Securing Immersive 360 Video Streams through Attribute-Based Selective Encryption

Mohammad Waquas Usmani
mohammadwaqu@umass.edu
University of Massachusetts Amherst
Massachusetts, USA

Susmit Shannigrahi
sshannigrahi@tntech.edu
Tennessee Technological University
Tennessee, USA

Michael Zink
zink@ecs.umass.edu
University of Massachusetts Amherst
Massachusetts, USA

Abstract

Delivering high-quality, secure 360° video content introduces unique challenges, primarily due to the high bitrates and interactive demands of immersive media. Traditional HTTPS-based methods, although widely used, face limitations in computational efficiency and scalability when securing these high-resolution streams. To address these issues, this paper proposes a novel framework integrating Attribute-Based Encryption (ABE) with selective encryption techniques tailored specifically for tiled 360° video streaming. Our approach employs selective encryption of frames at varying levels to reduce computational overhead while ensuring robust protection against unauthorized access.

Moreover, we explore viewport-adaptive encryption, dynamically encrypting more frames within tiles occupying larger portions of the viewer's field of view. This targeted method significantly enhances security in critical viewing areas without unnecessary overhead in peripheral regions. We deploy and evaluate our proposed approach using the CloudLab testbed, comparing its performance against traditional HTTPS streaming. Experimental results demonstrate that our ABE-based model achieves reduced computational load on intermediate caches, improves cache hit rates, and maintains comparable visual quality to HTTPS, as assessed by Video Multimethod Assessment Fusion (VMAF).

Keywords: 360° Video Streaming, DASH, Attribute-Based Encryption, Caching, ABR, Quality of Experience, CDN

1 Introduction

Digital Rights Management (DRM) has played a pivotal role in the widespread adoption and success of online video streaming by protecting content from unauthorized access and distribution. Traditional DRM approaches secure video content in transit and at rest separately, primarily relying on Hypertext Transfer Protocol Secure (HTTPS) coupled with Transport Layer Security (TLS). The success of video streaming is also based on the support by Content Distribution Networks (CDN) [2, 18], which make streaming scalable to hundreds of millions of users.

However, employing HTTPS with adaptive bitrate (ABR) streaming—commonly implemented through protocols such as Dynamic Adaptive Streaming over HTTP (DASH) on top of CDN distribution architectures—introduces significant computational overhead due to the necessity of decrypting and re-encrypting content at intermediate caching nodes. This limitation becomes particularly pronounced in bandwidth-intensive scenarios such as 360° video streaming, where the interactive nature and high bitrate require additional networking and computing resources.

To overcome these limitations, prior work [29] introduced Attribute-Based Encryption for DRM-enabled conventional video streaming.

This approach secures the data itself, removing the need for separate transport-layer encryption. It simplifies content distribution by encrypting data once with specific attributes, allowing caching nodes to serve encrypted video without additional cryptographic operations. Clients obtain attribute-based keys from a license server to decrypt and access the content, significantly reducing the computational load on caches and facilitating scalable streaming.

The work presented in this paper addresses the challenges of securely and efficiently streaming 360° videos. Due to their spherical nature, 360° videos are projected onto two-dimensional planes, tiled, and selectively streamed based on the user's viewport [35]. Depending on the viewer's head movement, their viewport might be composed of portions of different tiles. Consequently, the video frames of several tiles have to be streamed to the client in parallel, which is in contrast to conventional video, where there is only a maximum of one stream. Our approach introduces selective frame encryption strategies tailored explicitly to the tiled structure of 360° videos. This 2-dimensional method combines frame-selective and tile-selective encryption to reduce computational overhead while preserving content protection. Different frame types are encrypted at varying levels based on each tile's relevance to the viewport.

We evaluate our proposed approach in CloudLab [10], comparing it directly against conventional HTTPS-based streaming. Experimental results indicate that employing ABE substantially reduces computational overhead at intermediate caches, while preserving cache efficiency and maintaining comparable video quality, as assessed by the Video Multimethod Assessment Fusion (VMAF) metric. This work illustrates the benefits of integrating ABE into tiled 360° video streaming, and also establishes a promising pathway toward scalable and efficient secure immersive media distribution.

This paper makes the following contributions. We introduce an ABE-based two-dimensional selective encryption method, which is specifically designed for scalable streaming of 360° video. In addition, we conduct an in-depth evaluation, with two distinct experiment setups: a small-scale experimental and a large-scale multicache hierarchical experiment with strict bandwidth limitations. The results of this evaluation show that CPU usage at both the server and caches is significantly reduced with up to 63% reduction at the cache. In several cases, it also delivers improved hit rates, while maintaining QoE comparable to the TLS-based approach.

2 Background and Related Work

2.1 Fundamentals of 360° Video

360° videos offer immersive experiences by capturing a complete spherical view, enabling viewers to interactively explore the environment. As illustrated in Fig. 1, these videos provide three degrees

of freedom (3DoF): yaw (left/right), pitch (up/down), and roll (rotational movement), contributing to an engaging user experience.
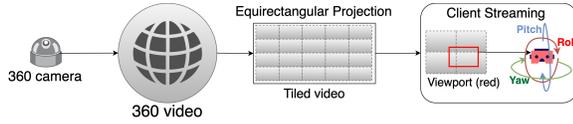

Figure 1: 360° Video Streaming Process

Equirectangular Projection: To encode spherical video data using standard video codecs, a spherical view must be projected onto a two-dimensional plane. The most widely used method is equirectangular projection [11, 33], mapping the spherical surface onto a rectangular grid. This allows compatibility with conventional video formats and codecs while preserving spatial relationships crucial for accurate spherical rendering during playback.

Viewport: A viewport represents the visible portion of the spherical content, typically defined by a horizontal and vertical field of view (FOV) ranging from 90 to 120 degrees. The viewport dynamically adjusts based on user interactions via VR headsets, mobile devices, or computer interfaces. By streaming primarily the viewport region, substantial bandwidth and computational savings are achieved without diminishing the immersive experience [11].

Tiling: Delivering high-quality streaming across an entire 360° video simultaneously is bandwidth-intensive and inefficient. To optimize performance, videos are spatially divided into smaller rectangular sections known as tiles. These tiles, derived from the equirectangular projection, allow adaptive streaming by enabling higher-quality delivery for tiles within the user's viewport, while peripheral tiles are streamed at lower quality or omitted. This technique significantly reduces bandwidth consumption and ensures smooth playback even under constrained network conditions [34].

Fig. 1 outlines the complete streaming workflow: capturing spherical video content, projecting it onto an equirectangular plane, partitioning it into tiles, and selectively streaming high-quality tiles based on the user's real-time viewport.

2.2 Attribute-Based Encryption (ABE)

ABE is a cryptographic technique designed to enable fine-grained access control for encrypted data. Unlike traditional symmetric encryption methods like AES, ABE uses attributes—descriptive elements assigned to users or data—to define access policies. In this scheme, data can be decrypted only if a user's attributes satisfy the access policy associated with the encrypted content.

