path does not.
- (2) **Update**: the object with the given path with the given value. A new object is created if the path does not exist.
- (3) **Remove**: the object with the given path and all its descendants.
- (4) **Merge**: Also known as the commit-time update operation [35], merge operation turns a read-modify-write operation into a single blind write operation that is performed at commit time. The given object value is a delta, rather than the new object value, and is applied to the existing object value at commit time. The operation has a precondition that the object already exists.

We describe the merge operation in more detail and highlight its importance. Merge operation is useful when objects are frequently

updated by commutative read-modify-write operations (e.g., counters), which may result in many conflicts. An important use case in TreeCat is applying changes in statistics that propagate up the hierarchy, which can result in contention on high-level statistics objects. We present an example of an update to table statistics that follows data insert:

Original Value: `{...size:1487..min:3...}`

Delta: `{...size:{op:+,val:124}..min:{op:min,val:0}...}`

Final Value: `{...size:1611..min:0...}`

Currently, the types of delta operations TreeCat supports include addition, subtraction, min, and max for a limited number of types.

4 SYSTEM ARCHITECTURE

We first outline the high level architecture of TreeCat, then elaborate on the core system components. We discuss the concurrency control mechanism in Section 5 in greater detail.

4.1 High Level Architecture

TreeCat is designed as a single-server database engine that handles incoming requests from remote client systems via the gRPC interface. Key components include:

- (1) **Backend Server:** Provides the API described in the previous section and schedules client tasks to the main thread pool.
- (2) **Executor:** Contains utilities for creating the execution plan of a read query and executing it.
- (3) **Transaction Manager:** Manages the overall concurrency mechanism. It keeps track of ongoing transaction states and launches a validation and commit process of the transaction upon receiving a commit request.
- (4) **Storage:** All catalog data is stored in the storage layer, which provides APIs to get an object by path and traverse all children of an object. The underlying storage engine is RocksDB [57].

This architecture allows a single TreeCat instance (possibly replicated, see Section 7) to handle the metadata for several different logical databases, and concurrently serve many execution engines while ensuring metadata consistency. Multiple TreeCat instances could be spun to serve different logical databases for administrative autonomy; however, any logical database (i.e., a collection of logically related tables or datasets) should ideally be served by a single TreeCat instance so that cross-table consistency can be enforced.

4.2 Backend Server

We use gRPC [33], a Remote Procedure Call (RPC) framework developed by Google, to provide the APIs listed in the previous section. gRPC is language-agnostic and offers portability across different languages. While REST API is another option with good portability, gRPC provides better performance due to its use of pre-compiled serialization and deserialization routines and support for server-side streaming, especially for large data loads (it is not uncommon for reads in this context to return MBs of data).

4.3 Storage Engine

When selecting the underlying storage engine, we considered several key requirements, including (1) optimizations for efficient traversal of hierarchical data, (2) efficient implementation of explicit

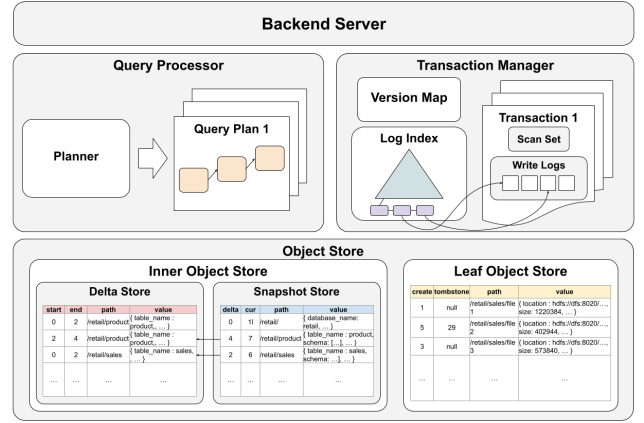


Figure 4: High-level architecture of TreeCat

versioning semantics, (3) fine-grained concurrency control with strong guarantees that can efficiently handle frequent updates to hierarchical data, and (4) embedded storage for low latency.

Numerous options exist for full-fledged database systems, ranging from RDBMS with XML [56, 63, 68] and graph [70] extensions, to graph databases [2, 18, 20, 31, 37, 40, 49, 51, 79], and document stores [2, 13, 46]. However, meeting all the requirements, especially (2) explicit versioning and (3) fine-grained concurrency control, proves to be difficult, as these systems are designed and implemented for more general use cases. As discussed before, lack of efficient support for serializable isolation level is especially problematic for graph databases and document stores.

After careful consideration, we built TreeCat on top of RocksDB, a write-optimized key-value store. Inspired by file systems built on top of write-optimized key-value stores [21, 39, 64], we adopt a storage layout where objects are sorted by the full object path, first by the path depth, then lexicographically. This layout allows fast listing of sibling objects by performing a range scan on the common prefix. It has advantages over storing the parent-child relation as a separate set of edges, which incurs the overhead of joining the objects (vertices) with edges during traversals. Although each range scan incurs a seek cost, the expense is amortized since objects typically have multiple children; as we show later, query cost is bounded by the query’s selectivity because most irrelevant objects are pruned early. This layout is also more suited for versioning and transactional updates compared to the scheme where an object’s children are stored together as a single list. With the latter, installing a new object version forces installation of an updated version of the parent’s children list, potentially cascading changes to the root and causing read-write conflicts that severely limit concurrency under serializable isolation. *This also imposes a fundamental limitation of Iceberg for supporting multi-table transactions.*

Object values are stored in BSON format [12], a binary JSON format that originated from MongoDB, is queryable, and can easily be converted to JSON format. While other efficient JSON-compatible binary formats exist (e.g., PostgreSQL’s JSONB), we chose BSON for its robust library support and rich feature set, which enables faster system development. We use two separate storage engines,

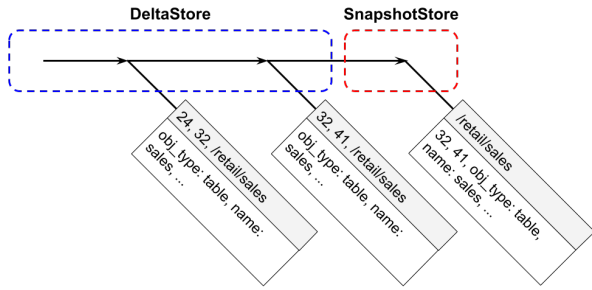


Figure 5: Example storage layout of the Sales table from Figure 3 in InnerObjectStore.

InnerObjectStore and **LeafObjectStore**, which store non-leaf objects and leaf objects respectively. **LeafObjectStore** stores one object value for each path and the *vid* range during which the object is visible. Unlike leaf objects, multiple past versions of non-leaf objects are persisted. For non-leaf objects, we adopt the “time-travel” storage scheme [77] where the most recent version is stored in the main **SnapshotStore** and the chain of past versions are stored in a separate **DeltaStore**, as illustrated in Figure 5. **SnapshotStore** has the following mapping:

path \rightarrow delta_vid, cur_vid, object value

where *cur_vid* is the *vid* from which the current snapshot is valid, and *delta_vid* is the *start_vid* of the previous object value. **DeltaStore** has the following mapping:

start_vid, end_vid, path \rightarrow object value

where the object value is visible between *start_vid* and *end_vid*. In **DeltaStore**, the objects are sorted by the object path, then by *start_vid* in descending order, so past object values can be traversed through a range scan. Retrieving the object value of a non-leaf object involves first checking the **SnapshotStore** for the current value and traversing the past values in **DeltaStore**, if necessary. Retrieving past versions incurs the cost of version traversal, but we expect most operations will access the most recent version. **LeafObjectStore** has the following mapping:

path \rightarrow create_vid, tombstone_vid, object value/primary path

where the object is visible between *create_vid* and *tombstone_vid*. Because leaf objects can have multiple parents, there could be a layer of indirection where the path could map to the primary path, which, in turn, maps to the object value.

