

High-throughput Publish/Subscribe on top of LSM-based Storage

Mohiuddin Abdul Qader

Vagelis Hristidis

Department of Computer Science & Engineering
University of California Riverside, California, USA

Email: {mabdu002,vagelis}@cs.ucr.edu

Abstract—State-of-the-art publish/subscribe systems are efficient when the subscriptions are relatively static – for instance, the set of followers in Twitter – or can fit in memory. However, now-a-days, many Big Data and IoT based applications follow a highly dynamic query paradigm, where both continuous queries and data entries are in the millions and can arrive and expire rapidly. In this paper we propose and compare several publish/subscribe storage architectures, based on the popular NoSQL Log-Structured Merge Tree (LSM) storage paradigm, to support high-throughput and highly dynamic publish/subscribe systems. Our framework naturally supports subscriptions on both historic and future streaming data, and generate instant notifications. We also extend our framework to efficiently support self-joining subscriptions, where streaming pub/sub records join with past pub/sub entries. Further, we show how hierarchical attributes, such as concept ontologies, can be efficiently supported; for example, a publication’s topic is “politics” whereas a subscription’s topic is “US politics.” We implemented and experimentally evaluated our methods on the popular LSM-based LevelDB system, using real datasets, for simple match and self-joining subscriptions on both flat and hierarchical attributes. Our results show that our approaches achieve significantly higher throughput compared to state-of-the-art baselines.

Index Terms—Log-Structured Merge Tree, LevelDB, Publish/Subscribe, Self-join subscription, Dewey, Big Data, Internet of Things, Continuous Lookup Queries



1 INTRODUCTION

In this age of big data and Internet of Things (IoT), large amounts of data are generated, stored, and used by a diverse set of entities – devices, vehicles, buildings, software, and sensors. It is challenging to efficiently ingest, manage, read and deliver the generated data to millions of users or entities in real time. Publish/Subscribe systems are used in many applications, such as social networks, messaging systems, and traffic alerting systems.

As a running example application, consider users who are driving and subscribe to nearby traffic or other incidents (accident, crime, roadwork, fire, natural disaster, protest etc). Every time a user moves to a new location (i.e., a geospatial cell) they need to subscribe to events/publications in the new location for a time duration starting from the near past to the near future – e.g., to know what happened in the last one minute and what will happen in the next one minute until the user moves to another cell. These subscriptions are *highly dynamic* as they come and go every few seconds. At the same time, users are publishing incidents. As millions of users are moving and subscribing to events, these large streams of subscriptions and publications must be stored and managed efficiently. As another application, an airplane continuously queries for data in its path such as turbulence, wind, air pressure, etc.

Challenges and requirements A key component of a Publish/Subscribe system is its storage module, which stores both the subscriptions and the publications. As described in detail in Section 2, the storage modules of existing Publish/Subscribe systems have several *limitations*, which make them inadequate for modern IoT applications. First, the number of subscriptions (continuous queries) is assumed to be relatively small to fit in main memory, which may not always be true. Second, the subscriptions

are viewed as relatively static, for example, a user infrequently modifies her following list in Twitter. This is not the case in applications such as traffic alerting or aviation where vehicles or airplanes subscribe to different cells of interest every few seconds. Third, most of the systems support only queries on future data. In contrast, an airplane may need to know the conditions in an area in the last few minutes and the next few seconds, so a combination of past and future data may be desired. The publications may also have an expiration time; for example, a snow storm warning might have a certain time limit.

A *baseline* solution to realize a Publish/Subscribe system, which we experimentally evaluate, is to maintain a list of subscriptions (queries) and periodically submit these queries on the publications database. A similar approach is currently used in the AsterixDB BAD system [1], where such subscriptions are referred as repetitive channels. A drawback with this approach is that one cannot get notification at the exact time of an incident. Also, this solution wastes resources as the same query will keep getting submitted even if no new matching publications exist. Finally, some publications with short expiration time may be completely missed.

To summarize, our goal is to design and build a Publish/Subscribe storage system with the following properties:

- 1) Scale to millions of dynamically changing subscriptions and publications per minute, that is, both subscriptions and publications arrive and expire at a rapid pace.
- 2) After a new publication, immediately identify and notify matching subscribers (as in traditional database triggers, discussed in Section 2), that is, not follow a periodic check paradigm.

- 3) Subscriptions or publications may have validity time periods. Subscriptions may request past data in addition to future data.
- 4) The subscriber should assume that all the matching publications should reach her, that is, there is no data loss which may occur with periodic check systems.

Proposed work To support these properties we propose an efficient storage framework on top of LSM-based NoSQL databases. We argue that NoSQL databases – such as Cassandra [2], Volde-mort [3], AsterixDB [4], MongoDB [5] and LevelDB [6] – provide the right primitives –fast write throughput and fast lookups on the primary key – on top of which we can design an efficient Publish/Subscribe storage system that satisfies the above properties.

We first propose an approach based on two separate databases (*TwoDB*), one for subscriptions (queries) and one for publications (data records). This approach is simpler to implement, and we show that it is efficient compared to a baseline based on repetitive queries (*RepQueries*) mentioned above, when queries (subscriptions) simply perform lookups, e.g., return all events (publications) inside my current cell id.

However, this two-database approach suffers from a high rate of random disk accesses when the subscriptions not only match publications but also past subscriptions – e.g., return publications that have not been already returned by a past subscription. To efficiently answer such complex subscriptions, which we refer as *self-joining subscriptions*, we propose a novel *DualDB* approach where both queries and data entries are stored in the same keyspace of the same database, which leads to much fewer random accesses.

We also propose techniques to extend these storage architectures to handle hierarchical attributes. For example, a publication may specify a large affected area (higher level bounding rectangle in a spatial hierarchy), whereas a subscription may only be interested in a smaller spatial cell. Similarly, a publication may specify a specific topic like “soccer”, whereas a subscription may specify a more general topic like “sports.”

In addition to the repetitive queries baseline, we compare our approaches to a state-of-the-art pub/sub system, Padres [7]. We chose Padres because it is open source and is easy to modify to fit our use-case. Padres supports querying on historic/past data, and uses the PostgreSQL database system. We show that Padres does not scale well with the number of subscriptions. Specifically, the system runs out of memory and crashes very quickly, because it uses an in-memory module to manage the subscriptions.

To summarize, we make the following contributions in this paper:

- We propose and implement two novel approaches built on top of on LSM-based storage systems, *DualDB* and *TwoDB*, to efficiently support high rates of dynamic subscriptions and publications (Section 4).
- We implemented our methods on LevelDB and conducted extensive experiments with various workloads with different subscription to publication ratios. For that, we extended LevelDB to handle lists of values per key (Section 6). Our experiments show that both approaches perform up to 1000% faster than baseline *RepQueries* (repetitive queries) and up to 3000% faster than a state-of-the-art pub/sub system Padres (Section 7).
- In addition to simple lookup subscriptions that return matching publications, we implemented methods to sup-

port subscription queries that join streaming publications/subscriptions with existing data/queries, which we refer as *self-joining subscriptions* (Section 5). We show that for such subscriptions *DualDB* performs better (up to 20%) compared to the *TwoDB* approach. Interestingly, *DualDB* performs 5% better than *TwoDB* even for simple match subscriptions (Section 7).

- To support hierarchical attributes, we propose a Dewey encoded representation of topics/cells and an efficient querying algorithm (Section 5). We show that our approach performs significantly better than reasonable baselines that handle multi-granular data (Section 7).
- We have published open-source version of our *DualDB* and *TwoDB* implementations on top of LevelDB [8].

The rest of the paper is organized as follows. Section 2 reviews the related work and background. Then, Section 3 presents the framework of the system, Finally, Section 8 concludes and presents future directions.

2 RELATED WORK AND BACKGROUND

2.1 Related Work

2.1.1 Publish/Subscribe Systems

Most of the academic work on pub/sub mainly has studied how to efficiently route the message through the distributed pub/sub network. Instead, in this paper we propose an efficient data storage management system for pub/sub systems. There exist many variations of pub/sub systems supporting content or topic based subscription [9] [10]. Very few works studied the storage architecture of pub/sub systems. Padres [7] is a popular open-source pub/sub system, which allow continuous lookup queries. It supports subscriptions on future and past data, using a PostgreSQL database inside each broker [11]. Pub/Sub systems integrated in relational database systems have also been studied [12]. However, they do not scale with the number of subscriptions as we show in our experiments.

