Handling Practical Video Analytics in a Distributed Cloud Environment

Sahil Jethva[1]
[1]Department of Computer and Engineering, CMR University, Bangalore.
Dr S Saravana Kumar[2]
[2]Professor and Head, Dept. of CSE, CMR University, Bangalore.
Dr Rubini P[3]
[3]Associate Professor, Dept. of CSE, CMR University, Bangalore.
[1]sahil.jethva2295@gmail.com, [2]saravanakumarmithun@gmail.com

Abstract

On the rise of distributed computing technologies, video big data analytics in the cloud have attracted researchers and practitioners' attention. The current technology and market trends demand an efficient framework for video big data analytics. However, the current work is too limited to provide an architecture on video big data analytics in the cloud, including managing and analysing video big data, the challenges, and opportunities. This study proposes a service-oriented layered reference architecture for intelligent video big data analytics in the cloud. Finally, we identify and articulate several open research issues and challenges, which have been raised by the deployment of big data technologies in the cloud for video big data analytics. This paper provides the research studies and technologies advancing video analyses in the era of big data and cloud computing. This is the first study that presents the generalized view of the video big data analytics in the cloud to the best of our knowledge. Video analytics has become an indispensable technology for various industries, including surveillance, healthcare, retail, and transportation. The increasing adoption of video-based applications has led to the generation of massive video data, creating challenges in terms of processing, storage, and analysis. To address these challenges, the use of a distributed cloud environment has gained popularity due to its scalability and resource flexibility. This research paper explores the practical implementation of video analytics in a distributed cloud, covering key aspects such as data pre-processing, distributed storage, real-time processing, and model deployment. Through this study, we aim to provide valuable insights into effectively harnessing the power of distributed cloud computing for handling video analytics applications.

Introduction

Overview of video analytics and its applications

Video analytics is a burgeoning field that involves the automated analysis of video content to extract meaningful insights, patterns, and information. It encompasses a range of techniques and technologies aimed at processing and interpreting video data, allowing for enhanced decision-making, security, and operational efficiency across various industries.

The applications of video analytics are diverse and continue to expand, transforming the way we interact with and understand visual data. In the realm of surveillance and security, video analytics play a crucial role in identifying suspicious activities, tracking individuals, and providing real-time alerts for potential threats. This is particularly valuable for public spaces, critical infrastructure, and retail environments.

The advent of deep learning and artificial intelligence has further accelerated the capabilities of video analytics, allowing for the identification of objects, faces, emotions, and actions within videos with remarkable accuracy. This has paved the way for applications in fields like marketing, where video content can be personalized based on viewer preferences and demographics.

In essence, video analytics is transforming the way we interpret visual data, enabling us to glean valuable insights from the ever-expanding pool of video content. Its applications span across diverse sectors, offering improvements in security, efficiency, and decision-making processes, ultimately enhancing our interaction with the visual world.

Advantages of using a distributed cloud environment for video analytics

Utilizing a distributed cloud environment for video analytics offers a multitude of advantages that significantly enhance the efficiency, scalability, and capabilities of video processing and analysis. Some of the key benefits include:

1) Scalability: Distributed cloud environments can dynamically scale resources up or down based on the demand. This is crucial for video analytics, where processing requirements can vary greatly depending on factors like data volume and complexity. This scalability ensures optimal performance during peak usage times and prevents over-provisioning during quieter periods.

2) Processing Power: Video analytics tasks, such as object detection, tracking, and recognition, require substantial processing power. Distributed cloud environments allow for the parallel processing of video data across multiple servers or nodes. This not only accelerates the analysis process but also supports the execution of complex algorithms and machine learning models.

3) Real-time Analysis: The distribution of processing tasks enables real-time video analysis, which is essential for applications like surveillance, security, and real-time decision-making. By leveraging the combined computational capabilities of multiple nodes, distributed cloud setups can achieve faster processing and response times.

4) Geographical Distribution: Distributed cloud environments consist of multiple data centre's located in different geographic regions. This enables data to be processed closer to its source, minimizing latency and ensuring timely analysis. For video analytics applications that require low-latency responses, this is a critical advantage.

5) Fault Tolerance and Reliability: Distributed cloud setups inherently offer greater resilience. If one node or data centre experiences an outage or failure, processing can be seamlessly shifted to other available nodes, ensuring uninterrupted service and preventing data loss.

6) Cost Efficiency: Distributed cloud environments often follow a pay-as-you-go model, allowing organizations to only pay for the resources they use. This cost-efficient approach eliminates the need for upfront infrastructure investments and ensures that resources are allocated where they are needed most.

