

# GRAPHONE: A Data Store for Real-time Analytics on Evolving Graphs

PRADEEP KUMAR, Department of Computer Science, William and Mary

H. HOWIE HUANG, Department of Electrical and Computer Engineering, George Washington University

There is a growing need to perform diverse set of real-time analytics (batch and stream analytics) on evolving graphs in order to deliver the values of big data to users. The key requirement from such applications is to have a data store to support their diverse data access efficiently, while concurrently ingesting fine-grained updates at a high velocity. Unfortunately, current graph systems, either graph databases or analytics engines, are not designed to achieve high performance for both operations, rather they excel in one area that keep a private data store in a specialized way to favor their operations only. To address this challenge, we have designed and developed GRAPHONE, a graph data store that abstracts the graph data store away from the specialized systems to solve the fundamental research problems associated with the data store design. It combines two complementary graph storage formats (edge list and adjacency list), and uses *dual versioning* to decouple graph computations from updates. Importantly, it presents a new data abstraction, *GraphView*, to enable data access at two different granularities of data ingestions (called *data visibility*) for concurrent execution of diverse classes of real-time graph analytics with only a small data duplication. Experimental results show that GRAPHONE is able to deliver 11.40× and 5.36× average speedup in ingestion rate against LLAMA and Stinger, the two state-of-the-art dynamic graph systems respectively. Further they achieve an average speedup of 8.75× and 4.14× against LLAMA, and by 12.80× and 3.18× against Stinger for BFS and PageRank analytics (batch version) respectively. GRAPHONE also gains over 2000× speedup against Kickstarter, a state-of-the-art stream analytics engine in ingesting the streaming edges and performing streaming BFS when treating first half as a base snapshot and rest as streaming edge in a synthetic graph. GRAPHONE also achieves an ingestion rate of two to three orders of magnitude higher than graph databases. Finally, we demonstrate that it is possible to run concurrent stream analytics from the same data-store.

CCS Concepts: • **Information systems** → **Data management systems**; **Main memory engines**; **Information storage systems**.

Additional Key Words and Phrases: Graph Systems, Graph Data Management, Unified Graph Data Store, Batch Analytics, Stream Analytics

## ACM Reference Format:

Pradeep Kumar and H. Howie Huang. 2019. GRAPHONE: A Data Store for Real-time Analytics on Evolving Graphs. *ACM Trans. Storage* 1, 1, Article 1 (January 2019), 40 pages. <https://doi.org/10.1145/3364180>

This article is an Extension of a FAST'19 Conference Paper by Kumar et al. [48]. There has been significant addition to the conference paper including presentation on concurrent analytics execution, comparison against two new systems: a snapshot based batch graph analytic system; and a stream graph analytics system. We also incorporate various new results on raw text data, evaluation of overhead on GraphView, many new evaluations on design parameters; and extended discussions using new diagrams and tables to improve the readability throughout the paper.

Authors' addresses: Pradeep Kumar, [pkumar@wm.edu](mailto:pkumar@wm.edu), Department of Computer Science, William and Mary, 251 Jamestown Road, Williamsburg, VA, 23185; H. Howie Huang, [howie@gwu.edu](mailto:howie@gwu.edu), Department of Electrical and Computer Engineering, George Washington University, 800 22nd Street NW, Washington, DC, 20052.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2019 Association for Computing Machinery.

1553-3077/2019/1-ART1 \$15.00

<https://doi.org/10.1145/3364180>

## 1 INTRODUCTION

We live in a world where information networks have become an indivisible part of our daily lives. A large body of research has studied the relationships in such networks, e.g., biological networks [35], social networks [22, 43, 49], and web [9, 33]. In these applications, graph queries and analytics are being used to gain valuable insights from the data, which can be classified into two broad categories: *batch analytics* (e.g. PageRank [69], graph traversal [11, 54, 56]) that analyzes a static snapshot of the data, and *stream analytics* (e. g. anomaly detection [8], topic detection [73]) that studies the incoming data over a time window of interest.

Generally speaking, batch analytics prefers a *base (data) store* that can provide indexed access on the non-temporal property of the graph such as the source vertex of an edge, and on the other hand, stream analytics needs a *stream (data) store* where data can be stored quickly and can be indexed by their arrival order for temporal analysis.

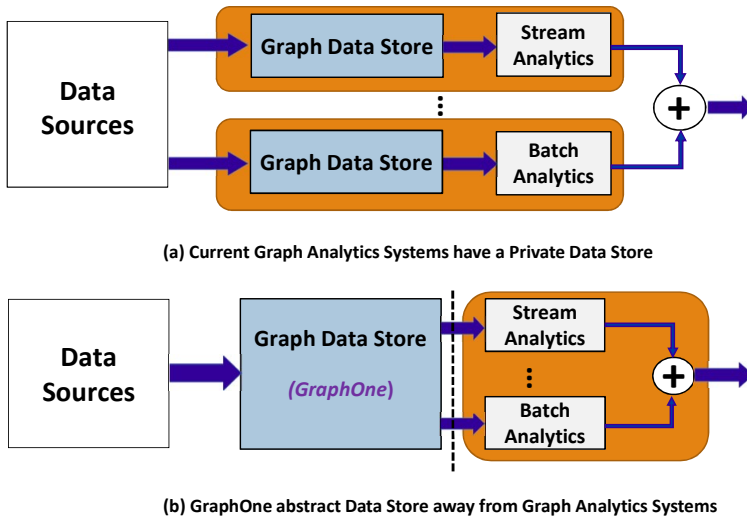
Increasingly, one needs to perform batch and stream processing together on evolving graphs [10, 77, 78, 93]. The key requirement here is to sustain a large volume of fine-grained updates at a high velocity, and simultaneously provide high-performance data access for various classes of real-time analytics and query support.

This trend poses a number of challenges to the underlying storage and data management system. First, batch and stream analytics perform different kinds of data access, that is, the former visits the whole graph while the latter focuses on the data within a time window. Second, each analytic has a different notion of real time, that is, data is visible to the analytics at different granularity of data ingestion (updates). For example, an iterative algorithm such as PageRank can run on a graph that is updated at a coarse granularity, but a graph query to output the latest shortest path requires data visibility at a much finer granularity. Third, such a system should also be able to handle a high arrival rate of updates, and maintain data consistency while running concurrent batch and stream processing tasks.

Table 1. A simplified classification of various graph data management systems depicting that each graph system only supports one type of data ingestion and access pattern.

Graph Systems	Ingestion Type Offered	Data Access Type Offered
Graph Database	Fine-grained (e.g., Neo4j [66])	Whole Data
Dynamic Graph System	Fine-grained (e.g., Stinger [23])	Whole Data
	Coarse-grained (e.g., Graphchi [50])	
Stream Graph System	Coarse-grained (e.g., Naiad [65])	Streaming Access of Whole Data
	Coarse-grained (e.g. GraphTau [34])	Streaming Access from Time-window

Unfortunately, current graph systems can neither provide diverse data access nor at different granularity of data ingestion in presence of a high data arrival rate as shown in Table 1. Many dynamic graph systems [50, 59] only support batched updates, and a few others [23, 79] offer data visibility at fine granularity of updates but with a weak consistency guarantee (atomic read consistency), which as a result may cause an analytic iteration to run on different data versions and produce undesired results. Relational and graph databases such as Neo4j [66] can handle fine-grained updates, but suffer from poor ingestion rate for the sake of strong consistency guarantee [62]. Also, such systems are not designed to support high-performance streaming data access over a time window. On the other hand, graph stream engines [34, 65, 76, 85] interleave incremental computation with data ingestion, i.e., graph updates are batched and not applied until the end of an iteration. In short, the existing systems manage a private data store in a way to favor their specialized analytics.



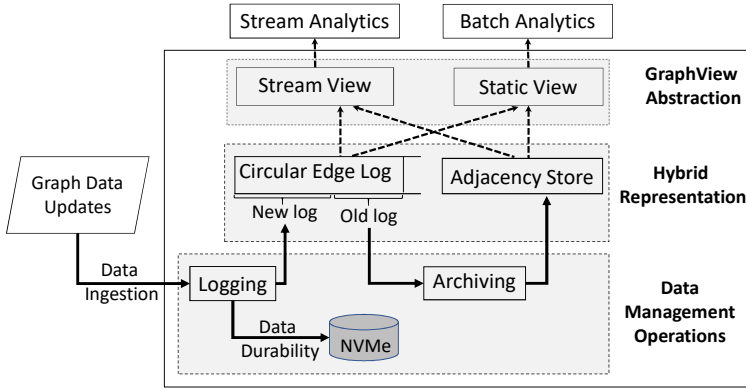
**Fig. 1:** Showing a high-level overview of graph analytics system. The evolving graph data arrives one edge at a time or in batches from the data sources, such as Apache Kafka, database change log, etc. Each of current graph analytics systems have a private data store optimized for one class of usecases that interface with data sources. GRAPHONE abstracts the data store away from the analytics systems to solve the problems associated with a general-purpose graph data store, and enables execution of diverse classes of analytics from the same data-store using GraphView instances.

In principle, one can utilize these specialized graph systems side-by-side to provide data management functions for dynamic graphs and support a wide spectrum of analytics and queries as shown in Fig. 1 (a). However, such an approach would be suboptimal [93], as it is only as good as the weakest component, in many cases the graph database with poor performance for streaming data. Worse, this approach could also lead to excessive data duplication, as each subsystem would store a replica of the same underlying data in their own format.

In this work, we have designed GRAPHONE, a unified graph data store abstracted away from the specialized graph systems, as shown in Fig. 1(b), in order to solve the fundamental problem of such data store design and management. GRAPHONE offers diverse data access at two granularity levels of data ingestions for various real-time analytics, so that many classes of analytics can be performed from the same data-store, while supporting data ingestion at a high arrival rate. Fig. 2 provides a high-level overview.

It leverages a *hybrid graph store* to combine a small circular edge log (henceforth *edge log*) and an *adjacency store* for their complementary advantages. Specifically, the edge log keeps the latest updates in the edge list format using a logging phase, and is designed to accelerate data ingestion. At the same time, the adjacency store holds the snapshots of the older data in the adjacency list format that is moved periodically from the edge log using a archiving phase, and is optimized for batch and streaming analytics. It is important to note that the graph data is not duplicated in two formats, although a small amount of overlapping is allowed to keep the original composition of the versions intact.

GRAPHONE enforces data ordering using the temporal nature of the edge log in the logging phase, and keeps the per-vertex edge arrival order intact in the adjacency store using a novel *edge sharding* in the archiving phase. A *dual versioning* technique then exploits the fine-grained versioning of



**Fig. 2:** High-level architecture of GRAPHONE. Solid and dotted arrows show the data management and access flow respectively.

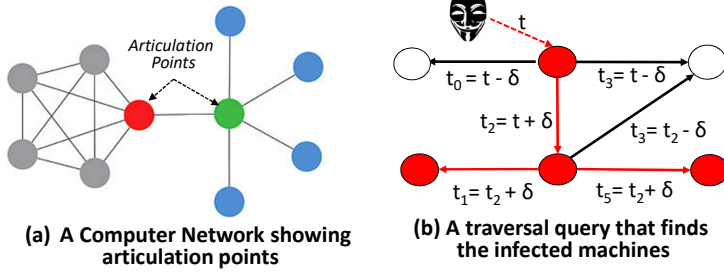
the edge list format and the coarse-grained versioning of the adjacency list format to create real-time versions to offer *snapshot isolation* data consistency. Further, GRAPHONE allows independent execution of analytics that run parallel to data management, and can fetch a new version at the end of its own incremental computation step in case of stream analytics. This is different from the current stream graph systems where the size of batched update is often determined by the heaviest compute, which hurts the timeliness of other concurrently executing analytics.

Additionally, we provide two optimization techniques, *cache line sized memory allocation* and *special handling of high degree vertices* of power-law graphs, to reduce the memory requirement of versioned adjacency store.

GRAPHONE simplifies the diverse data access by presenting a new data abstraction, *GraphView*, on top of the hybrid store. Two types of GraphView are supported as shown in Fig. 2 : (1) the *static view* offers real-time versioning of the latest data for batch analytics; and (2) the *stream view* supports stream analytics with the most recent updates. These views offers *visibility* of data updates to analytics at two levels of granularity where the edge log is used to offer it at the edge level, while the adjacency store provides the same at coarse granularity of updates. As a result, GRAPHONE provides high-level applications with the flexibility to choose the granularity of data visibility for a desired performance. In other words, the edge log can be accessed if fine-grained data visibility is required, which can be tuned (§7.4).

We have implemented GRAPHONE as an in-memory graph data-store with a limited durability guarantee on external non-volatile memory express solid-state drives (NVMe SSD). For comparison, we have evaluated it against many classes of in-memory graph systems: LLAMA [59] as a snapshot based coarse-grained graph batch analytics system; Stinger [23], a fine-grained dynamic graph system; Neo4j and SQLite, two graph data management systems; Kickstarter [85] as graph stream analytics. The experimental results show that GRAPHONE can support a high data ingestion rate, specifically it achieves  $11.40\times$  average speedup in ingestion rate than LLAMA,  $5.36\times$  higher ingestion rate than Stinger, and two to three orders of magnitude higher ingestion rate than graph databases.

In addition, GRAPHONE outperforms LLAMA by  $8.75\times$  and  $4.14\times$ , and Stinger by  $12.80\times$  and  $3.18\times$  for BFS and PageRank analytics (batch version) respectively. GRAPHONE also gains a speedup of over  $2000\times$  speedup over Kickstarter in ingesting the streaming edges and performing streaming BFS when treating first half as a base snapshot and rest as streaming edge in a synthetic graph. The stream processing in GRAPHONE runs parallel to data archiving phase (update to adjacency



**Fig. 3:** Graph traversal can locate possible infected nodes using real-time authentication graph if infected user and node are known

store) which offers 28.58% higher ingestion rate compared to the current practice of running the two in sequence. Due to this feature, GRAPHONE also enables concurrent stream analytics from the same dynamic adjacency store that is mostly absent in prior graph stream analytics engine.

To summarize, GRAPHONE makes three contributions:

- Unifies stream and base stores to manage the graph data in a dynamic environment;
- Provides batch and stream analytics through dual versioning, smart data management, and memory optimization techniques;
- Supports diverse data access for concurrent execution of various usecases with GraphView and data visibility abstractions.

The rest of the paper is organized as follows. We present a usecase in §2, opportunities and GRAPHONE overview in §3, the hybrid store in §4, data management internals and optimizations in §5, GraphView data abstraction in §6, evaluations in §7, related work in §8, and conclusion in §9.

**2 USE CASE: NETWORK ANALYSIS**

Graph analytics is a natural choice for data analysis on an enterprise network. Fig. 3(a) shows a graph representation of a simple computer network. Such a network can be analyzed in its entirety by calculating the diameter [52], and betweenness centrality [13] to identify the articulation points. This kind of batch analysis is very useful for network infrastructure management. In the meantime, as the dynamic data flow within the network captures the real-time behaviors of the users and machines, the stream analytics is used to identify security risks, e.g., denial of service, and lateral movement, which can be expressed in the form of path queries, parallel paths and tree queries on a streaming graph [18, 40].