AsterixDB BAD [1] supports complex continuous queries, where queries are executed repetitively on top of the LSM-based AsterixDB database. Due to its repetitive nature, it does not provide instant notification. On the other hand, BAD supports a richer set of subscription query types than this work, where we require that subscriptions and publications match based on a key attribute condition. XML Based pub/sub systems also use a relational database for content based subscriptions [13]. Elaps [14] is a Location aware pub/sub system focusing on efficient processing of continuous moving range queries against dynamic event streams. Their work focuses on the query optimization rather than an efficient storage architecture. Our focus is on millions of dynamically arriving and expiring subscriptions and publications, which requires an efficient specialized storage framework. A preliminary version of our work was recently published as a short paper [15], which does not include the *TwoDB* approach, hierarchical attributes’ support or self-joining subscriptions.

2.1.2 Triggers and Continuous Queries

Continuous queries on databases may be implemented using triggers [16] [17]. The TriggerMan project [18] proposes a method for implementing a scalable trigger system based on the assumption that some triggers may share a common structure. It uses a special selection predicate index and an in-memory trigger cache

to achieve scalability. However, they are only able to handle a very small number of triggers on a table [19], whereas we want our continuous queries to scale to millions. Further, triggers have a relatively high creation and deletion cost, which makes them inappropriate for dynamically changing subscriptions.

NiagraCQ [19] proposed techniques to group continuous queries with similar structure, to share common computation. TelegraphCQ [20] is another system that handles streams of continuous queries over high-volume, highly-variable data streams. Complex continuous queries on data streams have also been studied [21] [22]. Kafka [23] is a popular stream processing platform. But our storage systems also support subscriptions on historic data where both subscriptions and publications dynamically expire, which data streams cannot support. Gemfire [24] is one of the few NoSQL database systems where continuous queries are supported. They maintain a separate in-memory processing framework, which maintains subscriber channels and communicates with the storage. InfluxDB [25] supports continuous queries on a time series database and thus its application is limited. Microsoft supports efficient spatial continuous queries on SQL Server [26]. However, these works assume that the queries fit in memory and are relatively static, that is, they do not scale to arbitrary numbers of queries nor to rapidly added and expiring queries. Further, these models generally only support “future-only” queries, that is, queries that only return future data items.

2.1.3 Join-Indices

Our work to combine two databases into one dual database is also related to join indices, such as Oracle’s Bitmap join index [27], which creates a join index on attributes from two tables to facilitate faster joins. However, these are separate and specialized data structures that are generally expensive to maintain, in contrast to our proposed solution that does not add any redundant indexes and is efficient to maintain.

2.2 Background

2.2.1 LSM tree

An LSM tree generally consists of an in-memory component (a.k.a. Memtable) and several immutable on-disk components (a.k.a. SSTables). Each component consists of several data files and each data file consists of several data blocks. As depicted in Figure 1, all writes go to in-memory component (C0) first and then flush into the first disk component once the in-memory data is over the size limit of C0, and so on. The on disk components normally increase in size as shown in Figure 1 from C1 to CL. A background process (compaction) periodically merges the smaller components to larger components as the data grows. Deletion on LSM is achieved by writing a tombstone of the deleted entry to C0, which will later propagate to the component that stores this entry.

LSM storage is highly optimized for writes as a write only needs to update the in-memory component C0. This append-only style means that an LSM tree could contain different versions of the same entry (different key-value pairs with the same key) in different components. A read (GET) on an LSM tree starts from C0 and then goes to C1, C2 and so on until the desired entry is found, which makes sure the newest (valid) version of an entry will be returned. Hence, reads are slower than writes. The older versions of an entry (either updated or deleted) are obsolete and will be discarded during the merge (compaction) process.




Fig. 1: LSM tree components.

Because of LSM trees’ good performance on handling high write throughput, they are commonly used in NoSQL databases such as BigTable [28], HBase [29], Cassandra [2], LevelDB [6], RocksDB [30] and AsterixDB [4].

2.2.2 LevelDB Storage Architecture

As we will implement our proposed approaches on top of LevelDB, we present some storage details of LevelDB as background for ease of understanding. The SSTables of LevelDB are organized into a sequence of levels where Level- $(i + 1)$ is 10 times larger than level- i in LevelDB. Each level (component) consists of a set of disk files (SSTables), each having a size around 2 MBs. The SSTables in the same level may not contain the same key (except level-0¹). As described in the LSM storage, the lower level (smaller i) in LevelDB will contain the more recent key-value pairs. Recent Cassandra versions support this level merging (compaction) option used in LevelDB.

Data Block 1
Data Block 2
...
Filter Meta Blocks
Bloom filter for primary keys in data block 1,
Bloom filter for primary keys in data block 2,
...
Stats Meta Blocks
Other optional Meta blocks
Meta Index Block
Data Index Block
Footer

Fig. 2: LevelDB SSTable Structure.

Further, LevelDB partitions its SSTables into data blocks (normally tens of KB in size). The key-value pairs are stored in the *data blocks* sorted by their key. Meta blocks contain meta information of data blocks. Index blocks store the addresses of meta blocks and data blocks.

Compaction Policy: Similar to any LSM architecture, if any level L exceeds its limit, they compact it in a background thread. Level-0 files are just flushed memtables and that’s why they have overlapping key-ranges. But for all other level, each file contains non-overlapping key-ranges. Their policy picks a file from level L and all overlapping files from the next level $L+1$ to merge. This selection for a particular level rotates through the key space of the files (i.e. Round-Robin). Compaction also drops overwritten/deleted values which ensures one unique key per level.

TABLE 1: Set of Operations for Pub/Sub Storage Framework

Operation	Description
SUBSCRIBE (<i>ID</i> , <i>Subscription-JSON</i> , <i>SelfJoinFlag</i>)	Store the subscription with <i>ID</i> as primary key; return matching and valid past publications, published in $[T_{min}, T_{max}]$ time interval; if <i>SelfJoinFlag</i> is true, return matching list of subscriptions (see Section 5.1).
PUBLISH (<i>ID</i> , <i>Publication-JSON</i> , <i>SelfJoinFlag</i>)	store publication with <i>ID</i> as primary key; return matching and valid subscribers who subscribed to this <i>ID</i> (i.e. topic/region); if <i>SelfJoinFlag</i> is true, also return matching publications .

3 FRAMEWORK

To support pub/sub operations on top of a NoSQL data store, we define a basic API as shown in Table 1. To illustrate the functionality of this API, consider the motivating application described in Section , where Subscriptions S (API call *SUBSCRIBE* (*ID*, *Subscription-JSON*)) are generated by mobile users driving in their vehicles, who are interested to know about recent events, such as accidents or weather changes, close to their current location L_S . ‘‘Recent’’ may refer to events that happened in the time interval $[T_{min} = T_S - 10 \text{ sec}, T_{max} = T_S + 10 \text{ sec}]$, where T_S is the query (subscription) time. Note that the time interval associated with each subscription S can include both past and future time ranges. If the intervals only contained future times, no storage would be needed for the publications. Whenever there is a new publication (API call *PUBLISH* (*ID*, *Publication-JSON*)), i.e. an event occurred, we look for all the stored subscriptions to notify them if the *ID* (cell-id for moving objects or topic-id for topic-based pub/sub systems) of the publication matches that of the subscription. We emphasize that subscription time intervals span both the past and the future, we need a storage framework to store the subscriptions as well as the publications.

ID is the key that joins subscriptions and publications. *Subscription-JSON* and *Publication-JSON* are the JSON-formatted values in our Key-Value store. They hold the attributes shown in Table 2. Note that subscriptions may also specify additional conditions, such as keyword matching, for instance, return data close to me that contain the term ‘‘accident.’’ Such conditions can be supported by a postprocessing (filtering) layer on top of the location-constrained or topic-constrained results.

TABLE 2: Set of Attributes and Terminologies

<i>ID</i>	cell-id/topic-id/product-id
T_p, T_s	Execution time of a publication and subscription
T_{min}, T_{max}	The time interval in a subscription query. T_{max} is expiry time for subscription and publication.
<i>Sid, Pid</i>	Subscription and publication Identifier
<i>Uid</i>	User Identifier (One user can submit multiple subscriptions)
<i>Desc</i>	Text Description field in publication
<i>Subscription-JSON</i>	Body of a subscription in JSON $\{T_{min}, T_{max}, T_s, Sid, Uid\}$.
<i>Publication-JSON</i>	Body of a publication in JSON $\{T_p, T_{max}, Pid, Desc\}$

4 PROPOSED APPROACHES

We propose and implement several approaches to efficiently realize the API defined in Section 3, while providing instant response

1. which is C1 in Figure 1 as LevelDB does not number the memory component.

times to publications. Our first approach (*TwoDB*) maintains two databases for holding subscriptions and data records/publications whereas our second approach (*DualDB*) maintains a more optimized single database holding both datasets. To compare our approaches with reasonable baseline, we also propose a baseline approach *RepQueries* which perform queries via repetitive channel which can not provide instant response. For all of the approaches, we use the base storage described in following Section 6 which maintains a lazily updated value list instead of a single value associated with a key (i.e. instead of single (key,value) pair we have fragmented value list for a key throughout the storage).

