

VSS: A Storage System for Video Analytics

Brandon Haynes
brandon.haynes@microsoft.com
Gray Systems Lab, Microsoft

Amrita Mazumdar
amrita@cs.washington.edu
University of Washington

Maureen Daum
mdaum@cs.washington.edu
University of Washington

Magdalena Balazinska
magda@cs.washington.edu
University of Washington

Dong He
donghe@cs.washington.edu
University of Washington

Alvin Cheung
akcheung@cs.berkeley.edu
University of California, Berkeley

Luis Ceze
luisceze@cs.washington.edu
University of Washington

ABSTRACT

We present a new video storage system (VSS) designed to decouple high-level video operations from the low-level details required to store and efficiently retrieve video data. VSS is designed to be the storage subsystem of a video data management system (VDBMS) and is responsible for: (1) transparently and automatically arranging the data on disk in an efficient, granular format; (2) caching frequently-retrieved regions in the most useful formats; and (3) eliminating redundancies found in videos captured from multiple cameras with overlapping fields of view. Our results suggest that VSS can improve VDBMS read performance by up to 54%, reduce storage costs by up to 45%, and enable developers to focus on application logic rather than video storage and retrieval.

ACM Reference Format:

Brandon Haynes, Maureen Daum, Dong He, Amrita Mazumdar, Magdalena Balazinska, Alvin Cheung, and Luis Ceze. 2021. VSS: A Storage System for Video Analytics. In *Proceedings of the 2021 International Conference on Management of Data (SIGMOD '21)*, June 20–25, 2021, Virtual Event, China. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3448016.3459242>

1 INTRODUCTION

The volume of video data captured and processed is rapidly increasing: YouTube receives more than 400 hours of uploaded video per minute [52], and more than six million closed-circuit television cameras populate the United Kingdom, collectively amassing an estimated 7.5 petabytes of video per day [9]. More than 200K body-worn cameras are in service [24], collectively generating almost a terabyte of video per day [55].

To support this video data deluge, many systems and applications have emerged to ingest, transform, and reason about such data [19, 23, 25, 27, 28, 34, 43, 56]. Critically, however, most of these systems lack efficient storage managers. They focus on query execution for

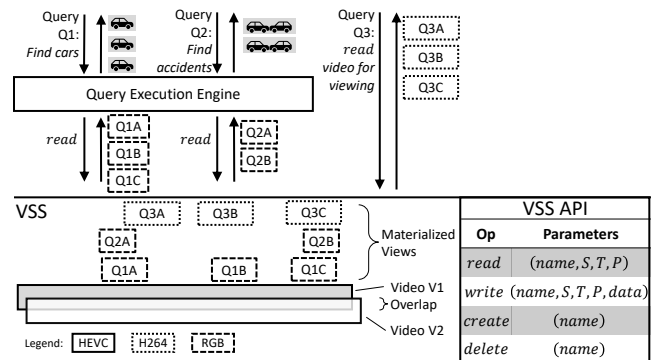


Figure 1: VSS overview & API. Reads and writes require specification of spatial (S ; resolution, region of interest), temporal (T ; start/end time, frame rate), and physical (P ; frame layout, compression codec, quality) parameters.

a video that is already decoded and loaded in memory [23, 27, 28] or treat video compression as a black box [25, 34, 56] (cf. [19, 43]). In practice, of course, videos are stored on disk, and the cost of reading and decompressing is high relative to subsequent processing [11, 19], e.g., constituting more than 50% of total runtime [29]. The result is a performance plateau limited by Amdahl's law, where an emphasis on post-decompression performance might yield impressive results in isolation, but ignores the diminishing returns when performance is evaluated end-to-end.

In this paper, we develop VSS, a *video storage system* designed to serve as storage manager beneath a video data management system or video processing application (collectively VDBMSs). Analogous to a storage and buffer manager for relational data, VSS assumes responsibility for storing, retrieving, and caching video data. It frees higher-level components to focus on application logic, while VSS optimizes the low-level performance of video data storage. As we will show, this decoupling dramatically speeds up video processing queries and decreases storage costs. VSS does this by addressing the following three challenges:

First, modern video applications commonly issue multiple queries over the same (potentially overlapping) video regions and build on each other in different ways (e.g., Figure 1). Queries can also vary video resolution and other characteristics (e.g., the SMOL system

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.

SIGMOD '21, June 20–25, 2021, Virtual Event, China

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8343-1/21/06...\$15.00

<https://doi.org/10.1145/3448016.3459242>

rescales video to various resolutions [29] and Chameleon dynamically adjusts input resolution [25]). Such queries can be dramatically faster with an efficient storage manager that maintains and evolves a cache of video data, each differently compressed and encoded.

Second, if the same video is queried using multiple systems such as via a VDBMS optimized for simple select and aggregate queries [27] and a separate vision system optimized for reasoning about complex scenes [48] (e.g., Figure 1), then the video file may be requested at different resolutions and frame rates and using different encodings. Having a single storage system that encapsulates all such details and provides a unified query interface makes it seamless to create—and optimize—such federated workflows. While some systems have attempted to mitigate this by making multiple representations available to developers [49, 54], they expensively do so for entire videos even if only small subsets (e.g., the few seconds before and after an accident) are needed in an alternate representation.

Third, many recent applications analyze large amounts of video data with overlapping fields of view and proximate locations. For example, traffic monitoring networks often have multiple cameras oriented toward the same intersection and autonomous driving and drone applications come with multiple overlapping sensors that capture nearby video. Reducing the redundancies that occur among these sets of physically proximate or otherwise similar video streams is neglected in all modern VDBMSs. This is because of the substantial difficulties involved: systems (or users) need to consider the locations, orientations, and fields of view of each camera to identify redundant video regions; measure overlap, jitter, and temporally align each video; and ensure that deduplicated video data can be recovered with sufficient quality. Despite these challenges, and as we show herein, deduplicating overlapping video data streams offers opportunities to greatly reduce storage costs.

VSS addresses the above challenges. As a storage manager, it exposes a simple interface where VDBMSs read and write videos using VSS’s API (see Figure 1). Using this API, systems write video data in any format, encoding, and resolution—either compressed or uncompressed—and VSS manages the underlying compression, serialization, and physical layout on disk. When these systems subsequently read video—once again in any configuration and by optionally specifying regions of interest and other selection criteria—VSS automatically identifies and leverages the most efficient methods to retrieve and return the requested data.

VSS deploys the following optimizations and caching mechanisms to improve read and write performance. First, rather than storing video data on disk as opaque, monolithic files, VSS decomposes video into sequences of contiguous, independently-decodable sets of frames. In contrast with previous systems that treat video as static and immutable data, VSS applies transformations at the granularity of these sets of frames, freely transforming them as needed to satisfy a read operation. For example, if a query requests a video region compressed using a different codec, VSS might elect to cache the transcoded subregion and delete the original.

As VSS handles requests for video over time, it maintains a per-video on-disk collection of materialized views that is populated passively as a byproduct of read operations. When a VDBMS performs a subsequent read, VSS leverages a minimal-cost subset of

these views to generate its answer. Because these materialized views can arbitrarily overlap and have complex interdependencies, finding the least-cost set of views is non-trivial. VSS uses a satisfiability modulo theories (SMT) solver to identify the best views to satisfy a request. VSS prunes stale views by selecting those least likely to be useful in answering subsequent queries. Among equivalently useful views, VSS optimizes for video quality and defragmentation.

Finally, VSS reduces the storage cost of redundant video data collected from physically proximate cameras. It does so by deploying a *joint compression* optimization that identifies overlapping regions of video and stores these regions only once. The key challenge lies in efficiently identifying potential candidates for joint compression in a large database of videos. Our approach identifies candidates efficiently without requiring any metadata specification. To identify video overlap, VSS incrementally fingerprints video fragments (i.e., it produces a feature vector that robustly characterizes video regions) and, using the resulting fingerprint index, searches for likely correspondences between pairs of videos. It finally performs a more thorough comparison between likely pairs.

In summary, we make the following contributions:

- We design a new storage manager for video data that leverages the fine-grained physical properties of videos to improve application performance (Section 2).
- We develop a novel technique to perform reads by selecting from many materialized views to efficiently produce an output while maintaining the quality of the resulting video data (Section 3).
- We develop a method to optimize the storage required to persist videos that are highly overlapping or contain similar visual information, an indexing strategy to identify such regions (Section 5), and a protocol for caching multiple versions of the same video (Section 4).

We evaluate VSS against existing video storage techniques and show that it can reduce video read time by up to 54% and decrease storage requirements by up to 45% (Section 6).