4.4 Query Execution

A read query is executed as a sequence of physical operators analogous to the correlated join operator from RDBMSs (also called CROSS APPLY or LATERAL JOIN), which passes each value from the outer query to a user-defined function or a subquery and evaluates it. Unlike other standard join operators, the correlated join operator allows for early pruning of objects in the inner query by only scanning objects that join with objects that satisfy the outer query. We use this mechanism to avoid scanning large number of leaf objects whose ancestors do not satisfy the preceding predicates in the path expression. For example, while executing `/[obj_id =`

`‘retail’]/[obj_id = ‘sales’]/**/*`, file objects that belong to the Customer table need not be scanned.

Given a path expression of depth n , the execution plan is constructed as a single chain of n execution nodes, each with the corresponding predicate, a fixed-sized output buffer, a pointer to its child node, and handles to the object store. The query plan follows a standard batch iterator model where each node invokes `next()` on its child node, retrieving a batch of context object paths. For each context object path, the executor scans and evaluates the predicate against every child object visible in the *read_vid*. If the object satisfies the predicate, its path (or its value if the last execution node) is appended to the output buffer. The execution of a simple example query is illustrated in Figure 6. The planner currently applies one optimization rule where the lower and/or upper bounds of the object store iterator are set if the predicate is only satisfied by a range of object ids (e.g. `[1 < obj_id < 100]`).

Assuming that the underlying data takes the form of a balanced tree T with a constant fan-out f and height h , we can compute the cost of a path query q of depth d where $selectivity(q) = (s_1, s_2, \dots, s_d)$ specifies the list of predicate selectivities. We also assume independent selectivities for simplicity. The two main cost components are the scan cost and seek cost of every range scan. We first compute the total number of scan operations. At i^{th} level of the query evaluation, the number of scanned objects increases by a factor of $s_{i+1}f$:

$$n_{scan} = f + s_1f^2 + \dots + \prod_{i=1}^{d-1} s_i f^d \leq \frac{s^d f^{d+1}}{sf - 1} \leq s^d f^{d+1}$$

where $s = \max_{i=1}^{d-1} s_i$ and $sf \geq 2$. The number of seek operations is equal to the number of scan operations of the preceding path sub-expression and can be computed to be $s^{d-1}f^d$ in a similar manner. Because there is no closed form formula for f in terms of $|T|$, we approximate f to $|T|^{1/h}$. Finally, the total cost can be approximated to $s^{d-1} \cdot |T|^{d/h} (s \cdot |T|^{1/h} \cdot c_{scan} + c_{seek})$ where c_{scan} and c_{seek} represent the cost of a single scan operation and that of a seek operation, respectively. We observe that the multiplier effect of selectivity s effectively bounds the number of scan and seek operations necessary for executing the given query, which aligns with the intuition about “early pruning” of objects. This contrasts with more popular join operators, such as sort-merge join or hash join that require at least one scan of the entire table, which can be costly for a large table size (analogously tree size).

4.5 Version Control Operations

The system supports globally consistent chronological versioning of the entire catalog data, using monotone increasing *vid* associated with every transaction. `snapshot()` operation is implemented by maintaining an additional mapping between the snapshot name and the *vid*, which can be used to retrieve the correct version. `clone()` operation copies the source object and all of its descendants (except for the leaf objects) at some *vid* to the destination path. For cloned leaf objects, the system adds a mapping from the new destination paths to the primary paths, so they are shared between the source and destination parents, saving the storage costs of duplication.

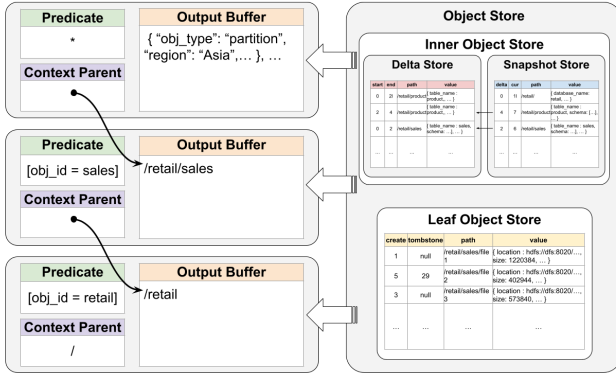


Figure 6: Query Execution of a Simple Example Query: `/[obj_id = retail]/[obj_id = sales]/*`. The output buffer in the top most node holds the final result set.

5 CONCURRENCY CONTROL

In this section, we present TreeCat’s concurrency control mechanism, which guarantees serializability. Besides our use case in TreeCat, the scheme is also of independent interest in a more general context of databases with a hierarchical data model.

5.1 Discussion

We first discuss the background and primary motivations, followed by an in-depth description of the implementation details. We adopt the MVOCC scheme where read-only transactions operate on a versioned snapshot of the data, and read-write transactions follow a protocol similar to the standard OCC mechanism. Because we implemented MVOCC on top of the existing version system, *vid* is synonymous with *timestamp* as commonly discussed in the concurrency control literature. For read-only transactions, we follow the common practice of assigning a *read_vid*, the version up to which data are safe for read, at the start of the transaction, and executing queries against the versions of objects visible at the assigned *read_vid*. Since this aspect of the mechanism aligns with other existing systems, our main focus is on the OCC, which utilizes a predicate-based method, often overlooked by standard schemes.

The important feature of TreeCat is that read queries involve **predicate read** operations. Unlike simple key-value read and write operations, predicate read operations can introduce *predicate dependencies* [1, 10], which are distinguished from the *item dependencies* (write-read, write-write, and read-write) that are commonly discussed in the concurrency control literature. The infamous *phantom read anomaly* is one example of a predicate anomaly that can occur if predicate dependencies are not handled correctly by the concurrency control mechanism [22]. Unlike item dependencies—which are based on conflicting operations on the same item—predicate dependencies occur when an update alters whether an object satisfies a given predicate. Specifically, if a write operation’s before-image satisfies the predicate of a predicate read operation but its after-image does not, or vice versa, a predicate dependency exists between the two operations [1]. Item-based approaches, including record locks in 2PL (2-Phase Locking), read and write sets in OCC,

and time stamps on individual items in TO (timestamp ordering), lack semantics for predicate reads, and cannot handle predicate dependencies without additional mechanisms.

The most common way to augment these schemes to handle predicate dependencies is to lock the entire scan range of the predicate read operation. If an index scan is used, the scan range can be limited to a range of values of the indexed attribute on which the predicate is defined. However, this optimization cannot be applied to arbitrary predicates, which may be a complex expression or involve attributes on which an index does not exist. The scheme is, at best, a coarse-grained approximation that may lock an irrelevant range of values, resulting in unnecessary conflicts.

We illustrate this point with a simple example. Let R be a logical table that is partitioned into multiple files. Transaction T_1 performs a predicate read operation on R ’s file metadata with the predicate p on the file statistics. Because there is no index on the file statistics, T_1 scans the entire file metadata of R . Meanwhile, T_2 inserts a file f with file metadata that does not satisfy p into R . There is no reason for the two operations to conflict, because the new file f is never visible to T_1 ’s predicate read, regardless of the order of operations. However, with scan range locking, T_1 would lock the entire range of R ’s file metadata. In a pessimistic scheme (S2PL), either of the two transactions would block the other. In an optimistic scheme (OCC), if T_2 is submitted for validation before T_1 , T_1 would be aborted.

