Register No : _____

THE KAVERY ENGINEERING COLLEGE

(An Autonomous Institution)

B.E/B.TECH DEGREE END SEMESTER THEORY EXAMINATIONS, NOV/DEC 2025

Fifth Semester

Computer Science Engineering

CCS368 – Stream Processing

(Common to Artificial Intelligence and Data Science & Information Technology)

Regulations - 2021

Duration : 3 Hrs *Maximum : 100 Marks*

*PART – A (ANSWER ALL QUESTIONS) (Marks 10*2=20)*

Q. No	*Questions*	*BLT*	*CO*	*Marks*
1.	Define stream processing and its significance.	L1	1	2
2.	List the key characteristics of real-time stream processing.	L2	1	2
3.	What is a stream processing engine? Give an example.	L1	2	2
4.	Summarize the differences between batch processing and stream processing.	L2	2	2
5.	Define the concept of windowing in stream processing.	L1	3	2
6.	State the significance of event time and processing time.	L2	3	2
7.	Identify different types of stateful stream processing.	L1	4	2
8.	Differentiate between stateless and stateful stream processing.	L2	4	2
9.	Describe the role of Apache Kafka in stream processing.	L1	5	2
10.	Compare fault tolerance mechanisms in stream processing frameworks.	L2	5	2

*PART – B (ANSWER ALL QUESTIONS) (Marks 5*16 = 80)*

Q.No			*Questions*	*BLT*	*CO*	*Marks*
11.	a	i	Explain the architecture of a typical stream processing system with a suitable diagram.	L2	1	8
		ii	Describe the role of distributed stream processing frameworks.	L2	1	8
			Or			
	b	i	Discuss the advantages and challenges of real-time stream processing.	L2	1	8
		ii	Explain the CAP theorem and its relevance to stream processing.	L2	1	8
12.	a	i	Compare batch processing with stream processing, illustrating with examples.	L2	2	8
		ii	Explain the need for time semantics in stream processing.	L2	2	8
			Or			

	b	i	Discuss the role of sliding, tumbling, and session windows in stream processing.	L2	2	8
		ii	Discuss the impact of late events and how they are handled.	L2	2	8
13.	a	i	Construct a NoSQL schema for an e-commerce application.	L3	3	8
		ii	Apply the relational model to represent a real-world dataset.	L3	3	8
			Or			
	b	i	Make use of Cypher Query Language to retrieve relationships in a graph database.	L3	3	8
		ii	Develop a SPARQL query to extract semantic relationships from a dataset.	L3	3	8
14.	a	i	Analyze the role of message brokers like Apache Kafka and RabbitMQ in stream processing.	L4	4	8
		ii	Examine the advantages of using Apache Flink for low-latency stream processing.	L4	4	8
			Or			
	b	i	Compare and contrast the features of Apache Storm, Spark Streaming, and Flink.	L4	4	8
		ii	How do data partitioning and parallelism analyze and optimize stream processing performance?	L4	4	8
15.	a	i	Evaluate strategies for managing backpressure in streaming applications.	L4	5	8
		ii	Analyze how exactly-once processing guarantees ensure data reliability in stream processing.	L4	5	8
			Or			
	b	i	Analyze how stream processing integrates with cloud platforms like AWS Kinesis and Google Dataflow.	L4	5	8
		ii	Evaluate the trade-offs between performance, consistency, and scalability in stream processing.	L4	5	8