4.1 Two-Database Approach (TwoDB)

As discussed in Section 2, previous works assume that the subscriptions are stored in memory, which is not realistic in our scenario. Figures 3a and 3b show the data flow for a new subscription (*SUBSCRIBE*) and a new publication of events/topics (*PUBLISH*), respectively, for a two-database algorithm, where separate databases (column families in Cassandra’s terminology) are maintained for, respectively, the queries and the published events/topics.

For every new continuous query, we must, in parallel, insert to the subscription database *SubscriptionDB* and query the data storage *PublicationDB* for matching events/topics. The former is needed when the T_{max} of the subscription is greater than current query time T_S while the latter is only needed if the T_{min} of the query is less than T_S . The query to the publication database *GetData*(*cell-id*, T_S , T_{min}) returns the list of events which is published within time interval $[T_{min}, T_S]$ and are valid (i.e.did not expire) on that time. The procedure of *GetData* operation is presented in Algorithm 1. Since the postings list could be scattered in different levels, *GetData* needs to merge them to get a complete list. For this, it checks the Memtable and then the SSTables, and moves down in the storage hierarchy one level at a time.

Algorithm 1 *GetData*(*ID*, T_S , T_{min}) operation

```

1: ListPB  $\leftarrow \emptyset$ 
   // Starts from Memtable (C0) and then moves to SSTable, C1 is LevelDB’s level-0 SSTable.
2: for j from 0 to L do
3:   if ID is not in Cj then
4:     NEXT
5:   List of Publications P  $\leftarrow$  the value of key ID in Cj
6:   for Publication in P do
7:     T  $\leftarrow$  Publication( $T_p$ )
8:     Texp  $\leftarrow$  Publication( $T_{max}$ )
9:     if T < TS and T > Tmin and T < Texp then
10:      ListPB.add(Publication)
11:   if T < Tmin then
12:     return ListPB
13: return ListPB

```

For each new publication of events/topics, we must, in parallel, check if an active subscription query matches it, and also insert it into the publication database. The matching function *GetSubscribers*(*ID*, T_p) returns a list of subscribers who subscribed to some events/topics on same *ID* and the query is still valid and if the publication time T_p is within time interval $[T_{min}, T_{max}]$ of the query. The procedure of *GetSubscribers* operation presented in Algorithm 2 is similar to *GetData* operation moving down level by level in storage to look for matching subscriptions.

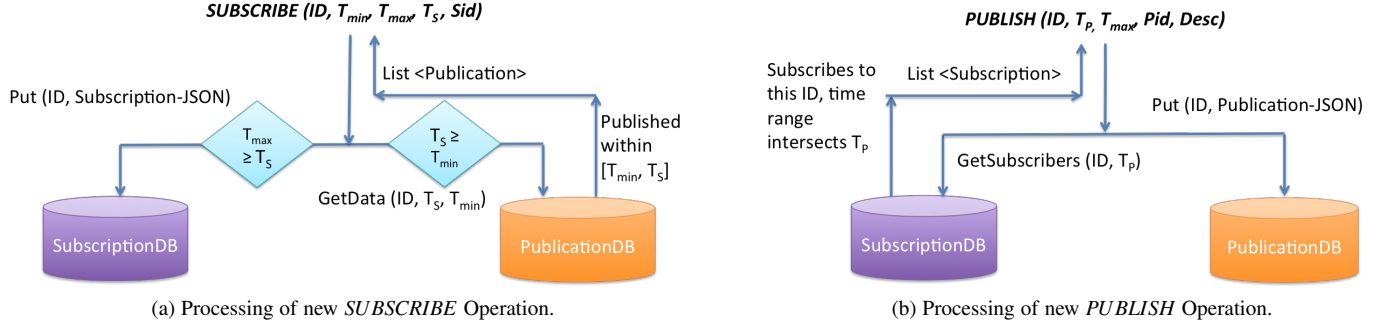


Fig. 3: TwoDB Approach

Algorithm 2 GetSubscribers(ID, T_P) operation

```

1: ListSB  $\leftarrow \emptyset$ 
   // Starts from Memtable (C0) and then moves to SSTable, C1 is LevelDB's level-0 SSTable.
2: for j from 0 to L do
3:   if ID is not in Cj then
4:     NEXT
5:   List of Subscriptions S  $\leftarrow$  the value of key ID in Cj
6:   for Subscription in S do
7:     Texp  $\leftarrow$  Subscription(Tmax)
8:     Tmin  $\leftarrow$  Subscription(Tmin)
9:     if Tp < Texp and Tp > Tmin then
10:      ListSB.add(Subscription)
11: return ListSB

```

4.1.1 Pruning search in SSTables

To answer a subscription that has a T_{min} in the past, we take advantage of an important characteristic of level-based LSM systems. In LevelDB, each level is ordered by time and older records go into lower levels. In our storage, each level can contain at most one list associated with a key. That means the fragmented lists associated with a key in different levels from top to bottom are ordered by time. Here, the publications in each lists are also ordered by execution time T_p . So for each *SUBSCRIBE* operation, when we perform *GetData* operation to look for all matching publications from past (T_{min}), we check for associated lists going level by level. For each matching list, we parse the list to check each publication record one by one. If we find one of them is older than T_{min} , we can safely assume that we neither need to go any further down the level nor further right to this current list, and hence we stop the search. Lines 11-12 in algorithm 1 show this optimization.

4.1.2 Cleaning up Invalid Records in Compaction

Our proposed system assumes that both the publication data and subscriptions can be highly dynamic and both may have an attribute in their value (T_{max}), which denotes their expiration time. The storage management system should efficiently remove the expired items from its storage. Our proposed storage is built on top of LSM based NoSQL databases which perform compaction in background. During compaction, we pick couple of SSTable files and compact them into a new set of files. During this time we sort and merge the list of publications or subscriptions associated with their ID, and check their expiration time against the current time to discard them if found invalid.

4.2 Single Database Approach (DualDB)

A key observation is that if *continuous subscription queries and data publications could be stored in the same key space, then a single database (in LevelDB terminology) could store both of them*. Then query and data insertions would only need to access a single database, which could reduce the number of disk accesses and more importantly improve the caching efficiency. Further, the compaction cost might be decreased, as we only have to compact a single database.

We propose to study the properties and performance of using a single database, which we call a *dual database (DualDB)*. The key idea is that the key-value data organization must be modified to accommodate a list of subscription and publication items in the value. That is, for a given ID (cell-id or topic-id) in our example, both the list of subscriptions and publications will be stored in the posting list of this ID. As LevelDB does not allow same key inside a component, we need to change the storage architecture to allow the disk and memory components holding at most two lists associated with the same key/cell-id. We assign a bit with the primary key (ID^S, ID^P) which will state whether it's a list of publications or subscriptions. Figures 4a and 4b show, respectively, how new query and new publication of events are processed in the proposed *DualDB*.

We can see that the key difference with the TwoDB approach is using single database against two and all the operations *Put*, *GetData* and *GetSubscribers* are performed with same *ID* with a assigned flag stating whether its a query or data. Here both *GetData* and *GetSubscribers* procedure follows the same algorithms 1 and 2 of TwoDB approach respectively. The optimization techniques for lookups described in Section 4.1.1 can also be applied for *DualDB* in similar fashion. *DualDB* also similarly does not allow the database to grow indefinitely and clean up the invalid records during compaction in background as described in Section 4.1.2.