2 VSS OVERVIEW

Consider an application that monitors an intersection for automobiles associated with missing children or adults with dementia. A typical implementation would first ingest video data from multiple locations around the intersection. It would then index regions of interest, typically by decompressing and converting the entire video to an alternate representation suitable for input to a machine learning model trained to detect automobiles. Many video query processing systems provide optimizations that accelerate this process [27, 35, 54]. Subsequent operations, however, might execute more specific queries only on the regions that have automobiles. For example, if a red vehicle is missing, a user might issue a query to identify all red vehicles in the dataset. Afterward, a user might request and view all video sequences containing only the likely candidates. This might involve further converting relevant regions to a representation compatible with the viewer (e.g., at a resolution compatible with a mobile device or compressed using a supported codec). We show VSS’s performance for this application in Section 6.

While today’s video processing engines perform optimizations for operations over entire videos (e.g., the indexing phase described above), their storage layers provide little or no support for subsequent queries over the results (even dedicated systems such as quFiles [49] or VStore [54] transcode entire videos, even when only a few frames are needed). Meanwhile, when the above application uses VSS to read a few seconds of low-resolution, uncompressed video data to find frames containing automobiles, it can delegate responsibility to VSS for efficiently producing the desired frames. This is true even if the video is streaming or has not fully been written to disk.

Critically, VSS automatically selects the most efficient way to generate the desired video data in the requested format and region of interest (ROI) based on the original video and cached representations. Further, to support real-time streaming scenarios, writes to VSS are non-blocking and users may query prefixes of ingested video data without waiting on the entire video to be persisted.

Figure 1 summarizes the set of VSS-supported operations. These operations are over *logical videos*, which VSS executes to produce or store fine-grained *physical video* data. Each operation involves a point- or range-based scan or insertion over a single logical video source. VSS allows constraints on combinations of temporal (T), spatial (S), and physical (P) parameters. Temporal parameters include start and end time interval ($[s, e]$) and frame rate (f); spatial parameters include resolution ($r_x \times r_y$) and region of interest ($[x_0..x_1]$ and $[y_0..y_1]$); and physical parameters P include physical frame layout (l ; e.g., YUV420, YUV422), compression method (c ; e.g., HEVC), and quality (to be discussed in Section 3.2).

Internally, VSS arranges each written physical video as a sequence of entities called *groups of pictures (GOPs)*. Each GOP is composed of a contiguous sequence of frames in the same format and resolution. A GOP may include the full frame extent or be cropped to some ROI and may contain raw pixel data or be compressed. Compressed GOPs, however, are constrained such that they are independently decodable and take no data dependencies on other GOPs.

Though a GOP may contain an unbounded number of frames, video compression codecs typically fix their size to a small, constant number of frames (30–300) and VSS accepts as-is ingested compressed GOP sizes (which are typically less than 512kB). For uncompressed GOPs, our prototype implementation automatically partitions video data into blocks of size ≤ 25 MB (the size of one RGB 4K frame), or a single frame for resolutions that exceed this threshold.

3 DATA RETRIEVAL FROM VSS

As mentioned, VSS internally represents a logical video as a collection of materialized physical videos. When executing a read, VSS produces the result using one or more of these views.

Consider a simplified version of the application described in Section 2, where a single camera has captured 100 minutes of 4K resolution, HEVC-encoded video, and written it to VSS using the name V . The application first reads the entire video and applies a computer vision algorithm that identifies two regions (at minutes 30–60 and 70–95) containing automobiles. The application then retrieves those fragments compressed using H264 to transmit to a device that only supports this format. As a result of these operations, VSS now contains the original video (m_0) and the cached versions of

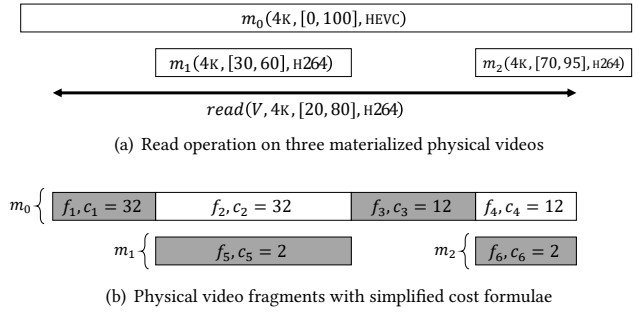


Figure 2: Figure 2(a) shows the query $read(V, 4K, [20, 80], H264)$, where VSS has materialized m_0 , m_1 , and m_2 . Figure 2(b) shows weighted fragments and costs. The lowest-cost result is shaded.

the two fragments (m_1, m_2) as illustrated in Figure 2(a). The figure indicates the labels $\{m_0, m_1, m_2\}$ of the three videos, their spatial configuration (4K), start and end times (e.g., $[0, 100]$ for m_0), and physical characteristics (HEVC or H264).

Later, a first responder on the scene views a one-hour portion of the recorded video on her phone, which only has hardware support for H264 decompression. To deliver this video, the application executes $read(V, 4K, [20, 80], H264)$, which, as illustrated by the arrow in Figure 2(a), requests video V at 4K between time $[20, 80]$ compressed with H264.

VSS responds by first identifying subsets of the available physical videos that can be leveraged to produce the result. For example, VSS can simply transcode m_0 between times $[20, 80]$. Alternatively, it can transcode m_0 between time $[20, 30]$ and $[60, 70]$, m_1 between $[30, 60]$, and m_2 between $[70, 80]$. The latter plan is the most efficient since m_1 and m_2 are already in the desired output format (H264), hence VSS need not incur high transcoding costs for these regions. Figure 2(b) shows the different selections that VSS might make to answer this read. Each *physical video fragment* $\{f_1, \dots, f_6\}$ in Figure 2(b) represents a different region that VSS might select. Note that VSS need not consider other subdivisions—for example by subdividing f_5 at $[30, 40]$ and $[40, 60]$ —since f_5 being cheaper at $[30, 40]$ implies that it is at $[40, 60]$ too.

To model these transcoding costs, VSS employs a *transcode cost model* $c_t(f, P, S)$ that represents the cost of converting a physical video fragment f into a target spatial and physical format S and P . The selected fragments must be of sufficient quality, which we model using a *quality model* $u(f, f')$ and reject fragments of insufficient quality. We introduce these models in the following two subsections.

3.1 Cost Model

We first discuss how VSS selects fragments for use in performing a read operation using its cost model. In general, given a *read* operation and a set of physical videos, VSS must first select fragments that cover the desired spatial and temporal ranges. To ensure that a solution exists, VSS maintains a *cover* of the initially-written video m_0 consisting of physical video fragments with quality equal to the original video (i.e., $u(m_0, f) \geq \tau$). Our prototype

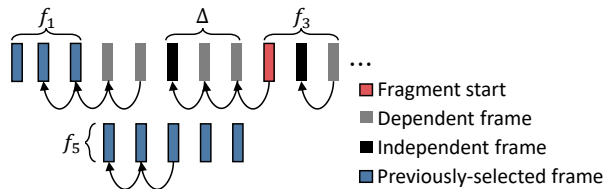


Figure 3: A simplified illustration based on Figure 2. VSS has decided to use f_1 and f_5 and is considering using f_3 starting at the red-highlighted frame. However, f_3 cannot be decoded without transitively decoding its dependencies shown by directed edges (labeled Δ).

sets a threshold $\tau = 40\text{dB}$, which is considered to be lossless. See Section 3.2 for details. VSS also returns an error for reads extending outside of the temporal interval of m_0 .

Second, when the selected physical videos temporally overlap, VSS must resolve *which* physical video fragments to use in producing the answer in a way that minimizes the total conversion cost of the selected set of video fragments. This problem is similar to materialized view selection [15]. Fortunately, a VSS read is far simpler than a general database query, and in particular is constrained to a small number of parameters with point- or range-based predicates.

We motivate our solution by continuing our example from Figure 2(a). First, observe that the collective start and end points of the physical videos form a set of *transition points* where VSS can switch to an alternate physical video. In Figure 2(a), the transition times include those in the set $\{30, 60, 70\}$, and we illustrate them in Figure 2(b) by partitioning the set of cached materialized views at each transition point. VSS ignores fragments that are outside the read’s temporal range, since they do not provide information relevant to the read operation.

Between each consecutive pair of transition points, VSS must choose exactly one physical video fragment. In Figure 2(b), we highlight one such set of choices that covers the read interval. Each choice of a fragment comes with a cost (e.g., f_1 has cost 32), derived using a cost formula given by $c_t(f, P, S) = \alpha(f_s, f_p, S, P) \cdot |f|$. This cost is proportional to the total number of pixels $|f|$ in fragment f scaled by $\alpha(S, P, S', P')$, which is the normalized cost of transcoding a single pixel from spatial and physical format (S, P) into format (S', P') . For example, using fragment m_1 in Figure 2 requires transcoding from physical format $P = \text{HEVC}$ to $P' = \text{H264}$ with no change in spatiotemporal format (i.e., $S = S'$).