Los Alamos Nation Laboratory (LANL) recently released a comprehensive data set [39] that captures a wide range of network information, including authentication events, process events, DNS lookups, and network flows. The LANL data covers over 1.5 billion events, 12,000 users, and 17,000 computers, and spans 58 consecutive days. For example, the network authentication data captures the login information that a user logs in to a network machine, and also from that machine to other machines. When the network defense system identifies a malicious user and node, it needs to find all the nodes that may have been infected. Instead of analyzing every node of the network, one can quickly run a path traversal query on the real-time authentication graph to identify the possible infected nodes, that is, find all the nodes whose login has originated from the chain of nodes that are logged in from the first infected machine [40] as shown in Fig. 3(b).

In summary, a high-performance graph store that captures dynamic data in the network, combined with user, machine information and network topology is advantageous in understanding the health of the network, accelerating network service, and protecting it against various attacks.

This clearly articulates the need to have a data store that can efficiently ingest the incoming data and make it available for analysis in a form that various analytics and query engine prefer, so that analytics writer can focus on the actual usecases rather than thinking about how to manage the incoming data, and providing right data consistency as well proving various data access types. To this end, GRAPHONE presents a graph data storage and management infrastructure that abstracts the data store away from the analytics engines to solve the basic data access problem associated with such diverse real-time analytics.

### 3 OPPORTUNITIES AND OVERVIEW

A graph can be defined as  $G = (V, E, W)$ , where  $V$  is the vertex set, and  $E$  is the edge set, and  $W$  is the set of edge weights. Each vertex may also have a label. In this section, graph formats and their traits are described as relevant for GRAPHONE, and then we present its high-level overview.

#### 3.1 Graph Representation: Opportunities

Fig. 4 shows three most popular data formats for a sample graph. First, the *edge list* is a collection of edges, a pair of vertices, and captures the incoming data in their arrival order. Second, the *compressed sparse row (CSR)* groups the edges of a vertex in an *edge array*. There is a metadata structure, *vertex array*, that contains the index of the first edge of each vertex. Third, the *adjacency list* manages the neighbors of each vertex in separate per-vertex edge arrays, and the vertex array stores a count (called degree) and pointer to indicate the length and the location of the corresponding edge arrays respectively. This format is better than the CSR for ingesting graph updates as it affects only one edge array at a time.

In the edge list, the neighbors of each vertex are scattered across, thus is not the optimal choice for many graph queries and batch analytics who prefer to get the neighboring edges of a vertex quickly [12, 31, 32, 36] etc. However, the edge list provides a natural support for fast update, as each update simply needs to be appended to the end of the list. On the other hand, the adjacency list format loses the temporal ordering as the incoming updates get scattered over the edge arrays, thus not suited directly for stream analytics. Also, creation of fine-grained snapshots are not easy on adjacency list, hence prior work has only attempted to create coarse-grained versioning [28, 59]. Stinger [23] directly updates the adjacency list format but can provide read committed only, a weaker data consistency, that does not provide any snapshot capability. Therefore different steps within an analytics will run on different versions of data, thereby producing unknown results in case of evolving graphs.

Given their advantages and disadvantages, neither format is ideally suited for supporting both the batch and stream analytics, and high ingestion of data on its own. We now identify two opportunities for this work:

**Opportunity #1: Utilize both the edge list and the adjacency list within a hybrid store.** The edge list format preserves the data arrival order and offers a good support for fast updates as each update is simply appended to the end of the list that requires sequential write only. On the other hand, the adjacency list keeps all the neighbors of a vertex indexed by the source vertex, which provides efficient data access for graph analytics. Thus it allows GRAPHONE to achieve high-performance graph computation while simultaneously supporting fine-grained updates.

**Opportunity #2: Fine-grained snapshot creation with the edge list format.** Graph analytics and queries require an immutable snapshot of the latest data for the duration of their execution. The edge list format provides a natural support for fine-grained snapshot creation without creating a physical snapshot due to its temporal nature, as tracking a snapshot is just remembering an offset in the edge list. Hence, the adjacency list format can use the fine-grained snapshot capability of the

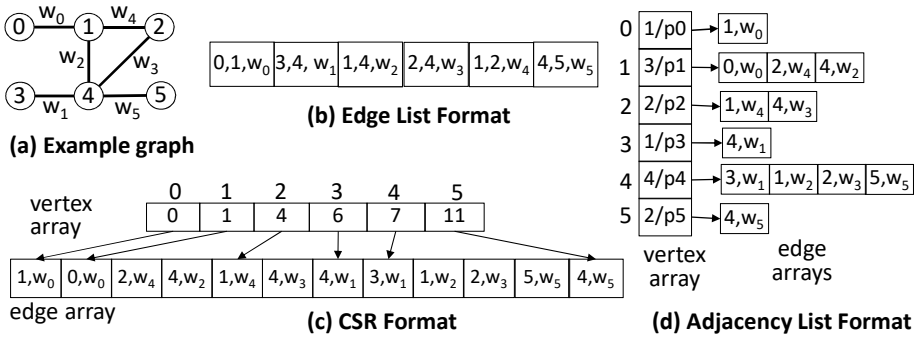


Fig. 4: Sample graph and its various storage format. GRAPHONE concentrates on the complementary properties of the edge list format and the adjacency list format.

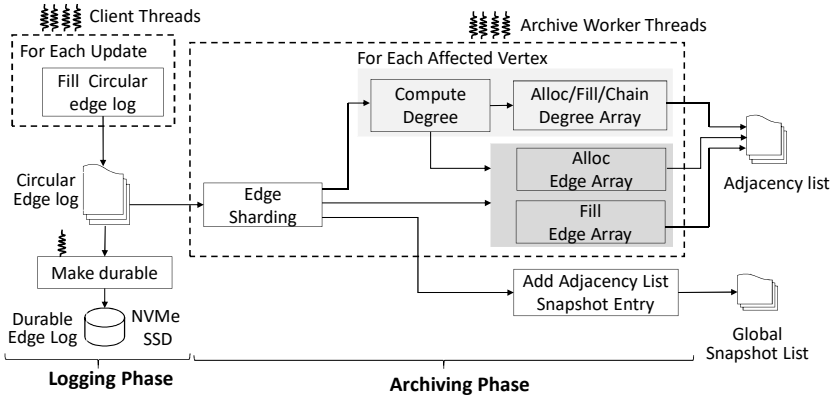
edge list to complement its coarse-grained snapshot capability [28, 59] to create real-time versions of graph data.

### 3.2 Overview

GRAPHONE utilizes a hybrid graph data store (discussed in §4) that consists of a small circular *edge log* and the *adjacency store*. Fig. 5 shows an high-level overview of GRAPHONE architecture. The hybrid store is managed in several phases (presented in §5). Specifically, during the *logging phase*, the edge log records the incoming updates in the edge list format atomically in their arrival order, and supports a high ingestion rate. We define *non-archived edges* as the edges in the edge log that are yet to be moved to the adjacency store. When their number crosses the *archiving threshold*, a parallel *archiving phase* begins, which merges the latest edges to the adjacency store to create a new adjacency list snapshot. This duration is referred to as an *epoch*. In the *durable phase*, the edge log is written to a disk in an asynchronous manner, and requires support from upstream buffering for complete recovery. In this paper, we use *data ingestion* to mean that newly arrived edges has been percolated to the adjacency store.

Thus, GRAPHONE is an in-memory graph data-store with a weaker durability guarantee and offers snapshot isolation as its data consistency which is slightly weaker model than full serializability that supports transactions, but is stronger than the read-committed isolation level offered by NoSQL data-stores such as Stinger [23] or Trinity [79, 90]. Due to logging phase, the non-archived edges move to adjacency store in batches, which enables a number of optimizations to improve the performance of archiving phase (§5.1.2), and to also reduce the memory consumption of adjacency store (§5.2). These techniques together bring the performance of GRAPHONE very close to static graph engines.

To efficiently create and manage immutable graph versions for real-time data analytics in presence of the incoming updates, we provide a set of GraphView APIs (discussed in §6) that data access for diverse classes of analytics. Specifically, *static view* API is for batch processing, while *stream view* API is for stream processing. Internally, the views utilize *dual versioning* technique where the versioning capability of both formats are exploited for faster creation of a snapshot of the current data at very low cost. For example, a real-time static view can readily be composed by using the latest coarse-grained version of the adjacency store, and the latest fine-grained version of non-archived edges. The *data overlap* technique allows GRAPHONE to keep the original composition of the created view intact, and thus guarantees its data access without interfering the data management operations.



**Fig. 5:** Architecture of GRAPHONE. Operations related to same data structures have been grayed out in archiving phase. Compaction Phase is not shown.

It is important to note that the GraphView also provides analytics with the flexibility to trade-off the granularity of *data visibility* of incoming updates for better analytics performance, i.e., an analytic gets to choose the type of data ingestion that it is interested in. For example, the analytics that prefer running only on batched updates will access the latest adjacency store, and thus avoid the cost associated with the access of the latest edges from the non-archived edges which are not indexed by source vertices.

Despite its simplicity, the abstraction of *data visibility* is very powerful, as it allows high-performance fine-grained ingestion, but offers the performance of batched updates to those analytics that are not interested in fine-grained ingestions. Put in other way, GRAPHONE removes the cost of fine-grained ingestion from data write path to data read path of those analytics that are really interested in running on fine-grained ingestion. All other data read path offers performance of a batched update system.

### 3.3 Major Differences

GRAPHONE is different than databases, graph batch analytics, graph stream analytics and key-value stores in some major way. In this sub-section, we provide an overview of those differences.

In databases, it is possible to provide functionality similar to adjacency list graph format. E.g., a database can create a simple schema to store graph edges containing source and destination vertices in an edge table. The data of edge table is similar to edge list format. Creating an index in the source vertex in a edge table results into a structure similar to adjacency list. In that case, databases updates the edge list log structure (for logging purpose), and update to the indexed structures (archiving to adjacency store) in the same thread context. Stinger [23] does not have any log structure, and maintains indexed structure only, but works same way, i.e. update of the indexed structure happens in the thread context of update thread.

On the other hand, GRAPHONE updates both of them in separate thread contexts, while delaying the update of the indexed structure for batching the edges to achieve better overall ingestion throughput. The major difference of GRAPHONE from the snapshot based graph batch analytics on evolving graphs, such as Kineograph [17], LLAMA [59] etc., is the ability of the GRAPHONE to create fine-grained snapshots, and controlling the data access of those snapshots by analytics using the data visibility abstraction.



Both the databases and snapshot based graph batch analytics and GRAPHONE run analytics in separate thread context than the data ingestion.

For graph stream analytics, update of the log structure happens separately than the analytics in GRAPHONE and prior stream engines both. The major difference is in the context of the indexed structure update (archiving phase). Prior graph stream engines perform archiving phase in the context of stream analytics using operators such map, join and group-by (as in Naiad [64], Graph-Tau [34] etc.) or using specific implementation (as in Kickstarter [85], GraphBolt [61] etc.), followed by actual stream computation. GRAPHONE, on the other hand, moves the archiving phase out of the analytics pipeline, i.e. archiving happens in parallel to actual stream computation that brings overall better performance.

Due to this overlapping nature of archiving and computation, GRAPHONE becomes the first ever system that can enable concurrent stream analytics from the same indexed structure. Even snapshot based stream analytics systems such as Kinerograph [17] can not execute two classes of stream analytics together: those that need to access whole data; and those that requires window access, as Kineograph propose a different snapshot technique to be implemented as future work for data-window access (see Section 6.3 in the paper [17]). Prior stream analytics can decides to deploy multiple instances of same stream analytics engine to run concurrent stream analytics but will also produce separate adjacency store.

Key-value store natively provides only three basic operations PUT, GET and DELETE operations. On the other hand, APPEND operation is a must-have requirement for graph systems in order to support the graph data ingestion. A new edge insertion potentially means appending a new entry in the neighbor list (the value) of source vertex (the key). Similarly a deletion of edge requires deleting one entry from neighbor list of the source vertex. Many prior graph analytics systems, such as Kineograph, Trinity etc., additionally implements a graph layer on top of a key-value store to support this APPEND operation. The append operation on top of key-value store works using copy-on-write in Trinity: an edge addition requires copy of the prior neighbor list (value) of the source vertex (key) from the memory, then appends a new neighbor to the neighbor list, and finally writes the whole neighbor list to a new memory location. Then, they provide some optimizations to reduce this migration on each append.

Clearly, key-value store is not natively suited for storing evolving graph data. On the other hand, GRAPHONE provides a native implementation of adjacency list with in-built support for APPEND operation without using copy-on-write approach.

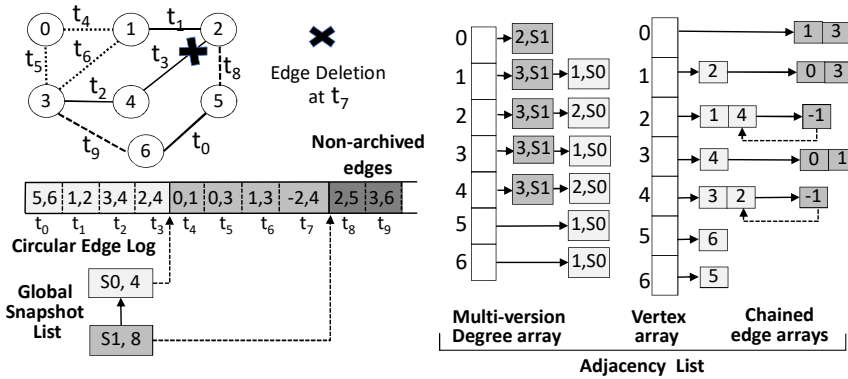
Finally, GRAPHONE data-store brings all different classes of analytics within a single system by allowing them to be performed concurrently from the same data-store using GraphView interface. In this paper, we concentrate on the data-store design and optimization, and demonstrate through discussions and experiments that its data-store abstraction and optimization techniques can be integrated to many classes of analytics engines to bring performance improvement. For this, we will be implementing actual analytics as suggested in prior works as much as possible for fair comparison while doing data-management our way.

## 4 HYBRID STORE AND VERSIONING

In this section, we first present internals of hybrid store followed by discussion on Versioning.

### 4.1 Hybrid Store

The hybrid store design presented in Fig. 6 consists of a small *circular edge log* that is used to record the latest updates in the edge list format. For deletion cases, we use tombstones, specifically the edge log adds a new entry but the most significant bit (MSB) of the source vertex ID of the edge is set to denote its deletion as shown in Fig 6 for deleted edge (2, 4) at time  $t_7$ .



**Fig. 6:** The hybrid store for the data arrived from time  $t_0$  to  $t_9$ : The vertex array contains pointers to the first and the last block of each edge array, while degree array contains deleted and added edge counts. However, only the pointer to the first block in the vertex array, and total count in the degree array are shown for brevity.

The *adjacency store* keeps the older data in the adjacency list format. The adjacency store is composed of *vertex array*, per-vertex *edge arrays*, and multi-versioned *degree array*. The *vertex array* contains a per-vertex flag and pointers to the first and last block of the edge arrays. Addition of a new vertex is done by setting a special bit in the per-vertex flag. Vertex deletion sets another bit in the same flag, and adds all of its edges as deleted edges to the edge log. These bits help GRAPHONE in garbage collecting the deleted vertex ID.