The operation of ABE begins with a setup phase conducted by a Trusted Authority (TA). During this phase, the TA generates a Master Key and a Public Key. The Data Owner encrypts content using the Public Key and defines an access policy based on attributes. The TA, using the Master Key, generates Private Keys for users, tailored to their specific attributes. When a user attempts to decrypt a ciphertext, the system verifies whether their attributes satisfy the encryption policy. If the attributes match, the decryption succeeds, granting access to the content. In our work, we utilized the **CPABE-Toolkit**, a command-line tool [6], which is based on the cryptographic framework outlined in [7]. Prior work [29] has utilized ABE to provide segment-level security for conventional video.

In this work, the authors apply ABE to individual video frames within each segment.

2.3 Selective Encryption of Video Streams

Selective encryption has been a widely researched topic in multimedia[1, 17, 32]. A recent study [1], explores the use of Advanced Encryption Standard (AES) for selectively encrypting H264/AVC videos. This system processes the H264 bitstream to identify *I*-frames and encrypts them using AES, based on the assumption that removing access to *I*-frames renders the dependent *P* and *B*-frames ineffective. While it is true that *I*-frames are self-contained and provide all the necessary information to display a complete image [25], this approach overlooks the fact that *P* and *B*-frames may still carry meaningful information. *P* and *B*-frames contain *I*-blocks, and when a sequence of these frames is correlated with their reference frames, they can still convey significant visual information [3]. Building on this observation, our work explores encrypting *P* and *B*-frames in addition to *I*-frames in H264 bitstreams, employing varying levels of encryption. This approach aims to address the limitations of solely encrypting *I*-frames by targeting the residual visual information carried by dependent frames, thereby enhancing security and reducing the risk of meaningful content leakage.

2.4 VMAF for 360° Video QOE Evaluation

Video Multi-method Assessment Fusion (VMAF) [16] is a video quality metric that combines human visual perception models with machine learning techniques to evaluate video quality. It assesses how closely a distorted or degraded video resembles a reference video, providing an objective measure of user experience. VMAF scores range from 0 to 100, with higher scores indicating better video quality and a more enjoyable viewing experience.

The metric is widely used in video streaming and encoding research to quantify quality degradation and optimize video delivery systems. Specifically, Orduna et al. [21] explore the application of VMAF for 360° videos. While VMAF was originally designed to evaluate conventional 2D videos, [21] demonstrated that it can also be effectively applied to 360° videos without any modifications. Their evaluation was based on leveraging a diverse dataset of 360° videos in equirectangular projection format.

2.5 HTTP-based Video Streaming and CDNs

DASH [27] and HLS [12] are popular ABR streaming standards that segment videos at various bitrates, allowing clients to adapt quality via algorithms like BOLA [28] or Pensieve [19]. HTTPS secures traffic via TLS [24], but its point-to-point encryption adds computational overhead, especially when 360° videos are cached.

Content Distribution Networks (CDNs) efficiently distribute digital content by caching it close to users, minimizing buffering and latency. Popular streaming services leverage CDNs extensively, employing caches globally to ensure high-quality service. However, the use of HTTPS complicates caching by necessitating TLS termination and re-encryption at CDN nodes, significantly adding overhead to streaming systems [29, 31].

3 System Architecture

The architecture we created for ABE-based 360° video distribution is inspired by the work presented in [29]. While previous work

has primarily focused on conventional video streaming, our approach is specifically tailored to 360° video content, requiring the design of an architecture that enables selective encryption for tiled 360° videos. Our proposed architecture is realized in a prototype implementation, which we use for the evaluation of our approach.
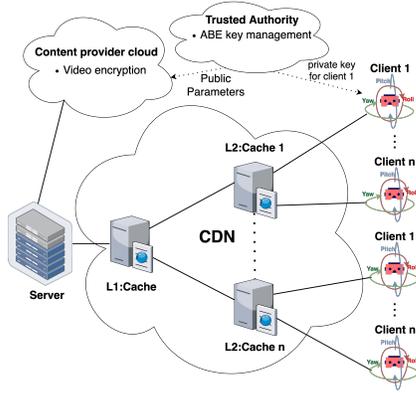

Figure 2: System Architecture

3.1 HTTP+ABE for Secured Streaming:

Figure 2 illustrates the system architecture of our ABE-based 360° video streaming framework, integrated into a typical DASH-based distribution pipeline.

The architecture consists of three key entities: Content Provider, the CDN, and a Trusted Authority. It includes three main components: origin servers, CDN edge servers (caches), and clients (e.g, head-mounted devices (HMD), laptops, phones). The Content Provider divides the equirectangular projection of the 360° videos into tiles and encodes each into multiple qualities, generates DASH metadata (e.g., MPD), and handles content encryption using ABE. All content is originally stored on the origin server.

Client requests are directed to edge servers based on CDN policies [18]. If the requested content is cached, the edge server delivers it directly; otherwise, it retrieves it from the origin server and decides whether to cache it for future use.

Integration into HTTP-based Streaming: To integrate ABE into standard DASH streaming, we made several modifications: A Trusted Authority (TA) generates a master key and a public key. The public key is used to encrypt video segments using an ABE policy, while the TA uses the master key to generate user-specific private keys embedded with access attributes. Unlike HTTPS-based approaches, where key exchange is handled via TLS, in the case of ABE, a separate mechanism is used where the TA issues decryption keys directly to clients. CDN caches remain agnostic to the decryption process—they store encrypted segments but do not perform any cryptographic operations.

Our system also introduces selective encryption of video streams (see Sect. 3.2). While prior work [22, 26] has explored selective encryption, modern HTTP-based streaming is typically binary: either fully encrypted (via HTTPS) or not encrypted (HTTP) at all. To support selective encryption in DASH, it is required that the server or cache can communicate to the client which portions of the video stream are encrypted. In our approach, this is achieved by

modifying the MPD to include "encryption level" for a segment that indicates what video frames are encrypted. On the client side, once a segment is downloaded, the DASH player reads this encryption level and invokes selective decryption logic. The client uses its private key to decrypt the targeted frames before playback. Additional implementation details are provided in Section 3.3.

All other components of standard DASH remain unchanged, allowing seamless integration into existing streaming infrastructures.