There are predicate-based methods that can precisely detect any predicate dependency. One scheme is to re-execute the scan operations to observe any changes in the visible objects [43]. While precise, re-executing scans can be costly for large ranges. Precision locking [41] is a form of predicate locking that evaluates the predicate directly against posted write operations. This technique detects predicate dependencies accurately and avoids the satisfiability test problem of conventional predicate locking. But in TreeCat, a read query is correlated whence an object satisfies the entire path expression only if all of its ancestors also satisfy the corresponding path sub-expressions. A naive implementation that only considers the predicates in a given query can, in fact, result in false conflicts.

In our scheme, we combine scan range locking with precision locking, leveraging the strengths of both methods. We first apply optimistic prefix-based scan range locking to detect candidate conflicts, i.e., write operations that fall within the scan range, but may not necessarily cause conflicts. Then, we use optimistic precision locking to confirm whether the candidate is actually a conflict. This scheme achieves the granularity of precision locking while improving efficiency by pruning unnecessary predicate evaluations with scan range locking. In addition, we make use of an in-memory index structure to quickly locate only the log records that fall in the scan range, rather than applying precision locking on the entire set.

5.2 Implementation

The concurrency control is handled by the **TransactionManager**, which keeps track of *read_vid*, the latest version up to which, the catalog can be read safely, and *commit_vid*, which is the monotone increasing vid assigned to a transaction before final commit. Key in-memory data structures include:

- (1) *ScanSet*: Maintained by each transaction. Stores pairs of the scan range (in the form of the context object path) and the

predicate of the predicate read operations executed as part of the transaction.

- (2) *LogBuffer*: Maintained by each transaction. Stores before- and after-images of all recent write operations.
- (3) *VersionMap*: Used by **TransactionManager** for validation. Maps object path to the most recent *vid* of its children.
- (4) *LogIndex*: Used by **TransactionManager** for validation. Index on the *LogBuffer*, mapping pairs of object path and *vid* to the log records. The index is implemented using a lock-free skip list and is sorted by the parent path, then *vid*.

When *startTransaction()* is called, a new transaction session is created and *read_vid* is loaded as the *vid* at which all read queries are executed in the session. During query execution, pairs of predicates and parent paths (prefix) of objects scanned for predicate evaluation are added to the *ScanSet*. Also, a private copy of the write set is constructed on the client side. When the transaction is submitted for the commit process, the write set is sent to the server and undergoes preprocessing, where it is copied to the *LogBuffer* and preconditions are checked. If the transaction passes these checks, the transaction is scheduled for the main commit process.

The main commit process consists of two phases: validation and batch write. For each phase, there is a dedicated thread pool of a fixed size, with each thread handling a single hash partition of the *ScanSet* and the write set. At the start of the validation, **TransactionManager** executes a *fetch_and_add()* operation on the *commit_vid* and assigns it to the transaction. Then, each worker thread validates the scan ranges in the *ScanSet* against the *VersionMap* to determine whether any objects scanned by predicate read operations have changed. If potential conflicts are detected, an iterator on *LogIndex* is set to scan and apply precision locks on corresponding recent writes with *vid* greater than the transaction's *read_vid* and fall in the scan range. If any conflict is detected through precision locking, the transaction is aborted. After all threads pass validation, they update the *VersionMap* and *LogIndex* based on their respective write set partitions for the transaction. The validation phase (omitting the initial *vid* assignment and latter updates to in-memory data structures) proceeds as per Algorithm 1. Note that for merge operation, before- and after-image may have to be constructed as they are not available until the final commit. However, once constructed, the images are cached so they are accessible during the validation of subsequent transactions.

Once validation is complete, the transaction proceeds to the write phase. During this phase, each thread processes its assigned write partition, evaluating the final values of merge operations (commit-time updates) and applying the writes as an atomic batch write operation to the underlying RocksDB instance. After every thread finishes its batch write, WAL records of the transaction are flushed to disk along with those of other transactions via *SyncWAL()* operation, which calls *fsync()* on the WAL files. The *read_vid* is also updated, making the committed changes visible for subsequent reads. Transactions are processed one at a time in each phase, but parallelism is achieved via hash-partitioning the workload and pipelining different phases of the commit process.

Old *LogBuffers* and entries in *VersionMap* and *LogIndex* that are no longer needed are asynchronously garbage collected by a background thread. For this purpose, the **TransactionManager**

Algorithm 1 Validation Algorithm

```

1: for (path, predicate)  $\in$  txn.ScanSet
2:   // scan range locking
3:   if VersionMap[path] > txn.read_vid
4:     iter  $\leftarrow$  LogIndex.newIterator()
5:     iter.seek(path, txn.read_vid)
6:     while iter.path.parent = path
7:       if iter.type = merge and iter.image.empty()
8:         iter.image.construct()
9:       // precision locking
10:      if iter.image.before  $\models$  predicate
11:        abort(txn)
12:      else if iter.image.after  $\models$  predicate
13:        abort(txn)

```

keeps track of *watermark*, the oldest *read_vid* among the current transactions. After each transaction commit, the *watermark* is updated and all entries in the aforementioned data structures with *vid* below the *watermark* are garbage collected.

5.3 Proof of Correctness

In this section, we prove that TreeCat's concurrency control mechanism guarantees serializable isolation.

Concurrency Control Model: We first formalize the concurrency control model. We follow Adya's model [1], where every object in the database has one or more versions. Every operation on an object is mapped to exactly one version of that object. Given an object x , x_i is a version with *vid* i , installed by transaction T_i . Every object is assumed to exist forever regardless of when they are inserted or deleted. The unborn state is represented as version x_{init} , while the dead state is represented as x_i , installed by the delete operation of T_i . Every object is uniquely identified by its full path.

It is important to discuss how TreeCat's query evaluation is represented in this model. Evaluation of a path expression involves performing a series of predicate range scans, each with respect to a context parent object. For example, given the path expression, */obj_id = "retail"/[name = "sales"]*, the evaluation of the second predicate, *name = "sales"*, is executed as a range scan on the children of */retail*. Each range scan can be modeled as a predicate read operation, where the predicate is a conjunction of the original predicate and the equivalence between the object parent path and the context parent path. Following the above example, evaluating *name = "sales"* against */retail*'s children objects can be modeled as a predicate read operation with the predicate, *name = "sales" \wedge parent = /retail*. The evaluation of a single predicate in a path query can correspond to multiple such predicate read operations, as there can be multiple context parent objects with which it is evaluated.

Main Proof: We now proceed with the main proof.

Lemma 1. *For any two committed transactions T_i and T_j where $commit_vid_i < commit_vid_j$, there is no dependency from T_j to T_i .*

PROOF. We prove the lemma by enumeration. First, there cannot be any ww-dependency from T_j to T_i because write operations are performed in the order of the *commit_vid* in the write phase.

Second, wr-dependency from T_j to T_i is also not possible. Given that $read_vid_i < commit_vid_i$ and $commit_vid_i < commit_vid_j$, $read_vid_i < commit_vid_j$ by transitivity. Because changes committed by T_j become visible only after the $read_vid$ has been incremented to $commit_vid_j$, T_i , which has a lower $read_vid$, cannot observe any changes installed by T_j . Lastly, we show by contradiction that there cannot be any rw-antidependency from T_j to T_i . Assume that there is a rw-antidependency from T_j to T_i . Consider the state when T_i passed validation, but T_j is about to enter validation. There must be a write operation w executed by T_i on an object o , whose before- or after-image satisfies the predicate of a predicate read operation P executed by T_j , which scans w 's before-image. Per our earlier discussion, the predicate of P is $p \wedge parent = c$, where p is the predicate in the path query and c is the context parent object. During the query execution of T_j , (c, p) was added to T_j 's *ScanSet*. Also, $(o.parent, i)$ was added to *VersionMap* and a mapping from o to w 's before- and after-images was added to *LogIndex* during the validation of T_i . During the validation of T_j , c is found in *VersionMap* because $c = o.parent$ holds true. Since $read_vid_j < commit_vid_i$, p is evaluated against o 's entries in *LogIndex*. Because the before- or after-image of w satisfies p , the rw-conflict is detected and T_j is aborted, leading to a contradiction. We have now considered all types of dependencies, so the lemma holds true. \square

Theorem 1. *Any schedule of committed transactions output by TreeCat is conflict serializable.*

PROOF. We prove the theorem by contradiction. Assume that there is a schedule s , output by TreeCat, that is not conflict serializable. Then, the corresponding precedence graph, $G(s) = (V, E)$ must contain at least one cycle. Consider a directed edge (T_i, T_j) in the cycle. Because there is a dependency from T_i to T_j , $commit_vid_i < commit_vid_j$ by Lemma 1. Because (T_i, T_j) is in a cycle there is a path from T_j to T_i . Using Lemma 1 and transitivity, we can prove by induction (which we omit for brevity) that $commit_vid_i > commit_vid_j$, resulting in a contradiction. \square

5.4 Related Work

We first discuss the concurrency control mechanisms of other catalog systems and then examine how predicate reads are handled by concurrency control mechanisms across more general data systems.