The *edge arrays* contains per-vertex edges of the adjacency list. It may contain many small edge blocks, each of which contains a count of the edges in the block and a memory pointer to the next block. The connection of edge blocks are referred to as *chaining*. An edge addition always happens at the end of the edge array of each vertex, which may require the allocation of a new edge block and linked to the last block. Fig. 6 shows chained edge arrays for the vertices with ID 1 to 4 for data updates that arrive in between  $t_4$  to  $t_7$ . The adjacency list treats an edge deletion as an addition but the deleted edge entry in the edge array keeps the negative position of the original edge, while the actual data is not modified at all, as shown for edge (2, 4). As a result, deletion never breaks the convergence of a previous computation as it does not modify the dataset of the computation.

The *degree array* contains the count of neighboring edges of each vertex, and is most important data structure to support multi-versioning of adjacency list. Since a degree array from an older adjacency store snapshot can identify the edges to be accessed even from the latest edge arrays due to the latter's append-only property. We exploit this property, and customize the degree structure to provide snapshot capability even in presence of data deletion. Hence, the degree array in GRAPHONE is *multi-versioned* to support adjacency store snapshots. It keeps the total added and deleted edge counts of each vertex. Both counts help in efficiently getting the valid neighboring edges, as a client can do the exact memory allocation (refer to the *get-nebrs-\**( $\ast$ ) API in Table 3). When an edge is added or deleted for a vertex, a new entry is added for this vertex in the degree array in each epoch. Two different versions S0 and S1 of the degree array are shown in Fig. 6 for two epochs  $t_0 - t_3$  and  $t_4 - t_7$ .

One can note that degree nodes are shared across epochs if there is no later activity in a vertex. For example, the same degree nodes for vertices with ID 5 and 6 are valid for both epochs in Fig. 6. The degree array nodes of an older versions are garbage collected when the corresponding adjacency store snapshot retires, i.e., not being used actively by any analytics, and is tracked using reference counting mechanism through the global snapshot list, which will be discussed shortly.

For example, if snapshot  $S_0$  is retired, then the degree nodes of snapshot  $S_0$  for vertices with ID  $1 - 4$  can be reused by later snapshots (e.g.  $S_2$ ). The garbage collection of degree node is part of archiving phase (see §5.1.2).

The *vertex array* is separated from the degree array for easier implementation of multi-versioning of the adjacency list, as the degree array nodes corresponding to older snapshots will be freed and recycled when the corresponding snapshot retires without affecting other data structures, such as edge arrays which are actively accessed by analytics.

The *global snapshot list* is a linked list of snapshot objects to manage the relationship between the edge log and adjacency store at each epoch. Each node contains an absolute offset to the edge log where the adjacency list snapshot is created, and a reference count to capture the number of views using this adjacency list snapshot. A new entry in the global snapshot list is created after each epoch, and it implies that the edge log data of the last epoch has been moved to the adjacency store atomically, and is now visible to the world.

**Weighted Graphs.** Edge weights are generally embedded in the edge arrays along with the destination vertex ID. Some graphs have static weights, e.g., an edge weight in an enterprise network can represent the network speed between the two nodes. A weight change is then treated internally as an edge deletion followed by an edge addition. On the other hand, if edge weights are dynamic, such as network data flow, then such weight changes are suited for various stream analytics if kept for a configurable time window, e.g., anomaly detection in the network flow. In this case GRAPHONE is configured to treat weight changes as a new edge to aid such analytics. Thus, streaming graph will not use the tombstones that we discussed in the beginning of this sub-section.

**An Example of Data Flow.** The data flow of GRAPHONE is presented in Fig. 6. There are addition and deletion of ten edges from time  $t_0$  to  $t_9$ , and two snapshots of the adjacency list are created. From time  $t_0$  to  $t_3$ , all the incoming edges are logged in the edge log. At  $t_4$  archiving starts moving these edge log data to the adjacency list in parallel to logging. During  $t_4$  to  $t_7$ , new edges are logged in the edge log. GRAPHONE again starts moving the last four updates to the adjacency list, while new arrivals are getting logged in the edge log from  $t_8$  onwards. At the end of  $t_9$ , the edge log contains two edges that are yet to move to the adjacency store.

## 4.2 Dual Versioning and Data Overlap

GRAPHONE uses *dual versioning* to create the instantaneous read-only graph views (snapshot isolation) for data analytics by exploiting both the *fine-grained* versioning property of the edge log, and the *coarse-grained* versioning capability of the adjacency list format. It should be noted that the adjacency list provides one version per epoch, while the edge log supports multiple versions per epoch, as many as the number of edges arrived during the epoch. So the dual versioning provides many versions within an epoch which is the basis for static views, and should not be confused with the adjacency list snapshots. In Fig. 6, static view at the time  $t_6$  would be adjacency list snapshot  $S_0$  plus the edges from  $t_4 - t_6$ .

A small amount of *data overlap* between the two stores keeps the composition of the view intact, and makes the view accessible even when the edge log data is moved to the adjacency store to create a new adjacency list version. Thus both stores have the copy of a few epochs of the same data. For one or more long running iterative analytics, we may use the durable edge log or a private copy of non-archived edges to provide data overlap, so that analytics can avoid interfering with data management operations of the edge log.




Fig. 7: Circular Edge log design showing various offset or markers. Markers for durable phase are similar to archiving and are omitted.

## 5 DATA MANAGEMENT AND OPTIMIZATIONS

Data management faces the key issues of minimizing the size of non-archived edges, providing atomic updates, data ordering, and cleaning of older snapshots. Addition and deletion of vertices and edges, and edge weight modification are all considered as an atomic update.

### 5.1 Data Management Phases

Fig. 5 depicts the internals of the data management operations. It consists of four phases: *logging*, *archiving*, *durable* and *compaction*. Client threads send updates, and the logging to the edge log happens in the same thread context synchronously. The archiving phase moves the non-archived edges to the adjacency store using many worker threads, and one of them assumes the role of the master, called the archive thread. The durable phase happens in a separate thread, while compaction is multi-threaded but happens much later.

A client thread wakes up the archive thread and durable thread to start the archiving and durable phases when the number of non-archived edges crosses a threshold, called *archiving threshold*. The logging phase continues as usual in parallel to them. Also, the archive thread and durable thread check if any non-archived edges are there at the end of each phase to repeat their process, or wait for work with a timeout.

The edge log has a distinct offset or marker, *head*, for logging, which is incremented every time an edge is ingested as shown in Fig. 7. For archiving, GRAPHONE manages a pair of markers, i.e. the archiving operation happens from the *tail archive marker* to the *head archive marker*, because the *head* will keep moving due to new updates. The durable phase also has a pair of markers to work with. Markers are always incremented and used with the modulo operator.

**5.1.1 Logging Phase.** The incoming update is converted to numerical identifiers, if required, and acquires an edge list format. The mapping between vertex label to vertex ID is managed using a hashmap. The reverse mapping to map the vertex ID to vertex name is managed internally. Many times, the incoming data may directly contain a unique pair of vertex ID in an edge. For such cases, user can configure GRAPHONE to directly use the existing ID space. In either case, a unique spot is claimed within the edge log by the atomic increment of the head, and the edge is written to a spot calculated using the modulo operation on the head. The source vertex also stores the operator (\$4), addition or deletion, along with the edges.

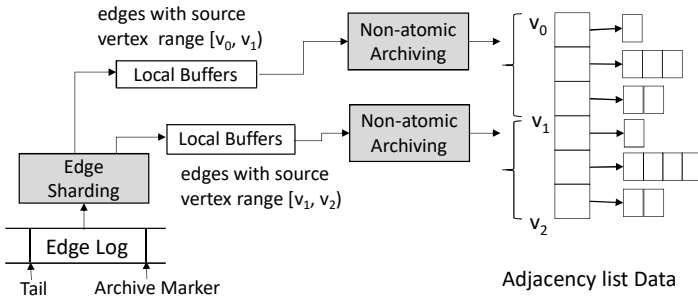
The atomicity of updates is ensured by the atomic increment of the head, and by writing the least significant bit of the edge log rewind count in the most significant bit of the destination vertex ID. This is written after writing the source vertex, and alternates between zero and one each time the edge log rewinds. Thus a reader of the edge log can figure out if it is reading the latest data, old data, or a partial update. Thus for each edge in the edge log, the MSB of source vertex contains the

deletion information while the MSB of destination vertex contains the rewind information, and are written in this order.

The edge log is automatically reused in the logging phase due to its circular nature, and thus is overwritten by newer updates. Hence the logging may get blocked occasionally if the whole buffer is filled as the archiving or durable phases may not be able to catch up. We keep sufficiently large edge log to avoid frequent blocking. In case of blocked client threads, they are woken up when the archiving or durable phases complete.

**5.1.2 Archiving Phase.** This phase moves the non-archived edges from the edge log to the adjacency store. This phase has many stages, and are discussed next.

**Edge Sharding.** A naive multi-threaded archiving, where each worker can directly work on a portion of non-archived edges, may not keep the data ordering intact. If a deletion comes after the addition of an edge within the same epoch, the edge may become alive or dead in the edge arrays depending on the archiving order of the two data points. To avoid this ordering problem, we propose *edge sharding* technique.



**Fig. 8:** Edge sharding separates the non-archived edges into many buffers based on their source vertex ID, so that the per-vertex edge arrays can keep the edge log arrival order, and enables non-atomic archiving.

An *edge sharding* stage in the archiving phase (Fig. 8) maintains per-vertex edges as per the edge log arrival to address the ordering problem. It shards the non-archived edges to multiple local buffers based on the range of their source vertex ID, i.e., each local buffer contains edges belonging to a few particular source vertices. Our implementation makes sure that the per-vertex edges in the local buffers are in the same arrival order as that of the edge log. This is implemented by scanning the edge log twice. In the first scan, each worker thread counts its contribution to all the local edge buffers, and in the second scan, the worker thread copies the edges to the corresponding slots in those local buffers.

For undirected graphs, the total edge count in the local buffer is twice of the non-archived edge count, as the ordering of reverse edges is also managed. For directed edges, both directions have their own local buffers. Hence each edge is stored twice, one in forward direction, another in reverse direction. However, GRAPHONE provides a way to avoid storing the edges twice, and are called *unigraph* in this paper.

**Non-Atomic Archiving.** Edge sharding has an additional advantage of avoiding the usage of atomic instructions for populating the edge arrays. i.e., the edges in each local buffer are archived in parallel without using any atomic instructions. A heuristic is required for workload distribution, as the equal division is not possible among threads, thereby the last thread may get more work assigned. To handle the workload imbalance among worker threads, we create a larger number of local buffers with smaller vertex range than the available threads, and assign different numbers of

local buffers to each thread so that each gets an approximately equal number of edges to archive. The idea here is to assign slightly more than equal work to each thread, so that all the threads are balanced while the last thread is either balanced or lightly loaded.

This stage allocates new degree nodes or can reuse the same from the older degree array versions if they are not being used by any analytics. We follow these rules for reusing the degree array from older versions. We track the degree array usage by analytics using reference counting per epoch [42], and can be reused if all static views created within that epoch have expired, i.e., the references are dropped to zero (not being used by any running analytics). It also ensures that a newly created view uses the latest adjacency list snapshot that should never be freed.

The stage then populates the degree array, and allocates memory for edge blocks that are chained before filling those blocks. If a previous edge array for a participating vertex exists, then the new edge block is chained to it, otherwise a new entry is created for the new edge block in the vertex array. We then create a new snapshot object, fill it up with relevant details, and add it atomically to the global snapshot list. At the end of the archiving phase, the archive thread sets the tail archive marker atomically to the value of the head archive marker, and wakes up any the blocked client threads.

*5.1.3 Durable Phase and Recovery.* The edge log data is periodically appended to a durable file in a separate thread context instead of logging immediately to the disk to avoid the overhead of IO system calls during each edge arrival. Also this will not guarantee durability unless `fsync()` is called. The logging uses buffered sequential write, and allows the buffer cache to work as spillover buffer for the access of non-archived edges if the edge log is over-written.

The durable edge log is a prefix of the whole ingested data, so GRAPHONE may lose some recent data in the case of an unplanned shutdown. The recovery depends on upstream backup that keep the latest data for some time, such as kafka [44], and replays it for the lost data, and creates the adjacency list on the whole data. Recovery is faster than building the data structures at an edge level, as only the archiving phase is involved working on bulk of data. Alternatively, persistent memory may be used for the edge log to provide durability at each update [47].

The durable phase also performs an incremental checkpointing of the adjacency store data from an old time-window, and frees the memory associated with it. This is useful for streaming data such as LANL network flow, where the old adjacency data can be checkpointed in disk, as the in-memory adjacency store within the latest time window is sufficient for stream analytics. By default, it is not enabled as it will break access of whole data in case of batch analytics or full access by streaming analytics. During checkpointing the adjacency store, the vertex ID and length of the edge array are persisted along with edge arrays so that data can be read easily later, if required.

*5.1.4 Compaction Phase.* The compaction of the edge arrays removes deleted data from per-vertex edge array blocks up to the latest retired snapshot identified via the reference counting scheme discussed in §5.1.2. The compaction needs a similar reference counting for the private static views (§6.1). For each vertex, it allocates new edge array block and copies valid data up to the latest retired snapshot from the edge arrays, and creates a link to the rest of the original edge array blocks. The newly created edge array block is then atomically replaced in the vertex array, while freeing happens later to ensure that cached references of the older data are dropped. This phase is generally clubbed with archiving phase where the degree array is updated to reflect the new combination.

## 5.2 Memory Overhead and Optimizations

The edge log and degree array are responsible for versioning. The edge log size is relatively small as it contains only the latest updates which moves quickly to the base store, e.g, the archiving threshold of  $2^{16}$  edges translates to only 1MB for a plain graph assuming 8 byte vertex ID. Thus the

edge log is only several MBs. The memory in degree arrays are also reused (§5.1.2). This leaves us to memory analysis of edge arrays which may consume a lot of memory due to excessive chaining in their edge blocks. For example, GRAPHONE runs archiving phase for  $2^{16}$  times for Kron-28 graph if the archiving threshold is  $2^{16}$ . In this case, the edge arrays would consume 148.73GB memory and have average 29.18 chain per-vertex. We will discuss the graph datasets used in this paper shortly. If all the edges were to be ingested in one archiving phase, this static system needs only an average 0.45 chain and 33.80GB memory. The chain count is less than one as 55% vertices do not have any neighbor.

GRAPHONE uses two memory allocation techniques, as we discuss next, to reduce the level of chaining to make the memory overhead of edge arrays modest compared to a static engine. The techniques work proactively, and do not affect the adjacency list versioning. Compaction further reduces the memory overhead to bring GRAPHONE at par with static analytics engine, but is performed less frequently.

Table 2. Impact of two optimizations on the chain count and memory consumption of the edge arrays of the adjacency store on the kroncker (Kron-28) graph.