3.2 Selective Encryption for Tiled 360° Videos

To balance content protection and computational efficiency, we propose a flexible, frame-based, two-dimensional selective encryption approach tailored to 360° video streaming. Rather than striving for complete secrecy, *our aim is to significantly impair unauthorized viewing by degrading video quality through partial encryption.* This strategy represents a trade-off between the extent of encryption and the resources consumed during en- and decryption. In addition, *our encryption approach is specifically designed for tile-based 360° streaming through the introduction of a scheme that adjusts the level of encryption based on the importance of a tile.*

In the first dimension, our method supports multiple levels of encryption, ranging from encrypting only *I*-frames to encrypting *P*-frames or all frames, providing scalability in protection and performance. Designed for H264/AVC encoded videos, we parse Network Abstraction Layer (NAL) units from DASH segments to identify *I, P,* and *B*-frames. Selected frames are encrypted using the CPABE toolkit with attribute-based policies, and reinserted into their original locations. We update frame size headers accordingly but intentionally avoid modifying the MP4 container's metadata and offsets, rendering the video unplayable unless decrypted correctly, adding an additional layer of access control.

Decryption mirrors this process: encrypted frame bytes are identified, decrypted using authorized keys, and reinserted into the segment. Frame size headers are restored to make the segment playable again. Figure 3 illustrates our selective encryption pipeline.
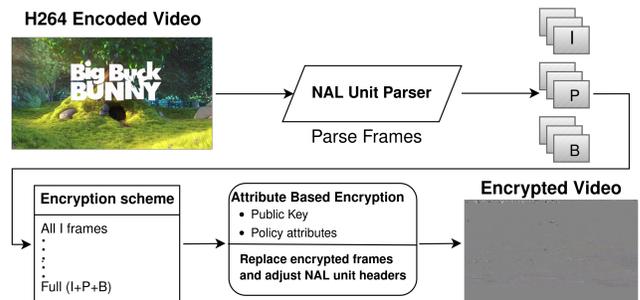

Figure 3: Frame-based selective encryption using ABE.

We extend this concept in the second dimension with **Tile-Based Selective Encryption**, which applies frame-based encryption selectively within tiles of a 360° video based on their relevance to the user's viewport. To illustrate our approach, we use the example of a 360° video that is tiled in a 3x3 grid. Without loss of generality, our approach can be easily extended to other tiling schemes. For tiles covering the major portion of the viewport—determined dynamically during playback—we aggressively encrypt both *I* and *P*-frames.
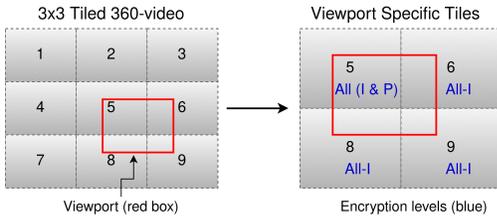
Figure 4: Tile-based selective frame encryption.

For peripheral viewport tiles, we encrypt only *I*-frames, achieving notable degradation in quality with reduced computational effort.

Figure 4 illustrates this process, where tile 5, occupying the core of the viewport, is fully encrypted (*I* and *P*-frames), while tiles 6, 8, and 9 are partially encrypted (*I*-frames only). This dual-layer strategy—spatial tiling and frame-based selection—achieves strong content protection while minimizing computational load.

3.3 Prototype Implementation

The evaluation of ABE for 360° video streaming is based on a prototype we implemented, making use of the following components:

CPABE Toolkit: We implement Attribute-Based Encryption (ABE) using the CPABE toolkit [6], based on the work in [7].

MPD: We modified the Media Presentation Description (MPD) to support our tiled 360° video streaming experiments using Attribute-Based Encryption (ABE). In our setup, each video segment consists of a group of four tiles representing the user's viewport. This structure is maintained across each adaptation set (representing different video bitrates), where each segment entry represents a list of four tile files. To enable adaptive bitrate streaming, the bitrate of each adaptation set is specified as the sum of the bitrates of the four tiles. Furthermore, each tile's filename includes a suffix that reflects the level of selective encryption (e.g., *allI* or *allI+P*) applied to it. This informs the client which frame types (*I* or *P*) are encrypted, guiding it to perform the corresponding ABE decryption during playback.

AStream-360: We extend the AStream DASH client, a Python-based DASH emulation tool [14, 15] supporting various ABR algorithms, to support 360° tiled video streaming with selective encryption. For each playback segment, AStream-360 downloads a set of four tiles corresponding to the user's current viewport and invokes the ABE selective decryption module to decrypt targeted frames within each tile. All four tiles are fetched at the same quality level, selected by the basic ABR algorithm based on the combined download and decryption time across all tiles. For simplicity, only viewport-relevant tiles are downloaded in our current implementation; however, this design can be easily extended to include non-viewport tiles to support seamless viewport transitions.

4 Evaluation

In this section, we first analyze the impact of our two-dimensional selective encryption technique under ABE for tiled 360° videos. We then describe the experimental setup and the metrics used to assess the performance of our attribute-based selective end-to-end encryption framework. Finally, we present experimental results comparing our approach to HTTPS-based 360° video streaming.

4.1 Impact of Selective Encryption under ABE

In this section, we apply frame-based selective encryption to tiled 360° videos using ABE under four different strategies. We evaluate each strategy in terms of the level of content unviewability it achieves and the associated computational trade-offs. For simplicity, our ABE policy used a single attribute.

Encryption Scheme: We evaluate four tile-based selective encryption strategies for 360° video streaming. Two of these employ viewport-aware, 2-dimensional selective encryption, where we encrypt frames in a tile segment based on the portion of viewport it covers: **(a) Major-allP:** All *I*- and *P*-frames are encrypted for the tile that predominantly covers the viewport (e.g., tile 5 in Fig. 4), while only *I*-frames are encrypted for the remaining minor tiles in the viewport. **(b) Major-allI:** Only *I*-frames from the major tile are encrypted, and all minor tiles in a viewport remain unencrypted.

To generate encrypted viewport videos, we identify the four tiles present in the viewport for each segment. Based on whether a tile is classified as major or minor, we apply the corresponding encryption level. These tiles are combined and cropped to form a viewport-specific clip for each segment. This process is repeated for all segments, and the clips are concatenated to produce the final encrypted viewport video.

The remaining strategies apply encryption independent of the viewport: **(c) Full:** All frames (*I, P, and B*) are encrypted in each tile. **(d) All I+P:** Only all *I*- and all *P*-frames are encrypted in each tile.

360-Video Dataset: Our evaluation uses four diverse 360° videos: (1) an action-packed short film, (2) a London tour with relatively still backgrounds, (3) a slower-paced documentary with frequent motion and dynamic background changes, and (4) a yoga session with minimal motion and a static indoor setting. These videos were selected to represent various genres and motion complexities.

Tiling and Encoding: Each video is divided into a 3x3 grid, resulting in nine tiles. With a resolution of 3840x1920, each tile measures 1280x640 pixels. Tiles are encoded using MP4-H264 and segmented for DASH streaming. A keyframe interval of 60 frames is used to support ABR, with additional keyframes inserted at scene changes using a scene threshold of 40.