As we are combining two databases into one database containing two lists, the storage components are changed accordingly to hold any such two heterogeneous data. Figures 5a and 5b show, respectively, snapshots of the entries in the corresponding databases for the Two-Database approach and the *DualDB* approach. In Figure 5a, we can see that the TwoDB approach holds posting lists of subscriptions (S1,S2,...) and publications (P2,P1,...) in two different databases and each list associated with a ID (C1,C2 etc.) is fragmented throughout different levels. Figure 5b shows the *DualDB* is holding lists of subscriptions and publications in same components with same key ID which holds a flag ($C1^S, C1^P$) that distinguishes between them. As in each component the records are

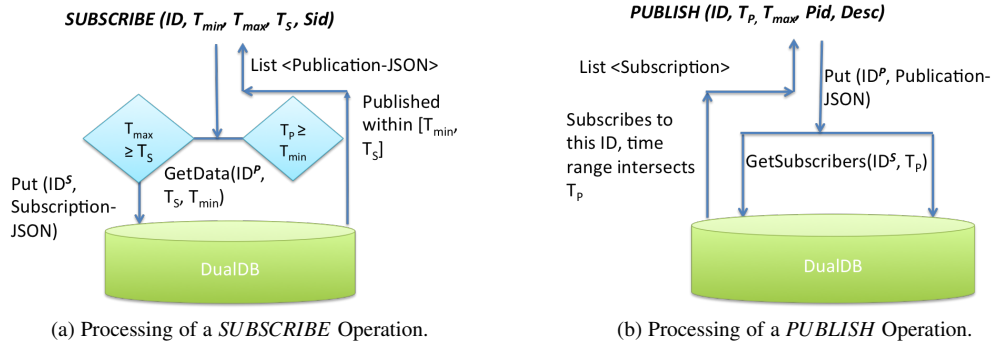


Fig. 4: DualDB Approach

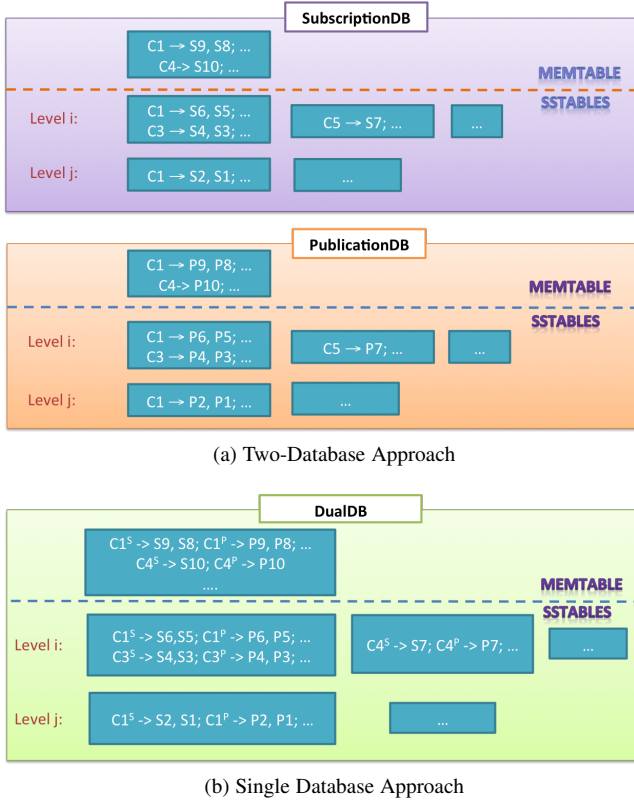


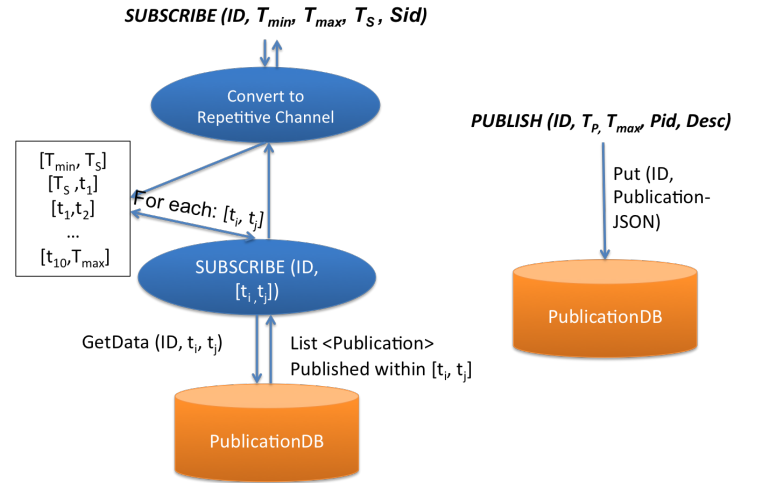
Fig. 5: Snapshot of storage components for our LevelDB-style LSM database for pub/sub System.

sorted by key, these two lists will always be co-located with each other. This should allow improvements in both LevelDB memory cache and the operating system page cache. These figures assume a leveled storage architecture (as in LevelDB and Cassandra); a stack-based architecture can be handled similarly.

4.3 Repetitive Queries Baseline (*RepQueries*)

Both of our proposed approaches discussed above support dynamic and massive amount of continuous queries on a storage framework holding lazily updated lists, which returns matching subscribers whenever an topic/event is published. This feature enables the system to support instance response to the user via continuous channel. To compare our systems, we also prepared a baseline *RepQueries* which relies on repetitive channel and can not guarantee instance response. This approach is used in many

pub-sub systems where they use a broker management as a middleware between database and the users. Here the user/broker is responsible for submitting the query repetitively to the publication database to find out the matching events/topics. Note that the user/broker is now responsible to detect duplicate data if some consecutive time range in the query intersects with each other resulting in return of duplicate publications. Also as the events are highly dynamic, some of them might expire in between the repetitive queries and user can potentially face missing data.

Fig. 6: **RepQueries Approach:** Processing of *SUBSCRIBE* and *PUBLISH* Operation

This approach only requires a database of publications. As periodic queries are issued from client application, and there is no need for instant response, we do not need to store the queries for future publications. To compare baseline with our proposed approaches, we converted each *SUBSCRIBE* operation into 10 repetitive query where we divide the time interval $[T_{min}, T_{max}]$ into 10 intervals. First it will issue a historic query on past data with time interval $[T_{min}, T_Q]$. Then it repetitively issue query and perform *GetData* operation after every $(T_{max} - T_Q)/10$ sec with time interval $[T_S, t_1]$, $[T_S, t_1]$, ..., and $[t_{10}, T_{max}]$. We consider this as equivalent to one *SUBSCRIBE* operation in instance-response approaches. The *GetData* operation follows the same procedure on publication database as algorithm 1.

Here the *PUBLISH* operation does not issue any read, it only writes the new publication to the database. So *PUBLISH* will be very fast and *SUBSCRIBE* will be very slow for *RepQueries* approach than our other two approaches. Figures 6 shows, how

a new subscription and a new publication of events are processed in the *RepQueries*.

5 EXTEND TO COMPLEX SUBSCRIPTIONS

5.1 Self-Joining Subscriptions

As discussed in Section 4, our *TwoDB* and *DualDB* approaches can answer simple queries where either a matching lists of subscribers (matched by topics or spatially and temporally) are returned for a new publication or matching lists of past publications are returned with a new subscription. A PUBLISH operation does not need to browse lists of publications and a SUBSCRIBE operation does not need to browse the lists of subscriptions. In this section, we handle more complex subscriptions where we also need to return the list of matching past events (publications) for a PUBLISH operation and the list of matching subscriptions for SUBSCRIBE operation.

Motivating Examples. Consider mobile users moving and subscribing to events on a region or cell. Imagine the scenario where a mobile user enters cell c_1 , then c_2 and then c_1 again, all within 10 seconds, and all subscriptions have a time interval of ± 10 sec. A desirable property for many applications is that data already returned when the user was in c_1 the first time should not be returned again during the second entry. For that, the system must efficiently identify what data has already been returned to a user. This is a classic use-case for using self-join subscriptions where for a new subscription, we need to access not only the lists of past matching publications, but also access the list of matching past subscriptions from the user in order to filter these results. As another application, consider a user who subscribes to a sensor temperature (or to a stock price) and wants to be notified if there has been a sudden change/drop in temperature/price. In this case, for a new publication, we need to traverse both the publications list to check past temperatures/prices in this location/stock, and the subscriptions lists to know who subscribes here.

That is, for these operations, a subscription (publication) must join with past subscriptions (publications), hence the name *self-joining* subscriptions. Our existing storage framework gives us flexibility to support such subscriptions efficiently. For both our approaches, if *SelfJoinFlag* is set we assume that it is a self-joining operation.

For a self-join subscription, we write our subscription and/or publication to the storage and retrieve a list of events and a list of queries from storage by issuing a *GetDataANDSubscribers*($ID, T_{P/S}, T_{min}$) operation. Here $T_{P/S}$ means execution time of either publication or subscription.

Algorithm 3 *GetDataANDSubscribers* ($ID, T_{P/S}, T_{min}$) operation on *TwoDB*

```

1:  $ListPB \leftarrow this.PublicationDB.GetData(ID, T_{P/S})$  // [Algorithm 1]
2:  $ListSB \leftarrow this.SubscriptionDB.GetSubscribers(ID, T_{P/S}, T_{min})$  //
   [Algorithm 2]
3: return  $ListPB, ListSB$ 

```

In the *TwoDB* approach, the implementation for self-join subscription is straightforward: we simply perform an extra *GetSubscribers* ($ID, T_{P/S}$) operation (Algorithm 2) on *PublicationDB* (along with *GetData* operation) to return the list of active matching subscribers for SUBSCRIBE operation and perform an extra *GetData* ($ID, T_{P/S}, T_{min}$) operation ((Algorithm 1)) on *PublicationDB* (along with *GetSubscribers* operation) to return the list of valid publications for PUBLISH operation.