Catalog Concurrency Control: Traditional RDBMSs store the catalog as a set of relations, relying on the same concurrency mechanism that is used for data. While data warehouses use proprietary mechanisms, Snowflake is known to use FoundationDB to persist transaction states and lock queues [47]. Even though resource locks at the level of micro-partitions are available, most single table DML queries require locking all micro-partitions [65]. Delta Lake uses snapshot-based OCC where, during the final commit, it tries to atomically write the next delta log record. If the underlying storage system does not support atomic file rename, it has to resort to another transactional system, such as DynamoDB [60]. Delta Lake currently only supports single-table transactions. However, there is early work on introducing a new commit protocol called *commit coordinator*, an external system that is dedicated to handling concurrency control, which may support multi-table transactions in the future [69]. Iceberg also uses snapshot-based OCC where,

after writing all the metadata files (manifest files, manifest list files, and metadata file), it does a final commit by switching the current pointer in the catalog to the new metadata file. If there is a conflict, the transaction undergoes a retry validation, so the entire operation need not be repeated from scratch [55]. Like Delta Lake, Iceberg's commit protocol only supports single-table transactions. Recently, multi-table transactions have been added to the Iceberg REST API. But as far as we know, Nessie [52], which supports git-like versioning, is the only compatible catalog that supports it. Even so, the transactions operate at the level of a single branch of the whole database. Lastly, Hudi also uses snapshot-based OCC and does conflict resolution based on the files modified by multiple writers [26]. Hudi requires an external lock provider for locking tables and can only support single-table transactions.

Predicate Read: To guarantee serializable isolation, a data system must account for any predicate dependencies [1]. The standard solution is to lock the entire scan range of the predicate read operation, an example of which is next key locking [45], which locks the current record and the gap until the next record. Since taking many record locks can cause excessive overheads, multiple granularity locking [32] is commonly used to escalate locks to more coarse-grained ones, such as table locks. Large scan range may cause unnecessary conflicts, so it is crucial to minimize the scan range. For simple predicates, performing an index scan can limit the scan range. A more advanced technique may be used to better approximate the minimum scan range for complex predicates. For example, Spanner uses a runtime data structure called a filter tree to compute the approximate scan range [7]. One way to precisely handle predicate dependencies is to validate the scan operation by re-executing it, which is adopted by Hekaton [19]. Another approach is to use predicate-based locking. Conventional predicate locking involves satisfiability test problem, which is known to be NP-complete [36]. However, precision locking avoids this problem by directly evaluating predicate against before- and after-images of posted writes [41]. The technique was not used in any real system for a long time, but was later adopted by Neumann et al. for implementing the MVOCC mechanism of Hyper [54].

6 EXPERIMENTAL EVALUATION

In this section, we present a comprehensive experimental evaluation of TreeCat. We set up a cluster of 4 physical machines on CloudLab [11], each equipped with a 10 core Intel Xeon Silver 4114 processor, 192GiB DDR4 memory, and a 64GB SATA SSD, interconnected by 10Gbps ethernet. We implemented a catalog API, including *getTable()* and *listFiles()* operations, in Apache Spark 3.4, allowing communication with a standalone instance of TreeCat.

6.1 Concurrency Control Mechanism

To evaluate the concurrency control mechanism, we implemented two standard concurrency control protocols in TreeCat. The first is OCC with optimistic scan range locking (OSL), which also primarily utilizes prefix-based scan range locks. However, instead of precision locking, it uses more fine-grained range locks for bounded range scans. For example, if only table partitions 1 to 100 are scanned, only the range of object id, [1, 100], rather than the entire range of table partitions, is optimistically locked. The second scheme implements

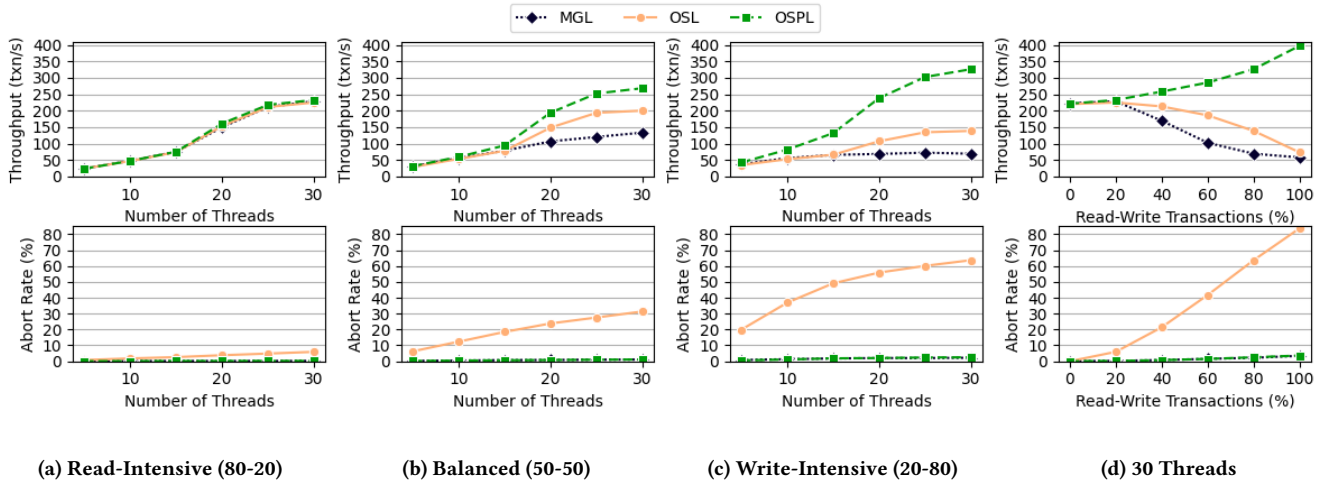


Figure 7: Throughput and abort rate across different number of client threads and different types of workloads.

S2PL with multiple granularity locking (MGL). All three schemes efficiently handle read-only transactions with an identical MVCC protocol. Our new proposed concurrency control mechanism, i.e., TreeCat’s default protocol, is referred to as OSPL (optimistic scan range and precision locking).

We used a microbenchmark largely derived from the TPC-DS benchmark. While TPC-DS proceeds as a single fixed sequence of phases, we design our benchmark so that multiple client threads randomly select queries and perform the corresponding metadata operations (against TreeCat) concurrently. Besides the original read queries and data maintenance operations, which involve data insertion and deletion on the fact tables, we introduce two new operations: an **optimize** operation, which selects and merges small files, and **dimension table insertion**. These new operations allow the simulation of real-world scenarios that require support for **multi-table transactions**. The relative distribution of write operations is adjusted with realistic assumptions about standard data warehouse workloads. For example, we assume that fact table insertion occur 12 times more frequently than dimension table insertion (every 5 minutes vs. every hour) and 144 times more frequently than optimize operations (every 5 minutes vs. every 12 hours). We adopted the schema of the TPC-DS benchmark, but used a custom data generator to generate relatively large amount of metadata. The amount of metadata per table, the cardinality of surrogate keys and business ids are in line with the TPC-DS specification for a scale factor of 100 terabytes. In our experiment, we deployed TreeCat on a single machine and client threads on the other three machines.