Viewport Simulation: We assume the viewport corresponds to the size of a single tile, approximately 120 degrees of the horizontal field of view (360/3). Since real head movement data was unavailable for most videos, we manually analyzed each video to identify visually engaging areas likely to attract user attention. A viewport spans portions of four different tiles and changes to different tiles at various points in the video. The viewport portions are cropped and stitched into a continuous "viewport video" for further analysis.

VMAF Evaluation: To assess quality degradation due to selective encryption, we compare each encrypted viewport video to its unencrypted counterpart using VMAF. In our context, a lower VMAF score indicates effective quality degradation, aligning with the goal of restricting high-quality video access while minimizing computational overhead.

Figure 5a presents the mean VMAF scores with 95% confidence intervals for the four selective encryption schemes. As expected, Full encryption results in the lowest VMAF, making content unviewable. Both All I+P and Major-P also achieve strong degradation,

keeping VMAF consistently below 5. In contrast, Major-allI performs the weakest, with VMAF reaching up to 30, as it encrypts only the *I*-frames of the major tile but leaves the rest unencrypted.

An interesting exception is the Yoga video, where VMAF remains near zero across all schemes, likely due to its static content and minimal motion. These findings suggest that encryption strategies can be content-aware, adapting frame selection based on video complexity to balance degradation and computational overhead.

Encryption & Decryption time overhead: Figure 5b shows the total encryption and decryption time for each scheme. Full encryption incurs the highest runtime, followed by All I+P, which reduces the cost to half by excluding *B*-frames. Major-allP cuts runtime even further—about half of All I+P—while still keeping VMAF below 5. Major-allI is the most efficient, with runtimes of only a few seconds, but it delivers the weakest degradation.

Total size increase overhead: Figure 5c, reports the size overhead introduced by each scheme. As with runtime, the Full scheme results in the largest overhead, most notably over 35% for the yoga video. All I+P keeps overhead moderate, under 10% only, while Major-allP limits it to around 5% by restricting encryption to select tiles and frames. Major-allI introduces negligible overhead, typically under 1%, aligning with its minimal computational cost.

4.2 360° Video Streaming Setup & Metrics

4.2.1 Setup **Videos:** For our streaming experiments, we used Google Spotlight Stories: HELP, the same action video referenced in Sect. 4.1. The 360-degree video, originally in 3840×1920 (4K) resolution, was divided into a 3×3 tile grid, resulting in nine tiles. Each tile was encoded using the H.264/AVC codec [13], packaged in MP4, and segmented into MPEG-DASH compliant .m4s files.

To enable adaptive bitrate streaming, we generated four quality levels per tile, each with distinct resolutions and bitrates (Tab. 1). Streams were encoded with a 60-frame keyframe interval, with additional keyframes inserted at scene changes (threshold = 40). We prepared two dataset versions with 2-second and 4-second segments for comparative evaluation. The 293-second video yields 147 segments for the 2-second version and 74 segments for the 4-second version per quality level. To create a more diverse dataset, we created four symbolic links for each tiled video stream, effectively simulating five distinct video copies.

A key reason for selecting this video was the availability of real head movement data from 48 viewers [30]. We parsed each user's head trace to extract the four tiles forming the viewport per segment, designating the tile covering the largest area as the Major tile and the rest of them as Minor tiles.

We averaged viewport coverage across users for both 2-second and 4-second segments to generate per-segment tile selections. As described in Sect. 3.3, these were encoded into the MPD, with each segment referencing four tiles—one major and three minor—based on real viewing data. The 2-second segmentation yielded finer viewport accuracy due to less averaging loss. We created 40 unique MPD files from 40 user head traces and evenly distributed five simulated video copies (via symbolic links) across them. To reflect realistic content popularity, clients requested videos according to a Zipf distribution ($s = 1.5$), ensuring more frequent access to certain videos—mirroring real-world streaming patterns.

Table 1: Bitrates and resolutions per video tile.

Stream	Bitrate (Mbps)	Resolution
1	0.5	480x240
2	1	640x320
3	2	960x480
4	3	1280x640

ABE: As part of our evaluation, we applied ABE encryption to all tiled video segments across all quality levels, using a single-attribute policy for simplicity. While prior work [23] shows that decryption time grows linearly with the number of attributes, its overall impact on system performance is minimal. Following encryption, we updated stream bitrates to account for overhead and revised the MPD files with the new bitrates and encrypted segment URLs. We evaluated two selective encryption schemes: (1) Major-P, which applies viewport-aware, two-dimensional encryption, and (2) All-I+P, which encrypts frames regardless of viewport coverage (see Sect. 4.1 for details).

We conducted our experiments on the **CloudLab** testbed [10] using two deployment scales.

(a) Small-Scale Setup: This setup consists of seven nodes: one origin server, one cache, and five client nodes. Each client node runs six DASH clients, totaling 30 clients, with each client using a unique headtrace-based MPD file. Bandwidth limits were configured using Linux's tc utility [9]. With each client streaming four tiles at up to 3 Mbps per tile (12 Mbps total), each node requires up to 72 Mbps, resulting in 360 Mbps across all clients. To avoid client-side contention, we allocated 72 Mbps between the cache and each client node. To stress the origin-cache link, we limited its bandwidth to 120 Mbps—one-third of the total client bandwidth.

(b) Large-Scale Hierarchical Setup: This setup spans eight nodes arranged in a two-level caching hierarchy. One node hosts the origin server, another acts as the level-1 (L1) cache, and two level-2 (L2) caches connect to the L1 cache. Each L2 cache connects to two client nodes, with each client node running 20 clients, totaling 80 clients. Since we had 40 unique headtraces, each was reused once. Bandwidth limits were enforced similar to the small-scale setup, between each L2 cache and its client nodes was capped at 240 Mbps (480 Mbps total). To simulate contention, we limited the L1-to-L2 links to 160 Mbps each (one-third of downstream capacity). The link between the origin and L1 cache was set to 320 Mbps, matching the combined bandwidth from L1 to both L2 caches.

Apache HTTP Server & Apache Traffic Server: We use the Apache HTTP Server [4] on the origin server and Apache Traffic Server (ATS) [5] as the caching proxy. Both were configured with TLS/SSL to support HTTPS streaming. In the small-scale experiment, cache sizes were varied across seven configurations: 0 MB (disabled), 100 MB, 250 MB, 500 MB, 1000 MB, 1500 MB, and 2000 MB. In the large-scale hierarchical setup, a minimum cache size of 10 MB was used, as 0 MB caching is not supported by ATS.