Algorithm 4 *GetDataANDSubscribers*($ID, T_{P/S}, T_{min}$) operation on *DualDB*

```

1:  $ListPB \leftarrow \emptyset$ 
2:  $ID^P \leftarrow Setflag(ID, E)$ 
3:  $ListSB \leftarrow \emptyset$ 
4:  $ID^S \leftarrow Setflag(ID, Q)$ 
   // Starts from Memtable (C0) and then moves to SSTable, C1 is LevelDB's level-0
   SSTable.
5: for  $j$  from 0 to  $L$  do
6:   if  $ID^P$  is not in  $C_j$  and  $ID^S$  not in  $C_j$  then
7:     NEXT
8:    $ListSB \leftarrow$  // Apply Line 5-10 in Algorithm 2 for key  $ID^S$  .
9:    $ListPB \leftarrow$  // Apply Line 5-12 in Algorithm 1 for key  $ID^P$  .
10: return  $ListPB, ListSB$ 

```

In *DualDB*, both the lists of publications and subscriptions associated with the same ID are stored in the same disk block next to each other. We convert the *GetSubscribers* and *GetData* operations performing on same Dual database into a single operation, which means we can get the both the lists of publications and subscriptions on single operation. We first convert the key ID to ID^P and ID^S by setting appropriate bits. Then we search in our *DualDB* starting from Memtable and moving to SSTables level by level. If we find a match for either ID^P or ID^S , we fetch the resulting block from the storage components which might contain both the lists. We extract both the lists if they are available in that particular fetched block. Here we can see that one disk I/O might sufficient against two if the list of publications and subscriptions are co-located inside same storage file block.

5.1.1 Cost Analysis (*TwoDB* Vs *DualDB* Vs *RepQueries*)

Table 3 summarizes the number of disk accesses for each PUBLISH/SUBSCRIBE operation for simple and self-joining subscriptions.

For simple subscriptions, only one I/O is needed to write the subscription and the publication in SUBSCRIBE and PUBLISH operation respectively. For both the operation, we issue a read (*GetData/GetSubscribers* operation) on the database which may involve up to L (number of levels) disk accesses (because in worst case the postings list is scattered across all levels). Here, at most one disk access is required for level except level-0. As level-0 contains files with overlapping key ranges, we need to traverse all files inside level-0. Let us assume l_0 is the number of level-0 files. As both *TwoDB* and *DualDB* approach issue one write and one read for their each operation, the total cost becomes $(L + l_0) + 1$ disk accesses for both of them. *RepQueries* performs 10 repetitive queries on publication database instead of a single subscription. But *RepQueries* publish operation only writes the publication record to the database.

For Self Joining subscriptions, similarly one write is needed for each operation. But for *TwoDB* approach, we need to submit two reads in two databases of publication and subscriptions (one in *PublicationDB* and one in *SubscriptionsDB*). This makes the total cost as $2 * (L + l_0) + 1$ disk accesses for each PUBLISH or SUBSCRIBE operation. But *DualDB* approach (Algorithm 4) only issues one read in dual database and the total cost becomes $(L + l_0) + 1$ disk accesses in the worst case. Note that we assume there is only one disk access per write operation. But in practice, LevelDB suffers from high write amplification factor where same data is written multiple times when moved from one level to another. This factor is excluded from our calculation.

TABLE 3: Number of disk accesses for PUBLISH/SUBSCRIBE operation for simple vs self-joining subscriptions.

Approach:Operation	Simple Subscription	Self-Join Subscription
<i>TwoDB : SUB</i>	$(L + l_0) + 1$	$2 * (L + l_0) + 1$
<i>DualDB : SUB</i>	$(L + l_0) + 1$	$(L + l_0) + 1$
<i>RepQueries : SUB</i>	$10 * (L + l_0)$,	Not Supported
<i>TwoDB : PUB</i>	$(L + l_0) + 1$	$2 * (L + l_0) + 1$
<i>DualDB : PUB</i>	$(L + l_0) + 1$	$(L + l_0) + 1$
<i>RepQueries : PUB</i>	1	Not Supported

5.2 Support for Hierarchical Attributes

In this section, we introduce an optimization in our storage system to efficiently support multi-granular data, which can be defined by a hierarchical model. For example, we may have a spatial hierarchy (e.g. Country->State->County->City->Block), where each publication or subscription may select spatial cells of different granularity. Similarly, topics from a topic hierarchy can be selected. We propose and compare two approaches to support such hierarchical attributes in a pub/sub storage system. Note that the choice of an approach is orthogonal to the choice of a storage architecture from the ones described in Section 4.

5.2.1 Fixed Granularity Approach

In this approach, both subscriptions and publications use the lowest level keys in the hierarchy. For example, if a subscription specifies a city which has 100 blocks, then we insert 100 copies of the subscription. Similarly we handle publications. A disadvantage is that multiple entries are needed for each subscription or publication. The key advantage is the simplicity, as this can be implemented as a preprocessing step on top of any storage framework.

5.2.2 Variable Granularity Approach

In this approach, we only have one entry per subscription or publication, and the storage framework manages the matchings. We propose to use Dewey encoding [31] (incremented integer numbers separated by dots) to represent each entity in the ontology, and this Dewey id becomes the key (ID) of the publication or subscription. Let us assume a hierarchy is defined by a quad tree. Figure 7a illustrates how each node in different depth of the quad tree can be represented by a Dewey encoding. Figure 7b illustrates how we partition the space using that quad tree. During a *SUBSCRIBE/PUBLISH* operation, we write the entry (publication/subscription) to the database and issue a range query from its Dewey to the Dewey of its next sibling.

For example given the quad tree partition in Figure 7, let a publication/subscription cover area (1.4) (or represent topic (1.4) if we use topics instead of locations as ID). The fixed granularity approach will issue PUBLISH/SUBSCRIBE operation for each areas/topics in the leaf level of the subtree of (1.4) [(1.4.1), (1.4.1.1), (1.4.1.2), (1.4.1.3), (1.4.1.4), (1.4.2), (1.4.3), (1.4.4)]. But in our optimized variable granularity approach, we will issue one write for this topic/area (1.4) and issue a range query from (1.4) to (1.5).

Next, we discuss how a range query based on Dewey IDs is efficiently supported, by modifying the approaches in Algorithm 1 and 2. We iterate from upper level to lower level, and for each level when we find the file that contains the start key, retrieve the value list associated with that key, and then instead of returning from there we continue iterating through that file (or next file if necessary) until the end key. As we know that the data in our

storage files (i.e.SSTables) are lexicographically sorted by their keys (i.e. Dewey numbers) and partitioned into fixed size blocks. So given a Dewey Number representing an internal node, all the nodes descended from that node by pre-order traversal should be stored. That makes this range query very efficient.

Figure 7c represents a snapshot of an SSTable in our DualDB, where the spatial hierarchy is represented by a quad tree in Figure 7a and 7b. Here we can see that how the list of publications or subscriptions in SSTable are sorted by its key lexicographically. If we issue a range query from (1,4) to (1,5), our algorithm will look for the file which contains (1,4) in each level, retrieve the disk block K of that SSTable and retrieve the list of subscriptions/publications associated with (1,4). Then it will retrieve the next records of that disk block subsequently [(1,4,1),(1,4,1,1),...] until it reaches the node (2). It might retrieve subsequent disk blocks in order to reach (1,5)/(2). In this figure, the last record of the block K is (1.4.1.2) and we need to retrieve block $(k + 1)$. Note that similarly to our simple point lookup algorithm in Section 4.1.1, we check the time predicates of all the publications/subscriptions in each list for validity, and return early if pruning can be utilized.

Algorithm 5 GetDataInRange ($StartID^P, EndID^P, T_p, T_{min}$)

```


1: ListPB ← ∅
   // Starts from Memtable (C0) and then moves to SSTable organized into L Levels.
   // C1 is LevelDB's level-0 SSTable.
2: for  $j$  from 0 to  $L$  do
3:   ListSSTs ← getOverlappingSSTablesInsideLevel(j, StartIDP)
4:   for  $f$  in ListSSTs do
5:     Iterator ← getIterator(Cf)
6:     Iterator ← Iterator.Seek(StartIDP)
7:     while Iterator is Valid do
8:       if Iterator.Key.getFlag() is Publication then
9:         List of Publications  $P$  ← Iterator.Value
10:      ListPB ← // Apply Line 6-9 with value of  $T_p$  and  $T_{min}$  in Algorithm 1 to Populate ListPB from  $P$ 
11:      Iterator ← Iterator.Next()
12:      if Iterator.Key >=  $EndID^P$  then
13:        return ListPB
14: return ListPB

```

Algorithm 5 summarizes how a single lookup *GetData* operation (Algorithm 1) for simple subscription queries can be extended to support a range operation *GetDataInRange* based on Dewey representation. Algorithm 2, 3 and 4 are extended similarly to support range operation on hierarchical attributes for simple and self-joining subscriptions.