We first conducted a scalability experiment to measure overall throughput, abort rate, and latency over a set time duration across different numbers of client threads and percentages of read-write transactions. For all three schemes, read-only transactions do not cause any conflict as they follow the MVCC protocol. Consequently, performance differences between the schemes are minimal for more read-intensive workloads, as shown in Figure 7. However, for write-intensive workloads, throughput diverges, with OSPL significantly

outperforming the other two as the number of client threads increases. The main source of contention is read-write conflicts on the customers table, which is one of the dimension tables to which data is inserted periodically. Inserting into a fact table requires joining the input data with dimension tables on the business ids (primary key in the source OLTP database) and should only conflict with dimension table insertion if the ranges of dimension table business IDs overlap. However, both OSL and MGL apply coarse-grained scan range lock on the dimension tables, leading to unnecessary conflicts. As a result, the abort rate of OSL increases almost linearly with the number of threads, reaching 31.3% for the balanced workload and 63.7% for the write-intensive workload at 30 client threads. Similarly, MGL maintains low throughput since transactions are blocked while waiting for lock acquisition. By contrast, OSPL avoids false conflicts through precision locking on the file statistics, achieving significantly higher throughput and a low abort rate. The latter two plots of Figure 7 show this trend more clearly as the percentage of read-write transactions increases for 30 client threads. A considerable proportion of read-only transactions scan entire fact tables and have longer latency than read-write transactions, which is why the throughput of OSPL increases with more write-intensive workloads.

Next, we compare latency across the different schemes for the balanced workload. Because read-only operations have longer latencies than read-write transactions and follow the same MVCC protocol in all three schemes, we isolate and analyze the latencies of read-write transactions. Figure 8 shows the 99th percentile latency of read-write transactions across different number of threads and for each operation type. While the latency remains consistent for the two OCC schemes, it increases almost linearly with the number of client threads for MGL as transaction lock queues build up. Categorizing transactions by operation type reveals that the latency of dimension table insertion is almost as twice as long as that of fact table insertion, indicating that the conflicts on the updated dimension table is the main source of contention for MGL. *The latency of optimize operation is significantly longer for OSPL compared to OSL*




Figure 8: 99th percentile operation latency of read-write transactions across different number of client threads and operation types for 30 threads, respectively. In the latter figure, we use abbreviations: Opt (optimize), Fact (fact table insertion), Dim (dimension table insertion), and Del (delete).

because higher number of small files accumulate with less aborts on insertion, resulting in larger optimize operation.

Lastly, we evaluate how two important structural properties of the metadata schema, breadth and depth, affect performance. These experiments use a separate microbenchmark designed to be more generalizable. The workload consists of a mix of read and read-write operations (50-50) on a single data set, which is range-partitioned into 100000 data files by a set of independent clustering attributes. Both types of operation perform a predicate read operation with fixed selectivity, using range predicates on randomly selected clustering attributes. Read-write operation additionally inserts a new data file. A partition level is constructed by grouping data files by a clustering attribute. Nested grouping by multiple attributes results in a multi-level hierarchy where the number of such attributes determines the overall partition level (depth) and the cardinality determines the fan-out (breadth). As shown in Figure 9, throughput is low across all schemes at low fan-out because of large scan ranges (e.g., all 100000 files have to be scanned at the partition level of 1). However, both OSL and MGL are further affected by contention from coarse-grained scan range locking. OSPL, on the other hand, avoids false conflicts as it uses a predicate-based method. Throughput converges with higher fan-out, as MGL and OSL benefit from more fine-grained lock ranges. A similar trend can be observed with different partition levels (depths) for the same reason. However, throughput is divergent even at the highest partition level because a subset of the clustering attributes do not form partitioning groups; a predicate read operation only on these attributes requires scanning all partition groups, resulting in higher contention for OSL and MGL. Although OSL outperforms MGL in throughput, it suffers from high abort rates.

6.2 Comparative Evaluation

Next, we compare TreeCat with HMS, Delta Lake, and Iceberg. An apples-to-apples comparison is difficult because TreeCat does not support a catalog API fully compatible with Spark and all the systems have different architectures and transactional guarantees. We nonetheless try our best to compare the relative costs of metadata operations and conduct an end-to-end experiment with an alternate database engine. We set up HDFS (Hadoop Distributed File System) and Spark on the 4 node cluster. We used Hive 2.3, Delta




Figure 9: Throughput for different breadths (fan-out) and depths (partition level) of the metadata schema.

Lake 2.4, and Iceberg 1.5, all of which are compatible with Spark 3.4. We used a single table with a schema identical to that of the Store Sales table in the TPC-DS, partitioned by the sales date, from 1998-01-01 to 2003-12-31. To test scalability of read performance, we populated the table with varying numbers of files and measured the latency of retrieving files with a range filter on the date. While our primary implementation of the catalog API uses a single thread to retrieve the query result via a single stream, both HMS and Iceberg leverage multiple cores on the client machine to retrieve the result over multiple streams where HMS calls `listFile()` and Iceberg retrieves manifest files in parallel. Client-side parallelism can be achieved with TreeCat quite easily by adding a partitioning predicate to the path query, so that multiple threads can concurrently submit queries and retrieve distinct result sets. Our implementation applies the `endwith()` operation to the partition value due to the limitations of the current query language. We used this parallel implementation for the following experiments.




Figure 10: Median latency of file retrieval across different number of files for 1 day and 365 days predicate ranges, respectively.

As shown in Figure 10, we first analyze the scalability of the file retrieval operation with high selectivity (date range of 1 day) and with low selectivity (date range of 365) across different numbers of files. For the 1-day predicate range, TreeCat consistently shows low latency as it efficiently traverses the hierarchical data on the server and only returns the file metadata in the given data partition. HMS also scales quite well as it calls `listPartitionsbyFilter()` to only retrieve the partitions in the given range from the thrift server. Although Delta Lake benefits from Spark’s distributed processing engine, the overhead of executing the job adds significant latency, resulting in a relatively high minimum latency of around 600 ms. This contrasts with previous experimental results [3, 38] and we suspect that performance could be significantly improved




Figure 11: Median latency of file retrieval across different predicate ranges and numbers of physical cores on client machine, resp. File count is 500000 for both while predicate range is 365 for the latter.

with a better-configured, larger Spark cluster. But the client-side implementation of getting the end result appears to be a major bottleneck. Although Iceberg utilizes multiple cores for processing manifest files, the manifest file size can get large, resulting in the longer latency for the higher file counts. TreeCat outperforms the other systems and scales well also for the 365-days predicate range. HMS beats both Delta Lake and Iceberg when the file count reaches 50000, which we mainly attribute to the low latency of HDFS.

Next, we analyze how the selectivity of the date range filter affects the latency for the file count of 500000. Both TreeCat and HMS effectively prune irrelevant date partitions, resulting in better performance for narrower predicate ranges. Iceberg, on the other hand, operates at the granularity of manifest files, so the latency remains relatively high. Delta Lake also does not scale too well as the number of retrieved files increases. We also evaluate how the number of physical cores on the client machine affects the overall performance, using 365-day date range and a file count of 500000. As expected, the performance of HMS, Iceberg, and TreeCat, which take advantage of client-side parallelism, degrades with limited number of cores while that of Delta Lake remains consistent.