DASH Clients: For client-side streaming, we use AStream-360 (see Sect.3.3), which includes ABE decryption for HTTP-ABE experiments. For HTTPS, the same client is used without decryption logic. Client session start times are scheduled using a Poisson distribution[20]. In the small-scale setup, we set $\lambda = 20$, resulting in an average interval of 20 seconds between client starts. For the
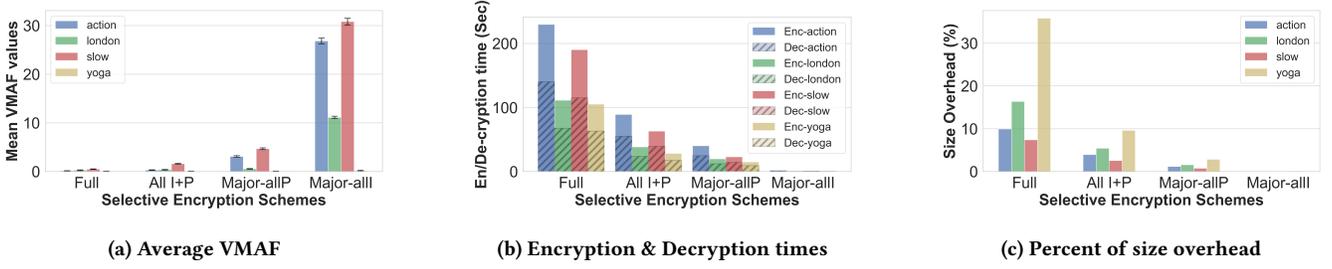
(a) Average VMAF

(b) Encryption & Decryption times

(c) Percent of size overhead

Figure 5: Impact of different selective encryption schemes under ABE.

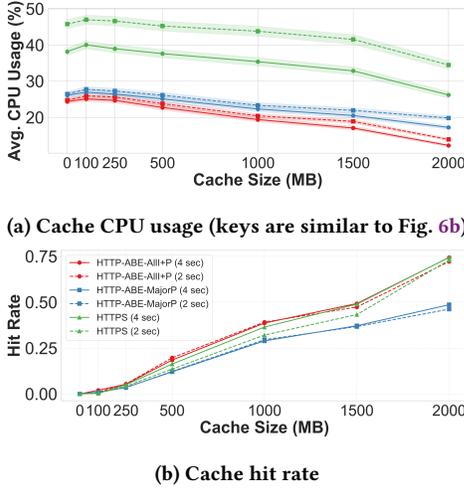(a) Cache CPU usage (keys are similar to Fig. 6b)

(b) Cache hit rate

Figure 6: Comparison of cache performance metrics for the small-scale experiment.

large-scale setup, we increase the rate to $\lambda = 10$ to reduce total experiment time, yielding an average interval of 10 seconds.

4.2.2 Evaluation Metrics, **CPU Load:** To compare the computational overhead of HTTP-ABE and HTTPS, we record CPU usage on both the cache and origin servers. We use Linux's pidstat tool [8] with a one-second refresh interval to monitor CPU utilization in real time throughout each experiment.

Hit Rate: We assess cache efficiency by comparing hit rates under HTTP-ABE and HTTPS across various cache sizes. Hit rates are computed using ATS's monitoring logs, which report the number of cache hits and misses during streaming sessions.

QoE-1: VMAF (Video Quality): To evaluate perceived video quality, we use VMAF [16], a metric developed by Netflix that combines human visual models with machine learning. VMAF scores range from 0 to 100, with higher scores indicating better visual quality. For each client, we compare the streamed quality of a tile segment to the corresponding tile segment from the highest-quality stream. We compute the average VMAF per client and then aggregate scores across all clients to assess overall video quality performance across cache sizes.

QoE-2: Rebuffering: Rebuffering measures playback interruptions caused by an empty buffer. We extract rebuffering durations from the AStream DASH client logs, which record segment-wise stall times. These metrics are aggregated per client and subsequently averaged across all clients to assess the impact of caching and protocol selection (HTTP-ABE vs. HTTPS) on playback smoothness.

4.3 HTTP-ABE vs HTTPS 360° video streaming

To evaluate system performance, we analyzed cache and origin server CPU load, cache hit rate, midgress traffic, and client QoE. Two experiments were conducted: a small-scale setup and a large-scale hierarchical multi-cache setup, as described in Sect. 4.2.1.

The small-scale setup included thirteen runs: one with 0 MB cache and two each for 100 MB, 250 MB, 500 MB, 1000 MB, 1500 MB, and 2000 MB. In the large-scale setup, where 0 MB was unsupported, we used 10 MB as the minimum, resulting in fourteen runs. Cache sizes were adjusted simultaneously across all three caches (L1 and both L2s). Each experiment included a warm-up run to populate the cache (no metrics collected), followed immediately by a second run that records metrics for comparison.

4.3.1 Small-Scale Experiment In this experiment, we evaluate the performance of our two ABE-based selective encryption schemes, MajorP and All I+P, under HTTP and compare them with HTTPS-based streaming. The evaluation is conducted using two segment durations: 2-second and 4-second video segments, to study the impact of segment granularity on system performance.

CPU Load on Cache: In Fig. 6a, we plot the average CPU usage throughout the streaming session, along with a 95% confidence interval, for various cache sizes. The results show that both HTTP-ABE approaches consume significantly less CPU resources than HTTPS, with HTTP-ABE-allI+P always outperforming Major-P.

A notable performance gap appears when comparing 4-second and 2-second segment lengths. With HTTPS, the CPU load increases substantially for 2-second segments, as the cache must perform TLS termination, decryption, and re-encryption for twice the number of segments compared to the 4-second setup. In contrast, both ABE-based approaches exhibit relatively stable CPU usage, with only a slight increase when moving to shorter segments—thanks to the absence of TLS termination. *This trend suggests that as segment duration decreases, HTTP-ABE becomes increasingly more efficient than HTTPS in terms of CPU usage.*

Specifically, for 4-second segments, HTTP-ABE-allI+P uses 36% less CPU at 0MB cache and up to 53% less at 2000MB, compared to HTTPS. For 2-second segments, the savings increase, ranging from 45% (at 0MB) up to 60% (at 2000MB). The HTTP-Major-P scheme shows similar trends, using 31% up to 34% less CPU than HTTPS for 4-second segments, and 42% up to 47% less for 2-second segments, as cache size increases from 0MB up to 2000MB.