During installation, VSS computes the domain of α by executing the vbench benchmark [31] on the installation hardware, which produces per-pixel transcode costs for a variety of resolutions and codecs. For resolutions not evaluated by vbench, VSS approximates α by piecewise linear interpolation of the benchmarked resolutions.

VSS must also consider the data dependencies between frames. Consider the illustration in Figure 3, which shows the frames within a physical video with their data dependencies indicated by directed edges. If VSS wishes to use a fragment at the frame highlighted in red, it must first decode all of the red frame’s *dependent frames*, denoted by the set Δ in Figure 3. This implies that the cost of transcoding a frame depends on *where* within the video it occurs, and whether its dependent frames are *also* transcoded.

To model this, we introduce a *look-back cost* $c_l(\Omega, f)$ that gives the cost of decoding the set of frames Δ on which fragment f depends *if they have not already been decoded*, meaning that they are not in the set of previously selected frames Ω . As illustrated in Figure 3, these dependencies come in two forms: independent frames $A \subseteq \Delta$ (i.e., frames with out-degree zero in our graphical representation) which are larger in size but less expensive to decode, and the remaining dependent frames $\Delta - A$ (those with outgoing edges) which are highly compressed but have more expensive decoding dependencies between frames. We approximate these per-frame costs using estimates from Costa et al. [10], which empirically concludes that dependent frames are approximately 45% more expensive than their independent counterparts. We therefore fix $\eta = 1.45$ and formalize look-back cost as $c_l(\Omega, f) = |A - \Omega| + \eta \cdot |(\Delta - A) - \Omega|$.

To conclude our example, observe that our goal is to choose a set of physical video fragments that cover the queried spatiotemporal range, do not temporally overlap, and minimize the decode and look-back cost of selected fragments. In Figure 2(b), of all the possible paths, the one highlighted in gray minimizes this cost. These characteristics collectively meet the requirements identified at the beginning of this section.

Generating a minimum-cost solution using this formulation requires jointly optimizing both look-back cost c_l and transcode cost c_t , where each fragment choice affects the dependencies (and hence costs) of future choices. These dependencies make the problem not solvable in polynomial time, and VSS employs an SMT solver [12] to generate an optimal solution. Our embedding constrains frames in overlapping fragments so that only one is chosen, selects combinations of regions of interest (ROI) that spatially combine to cover the queried ROI, and uses information about the locations of independent and dependent frames in each physical video to compute the cumulative decoding cost due to both transcode and look-back for any set of selected fragments. We compare this algorithm to a dependency-naïve greedy baseline in Section 6.1.

3.2 Quality Model

Besides efficiency, VSS must also ensure that the quality of a result has sufficient fidelity. For example, using a heavily *downsampled* (e.g., 32×32 pixels) or *compressed* (e.g., at a 1Kbps bitrate) physical video to answer a read requesting 4K video is likely to be unsatisfactory. VSS tracks quality loss from both sources using a quality model $u(f_0, f)$ that gives the expected quality loss of using a fragment f in a read operation relative to using the originally-written video f_0 . When considering using a fragment f in a read, VSS will reject it if the expected quality loss is below a user-specified cutoff: $u(f_0, f) < \epsilon$. The user optionally specifies this cutoff in the read’s physical parameters (see Figure 1); otherwise, a default threshold is used ($\epsilon = 40\text{dB}$ in our prototype).

The range of u is a non-negative peak signal-to-noise ratio (PSNR), a common measure of quality variation based on mean-squared error [22]. Values $\geq 40\text{dB}$ are considered to be lossless qualities, and $\geq 30\text{dB}$ near-lossless. PSNR is itself defined in terms of the mean-squared error (MSE) of the pixels in a frame relative to the corresponding pixels in a reference frame, normalized by the maximum pixel value.

As described previously, error in a fragment accumulates through two mechanisms—resampling and compression—and VSS uses the sum of both sources when computing u . We next examine how VSS computes error from each source.

Resampling error. First, for downsampled error produced through a resolution or frame rate change applied to f_0 , computing $MSE(f, f_0)$ is straightforward. However, VSS may transitively apply these transformations to a sequence of fragments. For example, f_0 might be downsampled to create f_1 , and f_1 later used to produce f_2 . In this case, when computing $MSE(f_0, f_2)$, VSS no longer has access to the uncompressed representation of f_0 . Rather than expensively re-decompressing f_0 , VSS instead bounds $MSE(f_0, f_n)$ in terms of $MSE(f_0, f_1), \dots, MSE(f_{n-1}, f_n)$, which are a single real-valued aggregates stored as metadata. We show a proof in [17].

Compression error. Unlike resampling error, tracking quality loss due to lossy compression error is challenging because it cannot be calculated without decompressing—an expensive operation—and comparing the recovered version to the original input. Instead, VSS estimates compression error in terms of mean bits per pixel per second (MBPP/S), which is a metric reported during (re)compression. VSS then estimates quality by mapping MBPP/S to the PSNR reported by the vbench benchmark [31], a benchmark for evaluating video transcode performance in the cloud. To improve on this estimate, VSS periodically samples regions of compressed video, computes exact PSNR, and updates its estimate.

4 DATA CACHING IN VSS

We now describe how VSS decides *which* physical videos to maintain, and which to evict under low disk space conditions. This involves making two interrelated decisions:

- When executing a read, should VSS admit the result as a new physical video for use in answering future reads?
- When disk space grows scarce, which existing physical video(s) should VSS discard?

To aid both decisions, VSS maintains a video-specific *storage budget* that limits the total size of the physical videos associated with each logical video. The storage budget is set when a video is created in VSS (see Figure 1) and may be specified as a multiple of the size of the initially written physical video or a fixed ceiling in bytes. This value is initially set to an administrator-specified default (10× the size of the initially-written physical video in our prototype). As described below, VSS ensures a sufficiently-high quality version of the original video can always be reproduced. It does so by maintaining a cover of fragments with sufficiently high quality (PSNR ≥ 40 dB in our prototype, which is considered to be lossless) relative to the originally ingested video.

The key idea behind VSS’s cache is to logically break physical videos into “pages.” That is, rather than treating each physical video as a monolithic cache entry, VSS targets the individual GOPs *within* each physical video. Using GOPs as cache pages greatly homogenizes the sizes of the entries that VSS must consider. VSS’s ability to evict GOP pages *within* a physical video differs from other variable-sized caching efforts such as those used by content delivery networks (CDNs), which make decisions on large, indivisible, and opaque entries (a far more challenging problem space with limited solutions [7]).

However, there are key differences between GOPs and pages. In particular, GOPs are related to each other; i.e., (i) one GOP might be a higher-quality version of another, and (ii) consecutive GOPs form a contiguous video fragment. These correlations make typical eviction policies like least-recently used (LRU) inefficient. In particular, naïve LRU might evict every other GOP in a physical video, decomposing it into many small fragments and increasing the cost of reads (which have exponential complexity in the number of fragments).

Additionally, given multiple, redundant GOPs that are all variations of one another, ordinary LRU would treat eviction of a redundant GOP the same as any other GOP. However, our intuition is that it is desirable to treat redundant GOPs different than singleton GOPs without such redundancy.

Given this intuition, VSS employs a modified LRU policy (LRU_{VSS}) that associates each fragment with a nonnegative sequence number computed using ordinary LRU offset by:

- **Position (p).** To reduce fragmentation, VSS increases the sequence number of fragments near the middle of a physical video, relative to the beginning or end. For a video with n fragments arranged in ascending temporal order, VSS increases the sequence number of fragment f_i by $p(f_i) = \min(i, n - i)$.
- **Redundancy (r).** VSS decreases the sequence number of fragments that have redundant or higher-quality variants. To do so, using the quality cost model u , VSS generates a u -ordering of each fragment f_i and all other fragments that are a spatiotemporal cover of f_i . VSS decreases the sequence number of f_i by its rank $r(f_i) \in \mathbb{Z}^{0+}$ in this order (i.e., $r(f_i) = 0$ for a fragment with no higher-quality alternatives, while $r(f_i) = n$ for a fragment with n higher-quality variants).
- **Baseline quality (b).** VSS never evicts a fragment if it is the only fragment with quality equal to the quality of the corresponding fragment m_0 in the originally-written physical video. To ensure this, given a set of fragments F in a video, VSS increases the sequence number of each fragment by (our prototype sets $\tau = 40$):

$$b(f_i) = \begin{cases} +\infty & \text{if } \nexists f_j \in F \setminus f_i. u(m_0, f_j) \geq \tau \\ 0 & \text{otherwise} \end{cases}$$

Using the offsets described above, VSS computes the sequence number of each candidate fragment f_i as $LRU_{VSS}(f_i) = LRU(f_i) + \gamma \cdot p(f_i) - \zeta \cdot r(f_i) + b(f_i)$. Here weights γ and ζ balance between position and redundancy, and our prototype weights the former ($\gamma = 2$) more heavily than the latter ($\zeta = 1$). It would be a straightforward extension to expose these as parameters tunable for specific workloads.

