

# Performance Benchmarking

## of Message Queues:

### Apache Kafka, Apache RabbitMQ, and Redis PubSub

Student: YangXingHao

Supervisor: Prof. Anwitaman Data

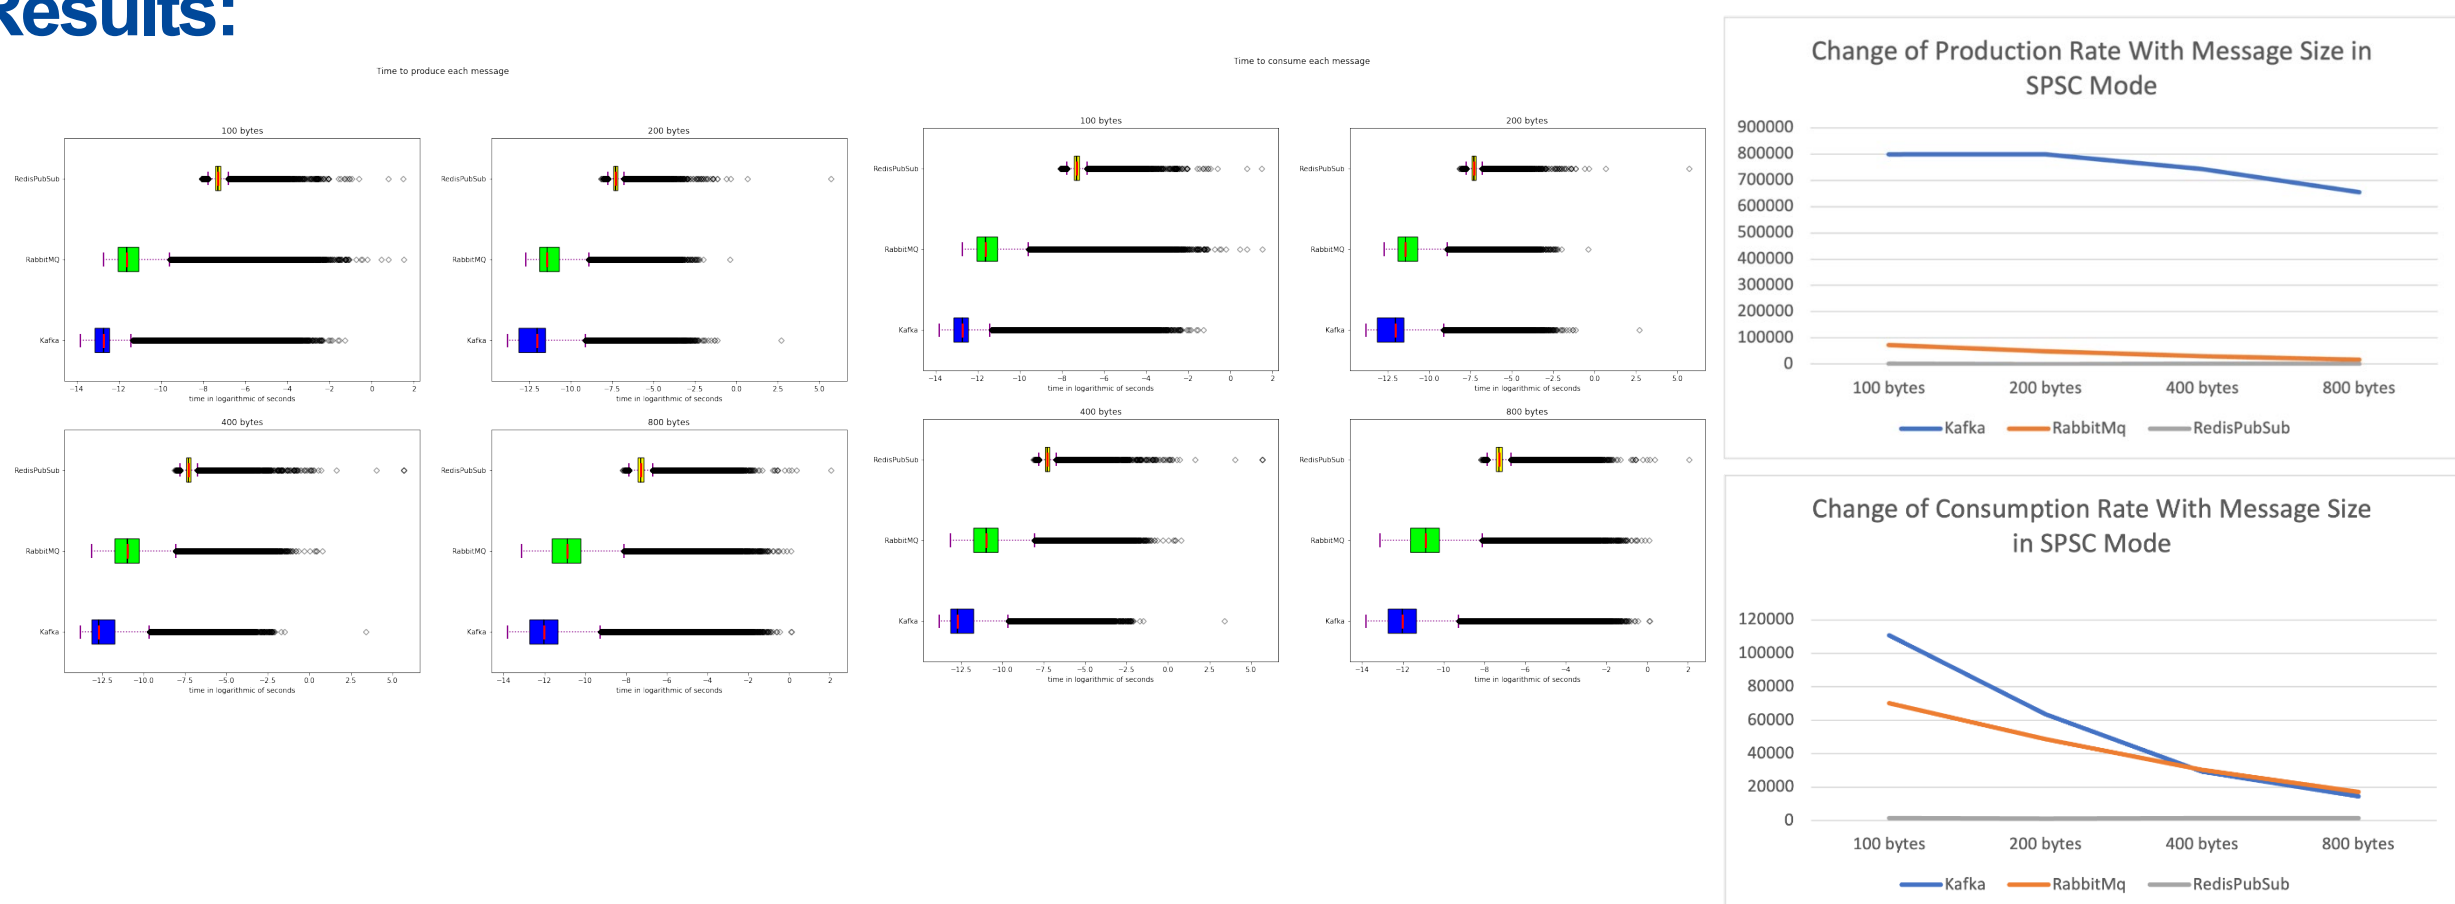
#### Objective:

Message queue (MQ) is a popular way for services to communicate asynchronously, especially in serverless and microservices architectures. There are many MQs available, each designed for different use cases, and selecting the right one depends on factors like throughput and latency. Unfortunately, existing benchmarking framework is limited to the total MQs that can be compared and cannot provide effective evidence explaining the variation in MQs' performance. Thus, we propose a new benchmarking framework that aims to address the issues of cross MQ compatibility and poor explainability of results, testing it with Apache Kafka, Apache RabbitMQ and Redis Pub/Sub.

#### Methodology:

We will be using logging to build the benchmarking framework. We will be measuring throughput and latency to understand the performance across message queues. Throughput is obtained by taking the difference between start and end of production/consumption time divided by the total number of messages published/consumed. Latency is calculated by finding the timestamp difference between each message that are logged into files. Evaluation is done by looking at the box plot and line graph of the distribution of message's production/consumption for each MQ.

#### Results:



The charts on the left is the latency and throughput comparison of the various message queues in single producer single consumer mode with varying message size. From the results, we can tell that performance wise, Kafka performs the best, followed by RabbitMQ and finally Redis PubSub. From observation, the bottle necks to the performance of the various MQs is often the rate of consumption, and its also obvious that when message size increases, there are significant rise in production/consumption time.