Optimizations	Chain Count		Memory Needed (GB)
	Average	Maximum	
Baseline System	29.18	65,536	148.73
+Cacheline memory	2.96	65,536	47.42
+Hub Vertex Handling	2.47	3,998	45.79
Static System	0.45	1	33.81

**Optimization #1: Cacheline Sized Memory Allocation.** Multiples of cacheline sized memory is allocated for the edge blocks. One cacheline (64 bytes) can store up to 12 neighbors for the plain graph of 32bit type, leaving the rest of the space for storing a count to track space usage in the block and a link to the next block. In this allocation method, the majority of the vertices will need only a few levels of chaining. For example, in a Twitter graph, 88.43% of the vertices will need at most 3 cachelines only, and so do 92.49% for Kron-28 graph. This optimization reduces the average chain count by 9.88 $\times$ , and memory consumption by 3.14 $\times$  in comparison to a baseline system as shown in Table 2. The baseline system uses a dynamic block size which is equivalent to the number of edges arrived during each epoch for each vertex.

**Optimization #2: Hub Vertex Handling.** A few vertices, called hub-vertices, have very high degree in a graph that follows power-law distribution [24]. They are very common in real-life graphs, such as for the twitter follower graph whose degree distribution is shown in Fig. 9. Such vertices are likely to participate in each archiving phase. Hence they will have a lot of chaining in their edge arrays, and the aforementioned memory management technique alone is not enough. In this case, we allocate in multiples of 4KB page-aligned memory for vertices that already have 8,192 edges or if the number of neighbors in any archiving phase crosses 256. The average chain count is reduced to 2.47, leading to a further reduction in memory utilization by 1.63GB as listed in Table 2. The reduction in memory overhead due to hub vertex handling are modest on top of cacheline sized memory allocations. This is because the number of hub vertices are very small in comparison to total vertex of the graph. Therefore, even if one varies the threshold to identify a hub vertex, the performance will remain similar to the cacheline sized memory (Fig. 18).

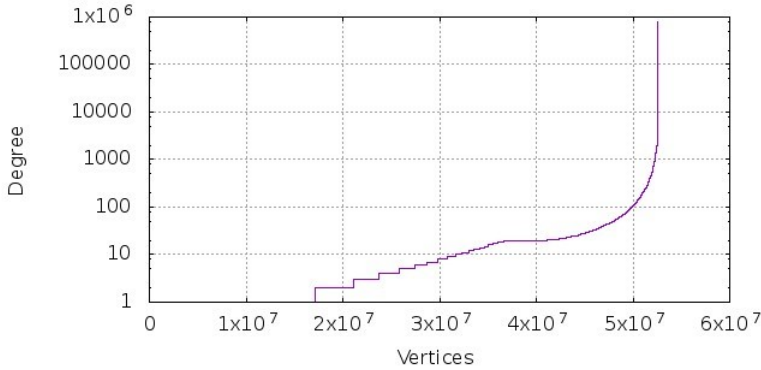


Fig. 9: Degree distribution of Twitter graph.

## 6 GRAPHVIEW AND DATA VISIBILITY

GraphView data abstraction hides the complexity of the hybrid store by embedding *data visibility* option and simplifying data access by providing APIs as shown in Tables 3, 4 and 5. The *static view* is suited for batch analytics and queries, while the *stream view* for stream processing. Both offer diverse data access at two granularities of data visibility of updates. The access of non-archived edges provides data visibility to analytics at the edge level granularity.




Fig. 10: GRAPHONE hybrid store illustrating various views with two adjacency store versions, S0 and S1, with a small edge log

A number of views may co-exist at any time without incurring much memory overhead, as the view data is composed of the same adjacency store and non-archived edges as shown in Fig. 10. Each analytics has its own view instance and executes in separate thread contexts and in parallel to any data-management phases and their threads. Generally, the registration of the batch or stream analytics triggers a thread creation that calls a function pointer specified in the create-\* or reg-\* APIs (Table 3, 4, and 5), and passing the handle. The function pointer is omitted for brevity from the API signature. Caller is responsible for all the computation inside the callback, including creation of more threads to achieve parallelism if required, and a return assumes that computation has been finished. In this section, we concentrate on data-access APIs, and leave the development or integration of a full-fledged analytics system to simplify some of the tasks of actual analytics as a future work.

Due to the cost of indexing the non-archived edges, GraphView provides an option to trade-off the granularity of data visibility to gain analytics performance by opting to compute *only* on the latest adjacency store snapshot. Further, one can use vertex-centric compute model [82] on the



adjacency list plus edge-centric compute model [45, 75, 96] on non-archived edges, so there is no need to index the latter as plotted later to find its optimal minimum size (Fig. 16).

### 6.1 Static View

Batch analytics and queries prefer snapshots for computation, which can be created in real-time using *create-static-view()* API, and then data may be accessed using other set of APIs as shown in Table 3. A static view is represented by an opaque handle that identifies the view composition, i.e., the non-archived edges and the latest adjacency list snapshot, and serves as input to other static view APIs. A created handle should be destroyed using *delete-static-view()*, otherwise, there would be memory leak as each static view may contain at least a degree array. However, the handle can be wrapped in a reference counting objects for auto deletion when the handle goes out of scope. Based on the input supplied to *create-static-view()* API, many types of static view are defined.

Table 3. Static View APIs

snap-handle	<i>create-static-view</i> (global-data, simple, private, stale)
status	<i>delete-static-view</i> (snap-handle)
count	<i>get-nebr-length</i> -{in/out}(snap-handle, vertex-id)
count	<i>get-nebrs</i> -{in/out}(snap-handle, vertex-id, ptr)
count	<i>get-nebrs-archived</i> -{in/out}(snap-handle, vertex-id, ptr)
count	<i>get-non-archived-edges</i> (snap-handle, ptr)

---

#### Algorithm 1 Traditional BFS using static view APIs

---

```

1: handle ← create-static-view(global-data, private=true, simple=true)
2: level = 1; active-vertex = 1; status-array[root-vertex] = level;
3: while active-vertex do
4:   active-vertex = 0;
5:   for vertex-type v = 0; v < vertex-count; v++ do ▷ Can be parallelized
6:     if status-array[v] == level then
7:       degree ← get-nebrs-out(handle, v, nebr-list);
8:       for j=0; j < degree; j++ do
9:         w ← nebr-list[j];
10:        if status-array[w] == 0 then
11:          status-array[w] ← level + 1; ++active-vertex;
12:        end if
13:      end for
14:    end if
15:  end for
16:  ++level;
17: end while
18: delete-static-view(handle)

```

---

**6.1.1 Basic Static View.** This is the most fundamental view of data, where the data is accessed directly from underlying storage. The main low-level API are: *get-nebrs-archived-\**() that returns the reference to the per-vertex edge array; and *get-non-archived-edges*() that returns the non-archived

edges. These two APIs are very useful for advanced users and for higher level library development which prefer more control and performance.

On the other hand, it also provides a high-level API, *get-nebrs-\**(), that returns the neighbor list of a vertex by combining the adjacency store and the non-archived edges in a user supplied memory buffer. It provides simplicity to analytics writers as they don't have to know the exact internal structures of the underlying data of GRAPHONE. However, there is a performance cost associated with usage of this API, as it will have to repeatedly scan the non-archived edges for finding the neighboring edges of each vertex, and also an additional step of data copying to user-supplied buffer. Though this API will be preferable by queries with high selectivity that only need to scan the non-archived edges for one or a few vertex, e.g. 1-hop query, and is not apt for long running analytics.

The implementation of *get-nebrs()* for the non-deletion case is a simple two step process: copy the per-vertex edge array to the user supplied buffer, followed by a scan of the non-archived edges to find and add the rest of the edges of the vertex to the buffer. For the deletion case, both the steps track the deleted positions in the edge arrays, and the last few edges from edge arrays and/or non-archived edge log are copied into those indexes of the buffer.

One can also pass *simple=true* in the *create-static-view()* to create a temporary in-memory adjacency list from the non-archived edges for optimizing the performance of *get-nebr-\**() API. Algorithm 1 shows a simplified BFS (push model) implementation using the high-level APIs of basic static view.

**6.1.2 Private Static View.** For long running analytics, keeping basic static views accessible have some undesirable impacts: (1) all the static views may have to use the durable edge log if the corresponding non-archived edges in the edge log has been overwritten; (2) the degree array cannot be reused in the archiving phase as it is still in use. To solve this, one can create a *private static view* by passing *private=true* in the *create-static-view()* API. In this case, a private copy of the non-archived edges and the degree array are kept inside the view handle with their global references dropped to make it independent from archiving.

The private static view approach of batch analytics is more flexible than static analytics engine where latter converts the whole data that takes much longer to complete [60] than the analytics run-time. In contrast, a private view only needs to copy a degree array whose cost will be much lower. This view is also better than other dynamic graph system that disallows the users to choose fine-grained control on snapshot creation.

Creation to many private static views may introduce memory overhead. To avoid this, a reference of the private degree array is kept in the snapshot object and is shared by other static views created within that epoch, and are locally reference counted for freeing. Thus, creating many private views within an epoch has overhead of just one degree array. However, creating many private static views across epochs may still cause the memory overhead, if older views are still being accessed by long running analytics. This also means that the machine is overloaded with computations, and they are not real-time in nature. In such a case, a user may prefer to copy the data to another machine to execute them.

**6.1.3 Stale Static View.** Many analytics are fine with data visibility at coarse-grained ingestions, thus some stale but consistent view of the data may be better for their performance. In this case, passing *stale=true* to *create-static-view()* API returns a handle corresponding to the snapshot of the latest adjacency list only. Thus, same set of APIs can be used by users to access the slightly stale data, but performance will automatically be of batched update case. This view can be combined with private static view where degree array will be copied.

## 6.2 Stream View

As we discussed in the beginning of this section, all stream computations runs in a separate thread contexts, i.e. one thread per stream analytics (can be parallelized to execute the core logic of stream analytics) while logging and archiving keep working in parallel in their own thread contexts. The stream view APIs as listed in Table 4 and 5 simplify the data access in presence of the data ingestion. For a stream analytics, *window-sz* is the size on which a stream analytics want the data-state (adjacency store) to be accessible for computation, and can stretch up to the beginning of the stream, if not specified. The *batch-sz* means stream analytics wants to perform incremental computation only if *batch-sz* count of edges have been accumulated since last incremental compute.

All stream analytics in GRAPHONE follow a pull method, i.e., the analytics *pulls* new snapshot by calling *update-\*-view()* API at the end of incremental compute to perform the next phase of incremental compute. If there is no new data arrived from last time the API was called or if the size of new data is less than *batch-sz* (*reg-\*-view()* API), this API will get blocked and will return only when new data equal or greater than *batch-sz* is available to continue the computation. Though, the focus of this sub-section is on data access for stream computations, one can pass appropriate flag value to *reg-\*-* APIs listed in Table 4 and 5 to configure data visibility of stream view.

Also, checkpointing the computation results and the associated data offset is the responsibility of the stream engine, so that the long running computation can be resumed from that point onwards in case of a fault. We now present data access patterns and corresponding APIs for two classes of stream analytics: *stateless* and *stateful* stream computations.

**6.2.1 Stateless Stream Processing.** A stateless computation, e.g. counting incoming edges (aggregation), only needs a batch of new edges. The APIs are listed in Table 4. The computation can be registered using the *reg-stream-view()* API, and the returned handle contains the batch of new edges. Algorithm 2 shows how one can use the API to do stateless stream computation. The handle also allows a pointer to point to analytics results to be maintained by the stream compute implementation. The implementation also needs to checkpoint only the edge log offset and the computation results as GRAPHONE keeps the edge log durable.

Table 4. Stateless Stream View APIs

stream-handle	<i>reg-stream-view(global-data, window-sz, batch-sz)</i>
status	<i>update-stream-view(stream-handle)</i>
status	<i>unreg-stream-view(stream-handle)</i>
count	<i>get-new-edges-length(stream-handle)</i>
count	<i>get-new-edges(stream-handle, ptr)</i>

An extension of the model is to process on a data window instead on the whole arrived data. For sliding window implementation, GRAPHONE manages a cached batch of edge data around the start marker of the data window in addition to the batch of new edges. The old cached data can be accessed by the analytics for updating the compute results, e.g., subtracting the value in aggregation over the data window. The cached data is fetched from the durable edge log, and shows sequential read due to the sliding nature of the window. A tumbling window implementation is also possible where the batch size of new edges is equal to the window size, and hence does not require older data to be cached. Additional checkpointing of the starting edge offset is required along with the edge log offset and computation results.

**6.2.2 Stateful Stream Processing.** A complex stream computation, such as graph coloring [76], is stateful that needs the streaming data and complete base store to access the computational

**Algorithm 2** A stateless stream compute skeleton

---

```

1: handle ← reg-stream-view(global-data, batch-sz=10s)
2: init-stream-compute(handle)
3: while true do
4:   if update-stream-view(handle) then
5:     count = get-new-edges(handle, new-edges)
6:     for j=0; j < count; j++ do
7:       | do-stream-compute(handle, new-edges[j])
8:     end for
9:   end if
10: end while
11: unreg-sstream-view(handle)

```

---

▶ Application specific  
 ▶ Or application specific exit criteria  
 ▶ Can be parallelized  
 ▶ Or any method

state of the neighbors of each vertex. A variant of static view is better suited for it because its per-vertex neighbor information eases the access of the computational state of neighbors. The APIs are listed in Table 5. The stateful computation is registered using *reg-sstream-view()*, and returns *sstream-handle*. For edge-centric computation, the handle also contains a batch of edges to identify the changed edges. For vertex-centric computation, the handle contains per-vertex one-bit status to denote the vertex with edge updates that can be identifies using the *has-vertex-changed()* API. This is updated during *update-sstream-view()* call that also updates the degree array. Algorithm 3 shows an example code snippet.

Table 5. Stateful Stream View APIs

<i>sstream-handle</i>	<i>reg-sstream-view</i> (global-data, window-sz, v-or-e-centric, simple, private, stale)
status	<i>update-sstream-view</i> ( <i>sstream-handle</i> )
status	<i>unreg-sstream-view</i> ( <i>sstream-handle</i> )
bool	<i>has-vertex-changed</i> ( <i>sstream-handle</i> , vertex-id)
count	<i>get-nebr-length</i> -{in/out}( <i>sstream-handle</i> , vertex-id)
count	<i>get-nebrs</i> -{in/out}( <i>sstream-handle</i> , vertex-id, ptr)
count	<i>get-nebrs-archived</i> -{in/out}( <i>sstream-handle</i> , vertex-id, ptr)
count	<i>get-non-archived-edges</i> ( <i>sstream-handle</i> , ptr)

As the degree array plays an important role for a stateful computation due to its association with the static view, using an additional degree array at the start marker of the data window eases the access of the data within the window from the adjacency store. The *sstream-handle* manages the degree array on behalf of the stream engine, and internally keeps a batch of cached edges around the start marker of the window to update the old degree array. The *get-nebrs-\**() function returns the required neighbors only. Checkpointing the computational results, the edge log offset at the point of computation, and window information is sufficient for recovery.