5 DATA COMPRESSION IN VSS

VSS employs two compression-oriented optimizations and one optimization that reduces the number of physical video fragments. Specifically, VSS (i) jointly compresses redundant data across multiple physical videos (Section 5.1); (ii) lazily compresses blocks of uncompressed, infrequently-accessed GOPs (Section 5.2); and (iii) improves the read performance by compacting temporally-adjacent video (Section 5.3).

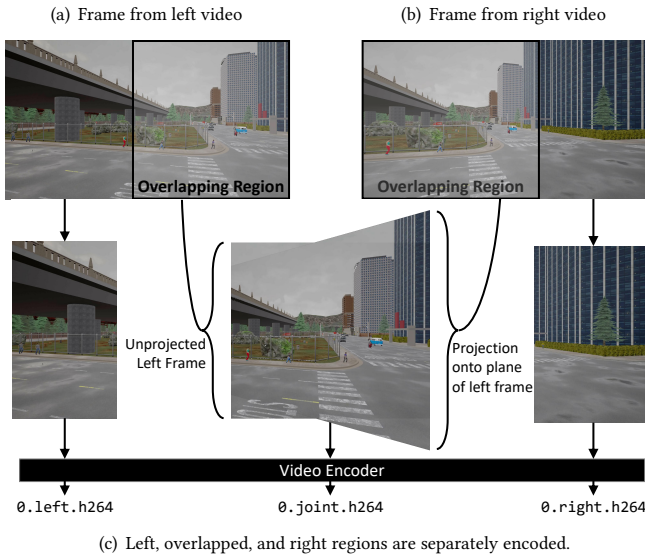


Figure 4: The joint compression process. VSS identifies overlap, combines it, and separately encodes the pieces.

5.1 Joint Physical Video Compression

Increasingly large amounts of video content is produced from cameras that are spatially proximate with similar orientations. For example, a bank of traffic cameras mounted on a pole will each capture video of the same intersection from similar angles. Although the amount of “overlapping video” being produced is difficult to quantify, it broadly includes traffic cameras (7.5PB per day in the United Kingdom [9]), body-worn cameras (>1TB per day [55]), autonomous vehicles (>15TB per vehicle per hour [21]), along with videos of tourist locations, concerts, and political events. Despite the redundant information that mutually exists in these video streams, most applications treat these video streams as distinct and persist them separately to disk.

VSS optimizes the storage of these videos by reducing the redundancy between pairs of highly-similar video streams. It applies this *joint compression* optimization to pairs of GOPs in different logical videos. VSS first finds candidate GOPs to jointly compress. Then, given a pair of overlapping GOP candidates, VSS recompresses them frame-by-frame (we describe this process in Section 5.1.1). For static cameras, once VSS compresses the first frame in a GOP, it can reuse the information it has computed to easily compress subsequent frames in the same GOP. We describe joint compression for dynamic cameras in Section 5.1.2. We finally describe the search process for overlapping GOPs in Section 5.1.3.

5.1.1 Joint frame compression. Figure 4 illustrates the joint compression process for two frames taken from a synthetic dataset (Visual Road-1K-50%, described in Section 6). Figures 4(a) and 4(b) respectively show the two frames with the overlap highlighted. Figure 4(c) shows the combined regions.

Because these frames were captured at different orientations, combining them is non-trivial and requires more than an isomorphic translation or rotation (e.g., the angle of the horizontal

sidewalk is not aligned). Instead, VSS estimates a homography between the two frames and a projection is used to transform between the two spaces. As shown in Figure 4(c), VSS transforms the right frame, causing its right side to bulge vertically. However, after it is overlaid onto the left frame, the two align near-perfectly.

Joint projection proceeds as follows (see [17] for formalization). First, VSS estimates a homography between two frames in the GOPs being compressed. Next, it applies a feature detection algorithm [32] that identifies features that co-occur in both frames. Using these features, it estimates the homography matrix used to transform between frame spaces.

With a homography estimated, VSS projects the right frame into the space of the left frame. This results in three distinct regions: (i) a non-overlapping “left” region of the left frame, (ii) an overlapping region, and (iii) a “right” region of the right frame that does not overlap with the left. VSS splits these into three distinct regions and uses an ordinary video codec to encode each region separately and write it to disk.

When constructing the overlapping region, VSS applies a *merge function* that transforms overlapping pixels from each unprojected region and outputs a merged, overlapping frame. An *unprojected merge* favors the unprojected frame (i.e., the left frame in Figure 4(c)), while a *mean merge* averages the pixels from both input frames. During reads, VSS reverses this process to produce the original frames. We show representative recovered frames in [17].

Some frames stored in VSS may be exact duplicates, however, for which the projection process described above introduces unnecessary computational overhead. VSS detects this case by checking whether the homography matrix would make a near-identity transform (specifically by checking $\|H - I\| \leq \epsilon$, where $\epsilon = \frac{1}{10}$ in our prototype). When this condition is met, VSS instead replaces the redundant GOP with a pointer to its near-identical counterpart.

5.1.2 Dynamic & mixed resolution cameras. For stationary and static cameras, the originally-computed homography is sufficient to jointly compress all frames in a GOP. For dynamic cameras, however, the homography quickly becomes outdated and, in the worst case, the cameras may no longer overlap. To guard against this, for each jointly compressed frame, VSS inverts the projection process and recovers the original frame. It then compares the recovered variant against the original using its quality model (see Section 3.2). If quality is too low (<24dB in our prototype), VSS re-estimates the homography and reattempts joint compression, and aborts if the reattempt is also of low quality.

For both static and dynamic cameras, VSS may occasionally poorly estimate the homography between two otherwise-compatible frames. The recovery process described above also identifies these cases. When detected (and if re-estimation is unsuccessful), VSS aborts joint compression for that pair of GOPs. We show an example of recovered frames where VSS produced an incorrect homography in an extended version of this paper [17].

VSS may also identify joint compression candidates that are at dissimilar resolutions. To handle this case, VSS first upscales the lower resolution fragment to that of the higher. It then applies joint compression as usual.

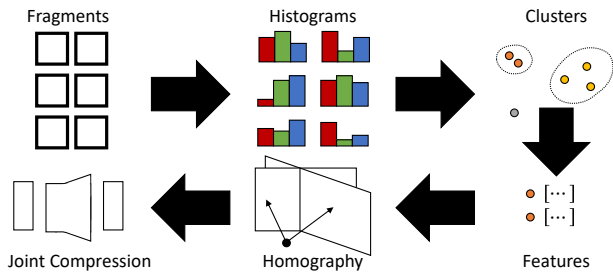


Figure 5: Joint compression fragment selection process: (i) compute and cluster fragment histograms, (ii) for the smallest cluster, compute features and search for fragments with many similar features, and (iii) for each pair, compute homography and (iv) compress.

5.1.3 Selecting GOPs for joint compression. Thus far we have discussed how VSS applies joint compression to a pair of GOPs, but not how the pairs are selected. Since the brute force approach of evaluating all $O(n^2)$ pairs is prohibitively expensive, VSS instead uses the multi-step process illustrated in Figure 5. First, to reduce the search space, VSS clusters all video fragments using their color histograms. Videos with highly distinct color histograms are unlikely to benefit from joint compression. The VSS prototype implementation uses the BIRCH clustering algorithm [57], which is memory efficient, scales to many data points, and allows VSS to incrementally update its clusters as new GOPs arrive.

Once VSS has clustered the ingested GOPs, it selects the cluster with the smallest radius and considers its constituents for joint compression. To do so, VSS applies a modified form of the homography computation described above. It begins by applying the feature detection algorithm [32] from Section 5.1.1. Each feature is a spatial histogram characterizing an “interesting region” in the frame (i.e., a keypoint). VSS next looks for other GOPs in the cluster that share a large number of interesting regions. Thus, for each GOP, VSS iteratively searches for similar features (i.e., within distance d) located in other GOPs within the cluster. A correspondence, however, may be ambiguous (e.g., if a feature in GOP 1 matches to multiple, nearby features in GOP 2). VSS rejects such matches.

When VSS finds m or more nearby, unambiguous correspondences, it considers the pair of GOPs to be sufficiently related. It then applies joint compression to the GOP pair as described above. Note that the algorithm described in Section 5.1.1 will abort if joint compressing the GOPs does not produce a sufficiently high-quality result. Our prototype sets $m = 20$, requires features to be within $d = 400$ (using a Euclidean metric), and disambiguates using Lowe’s ratio [33].