Distributed Cloud Architecture for Video Analytics

The rapid evolution of technology has paved the way for sophisticated video analytics applications that demand powerful computational resources, real-time processing capabilities, and scalable infrastructure. To address these requirements, a distributed cloud architecture emerges as a compelling solution, revolutionizing the landscape of video analytics. This architecture harnesses the potential of distributed computing and cloud technology to facilitate efficient and effective video data analysis.

Introduction to Distributed Cloud Architecture:

Distributed cloud architecture combines the advantages of cloud computing and distributed systems to create a dynamic and robust infrastructure for video analytics. In this setup, computational tasks are distributed across multiple interconnected nodes or data centres, working collaboratively to process and analyse video data. Each node contributes its processing power, memory, and storage, collectively forming a scalable and flexible network capable of handling the demands of complex video analytics applications.

Benefits of Distributed Cloud Architecture for Video Analytics:

- Enhanced Scalability: Accommodate changing workloads and data volumes seamlessly.
- Accelerated Processing: Parallel processing improves analysis speed and responsiveness.
- Real-time Insights: Enables timely decision-making through distributed real-time analysis.
- Reliability and Redundancy: Replicated data ensures fault tolerance and data integrity.

- Cost Savings: Pay-as-you-go model optimizes resource usage and reduces infrastructure costs.
- Global Accessibility: Enables collaborative efforts and remote monitoring across locations.

Scalability and resource management

In the era of data-driven decision-making, video analytics has emerged as a crucial tool across industries, enabling organizations to extract valuable insights from video data. To address the growing demands of processing and analysing massive volumes of video content, the integration of distributed cloud environments has become indispensable. This paper explores the strategic importance of scalability and resource management in handling practical video analytics within a distributed cloud setup.

1) Scalability in Distributed Cloud Environments:
   Scalability is pivotal in accommodating variable workloads and ensuring optimal performance in video analytics applications. Distributed cloud architecture allows organizations to scale resources horizontally by adding more nodes or vertically by allocating additional resources to existing nodes. This flexibility enables the architecture to handle fluctuations in video data volume and processing requirements seamlessly.

2) Efficient Resource Management:
   Efficiently managing resources is critical to ensuring optimal performance and cost-effectiveness in distributed cloud environments for video analytics.

3) Data Management and Latency Considerations:
   Managing the influx of video data and minimizing latency are paramount for effective video analytics in distributed cloud setups.

4) Hybrid Cloud Implementations:
   A hybrid cloud approach combines on-premises infrastructure with cloud resources, optimizing resource allocation for video analytics.

Data Pre-processing for Video Analytics

Data pre-processing is a critical phase that involves transforming raw video data into a format suitable for analysis. In distributed cloud environments, the complexities of pre-processing are compounded by the need for scalability and real-time processing.

1) Frame Extraction and Sampling: Due to the voluminous nature of video data, efficient frame extraction and sampling techniques are essential. This involves selecting representative frames for analysis, optimizing computational resources while retaining relevant information.

2) Resolution Normalization: Videos may have varying resolutions, impacting analysis accuracy. Normalizing resolution during pre-processing ensures uniformity, enabling consistent application of algorithms across the distributed architecture.

3) Data Cleansing: Removing noise, artefacts, and redundant frames enhances the quality of data for analysis. Parallelizing data cleansing tasks across nodes improves processing speed.

Efficient Data Ingestion

Data ingestion involves seamlessly bringing video data into the distributed cloud environment for analysis. Ensuring a smooth and optimized ingestion process is pivotal for video analytics.

1) Streaming vs. Batch Ingestion: Depending on the application, organizations may opt for streaming (real-time data ingestion) or batch (periodic data transfer) methods. Streaming is essential for applications requiring instant insights, such as surveillance, while batch ingestion suits applications with less time-sensitive analysis requirements.

2) Data Partitioning: Distributing video data across nodes for parallel processing requires efficient data partitioning. Strategies like range-based partitioning or hash-based partitioning ensure balanced workloads and optimal resource utilization.

3) Data Compression: Compressing video data before ingestion reduces storage and bandwidth requirements. However, it's essential to strike a balance between compression ratios and analysis accuracy.

Real-time Video Processing

The proliferation of video data in various domains, coupled with the demand for instant insights and interactivity, has spurred the development of real-time video processing technologies. Real-time video processing involves the rapid analysis and manipulation of video data as it is generated or ingested, enabling organizations to extract valuable insights, make informed decisions, and create interactive applications in time-critical scenarios.