**6.2.3 Discussion.** In this sub-section, we bring the differences to the existing graph stream analytics. The major difference arises for stateful computation, as computation requires the access of prior data in the form of adjacency store which is built in separate thread context. While GRAPHONE performs archiving in separate thread context than that of the analytics, prior steam analytics engine perform both of them in sequence in the same thread context. This difference can easily be explained using Algorithm 3. Prior stream engines will perform actual archiving phase (building/updating adjacency

---

**Algorithm 3** A stateful stream compute (vertex-centric) skeleton

---

```

1: handle ← reg-sstream-view(global-data, v-centric, stale=true)
2: init-sstream-compute(handle)
3: while true do
4:   if update-sstream-view(handle) then
5:     for v=0; v < vertex-count; v++ do
6:       if has-vertex-changed(handle, v) then
7:         do-sstream-compute(handle, v)
8:       end if
9:     end for
10:  end if
11: end while
12: unreg-sstream-view(handle)

```

▶ Application specific  
 ▶ Or application specific exit criteria  
 ▶ Can be parallelized  
 ▶ Application specific

---

store indexes) at line number 4 instead of calling *update-sstream-view(handle)*, while GRAPHONE does very lightweight metadata update at line number 4 by calling the above mentioned function so that view handle can point to the latest adjacency store as update to adjacency store happens in different thread context. Rest of the lines in the Algorithm 3 remains same. Thus, it will be possible to integrate many stream analytics systems with GRAPHONE data-store. This difference also enables concurrent execution of multiple stream analytics from the same adjacency store.

For stateless computation, there is no difference with the existing stream engines, as analytics don't require access to adjacency store at all. If these are the only analytics running in the system, GRAPHONE can be configured to not perform the archiving phase. Thus, prior streams systems will be identical to GRAPHONE, and both of them exactly looks like Algorithm 2.

### 6.3 Historic Views

GRAPHONE provides many views from recent past, but it is not designed for getting arbitrary historic views from the adjacency store. However, durable edge log can provide the same using *get-prior-edges()* API (Table 6) in edge list format as it keeps deleted data, behaving similar to existing data stores [14, 25].

Table 6. Historic View APIs

count	<i>get-prior-edges</i> (global-data, start, end, ptr)
-------	---

Moreover, in case of no deletion, one can create a degree array at a durable edge log offset by scanning the durable edge log, and the resultant degree array will serve older static or stream view from the adjacency store to gain insights from the historical data. For data access from a historical time-window in this case, one need to build two degree arrays at both the offsets of time-window of interest of the durable edge log.

## 7 EVALUATIONS

GRAPHONE is implemented in around 16,000 lines of C++ code including various analytics. It supports plain graphs and weighted graphs with either 4 byte or 8 byte vertex sizes. We store the fixed weights along with the edges, variable length weights in a separate weight store using indirection. Any type of value can be stored in place of weight such as integers, float/double, timestamps, edge-id or any custom weight as the code is written using C++ templates. So one

Table 7. Graph datasets showing vertex and edge counts in **millions**. For deletions see §7.3.

Graph Name	Graph Type	Vertex Counts (in Millions)	Edge Counts (in Millions)
LANL	Directed	0.16	1,521.19
Twitter	Directed	52.58	1,963.26
Friendster	Directed	68.35	2,586.15
Subdomain	Directed	101.72	2,043.20
Kron-28	Undirected	256	4,096
Kron-21	Undirected	2	32

can write a small plug-in describing the weight structures and other functions, and GRAPHONE would be ready to serve a custom weight. All experiments are run on a machine with 2 Intel Xeon CPU E5-2683 sockets, each having 14 cores with hyper-threading enabled. It has 512GB memory, Samsung NVMe 950 Pro 512GB, and CentOS 7.2. Prior results have also been performed on the same machine.

We choose data ingestion, BFS, PageRank and 1-Hop query to simulate the various real-time usecases to demonstrate the impact of GRAPHONE on analytics. BFS and PageRank are selected because many real-time analytics are iterative in nature, e.g. shortest path, and many prior graph systems readily implement them for comparison. 1-Hop query accesses the edges of random 512 non-zero degree vertices and sums them up to make sure we access them all. 1-Hop query simulates many small query usecases, such as listing one’s friends, or triangle completion to get friend suggestions in a social graph, etc.

**Datasets.** Table 7 lists the graph datasets. Twitter [3], Friendster [1] and Subdomain [4] are real-world graphs, while Kron-28 and Kron-21 are synthetic kronecker graphs generated using graph500 generator [27], all with 4 byte vertex size and without any weights. We also use these graphs for running comparison with LLAMA, that always expects weighted graphs. So, these graphs are also converted to weighted graph with 4 bytes of random weights assigned by a tool developed as part of LLAMA software suite. LANL network flow dataset [84] is a weighted graph where vertex and weight sizes are 4 bytes and 32 bytes respectively, and weight changes are treated as new streaming data. We run experiment on first 10 days of data. We test deletions on a weighted RMAT graph [15] generated with [62] where vertex and weight sizes are 8 bytes. It contains 4 million vertices, and 64 million edges, and a update file containing 40 million edges out of which 2,501,937 edges are for deletions.

These graphs have been taken from various domains, such as twitter and friendster from social media, subdomain from web domain, while LANL is captured from an enterprise cyber network apart from many synthetic graphs. We believe these are diverse enough to capture a wide variety of applications. We have two copies of each dataset, one in binary edge list format and another is raw text edge list format. During the ingestion, vertex name to vertex ID conversion was not needed for binary graph data as we directly used the vertex ID supplied with these datasets. This convention is also followed by other graph systems. We discuss two cases of raw text data in §7.2. All the edges will be stored twice in the adjacency list: in-edges and out-edges for directed graphs, and symmetric edges for undirected graphs. No compaction was running in any experiments unless mentioned.

Table 8. Different data management rates in **millions edges/s (M/s)** on various graph dataset. The results show that the ingestion rate would be upper and lower bounded by the logging and archiving rate. See §7.3 for results on datasets containing deletions.

Graph Name	Individual Phases (M/s)		In-Memory Rate (M/s)		Ext-Memory Rate (M/s)		Compaction Rate (M)
	Logging	Archiving	Ingestion	Recovery	Ingestion	Recovery	
LANL	35.98	28.91	26.99	30.23	25.26	29.48	41.85
Twitter	82.62	47.98	66.39	71.28	61.13	71.87	541.71
Friendster	82.85	49.32	60.40	95.78	58.35	95.44	520.65
Subdomain	82.86	43.43	68.25	180.75	61.54	151.96	444.84
Kron-28	79.23	43.68	52.39	116.18	49.70	107.61	798.91
Kron-21	78.91	78.40	58.31	90.44	57.02	66.66	1011.68

### 7.1 Data Ingestion Performance on Binary Data

In this sub-section, we present data ingestion when they already contain the data in binary edge format, i.e. GRAPHONE does not create any vertex ID from the incoming data. This will depend on the data source, however, this is a perfect scenario to test the performance of various stages of data management pipeline, and end-to-end ingestion rate as data is already in a ready state to be ingested in GRAPHONE.


**Logging and Archiving Rate.** Logging to edge log is naturally faster, while archiving rate depends upon the archiving threshold. Table 8 lists the logging rate of a thread, and archiving rate at the archiving threshold of  $2^{16}$  edges for our graph dataset. A thread can log close to 80 million edges per second, while archiving rate is only around 45 million edges second at the archiving threshold for most of the graphs. Both the rates are lower for LANL graph, as the weight size is 32 bytes, while others have no weights.

**Ingestion Rate.** It is defined as single threaded ingestion to the edge log at one edge at a time, and leaving the archive thread and durable phase to automatically change with the arrival rate. The number is reported when all the data are in the adjacency store, and persisted in the NVMe ext4 file. GRAPHONE achieves an ingestion rate of more than 45 million edges per second, except LANL graph. The ingestion rate is higher than archiving rate (at the archiving threshold) except in Kron-21, as edges more than the archiving threshold are archived in each epoch due to higher logging rate. This indicates that GRAPHONE can support a higher arrival rate as archiving rate can dynamically boost with increased arrival velocity. The Kron-21 graph is very small graph, and the thread communication cost affects the ingestion rate.

**Compaction Rate.** We run compaction as a separate benchmark after all the data has been ingested. The graph compaction rate is 345.53 million edges per second for the RMAT graph which has more than 2.5 million deleted edges out of total 104 million edges. Results for other graphs are shown in Table 8. The poor rate for LANL graph is due to long tail for compacting edge arrays of few vertices. As shown later in Fig. 19, the compaction improves the analytics performance where the static GRAPHONE serves compacted adjacency list as it had no link in its edge arrays.

**Durability.** The durable phase has less than 10% impact on the ingestion rate. Table 8 shows the in-memory ingestion rate and can be compared against that of GRAPHONE, which uses NVMe SSD for durability. This is because durable phase runs in a separate thread context, and exhibits only sequential write. The NVMe SSD can support up to 1500MB/s sequential write and that is sufficient for GRAPHONE as it only needs smaller write IO throughput, as shown in Fig. 11 for Friendster graph. This indicates that a higher logging rate can easily be supported by using a NVMe SSD.

**Recovery.** Recovery only needs to perform archiving phase at bulk of data. As we will show later in Fig. 16, the archiving is fastest when around  $2^{27}-2^{31}$  edges are cleaned together. Hence we take



**Fig. 11:** Write throughput for friendster in GRAPHONE comparing against average requirement and maximum available in an NVMe

the minimum of this size as recovery threshold to minimize the memory requirement of IO buffer and the recovery time, and also gets an opportunity to pipeline the IO read time of the data with recovery. Table 8 shows the total recovery time, including data read from NVMe SSD after dropping the buffer cache. Clearly, GRAPHONE hides the IO time when compared against in-memory recovery. The recovery rate varies a lot for different graph due to different distribution of the batch of graph data that has profound impact on parallelism and hence locality access of edge arrays.

## 7.2 Data Ingestion Performance on Text Data

In this sub-section, we expect the incoming data in raw text format without having any vertex ID, such as LANL network logs. Ingesting raw text data involves parsing the information, and identified key information will be converted to vertex ID using hashmap of GRAPHONE to create a binary edge. It will also need to parse the text to identify the edge weights that will be also be stored within the binary edge. Here, the additional cost of parsing the data and interaction with hashmap will reduce the performance of logging phase. So, we only demonstrate the performance of logging phase, as performance of other phase will remain same.

We pick two kinds of text data to ingest. The *first* kind is purely text format where the vertices of edges are in the string format, and inserting them to GRAPHONE additionally requires interaction with hashmap, where latter converts the strings to vertex ID. We pick LANL raw text data in this case. The hashmap is taken from open-sourced Intel Thread Building Block library [70]. *Second* kind is just a pair of vertex ID which only requires parsing the data but no conversion to vertex ID is required. Therefore interaction with hashmap is not required. We pick other graphs including Twitter, Friendster etc. in raw text graph to demonstrate the performance of logging phase. The second kind is selected to gain an intuition about the cost of vertex ID creation and lookup versus parsing cost.

Table 9. Impact of raw text data on Logging Phase. Logging rate is per-thread.

Graph Name	Logging Rate (in Million Edges/sec)	
	Binary Data	Raw Data
LANL	35.98	1.25
Twitter	82.62	3.10
Friendster	82.85	3.05
Subdomain	82.86	3.11
Kron-28	79.23	3.10
Kron-21	78.91	3.23



Table 9 shows that parsing the raw text takes most of time for Kron-28 graph. For LANL graph, the parsing and interacting with the hashmap takes most of the time. The numbers in Table 9 is for single threaded logging phase. We can confirm that they do show sub-linear scaling with increasing thread count as the hashmap can support parallel ingestion and retrieval. Specifically, with 8 threads for 8 LANL text files shows logging rate close to 5.55 million edges per second, a speedup of 4.44 $\times$ . The sub-linear scaling is due usage of the atomic instructions while claiming the spot in the edge log.

However, we feel that it is not the ideal setup for parallel ingestions in the real-world. For example, the actual data sources for GRAPHONE may not necessarily be the files, rather they could be the database log changes, Kafka etc. In case of Kafka, data is already sharded using shared-nothing partitioning technique, therefore, we believe logging data from different shards to the same edge log to create a global order is not worth, and rather separate edge logs would be best way to go forward. We confirm that batching in independent edge logs show linear scaling as per our testing.

### 7.3 Graph Systems Performance

We choose different classes of graph systems to compare against GRAPHONE. Stinger is a dynamic graph system, Neo4j and SQLite are graph databases, Galois and static version of GRAPHONE are static graph systems, and LLAMA [59] as a snapshot system. Except stream computations, all the analytics in this section are performed on private static view containing no non-archived edges as it is created at the end of the ingestion.

It should be noted that unlike these systems, the main contribution of GRAPHONE is its data-store. As the scope of this work is to show the advantage of our efficient data-store, we try to mimic the implementation strategy of prior work as much as possible when comparing the analytics performance, thus motivating the need to integrate such data-store in the existing real-time analytics engines empirically. We also compare the data ingestion performance to demonstrate the superior design of GRAPHONE.

**7.3.1 Snapshot Based System.** This section compares against graph batch analytics systems that create snapshots at bulk updates. We choose in-memory LLAMA [59] system for the comparison which is available in github as open-source. Since, LLAMA ingests data from file by using *fread* call which uses buffered IO system call API to batch the edges, we create *tmpfs* file-system so that the data read from files could be done faster to test the data ingestion. We also modified GRAPHONE to do ingestion in the same way by calling *fread* for each edge from files hosted in *tmpfs*. We use the same set of graphs as listed in Table 7 with few differences because of the way LLAMA is implemented as discussed next. Though Both of these difference are unfair to GRAPHONE, we could not find a better way of comparison.

*First*, LLAMA only takes weighted graph with 4 bytes of weight, but stores weights in a structure parallel to adjacency store. So when BFS and PageRank runs, it does not access the weight structure at all. GRAPHONE stores the 4 byte weights in the adjacency store directly, hence will be accessed while running BFS and PageRank even though not required. The required input binary file is created from text files using a utility provided by LLAMA that inserts a random 4 bytes weights for each edges. *Second*, LLAMA also creates one snapshot for a batch of input that are presented as a file. Interestingly, LLAMA's multi versioned vertex array consumes a lot of memory, therefore it can not create even 512 snapshots for most of the graphs except Kron-21 graph. So, we choose to create 256 snapshots, but Kron-28 still runs out of memory, so we drop its comparison. In GRAPHONE, we create one snapshot for each 65,536 batch of edges (the archiving threshold), which will be equal to 32,768 snapshots for Friendster graph.




Fig. 12: Data ingestion rate and analytics performance comparison of GRAPHONE against LLAMA

Despite the advantage of LLAMA over GRAPHONE in this setup, GRAPHONE achieves huge speedup over LLAMA as shown in Fig. 12 for data ingestion time, BFS and PageRank. GRAPHONE achieves an average speedup of 6.38 $\times$ , 1.66 $\times$ , 2.43 $\times$  for data ingestion, BFS and PageRank respectively for all graphs except Kron-21 for which the speedups are much higher at 31.47 $\times$ , 24.79 $\times$  and 1.23 $\times$  respectively. There is just one exception of BFS performance in Subdomain graph which has 140 levels in its BFS tree. The LLAMA has a highly optimized BFS implementation as it uses, in addition, a bitwise status for each active vertices which avoids reading the actual BFS algorithmic metadata for many vertices in so many iterations thereby GRAPHONE is slower by 33%.