5.2 Deferred Compression

Most video-oriented applications operate over decoded video data (e.g., RGB) that is vastly larger than its compressed counterpart (e.g., the VisualRoad-4K-30% dataset we describe in Section 6 is 5.2TB uncompressed as 8-bit RGB). Caching this uncompressed video quickly exhausts the storage budget.

To mitigate this, VSS adopts the following approach. When a video’s cache size exceeds a threshold (25% in our prototype), VSS activates its *deferred compression* mode. Thereafter when an uncompressed read occurs, VSS orders the video’s uncompressed cache entries by eviction order. It then losslessly compresses the *last* entry (i.e., the one least likely to be evicted). It then executes the read as usual.

Our prototype uses Zstandard for lossless compression, which emphasizes speed over compression ratio (relative to more expensive codecs such as PNG or HEVC) [13].

VSS performs two additional optimizations. First, Zstandard comes with a “compression level” setting in the range [1..19], with the lowest setting having the fastest speed but the lowest compression ratio (and vice versa). VSS linearly scales this compression level with the remaining storage budget, trading off decreased size for increased throughput. Second, VSS also compresses cache entries in a background thread when no other requests are being executed.

5.3 Physical Video Compaction

While caching, VSS persists pairs of cached videos with contiguous time and the same spatial and physical configurations. (e.g., entries at time [0, 90] and [90, 120]). Deferred compression may also create contiguous entries.

To reduce the number of videos that need to be considered during a read, VSS periodically and non-quiescently compacts pairs of contiguous cached videos and substitutes a unified representation. It does so by periodically examining pairs of cached videos and, for each contiguous pair, creating hard links from the second into the first. It then removes the second copy.

6 EVALUATION

We have implemented a prototype of VSS in Python and C++ using CUDA [40], NVENCODE [39], OpenCV [41], FFmpeg [6], and SQLite [45]. Our prototype adopts a no-overwrite policy and disallows updates. We plan on supporting both features in the future. Finally, VSS does not guarantee writes are visible until the file being written is closed.

Baseline systems. We compare against VStore [54], a recent storage system for video workloads, and direct use of the local file system. We build VStore with GPU support. VStore intermittently failed when operating on $>2,000$ frames and so we limit all VStore experiments to this size.

Experimental configuration. We perform all experiments using a single-node system equipped with an Intel i7 processor, 32GB RAM, and a Nvidia P5000 GPU.

Datasets. We evaluate using both real and synthetic video datasets (see Table 1). We use the former to measure VSS performance under real-world inputs, and the latter to test on a variety of carefully-controlled configurations. The “Robotcar” dataset consists of two highly-overlapping videos from vehicle-mounted stereo cameras [36]. The dataset is provided as 7,494 PNG-compressed frames at 30 FPS (as is common for datasets that target machine learning). We cropped and transcoded these frames into a H264 video with one-second GOPs.

Table 1: Datasets used to evaluate VSS

Dataset	Resolution	# Frames	Compressed Size (MB)
Robotcar	1280×960	7,494	120
Waymo	1920×1280	398	7
VisualRoad 1K-30%	960×540	108k	224
VisualRoad 1K-50%	960×540	108k	232
VisualRoad 1K-75%	960×540	108k	226
VisualRoad 2K-30%	1920×1080	108k	818
VisualRoad 4K-30%	3840×2160	108k	5,500

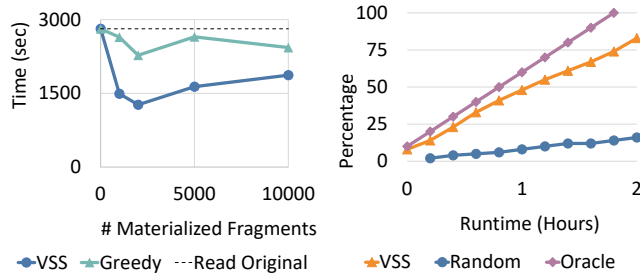


Figure 6: Time to select fragments and read video.

The “Waymo” dataset is an autonomous driving dataset [51]. We selected one segment (~20s) captured using two vehicle-mounted cameras. Unlike the Robotcar dataset, we estimate that Waymo videos overlap by ~15%.

Finally, the various “VisualRoad” datasets consist of synthetic video generated using a recent video analytics benchmark designed to evaluate the performance of video-oriented data management systems [20]. To generate each dataset, we execute a one-hour simulation and produce video data at 1K, 2K, and 4K resolutions. We modify the field of view of each panoramic camera in the simulation so that we could vary the horizontal overlap of the resulting videos. We repeat this process to produce five distinct datasets; for example, “VisualRoad-1K-75%” has two 1K videos with 75% horizontal overlap.

Because the size of the uncompressed 4K Visual Road dataset (~ 5TB) exceeds our storage capacity, we do not show results that require fully persisting this dataset uncompressed on disk.

6.1 Data Retrieval Performance

Long Read Performance. We first explore VSS performance for large reads at various cache sizes. We repeatedly execute queries of the form $read(VRoad-4K-30\%, 4k, [t_1, t_2], P)$, with parameters drawn at random. We assume an infinite budget and iterate until VSS has cached a given number of videos.

We then execute a maximal HEVC read ($t=[0-3600]$), which is different from the originally-written physical video (H264). This allows VSS to leverage its materialized fragments.

Figure 6 shows performance of this read. Since none of the other baseline systems support automatic conversion from H264 to HEVC, we do not show their runtimes for this experiment.

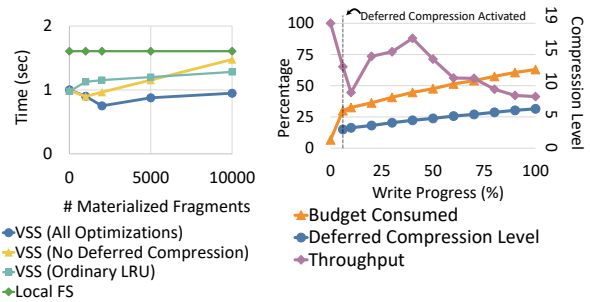


Figure 8: Selecting & Figure 9: Writes with reading short segments. deferred compression.

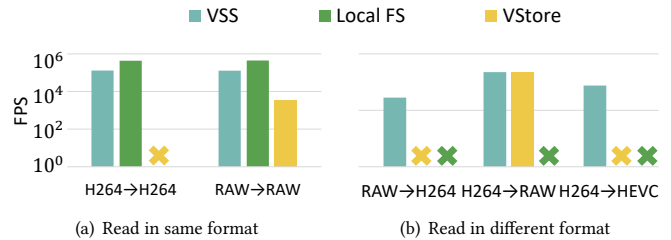


Figure 10: Read throughput. Each group $I \rightarrow O$ shows throughput reading in format I and outputting in format O . An \times indicates lack of support for a read type.

As we see in Figure 6, even a small cache improves read performance substantially—28% at 100 entries and up to a maximum improvement of 54%. Further, because VSS decodes fewer dependent frames, VSS’s solver-based fragment selection algorithm outperforms both reading the original video and a naïve baseline that greedily selects fragments.

Short Read Performance. We next examine VSS performance when reading small, one-second regions of video (e.g., to apply license plate detection only to regions of video that contain automobiles). In this experiment, we begin with the VSS state generated by the previous experiment and execute many short reads of the form $read(VisualRoad-4K-30\%, R, [t_1, t_2], P)$, where $0 \leq t_1 < 3600$ and $t_2 = t_1 + 1$ (i.e., random 1 second sequences). R and P are as in the previous experiment.

Figure 8 shows the result for VSS (“VSS (All Optimizations)”) versus reading from the original video from the local file system (“Local FS”). For this experiment, VSS is able to offer improved performance due to its ability to serve from a cache of lower-cost fragments, rather than transcoding the source video. We discuss the other optimizations in this plot in Section 6.3.

Read Format Flexibility. Our next experiment evaluates VSS’s ability to read video data in a variety of formats. To evaluate, we write the VRoad-1K-30% dataset to each system in both compressed (224MB) and uncompressed form (~328GB). We then read video from each system in various formats and measure throughput. Figure 10 shows read results for the same (10(a)) and different (10(b)) formats. Because the local file system does not support automatic

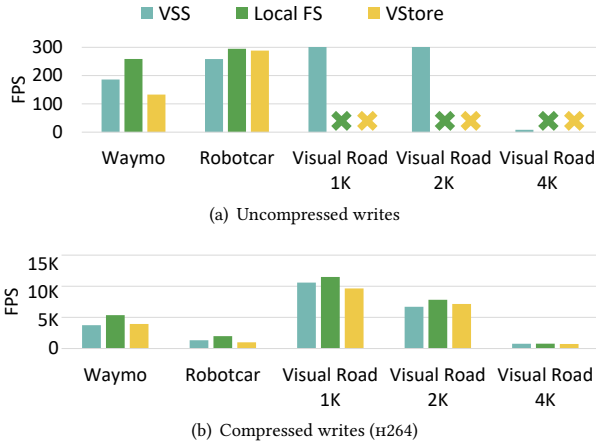


Figure 11: Throughput to write video.

transcoding (e.g., H264 to RGB), we do not show results for these cases. Additionally, VStore does not support reading some formats; we omit these cases.