Comparison of Hit Rate: Figure 6b compares the hit rates of the cache for different sizes for HTTP-ABE and HTTPS. An interesting observation is that the ABE-allI+P scheme achieves higher hit rates than ABE-MajorP for cache sizes 500MB onwards. This is attributed
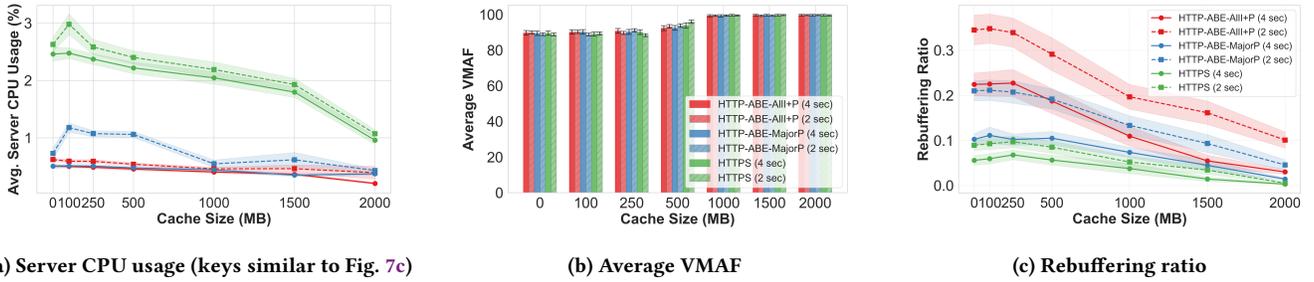
(a) Server CPU usage (keys similar to Fig. 7c)   **(b) Average VMAF**   **(c) Rebuffering ratio**

Figure 7: Comparison of origin server CPU usage and client side QoE metrics for small-scale experiment.

to the nature of two-dimensional selective encryption in MajorP: the same tile segment may be stored in the cache encrypted with only I-frames (as a minor tile in one request), but later requested with both I and P-frames encrypted (as a major tile in another request), resulting in a cache miss due to mismatch in encryption level. This fragmentation leads to lower cache reuse. In contrast, the ABE-allI+P scheme applies a uniform encryption level across all tiles, improving cacheability and hit rate consistency.

This also explains the higher CPU usage for ABE-MajorP compared to ABE-allI+P, as more cache misses require more frequent segment retrieval from the origin and additional writes to the cache.

When comparing segment lengths, both ABE schemes show stable hit rates between 2-second and 4-second segments. However, HTTPS experiences a noticeable drop in hit rates for 2-second segments, particularly at cache sizes between 500 MB and 1500 MB, where ABE-allI+P outperforms HTTPS. *Cache logs reveal that the ABE approach benefits from a higher number of Read While Write (RWW) hits, meaning the cache can immediately begin serving a segment to other clients while it is still being written.* In contrast, HTTPS incurs additional delays due to TLS termination, decryption, and re-encryption—even for cached segments.

This behavior becomes more pronounced with shorter segments, where client requests are more frequent, increasing the overhead of TLS operations in HTTPS. In contrast, HTTP-ABE avoids this overhead entirely, resulting in better hit rates and improved efficiency, especially with 2-second segmentation.

Origin Server: Figure 7a shows the origin server's CPU load with a 95% confidence interval. Across all cache sizes and for both 2-second and 4-second segments, both HTTP-ABE approaches consistently use less CPU than HTTPS. This is primarily due to the overhead of TLS encryption required for each client request under HTTPS. As cache size increases, the CPU usage gap between HTTPS and HTTP-ABE narrows, since more segments are served from the cache. However, overall CPU usage remains low for all configurations, with HTTPS peaking at approximately 3%.

Comparison of VMAF: We compared the client's QoE between HTTP-ABE and HTTPS by calculating the VMAF score for each tile segment and averaging the results per client. Figure 7b presents the average VMAF across 30 clients, along with a 95% confidence interval, for various cache sizes. The results show that all three approaches deliver comparable VMAF scores, regardless of cache size or segment duration (2-second or 4-second). Additionally, we observe that average VMAF scores are lower at smaller cache sizes

and increase with larger caches, as more segments are served directly from the cache. This reduces requests to the origin server and avoids potential congestion.

Comparison of Rebuffering ratio: Figure 7c shows the average rebuffering ratio over the duration of the entire video for 30 clients, with a confidence interval of 95% between varying cache sizes. We observe that HTTP-ABE-allI+P experiences significantly higher rebuffering compared to ABE-MajorP and HTTPS, particularly when using 2-second segments. The HTTPS approach consistently results in the lowest rebuffering, while MajorP performs similarly to HTTPS with 4-second segments but degrades noticeably with 2-second segments.

Rebuffering decreases for all approaches as cache size increases, since more segments are served directly from the cache. The higher rebuffering observed in HTTP-ABE approaches can be attributed to bandwidth congestion between the origin server and the cache, further exacerbated by the increased segment sizes introduced by ABE encryption overhead (as discussed in Sect. 4.1 and shown in Fig. 5c). This overhead is greater for the allI+P scheme compared to MajorP, which explains its poorer performance.

Similar findings were reported in [29], which shows that removing bandwidth constraints significantly reduces rebuffering for HTTP-ABE and HTTPS, limiting it to initial startup buffering.

*4.3.2 **Large Scale Hierarchical Experiment*** For the following evaluations, we focus on 2-second segments, which yielded more distinctive results in small-scale experiments.

CPU Load on L1 and L2 Caches: Figure 8 shows the average CPU usage for both L1 and L2 caches. Similar to the small-scale experiment, HTTPS consistently uses more CPU than both HTTP-ABE approaches. At the L1 cache, the CPU usage gap is most pronounced at smaller cache sizes. It narrows as cache size increases, eventually converging at 2000 MB, where most segments are served by the L2 caches, reducing the load on the L1 cache.

Between the two ABE schemes, ABE-allI+P continues to outperform MajorP in terms of CPU efficiency, as previously observed in the small-scale setup. The same trend applies to L2 caches: HTTPS consumes more CPU than both ABE schemes, and allI+P consistently uses less CPU than MajorP.

In case of the L1 cache, both ABE schemes consume up to 43% less CPU than HTTPS at 10 MB. At 2000 MB, the reduction drops to 17% for MajorP and 33% for allI+P. For the L2 caches, MajorP achieves CPU savings ranging from 30–57%, while allI+P performs even better, with reductions between 32–63% compared to HTTPS.

These results differ from prior findings [29], where in large-scale experiments HTTP-ABE was observed to use more cache CPU
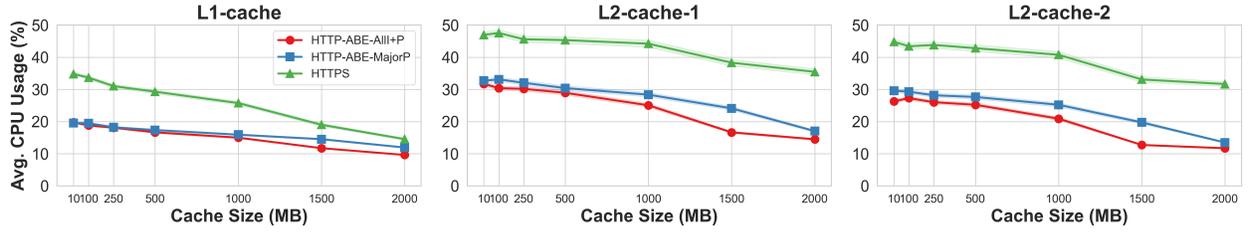
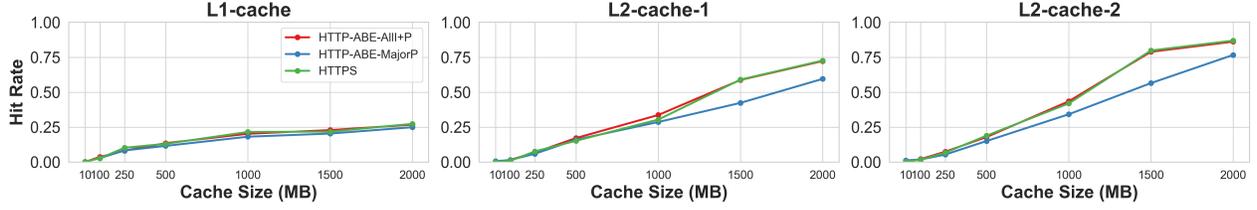Figure 8: CPU usage of L1 and L2 caches for different cache sizes.