Importance of Real-time Video Processing:

Real-time video processing plays a pivotal role in applications where timely responses are crucial. From surveillance and security to healthcare and entertainment, real-time insights enable immediate action, enhancing operational efficiency and user experiences.

Core Components of Real-time Video Processing:

Real-time video processing comprises several essential components that work in tandem to deliver seamless and responsive outcomes:

- Capture: Video data is captured from sources such as cameras, sensors, or streaming platforms.
- Pre-processing: Raw video data is pre-processed to enhance quality, reduce noise, and normalize formats.
- Analysis: Advanced algorithms are applied to analyse video content in real-time, detecting objects, patterns, and anomalies.
- Decision-making: Derived insights fuel timely decision-making and trigger responsive actions.
- Visualization: Processed video data is presented to users through visual interfaces or integrated into applications.

System Objectives

Figure 3.1 illustrates a high-level overview of the proposed real-time and scalable video analytics system in this work. As shown in the figure, the system is designed to carry out the set of video analytics jobs requested by the system user. The video analytics is performed by the system in a distributed and scalable manner and on any number of supplied video streams. The user-provided information is used by the system to authenticate with and retrieve the intended set of video streams.

As a practical matter, the above video streams comprise the footage from surveillance cameras that reside on the user/customer premises continuously monitoring a region of interest. The video analytics provide the system user with meaningful insights into varying set of events as they occur in each of the cameras' monitored regions. The proposed system is designed to scale on-demand and perform the required set of video analytics jobs on input video streams, hence, returning the obtained analytics results to the system's user in real-time.

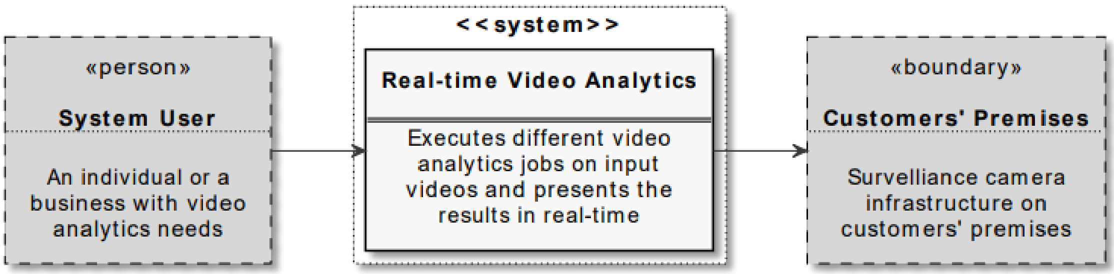

System Architecture Design

The architectural design of the proposed real-time video analytics system is carried out in this section, after due consideration of the above-stated set of requirements. Figure 3.2 illustrates at a high-level, the set of the required architectural components as identified in the preceding sections. The detailed description of each system architecture component is provided in sections that follow