We find that read performance *without* a format conversion from VSS is modestly slower than the local file system, due in part to the local file system being able to execute entirely without kernel transitions and VSS’s need to concatenate many individual GOPs. However, VSS can adapt to reads in *any* format, a benefit not available when using the local file system.

We also find that VSS outperforms VStore when reading uncompressed video and is similar when transcoding H264. Additionally, VSS offers flexible IO format options and does not require a workload to be specified in advance.

6.2 Data Persistence & Caching

Write Throughput. We next evaluate VSS write performance by writing each dataset to each system in both compressed and uncompressed form. For uncompressed writes, we measure throughput and show results in Figure 11(a).

For uncompressed datasets that fit on local storage, all systems perform similarly. On the other hand, no other systems have the capacity to store the larger uncompressed datasets (e.g., VisualRoad-4K-30% is >5TB uncompressed). However, VSS’s deferred compression allows it to store datasets no other system can handle (though with decreased throughput).

Figure 11(b) shows results for writing the *compressed* datasets to each store. Here all perform similarly; VSS and VStore exhibit minor overhead relative to the local file system.

Cache Performance. To evaluate the VSS cache eviction policy, we begin by executing 5,000 random reads to populate the cache, using the same parameters as in Section 6.1. In this experiment, instead of using an infinite storage budget, we limit it to multiples of the input size and apply either the least-recently used (LRU) or VSS eviction policy. This limits the number of physical videos available for reads. With the cache populated, we execute a final read for the entire video. Figure 12 plots runtimes for each policy and storage budget. This shows that VSS reduces read time relative to LRU.

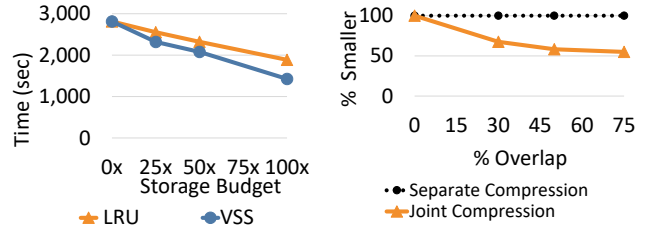


Figure 12: Read runtime by storage budget. Figure 13: Joint compressed cache eviction policy vs. separate video.

Table 2: Joint compression recovered quality

Dataset	Quality (PSNR)		Fragments
	Unprojected Left/Right	Mean Left/Right	Admitted (%) Unprojected / Mean
Robotcar	350 / 24	30 / 27	36 / 64
Waymo	352 / 29	32 / 30	39 / 68
VRoad-1K-30%	359 / 30	31 / 30	46 / 80
VRoad-1K-50%	358 / 28	29 / 29	41 / 72
VRoad-1K-75%	348 / 24	30 / 28	44 / 68
VRoad-2K-30%	352 / 30	30 / 30	52 / 82
VRoad-4K-30%	360 / 30	29 / 30	54 / 78

6.3 Compression Performance

Joint Compression Quality. We first examine the recovered quality of jointly-compressed physical videos. For this experiment we write various overlapping Visual Road datasets to VSS. We then read each video back from VSS and compare its quality—using peak signal-to-noise ratio (PSNR)—against its originally-written counterpart.

Table 2 gives the PSNR for recovered data compared against the written videos. Recall that a PSNR of ≥ 40 is considered to be lossless, and ≥ 30 near-lossless [22]. When applying the unprojected merge function during joint compression, VSS achieves almost perfect recovery for the left input (with PSNR values exceeding 300dB) and near-lossless quality for the right input. Loss in fidelity occurs when inverting the merge, i.e., performing the inverse projection on the right frame using left-frame pixels decreases the quality of the recovered frame. This merge function also leads to VSS rejecting approximately half of the fragments due to their falling below the minimum quality threshold. We conclude this merge function is useful for reducing storage size in video data that must maintain at least one perspective in high fidelity.

On the other hand, VSS attains balanced, near-lossless quality for *both* the left and right frames when applying the mean merge function during joint compression. Additionally, the number of fragments admitted by the quality model is substantially higher under this merge function. Accordingly, the mean merge function is appropriate for scenarios where storage size is paramount and near-lossless degradation is acceptable.

Joint Compression Throughput. We next examine read throughput with and without the joint compression optimization. First, we write the VisualRoad-1K-30% dataset to VSS, once with joint compression enabled and separately with it disabled. We then read

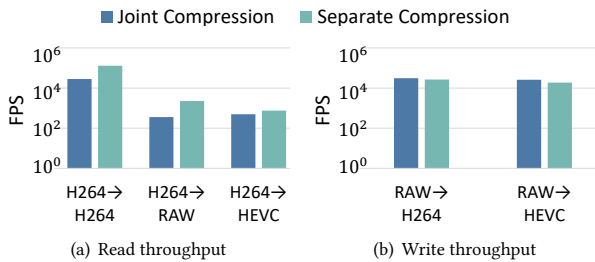


Figure 14: Joint compression throughput.

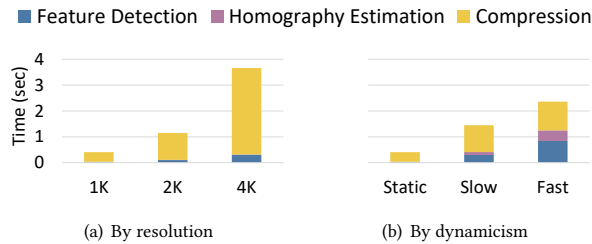


Figure 15: Joint compression overhead.

in various physical configurations over the full duration. Figure 14(a) shows the throughput for reads using each configuration. Our results indicate that read overhead when using joint compression is modest but similar to reads that are not co-compressed.

Joint compression requires several nontrivial operations, and we next evaluate this overhead by writing 1k, 2k, and 4k video and measuring throughput. Figure 14(b) shows the results. Joint writes are similar to writing each video stream separately. This speedup is due to VSS’s encoding the lower-resolution streams in parallel. Additionally, the overhead in feature detection and generating the homography is low. Figure 15 decomposes joint compression overhead into these subcomponents. First, Figure 15(a) measures joint compression overhead by resolution, where compression costs dominate for all resolutions. Figure 15(b) further shows VSS performance under three additional scenarios: a static camera, a slowly rotating camera that requires homography reestimation every fifteen frames, and a rapidly rotating camera that requires reestimation every five frames. In these scenarios non-compression costs scale with the reestimation period, and compression performance is loosely correlated since a keyframe is needed after every homography change.

We next evaluate VSS’s joint compression selection algorithm. Using VisualRoad-1K-30%, we count joint compression candidates using (i) VSS’s algorithm, (ii) an oracle, and (iii) random selection. Figure 7 shows performance of each strategy. VSS identifies 80% of the applicable pairs in time similar to the oracle and outperforms random sampling.

Joint Compression Storage. To show the storage benefit of VSS’s joint compression optimization, we separately apply the optimization to each of the Visual Road videos. We then measure the final on-disk size of the videos against their separately-encoded

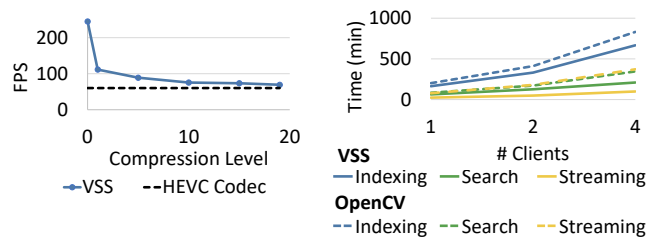


Figure 16: Throughput for reads over fragments with deferred compression. Figure 17: End-to-end application performance.

variants. Figure 13 shows the result of this experiment. These results show joint compression substantially reduces the storage requirements of overlapping video.

Deferred Compression Performance. We next evaluate deferred compression for uncompressed writes by storing 3600 frames of the VisualRoad-1K-30% dataset in VSS, leaving budget and deferred compression at their defaults.

The results are listed in Figure 9. The figure shows storage used as a percentage of the budget, throughput relative to writing without deferred compression activated, and compression level. Storage used exceeds the deferred compression threshold early in the write, and a slope change shows that deferred compression is moderating write size. Compression level scales linearly with storage cost. Throughput drops substantially as compression is activated, recovers considerably, and then slowly degrades as the level is increased.