We conclude with an end-to-end experiment, using DuckDB 1.2 as the execution engine. The workload involves repeatedly inserting data of constant size into an unpartitioned Store Sales table. A single operation proceeds in three steps: (1) Retrieve the table metadata, including the table directory and schema from the catalog; (2) Parse the schema, generate data, and write the output to a parquet file under the table directory, using DuckDB; (3) Invoke the commit protocol on the catalog with the metadata of the new data file. We measure both the total throughput over a set duration and the end-to-end latency of each invocation across different data file sizes, configured using the number of rows per file (powers of 10, starting from 100). The results are shown in Figure 12. We observe that the high latency of metadata operations can indeed cause a performance bottleneck, especially for smaller file sizes. As shown in the time breakdown of Figure 12, the main difference lies in the commit protocol: Delta Lake’s commit protocol takes over 1.5 seconds, and Iceberg takes around 250 ms, while TreeCat only takes around 10 ms. The commit protocol for lakehouse storage formats involves retrieving the latest table snapshot (which may require reading and processing metadata files), writing one or more metadata files, and invoking a CAS (compare-and-swap) operation for ACID guarantees. The entire process involves one or more operations on the shared storage system, which have non-trivial




Figure 12: Throughput of insert operations (log scale) for different write sizes, and time breakdown for an insert operation of 10000 row data file (≈ 238 KB), respectively.

overhead. Delta Lake incurs extra overhead from launching a Spark job for metadata processing. In contrast, TreeCat uses in-memory data structures for validation, applies writes to RocksDB, which is optimized for fast write, and is only delayed by *fsync()* operation on the WAL file, resulting in lower latency. The throughput eventually converges with larger file sizes as the cost of the data operation becomes more dominant. However, given applications such as streaming systems [4, 8] that generate small data files at high frequency, these results carry significant real-world implications.

7 CONCLUSION AND DISCUSSION

In this paper, we identified the primary use cases of a standalone operational catalog and justified the key design decisions for our catalog engine, TreeCat. Our experimental evaluation not only exposes the limitations of existing solutions but also validates the advantages of our design choices.

Looking ahead, we see a number of critical challenges in this space that need further research. Some of these challenges lie outside the scope of the catalog itself. For instance, **consistency between the data and the metadata** (e.g., ensuring that a pointer to a physical object is valid), is a major emerging challenge. Yet it is not clear where that responsibility lies. Similarly, **access control** is crucial but needs to be enforced by the layer that manages access to the data itself. **Garbage collection** of data objects and their metadata that are no longer visible may become necessary over time. A TreeCat instance can be periodically scanned to find and return leaf objects that are no longer reachable by any path from the root. These objects can be quarantined in temporary storage until the data layer invokes a final commit to remove them. Finally, **scalability and availability** are important concerns in a distributed setting. Given the relatively small volume of metadata and large numbers of read queries, we propose the adoption of a single-writer, multiple-reader architecture where all read-write transactions are handled by a single server, and logs are shipped to and replayed by read replicas [75]. This scheme would require minimum changes to the concurrency control mechanism because a single server will process all write operations. If the primary writer server fails, a new writer server can be elected via a quorum-based mechanism, ensuring availability.

ACKNOWLEDGMENTS

This research was supported in part by Dolby Laboratories.