The speedups of GRAPHONE are due to our memory management optimizations, overlap of logging and archiving, edge sharding technique, and avoiding the read of adjacency data through multi-versioned degree array. As we have demonstrated that memory management (§7.4.3) provides speedup to ingestion and analytics both, while edge sharding (§7.4.4) leads to further improvement in data ingestion. The multi-versioned array in LLAMA is the main structure responsible for providing snapshot capability in LLAMA, and is also used by analytics. On the other hand, GRAPHONE provides snapshot through multi-versioned degree array, but uses private degree array of static view for the analytics that avoids reading multi-versioned degree array completely. This technique apart from cacheline-/page- sized memory allocation further results into better analytics performance.

These advantages powers GRAPHONE to perform better than LLAMA despite creating an order of more adjacency list snapshots and reading the unneeded weights from the adjacency store for analytics. For results on Kron-21 graph, we believe that overlapping of logging and archiving plays a major role as each LLAMA snapshot is created using  $2^{18}$  batched edges only (total edges/snapshot count) due to its smaller size. For other graphs the batched edges per snapshot are much higher due to their larger sizes.

**7.3.2 Dynamic Graph System.** Stinger is an in-memory graph system that uses atomic instructions to support fine-grained updates. So it cannot provide semantically correct analytics if updates and computations are scheduled at the same time, as different iteration of the analytics will run on the different versions of the data. We used the benchmark developed in [62] to compare the results on the RMat graph.

Stinger is able to support 3.49 million updates/sec on the same weighted RMat graph, whereas GRAPHONE ingests 18.67 million edges/sec, achieving 5.36 $\times$  higher ingestion rate as shown in Fig. 13. Part of the reason for poor update rate of Stinger is that unlike GRAPHONE, it directly updates the adjacency store using atomic constructs.

We have also implemented PageRank and BFS in a similar approach as Stinger. The comparison is plotted in Fig. 13. Clearly, GRAPHONE is able to provide a better support for BFS and PageRank




Fig. 13: Comparison against Stinger for in-memory setup

achieving 12.76× and 3.18× speedup respectively. The reason behind the speedup is explicit optimization to reduce the chaining which removes a lot of pointer chasing, and better cache access locality due to cacheline sized edge blocks.

7.3.3 *Databases.* GRAPHONE is only designed for graph like data-structure, and do not offer any other indexing apart from those that are created on source and destination vertices at this moment. Also, GRAPHONE offers snapshot isolation and not the strict serializability required for transaction support, hence GRAPHONE cannot compete against workloads that require transaction support. The value proposition here is that the adjacency store design can be integrated with the existing databases where prime focus is management and analytics on graph data but without transaction support, such as analytics queries. E.g. the in-memory adjacency list in Neo4j [66] is optimized for read-only workloads, and new updates generally require invalidating and rebuilding those structures [72]. TitanDB [2], an open source graph analytics framework, is built on top of other storage engines such as HBase and BerkeleyDB, and thus does not offer adjacency list at the storage layer. Adjacency store of GRAPHONE can be very useful in these products. Therefore, even if GRAPHONE is limited in this sense, we expect it be beneficial by integrating it with graph databases, and expect that their query engine can be integrated using our GraphView APIs to support their query languages.




Fig. 14: Update rate comparison against SQLite and Neo4j

Therefore, we compare against SQLite 3.7.15.2, a relational database, Neo4j 3.2.3, a graph database for ingestion test only. SQLite and Neo4j support ACID transaction, and do not provide native support for graph analytics. It is known that higher update rate is possible by trading off the strict serializability of databases, however to measure the magnitude of improvement, it is necessary to conduct experiment. The in-memory configuration of SQLite can ingest 12.46K edges per second, while GRAPHONE is able to support 18.67 million edges per second in the same configuration for above dataset. Neo4j could not finish the benchmark after more than 12 hours, which is along the same line as observed in [62]. Hence we have tested on a smaller graph with 32K vertices, 256K edges, and 100K updates. Neo4j is configured to use disk to make it durable. Neo4j and GRAPHONE both use the buffer cache while persisting the graph data. Neo4j can ingest only 14.81K edges per second, whereas GRAPHONE ingests at 3.63M edges per second.

**7.3.4 Stream Graph Engines.** GRAPHONE runs stream computation and data ingestion to adjacency store (archiving phase) concurrently, while prior stream processing systems updates the adjacency store as part of stream computation where first operation is generally the update of the adjacency store as in Kickstarter [85], differential dataflow of Naiad [64], GraphBolt [61], GraphTau [34] etc. except Kineograph [17] that we discussed in §3.2. This results into lower ingestion rate in prior work. In this sub-section, we will evaluate both: the advantage of efficient data-store; as well as the design choice of having the data-store (adjacency store) getting updated in parallel to streaming analytics.

In this sub-section, we will first compare Kickstarter against the version of a GRAPHONE that updates adjacency store as first operation of streaming analytics workflow followed by actual analytics computation as done in Kickstarter. This serves a baseline version of GRAPHONE for stream analytics comparison, and then we demonstrate the advantage of GRAPHONE over the baseline GRAPHONE by following our design decision of updating the adjacency store in parallel to analytics.

We use an open-source binary executable file of Kickstarter’s single node multi-threaded implementation that has been shared by authors through github (source-code is not available yet), and contains only BFS and SSSP (single source shortest path) implementation. It only prints the time taken by the analytics, but not the time taken by the insertion, so we estimate it using wall clock (literally) by looking into some logs that are printed in the console. We choose BFS for the comparison and is naively implemented in GRAPHONE.

Kickstarter expects all the graph data in the text format, out of which the base snapshot should be provided as a adjacency format, while rest in edge list format along with the number of batches to be created out of this update data file. So, we put these two files in a *tmpfs* file system to avoid the disk IO in the path of data read. Interestingly, Kickstarter while performing streaming BFS very fast, it is very slow in updating the adjacency store. For Kron-21 graph, we supplied half of edges (16 million) as base snapshot in the adjacency format, while rest (16 million) are edge updates in the edge list format. The base snapshot creation and BFS are very fast, while the update insertion took slightly more than 210 minutes, and streaming BFS took 31 milliseconds. The process runs for total 210 minutes and 26.948 seconds. We then repeated the same steps for Twitter graph, but it did not finish running the process for over 18 hours. So we did not run any more experiments.

On the other hand, we supplied all the edges as edge list format in a text file hosted in *tmpfs* file system, also created base snapshot at half the edge count, and next snapshot after all the edges arrived. The edges are inserted to GRAPHONE at one edge at a time using logging phase, while the archiving is controlled by the steaming analytics. The baseline GRAPHONE (i.e. archiving and BFS done in sequence) took 6.226 second to update the adjacency list (including the wait time till all the edges arrive), and 36 millisecond for BFS, and total process time is 12.632 seconds. While our naive BFS is competitive to the Kickstarter’s implementation, the adjacency store update for streaming edges is more than 2000× faster than Kickstarter.

This experiment brings two observation: *first* is that our data-store is very optimized; and *second* is that such data-store can be easily integrated with Kickstarter to take the advantage of the state-of-the-art graph stream processing and our highly optimized data-store. We did not perform any additional experiment due to Kickstarter’s poor update rate and the fact that their analytics can be implemented on top of GRAPHONE.

**Advantage of Separating Archiving Phase from Stream Computation.** We now show the advantage of running the adjacency store update (archiving phase) in parallel to stream computation. We report the wall clock time of the whole process in GRAPHONE compared to baseline GRAPHONE. Here baseline means that archiving is run as part of streaming workflow, as done in Kickstarter etc. In both the setup, edges are always inserted one at a time using logging phase. We take the above

set up of providing the input data from a text file hosted in tmpfs containing data in the edge list format. This time BFS starts running as soon as first snapshot of adjacency store is available.

We first run GRAPHONE and figure out how many times streaming BFS was executed. This is because archiving runs in parallel to BFS, and when one streaming BFS converges in any snapshot (captured using stateful stream view), it immediately pulls the current snapshot by calling *update-sstream-view()* and performs streaming BFS again. Therefore streaming BFS may run less number of times than the archiving phase count. Hence, we print the count of how many times streaming BFS was executed. We then run same number of streaming BFS in the baseline GRAPHONE by creating exactly that number of adjacency list snapshots. Table 10 shows the wall clock timing of the two systems.

Table 10. Impact of separating archiving phase out of Analytics workflow. GRAPHONE, by this separation, decreases the overall run time of data ingestion and analytics as shown in third column of this table. The data is ingested from a text file hosted in tmpfs file system. All the numbers are in seconds.

Graph Name	Baseline GRAPHONE	GRAPHONE
Twitter	987.45	764.58
Friendster	1236.78	1032.35
Subdomain	975.74	817.90
Kron-28	2687.60	1990.68
Kron-21	19.37	13.87


The design decision to separate the archiving from the analytics workflow increases the overall average performance by 28.58% for graphs. It should be noted that GRAPHONE actually creates 29,954 adjacency store snapshot compared to 1,870 created by the baseline GRAPHONE, while both the versions run streaming BFS 1,870 times. The improvement in time despite creating more adjacency list snapshot in GRAPHONE comes due to continuous creation of adjacency store snapshot in parallel to stream analytics, which can schedule the stream analytics right at the time requested by pulling the latest snapshot using *update-sstream-view*. On the other hand, baseline GRAPHONE will incur huge delay in scheduling the stream analytics as it first has to create adjacency store on latest data.

The logging and archiving operations are examples of different categories of stream analytics: logging is a proxy to continuous stateless stream analytics, while archiving is same to the discrete stateful stream analytics. Thus, Table 8 is also an indication of their performance as well. We have also implemented a streaming weakly connected components using ideas from COST [63] using stateless stream view APIs and it can process 33.60 million stream edges/s on Kron-28 graph.

Another advantage of separating the archiving phase out of streaming analytics is that now concurrent stream analytics can be executed from the same data-store as we have discussed in §3.2 and §6.2. We explore this next.

**Concurrent Execution of Stream Analytics.** The separation of data-store as an abstraction provides benefit of executing many analytics together, such as multiple batch analytics or multiple stream analytics or any combination of both from the same data-store, and therefore many complex usecases can be implemented using GRAPHONE data-store and GraphView APIs. As we discussed in §3.2 and §6.2, prior graph analytics systems in theory can only execute multiple batch analytics but not multiple stream analytics from the same adjacency store. Even snapshot based stream analytics systems such as Kinerograph [17] can not execute together two classes of stream analytics: those that need to access whole data; and those that requires window access; as they propose a different snapshot technique to be implemented in future(see Section 6.3 in the paper [17]) for the latter

access. Prior stream analytics can decide to deploy multiple instances of same stream analytics engine to run concurrent stream analytics but will also produce separate adjacency store. Therefore, we choose multiple stream analytics to demonstrate the impact of their execution from the same adjacency store.



**Fig. 15:** Showing total time of GRAPHONE when different number of concurrent streaming BFS are executed while data is ingested in Kron-16 graph. All the data is ingested from a text file hosted in tmpfs file system.

We choose stream BFS starting from different source vertices, and each creates separate stateful stream view handle for execution. The choice of selection of streaming BFS is completely random, and we could have also picked up other stream analytics. We are only interested in demonstrating this capability that is enabled by GRAPHONE, while other consequences such as system overload due to multiple real-time analytics are outside the scope of the paper.

Fig. 15 shows that it is possible to execute more number of concurrent streaming BFS while the data is arrived in Kron-21 graph. The actual streaming BFS computation takes much lesser time as they utilize the prior computation results, and the graph is getting inserted from a text file hosted in tmpfs where parsing each text edge to convert it to binary edge list format slows down the data logging rate (see Table 9). This implies that system has sufficient CPU resources to run more streaming BFS. It is only when the count increases beyond 8, each component (logging, archiving and streaming BFS) start to fight for CPU and memory bandwidth resources that the overall system performance decreases.

It would also be possible to concurrently execute batch and streaming version analytics together as well, where batch version will simply run on the latest data at the time whenever they are called by creating static view and exiting after the computation is over, while streaming analytics will continue executing until unregistered.

**Discussion on Staleness.** It should be noted that the experiments in this section were run using *stale* stateful stream view. Though, including the latest edges in the computation will incur some penalty (see §7.4) if real-time nature is required, we claim that it is equally or more real-time in nature than Kickstarter or baseline GRAPHONE. This is because by the time the adjacency store is ready in those systems to execute the stream analytics, the adjacency store is more or equally stale than the one used in GRAPHONE as updating the adjacency store (archiving phase) takes some time (e.g. 6.22 seconds in Kron-21 for ingesting 16 million streaming edges to adjacency store), and then only actual analytics will start. On the other hand, GRAPHONE will supply the latest adjacency list that will be stale by at max non-archived edge count (which is always less than twice the archiving threshold) or will be no stale if slight lower analytics performance is acceptable.




Fig. 16: Algorithmic performance and archiving rate variation for different non-archived edge count

### 7.4 System Design Parameters

In this sub-subsection, we present the various results associated with various system design parameters, and individual techniques.

**7.4.1 Performance Trade-off in Hybrid Store.** We first characterize the behavior of the hybrid store for different number of non-archived edges. Fig. 16 shows the performance variation of archiving rate, BFS, PageRank, and 1-hop query for Kron-28 graph when the non-archived edge counts are increased, while the rest of the edges are kept in the adjacency store for Kron-28. The figure shows that up to 2<sup>17</sup> non-archived edges in the edge log brings negligible drop in the analytics performance. Hence, we recommend the value of archiving threshold as 2<sup>16</sup> edges as the logging overlaps with the archiving. GRAPHONE is able to sustain an archiving rate 43.68 million edges per second at this threshold. The 1-Hop query latency of all 512 queries together is only 53.766 ms, i.e. 0.105 ms for each query.

The archiving threshold of 2<sup>16</sup> edges is not unexpected as it is small compared to total edge count (2<sup>33</sup>) in Kron-28, and the analytics on non-archived edges are parallelized. Further, the parallelization cost dominates when the number of non-archived edges are small (2<sup>10</sup>). Thus the analytics cost remains same from 2<sup>10</sup> up to 2<sup>17</sup> and then drops as the number of non-archived edges becomes really large. This raises a question on whether there will be a performance drop in the analytics when we increase the edge count from 0 to 2<sup>10</sup>. Table 11 shows that the drop is only applicable for 1-hop query. This is because BFS and PageRank have already spawned multiple threads for execution from adjacency store, therefore they can easily process the edges in parallel in edge-centric processing (i.e. it only need to scan non-archived edges once each iteration), thus performance drop will be negligible. On the other hand, 1-Hop query is single threaded as getting neighbors of a vertex from adjacency list does not require multiple threads. But either the cost to create threads for parallel scan or doing it sequentially will lead to performance drop. The Table 11 and Fig. 16 shows the cost to process non-archived edges when we create multiple threads to take advantage of the parallelism offered by edge log scan (one time for each 1-Hop query). Scanning without parallelism will also reduce the performance for 1-Hop queries, but we have not measured it.

Fig. 16 also shows that higher archiving threshold leads to better archiving rate, e.g., a archiving threshold of 1,048,576 (2<sup>20</sup>) edges can sustain a archiving rate of 56.99 million edges/second. The drawback is that the analytics performance will be reduced as it will find more number of non-archived edges. On the contrary, archiving works continuously and tries to minimize the number of non-archived edges, so a smaller arrival rate will lead to frequent archiving, and thus fewer non-archived edges will be observed at any time. The drop in archiving rate at the tail is due to the

Table 11. Batch Analytics Performance (in Seconds) when there are 0 or 1024 non-archived edges in GRAPHONE for Kron-28 graph.