Similarly, Figure 16 shows throughput for reading fragments of raw video compressed at various levels. Though these reads have decreased performance and increased variance relative to uncompressed reads, at all levels Zstandard decompression remains much faster than using traditional video codecs.

Finally, Figure 8 explores the trade-offs between deferred compression performance and VSS’s cache eviction policy. In this experiment we variously disable deferred compression (“VSS (No Deferred Compression)”) and modify VSS to use ordinary LRU (“VSS (Ordinary LRU)”). The results show that VSS benefits from its eviction policy for small numbers of fragments (when deferred compression is off or at a low level) but offers increasingly large benefits as the cache grows. At large cache sizes as the storage budget is exhausted, deferred compression is increasingly important to mitigate eviction of fragments that are subsequently useful.

6.4 End-to-End Application Performance

Our final experiment evaluates the performance of the end-to-end application described in Section 2. In this scenario, VSS serves as the storage manager for an application monitoring an intersection for automobiles. It involves three steps: (i) an *indexing phase* that identifies video frames containing automobiles using a machine learning algorithm, (ii) a *search phase* that, given an alert for a missing vehicle, uses the index built in the previous step to query video frames containing vehicles with matching colors, and (iii) a *streaming content retrieval phase* that uses the search phase results to retrieve video clips containing vehicles of a given color.

We implement this application using VSS and a variant that reads video data using OpenCV and the local file system. For indexing, the application identifies automobiles using YOLOv4 [8] (both variants use OpenCV to perform inference using this model). For the search task, vehicle color is identified by computing a color histogram of the region inside the bounding box. We consider a successful detection to occur when the Euclidean distance between the largest bin and the search color is ≤ 50 . In the content retrieval phase, the application generates n video clips by retrieving contiguous frames containing automobiles of the search color.

We use as input four extended two-hour variants of the Visual Road 2k dataset. To simulate execution by multiple clients, we launch a separate process for each client. Both variants index automobiles every ten frames (i.e., three times a second). All steps exhaust all CPU resources at > 4 clients, and so we limit concurrent requests to this maximum.

Figure 17 shows the performance of each application step. The indexing step is a CPU-intensive operation that necessitates both video decoding and model inference, and because VSS introduces low overhead for reads, both variants perform similarly. Conversely, VSS excels at executing the search step, which requires retrieving raw, uncompressed frames that were cached during the indexing step. As such, it substantially outperforms the OpenCV variant. Finally, VSS’s ability to efficiently identify the lowest-cost transcode solution enables it to execute the streaming content retrieval step significantly faster than the OpenCV variant. We conclude that VSS’s performance greatly improves end-to-end application performance for queries that depend on cached video in multiple formats, and scales better with multiple clients.

7 RELATED WORK

Increased interest in machine learning and computer vision has led to the development of a number of systems that target video analytics, including LightDB [19], VisualWorldDB [18], Optasia [34], Chameleon [25], Panorama [58], Vaas [4], SurvQ [46], and Scanner [43]. These systems can be modified to leverage a storage manager like VSS. Video accelerators such as Blazelt [27], VideoStorm [56], Focus [23], NoScope [28], Odin [47], SQV [53], MIRIS [3], Tahoma[2], and Deluceva [50] can also benefit from VSS for training and inference.

Few recent storage systems target video analytics (although others have highlighted this need [14, 26]). VStore [54] targets machine learning workloads by staging video in pre-specified formats. However, VStore requires a priori knowledge of the workload and only makes preselected materializations available. By contrast, quFiles exploits data independence at the granularity of entire videos [49]. Others have explored on-disk layout of video for scalable streaming [30], and systems such as Haystack [5], AWS Serverless Image Handler [1], and VDMS [44] emphasize image and metadata operations.

Techniques similar to VSS’s joint compression optimization have been explored in the image and signal processing communities. For example, Melloni et al. develop a pipeline that identifies and aligns near-duplicate videos [38], and Pinheiro et al. introduce a fingerprinting method to identify correlations among near-duplicate videos [42]. However, unlike VSS, these techniques

assume that sets of near-duplicate videos are known a priori and they do not exploit redundancies to improve compression or read/write performance. Finally, the multiview extension to HEVC (MV-HEVC; similar extensions exist for other codecs) attempts to exploit spatial similarity in similar videos to improve compression performance [16]. These extensions are complementary to VSS, which could incorporate them as an additional compression codec for jointly-compressed video.

Finally, as in VSS, the database community has long exploited data independence to improve performance. Orthogonal optimizations could further improve VSS performance (e.g., perceptual compression [37] or homomorphic operators [19]).

8 CONCLUSION

We presented VSS, a video storage system that improves the performance of video-oriented applications. VSS decouples high-level operations (e.g., machine learning) from the low-level plumbing to read and write data in a suitable format. VSS automatically identifies the most efficient method to persist and retrieve video data. VSS reduces read time by up to 54%, and decreases the cost of persisting video by up to 45%.

As future work, we plan on extending VSS’s joint compression optimization to support more intelligent techniques for merging overlapping pixels. For example, VSS might intelligently detect occlusions and persist both pixels in these areas. This is important for cases where video must be maintained in its (near-)original form (e.g., for legal reasons).

Acknowledgments. This work is supported by the NSF through grants CCF-1703051, IIS-1546083, CCF-1518703, and CNS-1563788; DARPA award FA8750-16-2-0032 and RC Center grant G118518; DOE award DE-SC0016260; a Google Faculty Research Award; an award from the University of Washington Reality Lab; a CoMotion Commercialization Fellows grant; gifts from the Intel Science and Technology Center for Big Data, Intel Corporation, Adobe, Amazon, Facebook, Huawei, and Google; and by CRISP, one of six centers in JUMP, a Semiconductor Research Corporation program sponsored by DARPA.

REFERENCES

- [1] Amazon. 2020. Serverless Image Handler. <https://aws.amazon.com/solutions/implementations/serverless-image-handler>.
- [2] Michael R. Anderson, Michael J. Cafarella, German Ros, and Thomas F. Wenisch. 2019. Physical Representation-Based Predicate Optimization for a Visual Analytics Database. In *ICDE*. IEEE, 1466–1477.
- [3] Favyen Bastani, Songtao He, Arjun Balasingam, Karthik Gopalakrishnan, Mohammad Alizadeh, Hari Balakrishnan, Michael J. Cafarella, Tim Kraska, and Sam Madden. 2020. MIRIS: Fast Object Track Queries in Video. In *SIGMOD*. 1907–1921.
- [4] Favyen Bastani, Oscar R. Moll, and Samuel Madden. 2020. Vaas: Video Analytics at Scale. *VLDB* 13, 12 (2020), 2877–2880.
- [5] Doug Beaver, Sanjeev Kumar, Harry C. Li, Jason Sobel, and Peter Vajgel. 2010. Finding a Needle in Haystack: Facebook’s Photo Storage. In *OSDI*. 47–60.
- [6] Fabrice Bellard. 2018. FFmpeg. <https://ffmpeg.org>.
- [7] Daniel S. Berger, Nathan Beckmann, and Mor Harchol-Balter. 2018. Practical Bounds on Optimal Caching with Variable Object Sizes. *POMACS* 2, 2 (2018), 32:1–32:38.
- [8] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. 2020. YOLOv4: Optimal Speed and Accuracy of Object Detection. *arXiv CoRR abs/2004.10934* (2020).
- [9] Cloudview. 2018. Visual IoT: Where the IoT Cloud and Big Data Come Together. <http://www.cloudview.co/downloads/10>. (2018).