Figure 9: Hitrate of L1 and L2 caches for different cache sizes.

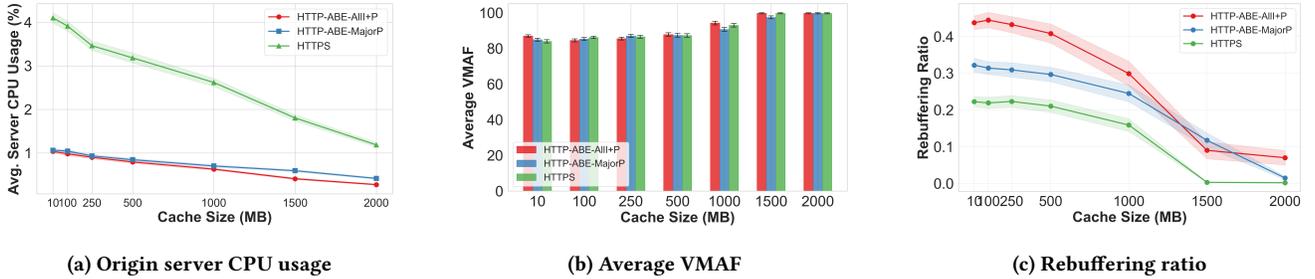(a) Origin server CPU usage   (b) Average VMAF   (c) Rebuffering ratio

Figure 10: Comparison of origin server CPU usage and client side QoE metrics for large-scale experiment.

than HTTPS at smaller cache sizes. In contrast, our results with $360°$ video streaming consistently show lower CPU usage for both ABE approaches compared to HTTPS across all cache sizes. This discrepancy can be attributed to the nature of $360°$ streaming, where four tiles are downloaded per segment, significantly increasing the load on the cache. Under HTTPS, this results in more frequent TLS termination, decryption, and re-encryption, contributing to higher CPU usage—an overhead avoided in the HTTP-ABE case.

Comparison of Hit Rate: Figure 9 shows that hit rates for both HTTP-ABE approaches and HTTPS are comparable at the L1 cache. However, at the L2 caches, only ABE-MajorP exhibits lower hit rates between 1000 MB and 2000 MB—consistent with the trend observed in the small-scale experiment.

Server CPU, VMAF and Rebuffering: Figure 10 presents the comparison of server CPU usage (10a), VMAF (10b), and rebuffering (10c) for the large-scale experiment. The observed trends closely align with those from the small-scale setup discussed in Sect. 4.3.1, and can be interpreted using the same reasoning.

5 Conclusions

This work presents a novel application of Attribute-Based Encryption (ABE) for HTTP-based $360°$ video streaming where we leverage selective frame-based encryption to secure the content rather than

the transport layer. Our approach enables efficient caching and targeted access control without the overhead of TLS.

We implemented and evaluated two ABE-based selective encryption schemes: ABE-allI+P, a uniform encryption strategy applied across all tiles, and MajorP, a viewport-aware, two-dimensional encryption scheme that prioritizes frames in the viewer's dominant field of view. We observe that HTTP-ABE significantly reduces cache CPU load up to 63% and improved hit rates in some cases, outperforming HTTPS in both efficiency and scalability. These benefits were especially pronounced in smaller segment durations, suggesting that ABE is well-suited for live streaming scenarios.

However, our evaluation also revealed important trade-offs. While both ABE schemes maintained comparable video quality (VMAF) to HTTPS, they incurred higher rebuffering.

We found MajorP offered a middle ground—with lower rebuffering than allI+P, smaller encryption and decryption overhead, and only slightly reduced content obfuscation. MajorP demonstrates the potential of viewport-adaptive encryption strategies for balancing efficiency and security. We will explore such trade-offs and their implications on live-streaming in a future work.

References

[1] Mohamed Abomhara, Omar M. Zakaria, Othman O. Khalifa, A. A. Zaidan, and B. B. Zaidan. 2022. Enhancing Selective Encryption for H.264/AVC Using Advanced Encryption Standard. *CoRR* abs/2201.03391 (2022). arXiv:2201.03391 https://arxiv.org/abs/2201.03391

[2] Vijay K. Adhikari, Yang Guo, Fang Hao, Volker Hilt, Zhi-Li Zhang, Matteo Varvello, and Moritz Steiner. 2015. Measurement Study of Netflix, Hulu, and a Tale of Three CDNs. *IEEE/ACM Transactions on Networking* 23, 6 (2015), 1984–1997. doi:10.1109/TNET.2014.2354262

[3] I. Agi and L. Gong. 1996. An empirical study of secure MPEG video transmissions. In *Proceedings of Internet Society Symposium on Network and Distributed Systems Security*. 137–144. doi:10.1109/NDSS.1996.492420

[4] Apache. 2024. *Apache HTTP Server Project*. Retrieved 2024 from https://httpd.apache.org/

[5] Apache. 2024. *Apache Traffic Server*. Retrieved 2024 from https://trafficserver.apache.org/

[6] John Bethencourt. 2011. *Using the cpabe Toolkit*. Retrieved 2024 from https://acsc.cs.utexas.edu/cpabe/tutorial.html

[7] John Bethencourt, Amit Sahai, and Brent Waters. 2007. Ciphertext-Policy Attribute-Based Encryption. In *2007 IEEE Symposium on Security and Privacy (SP '07)*. 321–334. doi:10.1109/SP.2007.11

[8] die.net. 2024. *pidstat - Report statistics for Linux tasks*. https://linux.die.net/man/1/pidstat

[9] die.net. 2024. *tc - show / manipulate traffic control settings*. https://linux.die.net/man/8/tc

[10] Dmitry Duplyakin, Robert Ricci, Aleksander Maricq, Gary Wong, Jonathon Duerig, Eric Eide, Leigh Stoller, Mike Hibler, David Johnson, Kirk Webb, Aditya Akella, Kuangching Wang, Glenn Ricart, Larry Landweber, Chip Elliott, Michael Zink, Emmanuel Cecchet, Snigdhaswin Kar, and Prabodh Mishra. 2019. The Design and Operation of CloudLab. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*. USENIX Association, Renton, WA, 1–14. https://www.usenix.org/conference/atc19/presentation/duplyakin

[11] Mario Graf, Christian Timmerer, and Christopher Mueller. 2017. Towards Bandwidth Efficient Adaptive Streaming of Omnidirectional Video over HTTP: Design, Implementation, and Evaluation. In *Proceedings of the 8th ACM on Multimedia Systems Conference* (Taipei, Taiwan) *(MMSys'17)*. Association for Computing Machinery, New York, NY, USA, 261–271. doi:10.1145/3083187.3084016

[12] Apple Inc. 2024. Apple HTTP Live Streaming. "https://developer.apple.com/streaming/".