REFERENCES

- [1] A. Adya, B. Liskov, and P. O’Neil. 2000. Generalized isolation level definitions. In *Proceedings of 16th International Conference on Data Engineering (Cat. No. 00CB37073)*. 67–78.
- [2] ArangoDB. 2025. ArangoDB. <https://arangodb.com/>. Accessed: July 8, 2025.
- [3] Michael Armbrust, Tathagata Das, Liwen Sun, Burak Yavuz, Shixiong Zhu, Mukul Murthy, Joseph Torres, Herman van Hovell, Adrian Ionescu, Alicja Łuszczak, Michał undefinedwitakowski, Michał Szafranski, Xiao Li, Takuya Ueshin, Mostafa Mokhtar, Peter Boncz, Ali Ghodsi, Sameer Paranjpye, Pieter Senster, Reynold Xin, and Matei Zaharia. 2020. Delta lake: high-performance ACID table storage over cloud object stores. *Proc. VLDB Endow.* 13, 12 (Aug. 2020), 3411–3424.
- [4] Michael Armbrust, Tathagata Das, Joseph Torres, Burak Yavuz, Shixiong Zhu, Reynold Xin, Ali Ghodsi, Ion Stoica, and Matei Zaharia. 2018. Structured Streaming: A Declarative API for Real-Time Applications in Apache Spark. In *Proceedings of the 2018 International Conference on Management of Data (Houston, TX, USA) (SIGMOD ’18)*. Association for Computing Machinery, New York, NY, USA, 601–613.
- [5] Michael Armbrust, Reynold S. Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K. Bradley, Xiangrui Meng, Tomer Kaftan, Michael J. Franklin, Ali Ghodsi, and Matei Zaharia. 2015. Spark SQL: Relational Data Processing in Spark. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Melbourne, Victoria, Australia) (SIGMOD ’15)*. Association for Computing Machinery, New York, NY, USA, 1383–1394.
- [6] Nikos Armenatzoglou, Sanuj Basu, Naga Bhanoori, Mengchu Cai, Naresh Chainani, Kiran Chinta, Venkatraman Govindaraju, Todd J. Green, Monish Gupta, Sebastian Hillig, Eric Hotinger, Yan Leshinsky, Jintian Liang, Michael McCreedy, Fabian Nagel, Ippokratis Pandis, Panos Parchas, Rahul Pathak, Orestis Polychroniou, Foyzur Rahman, Gaurav Saxena, Gokul Soundararajan, Sriram Subramanian, and Doug Terry. 2022. Amazon Redshift Re-invented. In *Proceedings of the 2022 International Conference on Management of Data (Philadelphia, PA, USA) (SIGMOD ’22)*. Association for Computing Machinery, New York, NY, USA, 2205–2217.
- [7] David F. Bacon, Nathan Bales, Nico Bruno, Brian F. Cooper, Adam Dickinson, Andrew Fikes, Campbell Fraser, Andrey Gubarev, Milind Joshi, Eugene Kogan, Alexander Lloyd, Sergey Melnik, Rajesh Rao, David Shue, Christopher Taylor, Marcel van der Holst, and Dale Woodford. 2017. Spanner: Becoming a SQL System. In *Proceedings of the 2017 ACM International Conference on Management of Data (Chicago, Illinois, USA) (SIGMOD ’17)*. Association for Computing Machinery, New York, NY, USA, 331–343.
- [8] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache flink: Stream and batch processing in a single engine. *The Bulletin of the Technical Committee on Data Engineering* 38, 4 (2015).
- [9] James Clark, Steve DeRose, et al. 1999. XML path language (XPath).
- [10] Jack Clark, Alastair F. Donaldson, John Wickerson, and Manuel Rigger. 2024. Validating Database System Isolation Level Implementations with Version Certificate Recovery. In *Proceedings of the Nineteenth European Conference on Computer Systems (Athens, Greece) (EuroSys ’24)*. Association for Computing Machinery, New York, NY, USA, 754–768.
- [11] CloudLab. 2025. CloudLab. <https://www.cloudbab.us/>. Accessed: July 8, 2025.
- [12] Creative Commons. 2025. BSON. <https://bsonspec.org/>. Accessed: July 8, 2025.
- [13] CouchDB. 2025. CouchDB. <https://couchdb.apache.org/>. Accessed: July 8, 2025.
- [14] Databricks. 2025. Catalogs and the REST catalog. <https://tabular.io/apache-iceberg-cookbook/getting-started-catalog-background/>. Accessed: July 8, 2025.
- [15] Databricks. 2025. Unity Catalog. <https://www.databricks.com/product/unity-catalog>. Accessed: July 8, 2025.
- [16] DataHub. 2025. DataHub. <https://datahubproject.io/>. Accessed: July 8, 2025.
- [17] DataHub. 2025. DataHub Architecture Overview. <https://datahubproject.io/docs/architecture/architecture/>. Accessed: July 8, 2025.
- [18] Dean De Leo and Peter Boncz. 2021. Tesoo and the analysis of structural dynamic graphs. *Proc. VLDB Endow.* 14, 6 (Feb. 2021), 1053–1066.
- [19] Cristian Diaconu, Craig Freedman, Erik Ismert, Per-Ake Larson, Pravin Mittal, Ryan Stonecipher, Nitin Verma, and Mike Zwilling. 2013. Hekaton: SQL server’s memory-optimized OLTP engine. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data (New York, New York, USA) (SIGMOD ’13)*. Association for Computing Machinery, New York, NY, USA, 1243–1254.
- [20] Ayush Dubey, Greg D. Hill, Robert Escriva, and Emin Gün Sirer. 2016. Weaver: a high-performance, transactional graph database based on refinable timestamps. *Proc. VLDB Endow.* 9, 11 (July 2016), 852–863.
- [21] John Esmet, Michael A. Bender, Martin Farach-Colton, and Bradley C. Kuszmaul. 2012. The TokuFS streaming file system. In *Proceedings of the 4th USENIX Conference on Hot Topics in Storage and File Systems (Boston, MA) (HotStorage’12)*. USENIX Association, USA, 14.
- [22] K. P. Eswaran, J. N. Gray, R. A. Lorie, and I. L. Traiger. 1976. The notions of consistency and predicate locks in a database system. *Commun. ACM* 19, 11 (Nov. 1976), 624–633.
- [23] The Apache Software Foundation. 2019. Apache Atlas. <https://atlas.apache.org/2.0.0/>. Accessed: July 8, 2025.
- [24] The Apache Software Foundation. 2022. Apache Hudi Technical Specification. <https://hudi.apache.org/tech-specs/>. Accessed: July 8, 2025.
- [25] The Apache Software Foundation. 2025. Apache Hudi. <https://hudi.apache.org/>. Accessed: July 8, 2025.
- [26] The Apache Software Foundation. 2025. Apache Hudi: Concurrency Control. https://hudi.apache.org/docs/concurrency_control/. Accessed: July 8, 2025.
- [27] The Apache Software Foundation. 2025. Apache Iceberg. <https://iceberg.apache.org/>. Accessed: July 8, 2025.
- [28] The Apache Software Foundation. 2025. Apache TinkerPop. <https://tinkerpop.apache.org/gremlin.html>. Accessed: July 8, 2025.
- [29] The Apache Software Foundation. 2025. Iceberg Table Spec. <https://iceberg.apache.org/spec/>. Accessed: July 8, 2025.
- [30] The Linux Foundation. 2025. Delta Lake. <https://delta.io/>. Accessed: July 8, 2025.
- [31] Per Fuchs, Domagoj Margan, and Jana Giceva. 2022. Sortedledon: a universal, transactional graph data structure. *Proc. VLDB Endow.* 15, 6 (Feb. 2022), 1173–1186.
- [32] J. N. Gray, R. A. Lorie, G. R. Putzolu, and I. L. Traiger. 1994. *Granularity of locks and degrees of consistency in a shared data base*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 181–208.
- [33] gRPC Authors. 2025. gRPC. <https://grpc.io/>. Accessed: July 8, 2025.
- [34] J.M. Hellerstein, M. Stonebraker, and J. Hamilton. 2007. *Architecture of a Database System*. Now Publishers. <https://books.google.com/books?id=aBPcm3C1avMC>
- [35] Yihe Huang, William Qian, Eddie Kohler, Barbara Liskov, and Liuba Shrira. 2020. Opportunities for optimism in contented main-memory multicore transactions. *Proc. VLDB Endow.* 13, 5 (Jan. 2020), 629–642.
- [36] Harry B. Hunt and Daniel J. Rosenkrantz. 1979. The complexity of testing predicate locks. In *Proceedings of the 1979 ACM SIGMOD International Conference on Management of Data (Boston, Massachusetts) (SIGMOD ’79)*. Association for Computing Machinery, New York, NY, USA, 127–133. <https://doi.org/10.1145/582095.582115>
- [37] HyperGraphDB. 2025. HyperGraphDB. <https://hypergraphdb.org/>. Accessed: July 8, 2025.
- [38] Paras Jain, Peter Kraft, Conor Power, Tathagata Das, Ion Stoica, and Matei Zaharia. 2023. Analyzing and Comparing Lakehouse Storage Systems. In *CIDR*.
- [39] William Jannen, Jun Yuan, Yang Zhan, Amogh Akshintala, John Esmet, Yizheng Jiao, Ankur Mittal, Prashant Pandey, Phaneendra Reddy, Leif Walsh, Michael Bender, Martin Farach-Colton, Rob Johnson, Bradley C. Kuszmaul, and Donald E. Porter. 2015. BetrFS: a right-optimized write-optimized file system. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies (Santa Clara, CA) (FAST’15)*. USENIX Association, USA, 301–315.
- [40] JanusGraph. 2025. JanusGraph. <https://janusgraph.org/>. Accessed: July 8, 2025.
- [41] J. R. Jordan, J. Banerjee, and R. B. Batman. 1981. Precision locks. In *Proceedings of the 1981 ACM SIGMOD International Conference on Management of Data (Ann Arbor, Michigan) (SIGMOD ’81)*. Association for Computing Machinery, New York, NY, USA, 143–147.
- [42] Oz Katz. 2021. Hive Metastore-It Didn’t Age Well. <https://lakefs.io/blog/hive-metastore-it-didnt-age-well/>. Accessed: July 8, 2025.
- [43] Per-Ake Larson, Spyros Blanas, Cristian Diaconu, Craig Freedman, Jignesh M. Patel, and Mike Zwilling. 2011. High-performance concurrency control mechanisms for main-memory databases. *Proc. VLDB Endow.* 5, 4 (Dec. 2011), 298–309.
- [44] Ajoy Majumdar and Zhen Li. 2018. Metacat: Making Big Data Discoverable and Meaningful at Netflix. <https://netflixtechblog.com/metacat-making-big-data-discoverable-and-meaningful-at-netflix-56fb36a53520>. Accessed: July 8, 2025.
- [45] C. Mohan. 1990. ARIES/KVL: a key-value locking method for concurrency control of multiaction transactions operating on B-tree indexes. In *Proceedings of the Sixteenth International Conference on Very Large Databases (Brisbane, Australia)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 392–405.
- [46] MongoDB. 2025. MongoDB. <https://www.mongodb.com/>. Accessed: July 8, 2025.
- [47] Ashish Motivala. 2018. How FoundationDB Powers Snowflake Metadata Forward. <https://www.snowflake.com/blog/how-foundationdb-powers-snowflake-metadata-forward/>. Accessed: July 8, 2025.
- [48] MySQL. 2016. MySQL 8: Improvements to information_schema. https://dev.mysql.com/blog-archive/mysql-8-0-improvements-to-information_schema/. Accessed: July 8, 2025.
- [49] NebbulaGraph. 2025. NebbulaGraph. <https://www.nebula-graph.io/>. Accessed: July 8, 2025.
- [50] Neo4j. 2025. Cypher Manual. <https://neo4j.com/docs/cypher-manual/>. Accessed: July 8, 2025.
- [51] Neo4j. 2025. Neo4j. <https://neo4j.com/>. Accessed: July 8, 2025.
- [52] Project Nessie. 2025. Project Nessie: Transactional Catalog for Data Lakes with Git-like semantics. <https://projectnessie.org/>. Accessed: July 8, 2025.
- [53] Netflix. 2025. Metacat. <https://github.com/Netflix/metacat>. Accessed: July 8, 2025.
- [54] Thomas Neumann, Tobias Mühlbauer, and Alfons Kemper. 2015. Fast Serializable Multi-Version Concurrency Control for Main-Memory Database Systems. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (Melbourne, Victoria, Australia) (SIGMOD ’15)*. Association for Computing Machinery, New York, NY, USA, 677–689.