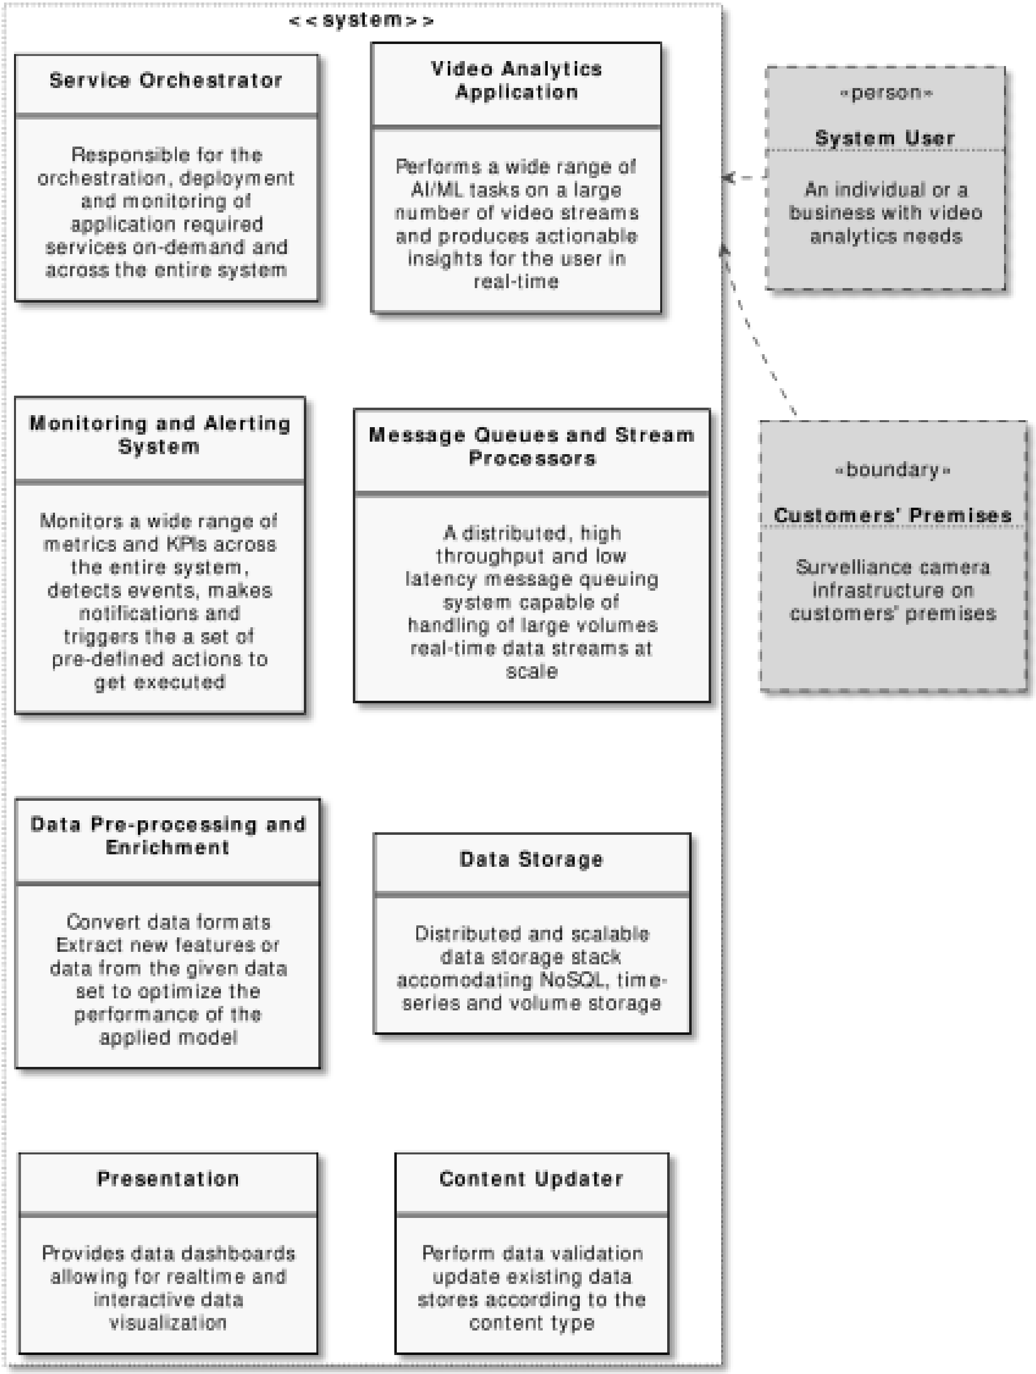

System Logical Design

The diagram represents at a broad level how the defined components in Chapter 3 will communicate and interface with each other. As illustrated in Figure 4.1, first the video analytics component of the proposed system securely connects the security cameras residing in the customer's premises and retrieves the real-time video footage captured therefrom. Therefore, the system performs the required video analytics jobs on the input video frames. The analytics job results are then forwarded to the proposed systems' Message Queues and Stream Processors component. In simple terms, the above component serves as the message broker that facilitates high-throughput and low-latency communication among the system components. Hence, the Data pre-processing and Enrichment component retrieves the published data streams in real time and performs the required data transformations or aggregations on each of the data streams. The processed data is ultimately forwarded to the Content Update component, responsible for updating the underlying Data Storage with data it receives in real-time. On the other hand, the Monitoring and Alerting system collects the events and performance metrics throughout the system and directly interfaced with the Data Storage component for long-term data persistence purposes.