5.2.3 Cost Analysis (Fixed Vs Variable Granularity Approaches)

We derive in Section 5.1.1 that a single PUBLISH or SUBSCRIBE operation can cost $O(L) = L + l_0 + 1$ number of disk accesses for both simple and self join subscriptions in *DualDB*. If we define the hierarchy as quad tree and the maximum depth as d , the number of leaf nodes in the worse case will be 4^d . So the cost for one operation will be $4^d * O(L)$ as we issue one operation for all leaf nodes separately. But for the variable granularity approach, we issue one write and a range query. Each range query similarly can hit $L + l_0$ number of files in each level and for each file, we may retrieve C consecutive disk blocks. C can be computed by dividing the size of the result set by the block size. The cost for each operation will be $C * (L + l_0) + 1$ number of disk accesses.



(a) A Quad Tree and Dewey Encoding for each Node

1.1	1.2	2	
1.3	1.4 1.1	1.4 1.2	1.4.2
	1.4 1.3	1.4 1.4	
3.1	3.2.1	3.2.2	4.1
	3.2.3	3.2.4	
3.3	3.4	4.3	4.4

(b) Partitioning Space using Quad Tree and Dewey Representation in Leaf Nodes

Data Block 1	(1.3) ^s -> (S12, ...)
Data Block 2	(1.3) ^p -> (P4, P7, ...)
...	...
Data Block K	(1.4) ^s -> (S23, ...)
	(1.4) ^p -> (P9, ...)
	(1.4.1) ^s -> (S2, ...)
	(1.4.1) ^p -> (P8, ...)
	(1.4.1.1) ^s -> (S5, ...)
	(1.4.1.1) ^p -> (P12, ...)
	(1.4.1.2) ^s -> (S6, S10, ...)
	(1.4.1.2) ^p -> (P47, ...)
	...
Data Block K+1	(1.4.1.3) ^s -> (S15, ...)
	(1.4.1.3) ^p -> (P5, P2, ...)
	(1.4.1.4) ^s -> (S13, ...)
	(1.4.1.4) ^p -> (P91, ...)
	(1.4.2) ^s -> (S23, ...)
	(1.4.2) ^p -> (P8, ...)
	(1.4.3) ^s -> (S33, S1, ...)
	(1.4.3) ^p -> (P10, P3, ...)
	(1.4.4) ^s -> (S39, S7, ...)
	(1.4.4) ^p -> (P5, ...)
	(2) ^s -> (S6, S17, S24, ...)
	(2) ^p -> (P18, P6, P11, ...)
	...
Data Block N-1	
Data Block N	
Stats Block	
Filter and Meta Index Blocks	

(c) Snapshot of an SSTable in DualDB containing lists in order of Dewey ID of the Key

Fig. 7: Variable Granularity Approach by Using Dewey Encoding

TABLE 4: Number of Disk Accesses for Hierarchical Attributes

Approach	Approximate I/O Cost
Fixed Granularity	$4^d * O(L)$
Variable Granularity	$C * (L + l_0) + 1$

6 ENABLE VALUE LISTS IN LEVELDB

The standard LevelDB implementation, as any key-value store, only supports the storage of (key,value) pairs. However, to realize our pub-sub storage algorithms, we need to be able to store (key, list of values) pairs, e.g., to store all publications for a single ID. In this section we discuss how we modify LevelDB to handle lists of values instead of a single value for every key. For example, the list of events for a single cell id is stored as a value list with the cell id as key.

LevelDB holds the LSM property shown in Figure 1 where data is first inserted in memtable which is flushed to disk components called SSTables later. SSTables like in Figure 2 hold records with unique key. We described how SSTables are organized into levels in LevelDB in Section 2.2.2. SSTables organized into levels in LevelDB and because of the way the compaction takes place, there can be one unique key in each level.

For our problem, we do not want uniqueness in primary key as we assume that for both *SUBSCRIBE* and *PUBLISH* operation, the location or cell-id is the primary key and we will match queries with events using this key. So, instead of one record associated with a key we will have a posting list of queries or events. So for each insertion with non-unique primary key/cell-id, we add the new record with the current list of records instead of dropping the old record. The naive way to do it will be an in-place update of the list by issuing a read on the current list, appending the new data to the list and writing back to the storage. Background compaction later will discard the old list. But here the main drawback is that

each write becomes very expensive, as we need to read a large list. High write throughput is one of the fundamental features of NoSQL databases which we can not compromise. Also we may have a lot of invalid large lists throughout all the levels in SSTables, which will waste a lot space.

To solve this problem we implemented a lazy update strategy on the postings list. When a new write/Put is issued on a record, we do not look for existing value lists associated with the same key/cell-id in the disk. We simply write them to the memtable. If there is a existing list in the memtable, we perform in-place update on the list in constant time. Memtable is flushed to SSTables when it is full and these SSTables are compacted later to move to lower levels. We modify this compaction strategy appropriately, where instead of discarding old records, we merge lists associated with same key and hence eventually large lists are created in lower levels. Original compaction ensures that there can be at most one value list associated with a key in each level except level-0 files as SSTables in non-zero levels contain non-overlapping keys in LevelDB. Now we have fragmented list associated with a key throughout different levels. When we issue a read on key, instead of returning the record on upper level, it continues to search level by level to collect the fragmented lists. Lets illustrate the lazy update strategy on posting lists with an example.

Assume the current state of a Events database right after the following sequence of publish operations $PUT(C1, P1)$, ..., $PUT(C1, P2)$, ..., $PUT(C2, P3)$, $PUT(C2, P4)$, ..., $PUT(C1, P5)$. Figure 8 depicts the state of the storage components (i.e. SSTables and Memtable) after these operation. We can see value list of C1 and C2 are fragmented in components of different levels. Then two operation $PUT(C1, P6)$, $PUT(C2, P7)$ are issued and we show in the Figure 8 how it affects the current lists and how these lists are lazily updated after compaction. We first see the P6 is added

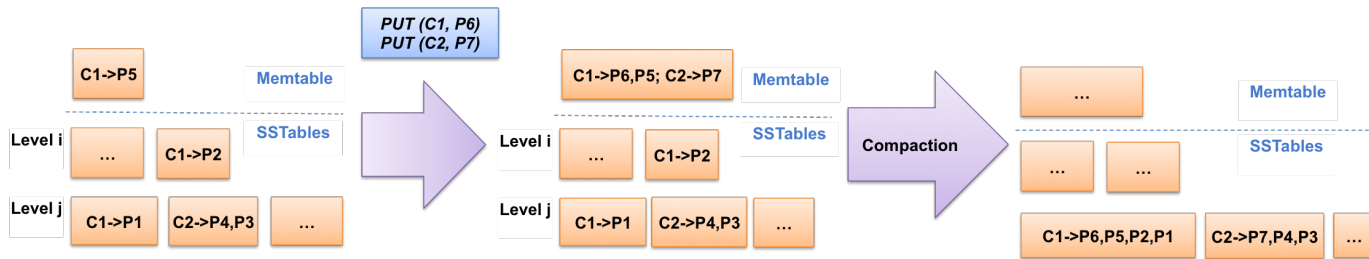


Fig. 8: Writing data to Storage containing lazy updated value lists and their compaction.

to the list of C1 with an in-place update and a new record C2->P7 was created inside memtable. Now after some compaction, these lists are moved to lower levels and we can see they are compacted and combined into larger lists.

7 EXPERIMENTS

7.1 Experimental Setup

We ran our experiments on a machine with the following configuration: Processor of AMD Phenom(tm) II X6 1055T and 8GB RAM, with Ubuntu version 15.04.

7.1.1 Data and Query Workload

We used the Twitter streaming API to collect 15 million geotagged tweets, taking 12 GB in JSON format, located within New York State. We use this dataset to generate our desired workloads for the experiment. We ran experiments using different Subscription-to-Publication ratios. As there are too many parameters, we set the time intervals for each query and expiration time of each event to a constant value: all subscriptions have T_{min} 10 seconds behind current time T_S and T_{max} 10 seconds ahead of T_S , and all the publications have expiration time (T_{max}) of 20s ahead of current time T_P .