- [55] Anton Okolnychiy, Chao Sun, Kazuyuki Tanimura, Russell Spitzer, Ryan Blue, SzeHon Ho, Yufei Gu, Vishwanath Lakkundi, and DB Tsai. 2024. Petabyte-Scale Row-Level Operations in Data Lakehouses. *Proc. VLDB Endow.* 17, 12 (Aug. 2024), 4159–4172.
- [56] Patrick O’Neil, Elizabeth O’Neil, Shankar Pal, Istvan Cseri, Gideon Schaller, and Nigel Westbury. 2004. ORDPATHs: insert-friendly XML node labels. In *Proceedings of the 2004 ACM SIGMOD International Conference on Management of Data* (Paris, France) (SIGMOD ’04). Association for Computing Machinery, New York, NY, USA, 903–908.
- [57] Meta Platforms. 2022. RocksDB. <https://rocksdb.org/>. Accessed: July 8, 2025.
- [58] PostgreSQL. 2025. Chapter 53. System Catalogs. <https://www.postgresql.org/docs/current/catalogs.html>. Accessed: July 8, 2025.
- [59] Sumedh Sakdeo. 2024. Open Sourcing OpenHouse: A Control Plane for Managing Tables in a Data Lakehouse. <https://www.linkedin.com/blog/engineering/open-source/open-sourcing-openhouse>. Accessed: July 8, 2025.
- [60] Denny Lee Scott Sandre and Mariusz Kryński. 2022. Multi-cluster writes to Delta Lake Storage in S3. <https://delta.io/blog/2022-05-18-multi-cluster-writes-to-delta-lake-storage-in-s3/>. Accessed: July 8, 2025.
- [61] Amazon Web Services. 2025. AWS Glue concepts. <https://docs.aws.amazon.com/glue/latest/dg/components-key-concepts.html>. Accessed: July 8, 2025.
- [62] Raghav Sethi, Martin Traverso, Dain Sundstrom, David Phillips, Wenlei Xie, Yutian Sun, Nezih Yegitbasi, Haozhun Jin, Eric Hwang, Nileema Shingte, and Christopher Berner. 2019. Presto: SQL on Everything. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. 1802–1813.
- [63] Jayavel Shanmugasundaram, Eugene Shekita, Jerry Kiernan, Rajasekar Krishnamurthy, Efstratios Viglas, Jeffrey Naughton, and Igor Tatarinov. 2001. A general technique for querying XML documents using a relational database system. *ACM Sigmod Record* 30, 3 (2001), 20–26.
- [64] Pradeep Shetty, Richard Spillane, Ravikant Malpani, Binesh Andrews, Justin Seyster, and Erez Zadok. 2013. Building workload-independent storage with VTrees. In *Proceedings of the 11th USENIX Conference on File and Storage Technologies* (San Jose, CA) (FAST’13). USENIX Association, USA, 17–30.
- [65] Snowflake. 2023. How Snowflake acquires resource locks while execution DML queries. <https://community.snowflake.com/s/article/Understanding-Lock-Granularity-in-Snowflake>. Accessed: July 8, 2025.
- [66] Snowflake. 2024. Introducing Polaris Catalog: An Open Source Catalog for Apache Iceberg. <https://www.snowflake.com/en/blog/introducing-polaris-catalog/>. Accessed: July 8, 2025.
- [67] Snowflake. 2024. Micro-partitions & Data Clustering. <https://docs.snowflake.com/en/user-guide/tables-clustering-micropartitions>. Accessed: July 8, 2025.
- [68] Igor Tatarinov, Stratis D. Viglas, Kevin Beyer, Jayavel Shanmugasundaram, Eugene Shekita, and Chun Zhang. 2002. Storing and querying ordered XML using a relational database system (SIGMOD ’02). Association for Computing Machinery, New York, NY, USA, 204–215.
- [69] Scott Sandre Tathagata Das, Allison Portis, Susan Pierce, and Carly Akerly. 2024. Delta Lake 4.0 Preview. <https://delta.io/blog/delta-lake-4-0/>. Accessed: July 8, 2025.
- [70] Daniel ten Wolde, Tavneet Singh, Gábor Szárnyas, and Peter Boncz. 2023. DuckPGQ: Efficient property graph queries in an analytical RDBMS. In *Proceedings of the Conference on Innovative Data Systems Research*.
- [71] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Ning Zhang, Suresh Antony, Hao Liu, and Raghotham Murthy. 2010. Hive - a petabyte scale data warehouse using Hadoop. In *2010 IEEE 26th International Conference on Data Engineering (ICDE 2010)*. 996–1005.
- [72] TigerGraph. 2025. GSQL Language. <https://docs.tigergraph.com/gsql-ref/4.2/querying/>. Accessed: July 8, 2025.
- [73] Win Vahlkamp. 2024. Catalog ONTAP data to unleash its power. <https://www.netapp.com/blog/ontap-business-technical-catalogs/>. Accessed: July 8, 2025.
- [74] Ben Vandiver, Shreya Prasad, Pratibha Rana, Eden Zik, Amin Saeidi, Pratyush Parimal, Styliani Pantela, and Jaimin Dave. 2018. Eon Mode: Bringing the Vertica Columnar Database to the Cloud. In *Proceedings of the 2018 International Conference on Management of Data* (Houston, TX, USA) (SIGMOD ’18). Association for Computing Machinery, New York, NY, USA, 797–809.
- [75] Alexandre Verbitski, Anurag Gupta, Debanjan Saha, Murali Brahmadesam, Kamal Gupta, Raman Mittal, Sailesh Krishnamurthy, Sandor Maurice, Tengiz Kharatishvili, and Xiaofeng Bao. 2017. Amazon Aurora: Design Considerations for High Throughput Cloud-Native Relational Databases. In *Proceedings of the 2017 ACM International Conference on Management of Data* (Chicago, Illinois, USA) (SIGMOD ’17). Association for Computing Machinery, New York, NY, USA, 1041–1052.
- [76] W3C. 2017. XQuery 3.1. <https://www.w3.org/TR/xquery-31/>. Accessed: July 8, 2025.
- [77] Yingjun Wu, Joy Arulraj, Jiexi Lin, Ran Xian, and Andrew Pavlo. 2017. An empirical evaluation of in-memory multi-version concurrency control. *Proc. VLDB Endow.* 10, 7 (March 2017), 781–792.
- [78] Jingyu Zhou, Meng Xu, Alexander Shraer, Bala Namasivayam, Alex Miller, Evan Tschannen, Steve Atherton, Andrew J. Beamon, Rusty Sears, John Leach, Dave Rosenthal, Xin Dong, Will Wilson, Ben Collins, David Scherer, Alec Grieser, Young Liu, Alvin Moore, Bhaskar Muppana, Xiaoge Su, and Vishesh Yadav. 2021. FoundationDB: A Distributed Unbundled Transactional Key Value Store. In *Proceedings of the 2021 International Conference on Management of Data* (Virtual Event, China) (SIGMOD ’21). Association for Computing Machinery, New York, NY, USA, 2653–2666.
- [79] Xiaowei Zhu, Guanyu Feng, Marco Serafini, Xiaosong Ma, Jiping Yu, Lei Xie, Ashraf Aboulmaga, and Wenguang Chen. 2020. LiveGraph: a transactional graph storage system with purely sequential adjacency list scans. *Proc. VLDB Endow.* 13, 7 (March 2020), 1020–1034.