Non-archived Edge Count	BFS	PageRank (5 Iterations)	1-Hop (512 Queries)
0	1.3101	32.1611	$8.32 \times 10^{-4}$
1024	1.3255	32.2803	$5.47 \times 10^{-2}$

impact of large working set size that leads to more last-level-cache transactions and misses while filling the edge arrays.

**7.4.2 Scalability.** The edge sharding stage removes the need of atomic instruction or locks completely in the archiving phase, which results into better scaling of archiving rate with increasing number of threads as plotted in Fig. 17. There is some super-linear behavior when thread count is increased from 16 to 32. This is due to the second socket coming into picture with its own hardware caches, and non-atomic behavior makes it to scale super-linear. This observation is confirmed by running the archiving using 16 threads spread equally across two sockets, and achieves higher archiving rate compared to the case when the majority of threads belong to one socket.




Fig. 17: Archiving rate scaling with thread count

**7.4.3 Memory Allocation.** Fig. 18 shows that the cacheline sized memory allocation and special handling of hub-vertices improve the performance of the archiving and analytics. The cacheline sized memory optimization improves the archiving rate at the archiving threshold by 2.20 $\times$  for Kron-28 graph, while speeding up BFS, PageRank and 1-Hop query performance by 1.37 $\times$ , 3.11 $\times$  and 8.82 $\times$ .




Fig. 18: Caheline sized memory allocation brings huge performance gain, while hub vertex handling on top of cacheline size memory allocation improves the query performance only.

Hub vertices handling additionally improves the query performance (by 7.5%). This was expected as the number of hub-vertices in any power-law graph are very small, so performance improvement would be limited.



**GRAPHONE Without any Memory Chain.** Memory allocation techniques tries to minimize the memory link or chain counts, but it cannot reduces it to just one due to evolving nature of input graph data. In this part, we show that the how much GRAPHONE suffers from such an ideal system if it can achieve such a goal. We design an ideal system, called static GRAPHONE that have only one memory link.




Fig. 19: Graph analytics performance in GRAPHONE compared to its static version that have no chaining requirement.

Fig. 19 shows this performance drop, specifically trading off just 17% average performance for real-world graphs (26% for all the graphs plotted) from the static system, one can support high arrival velocity of fine-grained updates. However, the performance drop is only temporary as the compaction process will remove the chaining in the background.

**7.4.4 Advantage of Edge Sharding.** Fig. 20 shows the performance of the archiving phase with edge sharding stage on Kron-28 graph in comparison to an alternate approach where each worker thread can directly work on the edge log by dividing the number of edges equally among them. This will need to use atomic instructions when two threads compete to add the edges to the adjacency list of the same vertex. Also, the order of processing will not be the same as the order of arrival of updates, thus the alternate method will suffer from ordering issues that we have outlined in §5.




Fig. 20: Impact of the edge sharding on archiving phase performance

The edge sharding stage improves the whole archiving phase by 2.59× for Kron-28 graph as shown in Fig. 20. It should be noted that the degree array computation also shows speedup as we removed the atomic instruction usage from this stage as well.

**7.4.5 Edge Log Size.** Fig. 21 shows the effect of edge log size on overall ingestion rate on Kron-28 graph. A smaller edge log would block the logging thread frequently, as the archiving rate is generally slower than the logging rate. Therefore, a larger edge log size can accommodate more edges, which will than force the archiving to run on large non-archived edges, which results in higher archiving rate. Clearly, an edge log size greater than 4 million edges (32 MB) does not have any impact on overall ingestion rate.




Fig. 21: Showing Ingestion rate when edge log size increases.

**7.4.6 Degree Array Snapshot Counts.** GRAPHONE keeps many recent degree values of each vertex, that are recycled after their number increases beyond a number. Keeping more degree array in GRAPHONE provides a higher number of adjacency list snapshots, but managing them incurs more memory cost, and results into slower archiving rate as recycling of old degree nodes requires more pointer chasing in degree array structure. Fig. 22 shows this impact on archiving rate for Kron-28 graph. Clearly, having less number of degree array versions are better for data ingestions. Therefore, GRAPHONE keeps only 3 degree array versions, and recommends using private views for long running analytics.




Fig. 22: Impact of the degree node count on archiving phase performance

## 7.5 GraphView

In this subsection, we measure the cost of creating different kind of GraphViews, and measure the performance impact of different parameters.

**Cost of High-level API of GraphView.** As we discussed in §6, we provide two levels of API: the low-level API, *get-nebrs-archived-\**(*i*), that directly allows access of references of the internal edge array blocks; the high-level API, *get-nebr-\**, simplifies the job of programmer to implement the vertex-centric implementation, and also make the programmer not to traverse all the links in the edge array. However, the high-level API has a cost as it internally traverses and copies all the edge array blocks data for each vertex to a user-supplied buffer.

Table 12. BFS performance when using low-level vs high-level GraphView API.

Graph Name	Low-Level API		High-Level API	
	BFS	PageRank	BFS	PageRank
Kron-28	15.09	58.62	15.51	60.17
Twitter	3.66	14.39	3.67	14.48

Table 12 shows the cost of BFS (push based) when using high-level APIs over low-level APIs. Clearly, simplifying the programmer’s job does not have much cost. However, there could be a cost

with high-level APIs in some cases, such as when implementing BFS using direction optimized technique [11] (i.e., using push+pull both). In this case, few levels of BFS are computed using pull method where not every neighbors of active vertices are visited. Thus, high-level API will have obvious disadvantage as it always copies whole neighbor list to user supplied buffer. Clearly, this is a limitation of GraphView API in its current form, as for direction-optimized BFS on kron-28 graph, high-level API is around 140% overhead. However, we do not see any limitation on extending the set of APIs, such as getting only last  $k$  neighbors of a vertex, where  $k$  would be a user-supplied value. We leave this work as a future work.

**Cost of Private Static View Creation.** Passing *private* flag in the view creation API copies the degree array and non-archived edges to a private buffer. The copy takes some time, but improves the working of GRAPHONE as the copying makes the data management independent of analytics as former can recycle the degree array easily. Table 13 shows the cost of creating a private static view which requires copying the degree array is in milliseconds, much less than BFS.

Table 13. Time spent (in milliseconds) in GraphView Creation.

Graph Name	Kron-28	Kron-21	Twitter
View Creation Time	215.53	1.77	111.66

## 8 RELATED WORK

Static graph analytics systems [7, 16, 29, 30, 38, 45, 46, 53, 55, 57, 58, 67, 74, 75, 87, 89, 91, 94–96] support only batch analytics where pre-processing consumes much more time than the computation itself [60], and are clearly not designed for supporting real-time analytics.

Graphchi [50] and other snapshot based systems [28, 37, 41, 59, 71, 86] support bulk updates only. However, in Graphchi, the updates are not made visible to the graph until the currently running analytics finishes. This implies that updates can be ingested only in long intervals, and execution of another analytics is not supported while the first one is running, as the updates will hardly get a chance to be merged to the graph if one or the other analytics is always running. LLAMA [59] uses multi-versioned vertex array and reverse chaining of delta adjacency list to create multiple delta snapshots. However, it supports only bulk updates. Clearly, creating a new delta snapshot using adjacency list at each fine-grained update will result into thousands of snapshots of vertex array that must be kept for some time as deletion can only happen once computation finishes on the snapshot. Thus soon the system will run out of memory and storage resources.

Chronos [28], e.g., also supports snapshots for bulk updates only, and uses a bitmap for each edge in the edge array instead of multi-version vertex array to identify whether an edge is part of a particular snapshot or not. The bit length assigned to each bitmap decides how many snapshot Chronos can manage, which clearly cannot support snapshots at fine-grained updates.

Naiad [65], a timely dataflow framework, supports iterative and incremental compute but does not offer the data window on the graph data. Other stream analytics systems [17, 34, 65, 81] support stream processing and snapshot creation, some offering data window but mostly at bulk updates only. Stream databases [5, 6] provide only stream processing. TIDE [88] introduces probabilistic edge decay that samples data from base store.

Graph data sharding [20, 45, 50, 51, 57, 92] is well-studied problem. Our edge sharding is different as our aim is to keep ordering intact in adjacency store, and remove usage of atomic instructions from archiving phase. Nonetheless, prior sharding techniques may be very useful for dividing the GRAPHONE data in many machines, which is left as a future work.

Prior works [26, 78] follow integrated graph system model that manage online updates and queries in the database, and replicate data in an offline analytics engines for long running graph analytics tasks. As we have identified in §1, they suffers from excessive data duplication and weakest component problem. Zhang et al [93] also argue that such composite design is not optimal. GraPU [80] proposes to pre-processes the buffered updates instead of making them available to compute as in GRAPHONE. Trading-off granularity of data visibility is similar to Lazybase [19], but we additionally tune the access of non-archived edges to reduce performance drop in our setup and offer diverse data views.

Many dynamic graph systems manage graph data at read-committed isolation level only [21, 23, 79] using atomic update construct without providing a streaming access or creation of snapshots. In other words, these systems modify the per-vertex edge arrays, while an analytics task is running. As a result, different iterations of the analytics may run on different versions of the dataset. Hence the result can be semantically incorrect, or the convergence may get delayed [85].

The in-memory adjacency list in Neo4j [66] is optimized for read-only workloads, and new updates generally require invalidating and rebuilding those structures [72]. Titan [2] (also known as Janus Graph), an open source graph analytics framework, is built on top of other storage engines such as HBase and BerkeleyDB, and thus does not offer adjacency list at the storage layer. OrientDB [68] builds a graph model using document store model where all information of a vertex including all of its neighbor are stored in a document.

Key value stores are basic building blocks for many NoSQL data storage, that supports atomic updates only. However, append operation or partial update of values are generally not possible in key-value store, a requirement to store graph data in adjacency list format. Therefore, systems must develop a graph data layer on top of key value store, such as trinity [79] which defined append operation using migration of old value (neighbor list) of a source vertex (key), which may trigger a lot of migration. Therefore, Facebook also uses a specialized data-store, called TAO [14], for storing its social graph by moving away from memcache key-value architecture. Hence, in absence of a higher layer, a key-value store is not natively suited for a graph eco-system, while GRAPHONE provides a native support of adjacency list format with an in-built support for append operation.

Timestamp is a known method to provide snapshots in adjacency list format [83]. As we have discussed, it is inefficient to maintain adjacency list snapshots at the edge level granularity and cannot provide streaming view. Also, additional storage for storing time-stamp will double the memory cost, and will introduce additional complexity and performance degradation as each analytics and query need to be made time-stamp aware.

## 9 CONCLUSION

We have presented GRAPHONE, a unified graph data store abstraction that offers diverse data access at different granularity for various real-time analytics and queries at high-performance, while simultaneously supporting high arrival velocity of fine-grained updates.

## ACKNOWLEDGMENTS

The authors thank our FAST'19 Shepherd Ashvin Goel, Usenix FAST'19 and ACM Transaction on Storage reviewers for their suggestions. Pradeep Kumar did part of this work as a graduate research assistant at the George Washington University. This work was supported in part by National Science Foundation CAREER award 1350766 and grants 1618706 and 1717774.

## REFERENCES

- [1] Friendster Network Dataset – KONECT. <http://konect.uni-koblenz.de/networks/friendster>.
- [2] Titan Graph Database. <https://github.com/thinkaurelius/titan>.

- [3] Twitter (MPI) Network Dataset – KONECT. [http://konect.uni-koblenz.de/networks/twitter\\_mpi](http://konect.uni-koblenz.de/networks/twitter_mpi).
- [4] Web Graphs. <http://webdatacommons.org/hyperlinkgraph/2012-08/download.html>.
- [5] D. J. Abadi, Y. Ahmad, M. Balazinska, U. Cetintemel, M. Cherniack, J.-H. Hwang, W. Lindner, A. Maskey, A. Rasin, E. Ryvkina, et al. The Design of the Borealis Stream Processing Engine. In *Cidr*, volume 5, pages 277–289, 2005.
- [6] D. J. Abadi, D. Carney, U. Cetintemel, M. Cherniack, C. Convey, S. Lee, M. Stonebraker, N. Tatbul, and S. Zdonik. Aurora: A New Model and Architecture for Data Stream Management. *VLDB Journal*, 12(2):12039, 2003.
- [7] Z. Ai, M. Zhang, Y. Wu, X. Qian, K. Chen, and W. Zheng. Squeezing out All the Value of Loaded Data: An out-of-core Graph Processing System with Reduced Disk I/O. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, (Santa Clara, CA), pages 125–137, 2017.
- [8] L. Akoglu, H. Tong, and D. Koutra. Graph based anomaly detection and description: a survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, May 2015.
- [9] R. Albert, H. Jeong, and A.-L. Barabási. Internet: Diameter of the World-Wide Web. *Nature*, 401(6749):130–131, Sept. 1999.
- [10] D. Anicic, P. Fodor, S. Rudolph, and N. Stojanovic. EP-SPARQL: A Unified Language for Event Processing and Stream Reasoning. In *Proceedings of the 20th international conference on World wide web*, pages 635–644. ACM, 2011.
- [11] S. Beamer, K. Asanovic, and D. Patterson. Direction-Optimizing Breadth-First Search. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2012.
- [12] B. Bhattarai, H. Liu, and H. H. Huang. CECI: Compact Embedding Cluster Index for Scalable Subgraph Matching. In *Proceedings of the 2019 International Conference on Management of Data*, SIGMOD '19, 2019.
- [13] U. Brandes. A faster algorithm for betweenness centrality\*. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.
- [14] N. Bronson, Z. Amsden, G. Cabrera, P. Chakka, P. Dimov, H. Ding, J. Ferris, A. Giardullo, S. Kulkarni, H. C. Li, et al. TAO: Facebook’s Distributed Data Store for the Social Graph. In *USENIX Annual Technical Conference*, pages 49–60, 2013.
- [15] D. Chakrabarti, Y. Zhan, and C. Faloutsos. R-MAT: A Recursive Model for Graph Mining. In *SDM*, 2004.
- [16] R. Chen, J. Shi, Y. Chen, and H. Chen. PowerLyra: Differentiated Graph Computation and Partitioning on Skewed Graphs. In *Proceedings of the Tenth European Conference on Computer Systems*, 2015.
- [17] R. Cheng, J. Hong, A. Kyröla, Y. Miao, X. Weng, M. Wu, F. Yang, L. Zhou, F. Zhao, and E. Chen. Kineograph: Taking the Pulse of a Fast-Changing and Connected World. In *Proceedings of the 7th ACM european conference on Computer Systems*, 2012.
- [18] S. Choudhury, L. B. Holder, G. Chin, K. Agarwal, and J. Feo. A Selectivity based approach to Continuous Pattern Detection in Streaming Graphs. In *18th International Conference on Extending Database Technology (EDBT)*, 2015.
- [19] J. Cipar, G. Ganger, K. Keeton, C. B. Morrey III, C. A. Soules, and A. Veitch. LazyBase: trading freshness for performance in a scalable database. In *Proceedings of the 7th ACM european conference on Computer Systems*, pages 169–182. ACM, 2012.
- [20] D. Dai, W. Zhang, and Y. Chen. IOGP: An Incremental Online Graph Partitioning Algorithm for Distributed Graph Databases. In *Proceedings of the 26th International Symposium on High-Performance Parallel and Distributed Computing*, pages 219–230. ACM, 2017.
- [21] L. Dhulipala, G. E. Blelloch, and J. Shun. Low-latency graph streaming using compressed purely-functional trees. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 918–934. ACM, 2019.
- [22] D. Easley and J. Kleinberg. *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press, 2010.
- [23] D. Ediger, R. McColl, J. Riedy, and D. A. Bader. Stinger: High Performance Data Structure for Streaming Graphs. In *High Performance Extreme Computing (HPEC), 2012 IEEE Conference on*, pages 1–5. IEEE, 2012.
- [24] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On Power-Law Relationships of the Internet Topology. In *ACM SIGCOMM computer communication review*, volume 29, pages 251–262. ACM, 1999.
- [25] FlockDB. [https://blog.twitter.com/engineering/en\\_us/a/2010/introducing-flockdb.html](https://blog.twitter.com/engineering/en_us/a/2010/introducing-flockdb.html), 2010.
- [26] Graph Compute with Neo4j. <https://neo4j.com/blog/graph-compute-neo4j-algorithms-spark-extensions/>, 2016.
- [27] Graph500. <http://www.graph500.org/>.
- [28] W. Han, Y. Miao, K. Li, M. Wu, F. Yang, L. Zhou, V. Prabhakaran, W. Chen, and E. Chen. Chronos: A Graph Engine for Temporal Graph Analysis. In *EuroSys*, 2014.
- [29] W.-S. Han, S. Lee, K. Park, J.-H. Lee, M.-S. Kim, J. Kim, and H. Yu. TurboGraph: A Fast Parallel Graph Engine Handling Billion-scale Graphs in a Single PC. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2013.
- [30] S. Hong, H. Chafi, E. Sedlar, and K. Olukotun. Green-Marl: a DSL for Easy and Efficient Graph Analysis. In *ACM SIGARCH Computer Architecture News*, 2012.