There are several public workloads for key-value store databases available online such as YCSB [32]. However, to simulate a highly dynamic publish/subscribe model, we need to generate a mix of dynamically expiring subscriptions and publications. We decided to write our own script to generate different workloads from the real twitter dataset. Our workload generator considers each tweet as either a subscription or a publication depending on the Subscription-to-Publication ratio.

As all our tweets were collected within New York State, we use the bounding box rectangle around New York State and partition it to generate 500×500 uniform sized cells each having a unique identifier cell-id. We map the Geo-location of the tweet to appropriate cell-id and use this cell-id as a primary key for input. If the tweet is an event (publication), the text is considered as event description. Tweet ID is used as either subscriber ID or publication ID. The time intervals are set according to the execution time of that particular operation as discussed above. In our dataset, as we have about 0.25 million cells and 15.3 million tweets, average number of tweets per cell is about 60 and as described, these tweets are converted into events and subscriptions.

7.1.2 Padres Baseline

In addition to the repetitive queries baseline (RepQueries) described in Section 4.3, we also compare to a popular Pub/Sub system, Padres. We understand that Padres has a client-server architecture, which may incur additional overhead in delivering a subscription result, but we show that the performance difference

is quite dramatic. First, we have to express our problem setting using Padres' model. For that, we convert our workload into Padres subscription and publication operations. Padres supports *historic subscriptions* on past queries. So, each subscriptions in our dataset is equivalent to one historic subscription and one regular subscription in Padres. Each event publication is also converted to a publish operation in Padres. We installed Padres locally containing one broker and two clients connecting to the broker. One client is subscribing queries and the other client is publishing events. For clarification, let's illustrate the dataset conversion with an example. Suppose at T_S , a subscription query with cid as cell-id and interval $[T_{min}, T_{max}]$ is issued. We then convert it to a composite subscription (Expression 1) and a regular subscription (Expression 2).

$$CS[class, eq, historic], [subclass, eq, events], [T, <, T_S], \dots, [cellid, eq, cid] \& [T, >, T_{min}], \dots, [cellid, eq, cid] \quad (1)$$

$$S[class, eq, events], [T, <, T_{max}], \dots, [cellid, eq, cid] \quad (2)$$

A new generated event publication at T_P with cid as cell-id and T_{max} as expiration time will be converted to the following publication operation (Expression 3).

$$P[class, events], [T, T_P], [description, any], [cellid, cid] \quad (3)$$

7.2 Experimental Results

We conduct our experiments on workloads for both simple matching subscription and more advanced self-joining and multi granular subscriptions for different subscription-to-publication ratios, which represent different use cases.

7.2.1 Simple Subscriptions

Figures 9 and 10 show the overall, SUBSCRIBE and PUBLISH performance of all systems subscription heavy ($\frac{Subscription}{Publication} = 3$) and publication heavy ($\frac{Subscription}{Publication} = \frac{1}{3}$) workloads, respectively. In all figures, we record the performance once per million operations. We display the cumulative total time taken for both PUBLISH and SUBSCRIBE operation separately and also collectively and calculate average time per operation in every million operations.

Figures 9 and 10 show that if the system relies on repetitive queries instead of instant response queries, it can not scale to millions of operations. As *RepQueries* does not issue any read after each publication, and only issues a write to a single database, PUBLISH has very good performance as expected. But SUBSCRIBE has bad performance as *RepQueries*. The overall performance is far worse than our proposed instant-response variants. Both *TwoDB* and *DualDB* approaches have about 1000% better performance than *RepQueries* for subscription heavy workload, and 300% better for publication heavy workload.




Fig. 9: Performance of different storage variants for simple subscription queries on subscription heavy workload




Fig. 10: Performance of different storage variants for simple subscription queries on publication heavy workload

According to the theoretical analysis in Section 5.1.1, both *TwoDB* and *DualDB* approaches cost the same number of disk accesses for each PUBLISH and SUBSCRIBE in the worst case. This is confirmed experimentally in Figure 10, where we find that they have identical performance for publication heavy workload.

Figure 9 shows the results for subscription heavy workload, where *DualDB* is better than *TwoDB* by about 5%. This behavior is expected in *DualDB*, although it does not perform any extra disk access than *TwoDB* theoretically. The reason is that *DualDB* holds both query and data associated with a key inside the same disk block or in a neighbor disk block, which leads to better cache performance for both OS page cache and LevelDB block cache.

We see in Figures 9 and 10 that *Padres* performs poorly not only compared to our two approaches, but also to the repetitive baseline. Also, in our experiments *Padres* is not able to cope with the increasing number of subscriptions, and runs out of memory very quickly (even before 2 million operations) and the system crashes. This is because it relies on an in-memory data structure to manage subscriptions and fails to cope with a very large number of queries. Note that we allocated maximum memory for *Padres* (i.e. 8GB) and it still runs out of memory. Here we can see that even if we have sufficient memory to support small number of subscriptions, the use of traditional SQL-like database perform

much worse (up to 300%) than even our baseline *RepQueries* approach.

7.2.2 Self-Joining Subscriptions

Similar to the simple matching subscription, we observe the performance for complex self-joining subscriptions. Figures 11 and 12 show the overall, SUBSCRIBE and PUBLISH performance of *DualDB* and *TwoDB* approaches on both subscription and publication heavy workloads, respectively. Here we do not consider baseline *RepQueries* approach, which relies on repetitive queries on Events database, because it does not store the subscriptions in a database and hence it cannot perform self-joining queries.

Figure 11 shows that *DualDB* is overall 20% (About 6% on PUBLISH and About 25% on SUBSCRIBE) faster than *TwoDB* for subscription heavy workload. We see similar performance gain of 12% overall improvement for *DualDB* over *TwoDB* in publication heavy workload. Note that here PUBLISH performs much better than subscription heavy workload which is desired in a publication heavy workload.




Fig. 11: Performance of different storage variants for Self-joining subscriptions on subscription heavy workload




Fig. 12: Performance of different storage variants for Self-joining subscriptions on publication heavy workload

7.2.3 Subscriptions for Multi Granular Data

To simulate the environment for multi granular attributes, we partitioned the space by a balanced quad tree of depth 5 instead of a uniform grid. Here each cell represents a node in the tree and each cell-id is encoded by Dewey Numbering (described in Section 5.2 and Figure 7). We set subscription-to-publication ratio as 1. This workload generates a mix of subscriptions and publications from different levels. We argue that the smaller cells will likely be accessed more than larger cells up in the hierarchy. For that, we design the workload generator in such way that a cell in level i has probability proportional to $\frac{P}{2^{d-i}}$ to be selected, where d is the maximum depth of the tree. In parallel, we also generated the workload for the fixed granularity approach by translating each operation into a multi operation (one operation for all the lowest level cells).

Figure 13 shows the overall, SUBSCRIBE and PUBLISH performance of *DualDB* for this workload of variable granularity against the fixed granularity approach. We see that our intelligent range query based variable granularity approach performs about 300% better for both PUBLISH and SUBSCRIBE operation. Figure 15 shows the performance for both approaches in different levels. It shows the average time per operation after performing 15

million operation. We see that our variable granularity approach performs even better for upper levels in the hierarchy. These results justify the cost analysis in Section 5.2.3.

In the above experiment, we assume that a subscription or publication must pick one of the existing geospatial regions of the quad tree. Next, we show how arbitrary spatial regions, specifically circles can be supported, represented by a geo-location point and a radius. We generated a workload from our Twitter dataset for this application, where we select a radius for each operation from a set of five radii ($4 * R$, $2 * R$, R , $R/2$, $R/4$). Here, R is the length of the smallest cell in the quad tree. For the variable granularity approach, each operation's geo-location and a randomly selected radius (from the pre-defined set) is mapped to the cell in our quad tree whose area fully contains the circle region. For fixed granularity approach, this circle is mapped to all the cells in the lowest level whose rectangular region intersect with this circle. We perform one operation for all of these cells separately. Figure 14 shows that our variable granularity approach performs about 200% better for both PUBLISH and SUBSCRIBE operations. Figure 16 shows the performance for both approaches for different radii. We see that the performance gain of our variable granularity approach becomes greater for larger radii.