[13] ITU-T. 2021. *H.264 : Advanced video coding for generic audiovisual services*. https://handle.itu.int/11.1002/1000/14659

[14] Parikshit Juluri, Venkatesh Tamarapalli, and Deep Medhi. 2015. SARA: Segment aware rate adaptation algorithm for dynamic adaptive streaming over HTTP. In *2015 IEEE International Conference on Communication Workshop (ICCW)*. 1765–1770. doi:10.1109/ICCW.2015.7247436

[15] Umakant Kulkarni, Yufeng Chen, Patrick Melampy, and Sonia Fahmy. 2023. Toward QoE-based Routing Path Selection. In *2023 IEEE 24th International Conference on High Performance Switching and Routing (HPSR)*. 114–119. doi:10.1109/HPSR57248.2023.10147938

[16] Zhi Li, Anne Aaron, Ioannis Katsavounidis, Anush Moorthy, and Megha Manohara. 2016. *Toward A Practical Perceptual Video Quality Metric*. Retrieved 2024 from https://netflixtechblog.com/toward-a-practical-perceptual-video-quality-metric-653f208b9652

[17] Zheng Liu and Xue Li. 2004. Motion Vector Encryption in Multimedia Streaming. In *Proceedings of the 10th International Multimedia Modelling Conference (MMM '04)*. IEEE Computer Society, USA, 64.

[18] Bruce M Maggs and Ramesh K Sitaraman. 2015. Algorithmic nuggets in content delivery. *ACM SIGCOMM Computer Communication Review* 45, 3 (2015), 52–66.

[19] Hongzi Mao, Ravi Netravali, and Mohammad Alizadeh. 2017. Neural adaptive video streaming with pensieve. In *Proceedings of the conference of the ACM special interest group on data communication*. 197–210.

[20] NumPy. 2024. *numpy.random.poisson*. Retrieved 2024 from https://numpy.org/doc/stable/reference/random/generated/numpy.random.poisson.html

[21] Marta Orduna, César Díaz, Lara Muñoz, Pablo Pérez, Ignacio Benito, and Narciso García. 2020. Video Multimethod Assessment Fusion (VMAF) on 360VR Contents. *IEEE Transactions on Consumer Electronics* 66, 1 (2020), 22–31. doi:10.1109/TCE.2019.2957987

[22] Lintian Qiao and Klara Nahrstedt. 2001. A New Algorithm for MPEG Video Encryption. (08 2001).

[23] David Reddick, Justin Presley, Frank Alex Feltus, and Susmit Shannigrahi. 2022. WiP: AABAC-Automated Attribute Based Access Control for Genomics Data. In *Proceedings of the 27th ACM on Symposium on Access Control Models and Technologies*. 217–222.

[24] Eric Rescorla. 2018. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446. doi:10.17487/RFC8446

[25] D. Salomon. 2007. *Data compression: The complete reference*. Springer.

[26] Changgui Shi and Bharat Bhargava. 1998. A Fast MPEG Video Encryption Algorithm. In *Proceedings of the Sixth ACM International Conference on Multimedia* (Bristol, United Kingdom) *(MULTIMEDIA '98)*. Association for Computing Machinery, New York, NY, USA, 81–88. doi:10.1145/290747.290758

[27] Iraj Sodagar. 2011. The mpeg-dash standard for multimedia streaming over the internet. *IEEE multimedia* 18, 4 (2011), 62–67.

[28] Kevin Spiteri, Rahul Urgaonkar, and Ramesh K Sitaraman. 2020. BOLA: Near-optimal bitrate adaptation for online videos. *IEEE/ACM Transactions on Networking* 28, 4 (2020), 1698–1711.

[29] Mohammad Waquas Usmani, Susmit Shannigrahi, and Michael Zink. 2025. Secure the Stream, Not the Hosts: Attribute-Based Encryption for DRM Enabled Video Streaming. In *Proceedings of the 16th ACM Multimedia Systems Conference* (Stellenbosch, South Africa) *(MMSys '25)*. Association for Computing Machinery, New York, NY, USA, 190–200. doi:10.1145/3712676.3714450

[30] Chenglei Wu, Zhihao Tan, Zhi Wang, and Shiqiang Yang. 2017. A Dataset for Exploring User Behaviors in VR Spherical Video Streaming. In *Proceedings of the 8th ACM on Multimedia Systems Conference* (Taipei, Taiwan) *(MMSys'17)*. Association for Computing Machinery, New York, NY, USA, 193–198. doi:10.1145/3083187.3083210

[31] Rui Xin, Shihan Lin, and Xiaowei Yang. 2023. Quantifying User Password Exposure to Third-Party CDNs. In *Passive and Active Measurement: 24th International Conference, PAM 2023, Virtual Event, March 21–23, 2023, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 652–668. doi:10.1007/978-3-031-28486-1_27

[32] Siu-Kei Au Yeung, Shuyuan Zhu, and Bing Zeng. 2009. Partial video encryption based on alternating transforms. *IEEE Signal Processing Letters* 16, 10 (2009), 893–896.

[33] Alireza Zare, Alireza Aminlou, Miska M. Hannuksela, and Moncef Gabbouj. 2016. HEVC-compliant Tile-based Streaming of Panoramic Video for Virtual Reality Applications. In *Proceedings of the 24th ACM International Conference on Multimedia* (Amsterdam, The Netherlands) *(MM '16)*. Association for Computing Machinery, New York, NY, USA, 601–605. doi:10.1145/2964284.2967292

[34] Ali Zeynali, Mohammad H. Hajiesmaili, and Ramesh K. Sitaraman. 2024. BOLA360: Near-optimal View and Bitrate Adaptation for 360-degree Video Streaming. In *Proceedings of the 15th ACM Multimedia Systems Conference* (Bari, Italy) *(MMSys '24)*. Association for Computing Machinery, New York, NY, USA, 12–22. doi:10.1145/3625468.3647607

[35] Michael Zink, Ramesh Sitaraman, and Klara Nahrstedt. 2019. Scalable 360° Video Stream Delivery: Challenges, Solutions, and Opportunities. *Proc. IEEE* 107, 4 (2019), 639–650. doi:10.1109/JPROC.2019.2894817