- [31] Y. Hu, P. Kumar, G. Swope, and H. H. Huang. Trix: Triangle Counting at Extreme Scale. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2017.
- [32] Y. Hu, H. Liu, and H. H. Huang. TriCore: Parallel Triangle Counting on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, page 14. IEEE Press, 2018.
- [33] B. A. Huberman and L. A. Adamic. Internet: Growth dynamics of the World-Wide Web. *Nature*, 1999.
- [34] A. P. Iyer, L. E. Li, T. Das, and I. Stoica. Time-Evolving Graph Processing at Scale. In *Proceedings of the Fourth International Workshop on Graph Data Management Experiences and Systems*, page 5. ACM, 2016.
- [35] H. Jeong, B. Tombor, R. Albert, Z. N. Oltvai, and A.-L. Barabási. The large-scale organization of metabolic networks. *Nature*, 2000.
- [36] Y. Ji, H. Liu, and H. H. Huang. iSpan: Parallel Identification of Strongly Connected Components with Spanning Trees. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, page 58. IEEE Press, 2018.
- [37] X. Ju, D. Williams, H. Jamjoom, and K. G. Shin. Version Traveler: Fast and Memory-Efficient Version Switching in Graph Processing Systems. In *2016 USENIX Annual Technical Conference (USENIX ATC 16)*, pages 523–536. USENIX Association, 2016.
- [38] U. Kang, H. Tong, J. Sun, C.-Y. Lin, and C. Faloutsos. GBASE: A Scalable and General Graph Management System. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2011.
- [39] A. D. Kent. Comprehensive, Multi-Source Cyber-Security Events. Los Alamos National Laboratory, 2015.
- [40] A. D. Kent, L. M. Liebrock, and J. C. Neil. Authentication graphs: Analyzing user behavior within an enterprise network. *Computers & Security*, 48:150–166, 2015.
- [41] U. Khurana and A. Deshpande. Efficient Snapshot Retrieval over Historical Graph Data. In *Data Engineering (ICDE), 2013 IEEE 29th International Conference on*, pages 997–1008. IEEE, 2013.
- [42] K. Kim, T. Wang, R. Johnson, and I. Pandis. Ermia: Fast Memory-Optimized Database System for Heterogeneous Workloads. In *Proceedings of the 2016 International Conference on Management of Data*, pages 1675–1687. ACM, 2016.
- [43] D. Knoke and S. Yang. *Social network analysis*, volume 154. Sage, 2008.
- [44] J. Kreps, N. Narkhede, J. Rao, et al. Kafka: a Distributed Messaging System for Log Processing. In *Proceedings of the NetDB*, pages 1–7, 2011.
- [45] P. Kumar and H. H. Huang. G-Store: High-Performance Graph Store for Trillion-Edge Processing. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2016.
- [46] P. Kumar and H. H. Huang. Falcon: Scaling IO Performance in Multi-SSD Volumes. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*. USENIX Association, pages 41–53, 2017.
- [47] P. Kumar and H. H. Huang. SafeNVM: A Non-Volatile Memory Store with Thread-Level Page Protection. In *IEEE International Congress on Big Data (BigData Congress), 2017*, pages 65–72. IEEE, 2017.
- [48] P. Kumar and H. H. Huang. GraphOne: A Data Store for Real-time Analytics on Evolving Graphs. In *Proceedings of the 17th USENIX Conference on File and Storage Technologies (FAST'19)*, 2019.
- [49] H. Kwak, C. Lee, H. Park, and S. Moon. What is Twitter, a social network or a news media? In *WWW*, 2010.
- [50] A. Kyrola, G. Blleloch, and C. Guestrin. Graphchi: Large-Scale Graph Computation on Just a PC. In *OSDI*, 2012.
- [51] K. Lakhotia, R. Kannan, S. Pati, and V. Prasanna. GPoP: A Cache and Memory-efficient Framework for Graph Processing over Partitions. In *Proceedings of the 24th Symposium on Principles and Practice of Parallel Programming, PPOPP '19*, pages 393–394, New York, NY, USA, 2019. ACM.
- [52] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining*, pages 177–187. ACM, 2005.
- [53] H. Lin, X. Zhu, B. Yu, X. Tang, W. Xue, W. Chen, L. Zhang, T. Hoefler, X. Ma, X. Liu, et al. ShenTu: Processing Multi-Trillion Edge Graphs on Millions of Cores in Seconds. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis*, page 56. IEEE Press, 2018.
- [54] H. Liu and H. H. Huang. Enterprise: Breadth-First Graph Traversal on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2015.
- [55] H. Liu and H. H. Huang. Graphene: Fine-Grained IO Management for Graph Computing. In *Proceedings of the 15th USENIX Conference on File and Storage Technologies (FAST)*, 2017.
- [56] H. Liu, H. H. Huang, and Y. Hu. iBFS: Concurrent Breadth-First Search on GPUs. In *Proceedings of the SIGMOD International Conference on Management of Data*, 2016.
- [57] Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein. Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud. *Proceedings of the VLDB Endowment (VLDB)*, 2012.
- [58] S. Maass, C. Min, S. Kashyap, W. Kang, M. Kumar, and T. Kim. Mosaic: Processing a Trillion-Edge Graph on a Single Machine. In *Proceedings of the Twelfth European Conference on Computer Systems, EuroSys '17*, 2017.

- [59] P. Macko, V. J. Marathe, D. W. Margo, and M. I. Seltzer. LLAMA: Efficient Graph Analytics Using Large Multiversioned Arrays. In *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*, pages 363–374. IEEE, 2015.
- [60] J. Malicevic, B. Lepers, and W. Zwaenepoel. Everything you always wanted to know about multicore graph processing but were afraid to ask. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, pages 631–643, Santa Clara, CA, 2017. USENIX Association.
- [61] M. Mariappan and K. Vora. GraphBolt: Dependency-Driven Synchronous Processing of Streaming Graphs. In *Proceedings of the Fourteenth EuroSys Conference 2019*, page 25. ACM, 2019.
- [62] R. C. McColl, D. Ediger, J. Poovey, D. Campbell, and D. A. Bader. A Performance Evaluation of Open Source Graph Databases. In *Proceedings of the First Workshop on Parallel Programming for Analytics Applications, PPAA '14*, pages 11–18, New York, NY, USA, 2014. ACM.
- [63] F. McSherry, M. Isard, and D. G. Murray. Scalability! But at what COST? In *15th Workshop on Hot Topics in Operating Systems (HotOS XV)*, 2015.
- [64] F. McSherry, D. G. Murray, R. Isaacs, and M. Isard. Differential Dataflow. In *Innovative Data Systems Research (CIDR)*, 2018.
- [65] D. G. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, and M. Abadi. Naiad: A Timely Dataflow System. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles (SOSP)*, 2013.
- [66] Neo4j Inc. <https://neo4j.com/>, 2016.
- [67] D. Nguyen, A. Lenharth, and K. Pingali. A Lightweight Infrastructure for Graph Analytics. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, 2013.
- [68] Orient Database. <https://orientdb.com/>, 2016.
- [69] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: bringing order to the Web. 1999.
- [70] J. Reinders. *Intel Threading Building Blocks*. O'Reilly & Associates, Inc., Sebastopol, CA, USA, first edition, 2007.
- [71] C. Ren, E. Lo, B. Kao, X. Zhu, and R. Cheng. On Querying Historical Evolving Graph Sequences. *Proceedings of the VLDB Endowment*, 4(11):726–737, 2011.
- [72] I. Robinson, J. Webber, and E. Eifrem. *Graph Databases*. O'Reilly Media, 2013.
- [73] D. M. Romero, B. Meeder, and J. Kleinberg. Differences in the mechanics of information diffusion across topics: idioms, political hashtags, and complex contagion on twitter. In *Proceedings of the 20th international conference on World wide web*, pages 695–704. ACM, 2011.
- [74] A. Roy, L. Bindschaedler, J. Malicevic, and W. Zwaenepoel. Chaos: Scale-out Graph Processing from Secondary Storage. In *SOSP*. ACM, 2015.
- [75] A. Roy, I. Mihailovic, and W. Zwaenepoel. X-stream: Edge-centric Graph Processing using Streaming Partitions. In *SOSP*. ACM, 2013.
- [76] S. Sallinen, K. Iwabuchi, S. Poudel, M. Gokhale, M. Ripeanu, and R. Pearce. Graph Colouring as a Challenge Problem for Dynamic Graph Processing on Distributed Systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, page 30. IEEE Press, 2016.
- [77] J. Seo, J. Park, J. Shin, and M. S. Lam. Distributed SocialLite: A Datalog-Based Language for Large-Scale Graph Analysis. *Proceedings of the VLDB Endowment*, 6(14):1906–1917, 2013.
- [78] M. Sevenich, S. Hong, O. van Rest, Z. Wu, J. Banerjee, and H. Chafi. Using Domain-specific Languages for Analytic Graph Databases. *Proc. VLDB Endow.*, 9(13):1257–1268, Sept. 2016.
- [79] B. Shao, H. Wang, and Y. Li. Trinity: A Distributed Graph Engine on a Memory Cloud. In *Proceedings of the SIGMOD International Conference on Management of Data*, 2013.
- [80] F. Sheng, Q. Cao, H. Cai, J. Yao, and C. Xie. GraPU: Accelerate Streaming Graph Analysis Through Preprocessing Buffered Updates. In *Proceedings of the ACM Symposium on Cloud Computing, SoCC '18*, 2018.
- [81] X. Shi, B. Cui, Y. Shao, and Y. Tong. Tornado: A System For Real-Time Iterative Analysis Over Evolving Data. In *Proceedings of the 2016 International Conference on Management of Data*, pages 417–430. ACM, 2016.
- [82] J. Shun and G. E. Blelloch. Ligra: A Lightweight Graph Processing Framework for Shared Memory. In *Proceedings of the 18th ACM SIGPLAN symposium on Principles and practice of parallel programming (PPoPP)*, 2013.
- [83] M. Then, T. Kersten, S. Günemann, A. Kemper, and T. Neumann. Automatic algorithm transformation for efficient multi-snapshot analytics on temporal graphs. *Proceedings of the VLDB Endowment*, 10(8):877–888, 2017.
- [84] M. J. M. Turcotte, A. D. Kent, and C. Hash. Unified Host and Network Data Set. *ArXiv e-prints*, Aug. 2017.
- [85] K. Vora, R. Gupta, and G. Xu. Kickstarter: Fast and Accurate Computations on Streaming Graphs via Trimmed Approximations. In *Proceedings of the Twenty-Second International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 237–251. ACM, 2017.
- [86] G. Wang, W. Xie, A. J. Demers, and J. Gehrke. Asynchronous Large-Scale Graph Processing Made Easy. In *CIDR*, volume 13, pages 3–6, 2013.
- [87] M. Wu, F. Yang, J. Xue, W. Xiao, Y. Miao, L. Wei, H. Lin, Y. Dai, and L. Zhou. GRAM: Scaling Graph Computation to the Trillions. In *Proceedings of the Sixth ACM Symposium on Cloud Computing*, 2015.

- [88] W. Xie, Y. Tian, Y. Sismanis, A. Balmin, and P. J. Haas. Dynamic interaction graphs with probabilistic edge decay. In *Data Engineering (ICDE), 2015 IEEE 31st International Conference on*, pages 1143–1154. IEEE, 2015.
- [89] D. Yan, Y. Huang, M. Liu, H. Chen, J. Cheng, H. Wu, and C. Zhang. GraphD: Distributed vertex-centric graph processing beyond the memory limit. *IEEE Transactions on Parallel and Distributed Systems*, 29(1):99–114, 2018.
- [90] K. Zeng, J. Yang, H. Wang, B. Shao, and Z. Wang. A Distributed Graph Engine for Web Scale RDF Data. In *Proceedings of the 39th international conference on Very Large Data Bases (PVLDB'13)*. VLDB Endowment, 2013.
- [91] K. Zhang, R. Chen, and H. Chen. NUMA-Aware Graph-Structured Analytics. *ACM SIGPLAN Notices*, 50(8):183–193, 2015.
- [92] W. Zhang, Y. Chen, and D. Dai. AKIN: A Streaming Graph Partitioning Algorithm for Distributed Graph Storage Systems. In *Proceedings of the 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, pages 183–192. IEEE Press, 2018.
- [93] Y. Zhang, R. Chen, and H. Chen. Sub-millisecond Stateful Stream Querying over Fast-evolving Linked Data. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 614–630. ACM, 2017.
- [94] D. Zheng, D. Mhembere, R. Burns, J. Vogelstein, C. E. Priebe, and A. S. Szalay. FlashGraph: Processing Billion-Node Graphs on an Array of Commodity SSDs. In *Proceedings of the 13th USENIX Conference on File and Storage Technologies (FAST)*, 2015.
- [95] X. Zhu, W. Chen, W. Zheng, and X. Ma. Gemini: A Computation-Centric Distributed Graph Processing System. In *OSDI*, pages 301–316, 2016.
- [96] X. Zhu, W. Han, and W. Chen. GridGraph: Large-scale Graph Processing on a Single Machine Using 2-level Hierarchical Partitioning. In *Proceedings of the USENIX Conference on Usenix Annual Technical Conference*, 2015.