Fig. 13: Performance of Pub/Sub in DualDB for high granular data in Different levels of hierarchy for whole workload




Fig. 14: Performance of Pub/Sub in DualDB for high granular spatial data covering area of different radii




Fig. 15: Performance comparison in different levels of Hierarchy

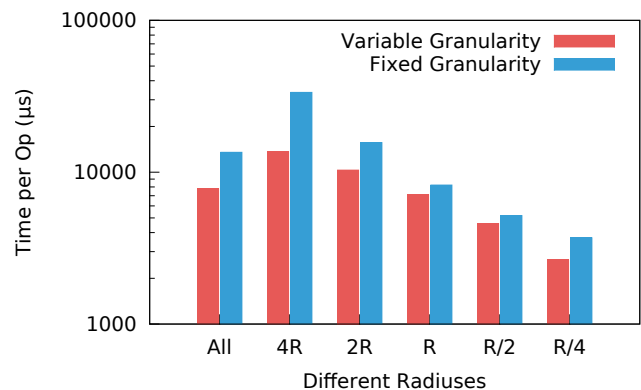


Fig. 16: Performance comparison for different radii

8 CONCLUSIONS AND FUTURE WORK

In this paper we present efficient storage and indexing approaches to achieve high throughput publish/subscribe on LSM-based databases where both the number of subscriptions and publications are massive in scale and one or both of them can arrive and expire

with time. Our approaches support instant notifications. We also consider a baseline approach that relies on repetitive queries. We also show how hierarchical attributes in multi granular data can be supported efficiently by a Dewey encoded representation.

We implement these storage frameworks on top of the popular

LSM-based LevelDB system and conduct extensive experiments using real datasets. The experimental results show that our proposed approaches outperform the state-of-the-art *Padres* pub-sub system (by up to 3000%) and also outperform the repetitive baseline *RepQueries* (by up to 1000%). We also show that our *DualDB* approach outperforms the *TwoDB* approach for simple subscriptions by a small margin (up to 5%) and for self-joining complex subscriptions by larger margin (up to 20%). For hierarchical attributes, we show that our variable granularity approach based on Dewey encoded IDs performs much better (up to 300%) than the fixed granularity approach.

In the future, we plan to extend this work to a distributed environment and allow more complex subscription queries that do not match based on a primary key conditions.

ACKNOWLEDGMENTS

This project is partially supported by NSF grants IIS-1447826 and IIS-1619463-.

REFERENCES

- [1] M. J. Carey, S. Jacobs, and V. J. Tsotras, "Breaking bad: a data serving vision for big active data," in *Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems*. ACM, 2016, pp. 181–186.
- [2] A. Lakshman and P. Malik, "Cassandra: A decentralized structured storage system," *SIGOPS Oper. Syst. Rev.*, vol. 44, no. 2, pp. 35–40, apr 2010.
- [3] A. Feinberg, "Project voldemort: Reliable distributed storage," in *Proceedings of the 10th IEEE International Conference on Data Engineering*, 2011.
- [4] S. Alsubaiee, A. Behm, V. Borkar, Z. Heilbron, Y.-S. Kim, M. J. Carey, M. Dreseler, and C. Li, "Storage management in asterixdb," *Proceedings of the VLDB Endowment*, vol. 7, no. 10, 2014.
- [5] "Mongodb," <https://www.mongodb.com>.
- [6] "Leveldb," <http://leveldb.org/>.
- [7] E. Fidler, H.-A. Jacobsen, G. Li, and S. Mankovski, "The padres distributed publish/subscribe system," in *FIW*, 2005, pp. 12–30.
- [8] "Project website for open source code." <http://dblab.cs.ucr.edu/projects/PubSub-Store/>.
- [9] P. T. Eugster, P. A. Felber, R. Guerraoui, and A.-M. Kermarrec, "The many faces of publish/subscribe," *ACM Computing Surveys (CSUR)*, vol. 35, no. 2, pp. 114–131, 2003.
- [10] A.-M. Kermarrec and P. Triantafyllou, "Xl peer-to-peer pub/sub systems," *ACM Computing Surveys (CSUR)*, vol. 46, no. 2, p. 16, 2013.
- [11] H.-A. Jacobsen, V. Muthusamy, and G. Li, "The padres event processing network: Uniform querying of past and future eventsdas padres ereignisverarbeitungsnetzwerk: Einheitliche anfragen auf ereignisse der vergangenheit und zukunft," *it-Information Technology Methoden und innovative Anwendungen der Informatik und Informationstechnik*, vol. 51, no. 5, pp. 250–260, 2009.
- [12] N. Bhatt, D. Gawlick, E. Soylemez, and R. Yaseem, "Content based publish-and-subscribe system integrated in a relational database system," Jun. 11 2002, uS Patent 6,405,191.
- [13] F. Tian, B. Reinwald, H. Pirahesh, T. Mayr, and J. Myllymaki, "Implementing a scalable xml publish/subscribe system using relational database systems," in *Proceedings of the 2004 ACM SIGMOD international conference on Management of data*. ACM, 2004, pp. 479–490.
- [14] L. Guo, D. Zhang, G. Li, K.-L. Tan, and Z. Bao, "Location-aware pub/sub system: When continuous moving queries meet dynamic event streams," in *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. ACM, 2015, pp. 843–857.
- [15] M. A. Qader and V. Hristidis, "Dualdb: An efficient lsm-based publish/subscribe storage system," in *Proceedings of the 29th International Conference on Scientific and Statistical Database Management (SS-DBM)*, 2017.
- [16] J. Widom and S. J. Finkelstein, "Set-oriented production rules in relational database systems," in *ACM SIGMOD Record*, vol. 19, no. 2. ACM, 1990, pp. 259–270.
- [17] U. Schreier, H. Pirahesh, R. Agrawal, and C. Mohan, "Alert: An architecture for transforming a passive dbms into an active dbms," in *Proceedings of the 17th International Conference on Very Large Data Bases*. Morgan Kaufmann Publishers Inc., 1991, pp. 469–478.
- [18] E. N. Hanson, C. Carnes, L. Huang, M. Konyala, L. Noronha, S. Parthasarathy, J. Park, and A. Vernon, "Scalable trigger processing," in *Data Engineering, 1999. Proceedings., 15th International Conference on*. IEEE, 1999, pp. 266–275.
- [19] J. Chen, D. J. DeWitt, F. Tian, and Y. Wang, "Niagaracq: A scalable continuous query system for internet databases," in *ACM SIGMOD Record*, vol. 29, no. 2. ACM, 2000, pp. 379–390.
- [20] S. Chandrasekaran, O. Cooper, A. Deshpande, M. J. Franklin, J. M. Hellerstein, W. Hong, S. Krishnamurthy, S. R. Madden, F. Reiss, and M. A. Shah, "Telegraphcq: continuous dataflow processing," in *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*. ACM, 2003, pp. 668–668.
- [21] S. Babu and J. Widom, "Continuous queries over data streams," *ACM Sigmod Record*, vol. 30, no. 3, pp. 109–120, 2001.
- [22] S. Madden, M. Shah, J. M. Hellerstein, and V. Raman, "Continuously adaptive continuous queries over streams," in *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*. ACM, 2002, pp. 49–60.
- [23] N. Garg, *Apache Kafka*. Packt Publishing Ltd, 2013.
- [24] "Gemfire continuous querying," https://pubs.vmware.com/vfabric5/index.jsp?topic=/com.vmware.vfabric.gemfire.6.6/developing/continuous_querying/how_continuous_querying_works.html.
- [25] "Influxdb," <https://www.influxdata.com/>.
- [26] A. M. Hendawi, J. Gupta, Y. Shi, H. Fattah, and M. Ali, "The microsoft reactive framework meets the internet of moving things."
- [27] "Oracle bitmap indexes," https://docs.oracle.com/cd/B10500_01/server.920/a96520/indexes.htm.
- [28] F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber, "Bigtable: A distributed storage system for structured data," *TOCS*, vol. 26, no. 2, p. 4, 2008.
- [29] L. George, *HBase: the definitive guide*. O'Reilly Media, Inc., 2011.
- [30] "Rocksdb," <http://rocksdb.org/>.
- [31] I. Tatarinov, S. D. Viglas, K. Beyer, J. Shanmugasundaram, E. Shekita, and C. Zhang, "Storing and querying ordered xml using a relational database system," in *Proceedings of the 2002 ACM SIGMOD international conference on Management of data*. ACM, 2002, pp. 204–215.
- [32] B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears, "Benchmarking cloud serving systems with ycsb," in *Proceedings of the 1st ACM symposium on Cloud computing*, 2010, pp. 143–154.



Mohiuddin Abdul Qader is a PhD Candidate in Computer Science at UC Riverside. He received his B.Sc. degree in Computer Science and Engineering from Bangladesh University of Engineering and Technology (BUET) in 2012 and received his M.S. in Computer Science from UC Riverside in 2015. His main research interests include Big Data Management Systems, Web Search Engine and Information Retrieval.



Vagelis Hristidis is a Professor of Computer Science at UC Riverside. His key areas of expertise are Databases, Information Retrieval, and particularly the intersection of these two areas. His work in these areas has received more than 4,000 citations according to Google Scholar. His key achievements also include the NSF CAREER award, a Google Research Award, an IBM Award, and a Kauffmann Entrepreneurship Award.