- [10] Victor H. Costa, Pedro A. Amado Assunção, and Paulo J. Cordeiro. 2018. A pixel-based complexity model to estimate energy consumption in video decoders. In *ICCE*. 1–5.
- [11] Maureen Daum, Brandon Haynes, Dong He, Amrita Mazumdar, Magdalena Balazinska, and Alvin Cheung. 2020. TASM: A Tile-Based Storage Manager for Video Analytics. arXiv:2006.02958
- [12] Leonardo Mendonça de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *TACAS*. 337–340.
- [13] Facebook. [n.d.]. Zstandard real-time compression algorithm. <https://facebook.github.io/zstd>.
- [14] Vishakha Gupta-Cledat, Luis Remis, and Christina R. Strong. 2017. Addressing the Dark Side of Vision Research: Storage. In *HotStorage*.
- [15] Alon Y. Halevy. 2001. Answering queries using views: A survey. *VLDB* 10, 4 (2001), 270–294.
- [16] Miska M. Hannuksela, Ye Yan, Xuehui Huang, and Houqiang Li. 2015. Overview of the multiview high efficiency video coding (MV-HEVC) standard. In *ICIP*. 2154–2158.
- [17] Brandon Haynes, Maureen Daum, Dong He, Amrita Mazumdar, Magdalena Balazinska, Alvin Cheung, and Luis Ceze. 2021. VSS: A Storage System for Video Analytics [Technical Report]. arXiv:2103.16604
- [18] Brandon Haynes, Maureen Daum, Amrita Mazumdar, Magdalena Balazinska, Alvin Cheung, and Luis Ceze. 2020. VisualWorldDB: A DBMS for the Visual World. In *CIDR*.
- [19] Brandon Haynes, Amrita Mazumdar, Armin Alaghi, Magdalena Balazinska, Luis Ceze, and Alvin Cheung. 2018. LightDB: A DBMS for Virtual Reality Video. *PVLDB* 11, 10 (2018), 1192–1205.
- [20] Brandon Haynes, Amrita Mazumdar, Magdalena Balazinska, Luis Ceze, and Alvin Cheung. 2019. Visual Road: A Video Data Management Benchmark. In *SIGMOD*. 972–987.
- [21] Stephan Heinrich and Lucid Motors. 2017. Flash memory in the emerging age of autonomy. *Flash Memory Summit* (2017).
- [22] Alain Horé and Djemel Ziou. 2010. Image Quality Metrics: PSNR vs. SSIM. In *ICPR*. 2366–2369.
- [23] Kevin Hsieh, Ganesh Ananthanarayanan, Peter Bodík, Shivaram Venkataraman, Paramvir Bahl, Matthai Philipose, Phillip B. Gibbons, and Onur Mutlu. 2018. Focus: Querying Large Video Datasets with Low Latency and Low Cost. In *OSDI*. 269–286.
- [24] Shelley S Hyland. 2018. Body-worn cameras in law enforcement agencies, 2016. *Bureau of Justice Statistics Publication No. NCJ251775* (2018).
- [25] Junchen Jiang, Ganesh Ananthanarayanan, Peter Bodík, Siddhartha Sen, and Ion Stoica. 2018. Chameleon: scalable adaptation of video analytics. In *SIGCOMM*. 253–266.
- [26] Junchen Jiang, Yuhao Zhou, Ganesh Ananthanarayanan, Yuanhao Shu, and Andrew A. Chien. 2019. Networked Cameras Are the New Big Data Clusters (*HotEdgeVideo'19*). 1–7.
- [27] Daniel Kang, Peter Bailis, and Matei Zaharia. 2019. BlazeIt: Optimizing Declarative Aggregation and Limit Queries for Neural Network-Based Video Analytics. *PVLDB* 13, 4 (2019), 533–546.
- [28] Daniel Kang, John Emmons, Firas Abuzaid, Peter Bailis, and Matei Zaharia. 2017. NoScope: Optimizing Deep CNN-Based Queries over Video Streams at Scale. *PVLDB* 10, 11 (2017), 1586–1597.
- [29] Daniel Kang, Ankit Mathur, Teja Veeramacheni, Peter Bailis, and Matei Zaharia. 2020. Jointly Optimizing Preprocessing and Inference for DNN-based Visual Analytics.
- [30] Sooyong Kang, Sungwoo Hong, and Youjip Won. 2009. Storage technique for real-time streaming of layered video. *MMSys* 15, 2 (2009), 63–81.
- [31] Andrea Lottarini, Alex Ramirez, Joel Coburn, Martha A. Kim, Parthasarathy Ranganathan, Daniel Stodolsky, and Mark Wachler. 2018. vbench: Benchmarking Video Transcoding in the Cloud. In *ASPLOS*. 797–809.
- [32] David G. Lowe. 1999. Object Recognition from Local Scale-Invariant Features. In *ICCV*. 1150–1157.
- [33] David G. Lowe. 2004. Distinctive Image Features from Scale-Invariant Keypoints. *IJCV* 60, 2 (2004), 91–110.
- [34] Yao Lu, Aakanksha Chowdhery, and Srikanth Kandula. 2016. Optasia: A Relational Platform for Efficient Large-Scale Video Analytics. In *SoCC*. 57–70.
- [35] Yao Lu, Aakanksha Chowdhery, Srikanth Kandula, and Surajit Chaudhuri. 2018. Accelerating Machine Learning Inference with Probabilistic Predicates. In *SIGMOD*. 1493–1508.
- [36] Will Maddern, Geoff Pascoe, Chris Linegar, and Paul Newman. 2017. 1 Year, 1000km: The Oxford RobotCar Dataset. *IJRR* 36, 1 (2017), 3–15.
- [37] Amrita Mazumdar, Brandon Haynes, Magda Balazinska, Luis Ceze, Alvin Cheung, and Mark Oskin. 2019. Perceptual Compression for Video Storage and Processing Systems. In *SoCC*. 179–192.
- [38] Andrea Melloni, S. Lameri, Paolo Bestagini, Marco Tagliasacchi, and Stefano Tubaro. 2015. Near-duplicate detection and alignment for multi-view videos. In *ICIP*. 2444–2448.
- [39] nvec 2020. Nvidia Video codec. <https://developer.nvidia.com/nvidia-video-codec-sdk>.
- [40] NVIDIA Corporation. 2007. *NVIDIA CUDA Compute Unified Device Architecture Programming Guide*. NVIDIA Corporation.
- [41] OpenCV. 2018. Open Source Computer Vision Library. <https://opencv.org>.
- [42] Giuliano Pinheiro, Marcos Cirne, Paolo Bestagini, Stefano Tubaro, and Anderson Rocha. 2019. Detection and Synchronization of Video Sequences for Event Reconstruction. In *ICIP*. 4060–4064.
- [43] Alex Poms, Will Crichton, Pat Hanrahan, and Kayvon Fatahalian. 2018. Scanner: efficient video analysis at scale. *TOG* 37, 4 (2018), 138:1–138:13.
- [44] Luis Remis, Vishakha Gupta-Cledat, Christina R. Strong, and Ragaad Altarawneh. 2018. VDMS: An Efficient Big-Visual-Data Access for Machine Learning Workloads. *CoRR* abs/1810.11832 (2018).
- [45] SQLite Consortium. 2020. SQLite. <https://www.sqlite.org>.
- [46] Michael Stonebraker, Bharat Bhargava, Michael Cafarella, Zachary Collins, Jenna McClellan, Aaron Sipsper, Tao Sun, Alina Nesen, KMA Solaiman, Ganapathy Mani, et al. 2020. Surveillance Video Querying with a Human-in-the-Loop. (2020).
- [47] Abhijit Suprem, Joy Arulraj, Calton Pu, and João Eduardo Ferreira. 2020. ODIN: Automated Drift Detection and Recovery in Video Analytics. *VLDB* 13, 11 (2020), 2453–2465.
- [48] Yao-Hung Hubert Tsai, Santosh Kumar Divvala, Louis-Philippe Morency, Ruslan Salakhutdinov, and Ali Farhadi. 2019. Video Relationship Reasoning Using Gated Spatio-Temporal Energy Graph. In *CVPR*. 10424–10433.
- [49] Kaushik Veeraraghavan, Jason Flinn, Edmund B. Nightingale, and Brian Noble. 2010. quFiles: The Right File at the Right Time. In *FAST*. 1–14.
- [50] Jingjing Wang and Magdalena Balazinska. 2020. Deluceva: Delta-Based Neural Network Inference for Fast Video Analytics. In *SSDBM*. 14:1–14:12.
- [51] waymo 2020. Waymo Open Dataset. <https://waymo.com/open>.
- [52] Susan Wojcicki. 2018. The Potential Unintended Consequences of Article 13. <https://youtube-creators.googleblog.com/2018/11/i-support-goals-of-article-13-i-also.html>.
- [53] Ioannis Xarchakos and Nick Koudas. 2019. SVQ: Streaming Video Queries. In *SIGMOD*. 2013–2016.
- [54] Tiantu Xu, Luis Materon Botelho, and Felix Xiaozhu Lin. 2019. VStore: A Data Store for Analytics on Large Videos. In *EuroSys*. 16:1–16:17.
- [55] Billy Yates. 2018. Body Worn Cameras: Making Them Mandatory. (2018).
- [56] Haoyu Zhang, Ganesh Ananthanarayanan, Peter Bodík, Matthai Philipose, Paramvir Bahl, and Michael J. Freedman. 2017. Live Video Analytics at Scale with Approximation and Delay-Tolerance. In *NSDI*. 377–392.
- [57] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. 1996. BIRCH: An Efficient Data Clustering Method for Very Large Databases. In *SIGMOD*. 103–114.
- [58] Yuhao Zhang and Arun Kumar. 2019. Panorama: A Data System for Unbounded Vocabulary Querying over Video. *VLDB* 13, 4 (2019), 477–491.