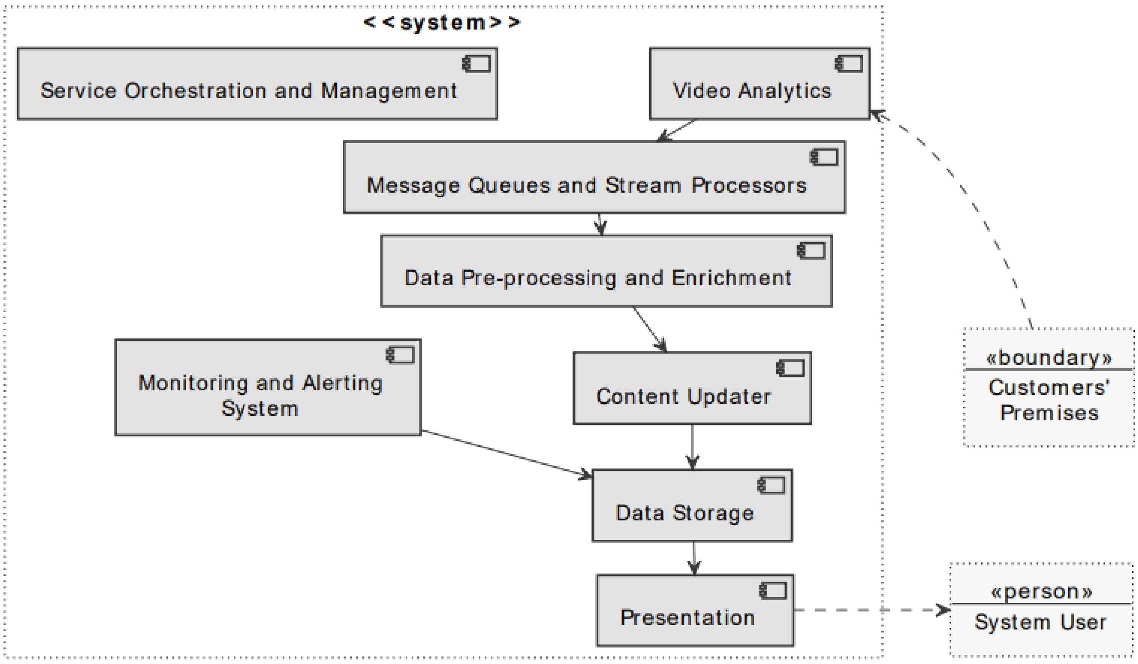

Conclusion

In the recent past, the number of public surveillance cameras has increased significantly, and an enormous amount of visual data is produced at an alarming rate. Such large-scale video data pose the characteristics of big data. Video big data offer opportunities to the video surveillance industry and permits them to gain insights in almost real-time.
This work consisted of designing, deploying, and evaluating a distributed, scalable and real-time video analytics system. This was accomplished by first deriving the functional and non-functional requirements of a typical video analytics application using the set of most commonly used requirements of users/businesses in modern video analytics applications. Therefore, complete system design with due consideration of the foregoing was devised. This paper provides an extensive study on intelligent video big data in the cloud. First, we define basic terminologies and establish the relation between video big data analytics and cloud computing. A comprehensive layered architecture has been proposed for intelligent video big data analytics in the cloud under the aaS model called L-CVAS.

Furthermore, to show the significance and recent research trends of IVA in the cloud, a broad literature review has been conducted. The research issues, opportunity, and challenges being raised by the uniqueness of the proposed L-CVAS, and the triangular relation among video big data analytics, distributed computing technologies, and cloud has been reported.

Reference

1. He Liu, Amy R Reibman, and Jacquelyn P Boerman. A cow structural model for video analytics of cow health. March 2020.
2. M. INC., “Video analytics market,” 2012, accessed: 2019-12-20. [Online]. Available: https://github.com/apache/rocketmq
3. Facebook.Com, “Facebook,” 2020, accessed: 2019-12-20. [Online]. Available: https://www.facebook.com
4. youtube.com, “Youtube statistics,” 2019, accessed: 2019-12-20. [Online]. Available: https://www.youtube.com/about/press/
5. Netflix.Com, “Netflix,” 2020, accessed: 2019-12-20. [Online]. Available: https://www.netflix.com/
6. S. Pouyanfar, Y. Yang, S.-C. Chen, M.-L. Shyu, and S. Iyengar, “Multimedia big data analytics: A survey,” ACM Computing Surveys (CSUR), vol. 51, no. 1, p. 10, 2018.
7. T. Huang, “Surveillance video: The biggest big data,” Computing Now, vol. 7, no. 2, pp. 82–91, 2014.
8. P. Mell, T. Grance et al., “The nist definition of cloud computing,” 2011.
9. J.-P. Dijcks, “Oracle: Big data for the enterprise,” Oracle white paper, p. 16, 2012.
10. J. R. Mashey, “Big data and the next wave of infras-tress,” in Computer Science Division Seminar, University of California, Berkeley, 1997.