

NIRMA UNIVERSITY

| | |
|------------------------------|---|
| Institute: | Institute of Technology, School of Technology |
| Name of Programme: | MTech CSE (Data Science) |
| Course Code: | 6CS302CC25 |
| Course Title: | Data-Science System Design |
| Course Type: | Core |
| Year of Introduction: | 2025-26 |

| L | T | Practical Component | | | | C |
|---|---|---------------------|----|---|---|---|
| | | LPW | PW | W | S | |
| 2 | 0 | 2 | - | - | - | 3 |

Course Learning Outcomes (CLO):

At the end of the course, the students will be able to:

1. identify the appropriate infrastructure needs for the ML system (BL3)
2. apply suitable tools to implement, test and deploy the applications (BL3)
3. analyse data-science applications with respect to fault-tolerance and scalability (BL4)
4. design the high-level architecture of the ML system. (BL6)

| Unit | Contents | Teaching Hours (Total 30) |
|----------|--|------------------------------|
| Unit-I | Introduction to Data-Science (DS) Systems Design: Requirements for DS systems, Framing DS Problems, ML Project Lifecycle, DS System Architecture | 03 |
| Unit-II | Network and Distributed Systems: IP and Addressing, TCP/UDP, HTTP, Concurrency and Synchronization, Scheduling and Logging, Rate Limiter, Leader Election, Clustering, Availability, Scalability | 05 |
| Unit-III | Communication Approaches: REST API and API Gateway, GraphQL, gRPC, Websockets, Long-Polling, Server Sent Events, Message Queues, Message Brokers, Publish and Subscribe, Distributed Queues | 05 |
| Unit-IV | Data Processing and Storage: SQL and NoSQL Databases, Data Models and Data Flow, Indexing, Searching, Normalization, Replication and Sharding, Consistency models, Distributed Transactions, Data Preprocessing, Data Pipelines, Scaling Data Storage and Data Processing, CDN, Blob Storage and S3 | 06 |
| Unit-V | Performance Aspects and Resiliency: Load Balancing, Caching, Hashing, Service Discovery, Circuit Breakers, Disaster Recovery, Testing and Monitoring, Securing the System | 04 |

| | | |
|---------|--|----|
| Unit-VI | Architecture, Infrastructure and Case Studies: Microservices Architecture, Cloud Infrastructure, Serverless Computing, Virtual Machines, Containers and Orchestration Tools, Responsible ML Engineering, Case Studies: Recommendation Engine, Ad Click Prediction, Visual Search, Twitter Feed, Food Delivery, Search Ranking. | 07 |
|---------|--|----|

Self-Study:

The self-study contents will be declared at the commencement of the semester. Around 10% of the questions will be asked from self-study content.

Suggested Readings/ References:

1. Chip Huyen, Designing Machine Learning Systems: An Iterative Process for Production-Ready Applications, O'Reilly
2. Martin Kleppmann, Designing Data-Intensive Applications, O'Reilly
3. Roberto Vitillo, Understanding Distributed Systems: What every developer should know about large distributed applications, Shroff Publishers
4. Emmanuel Ameisen, Building Machine Learning Powered Applications: Going from Idea to Product, O'Reilly
5. Machine Learning in Production: From Models to Products, Christian Kästner, CMU Press
6. Valerii Babushkin and Arseny Kravchenko, Machine Learning System Design, Manning Publications.

Suggested List of Experiments:

| Sr. No. | Name of Experiments/Exercises | Hours |
|---------|--|-------|
| 1 | Design REST API based server with CRUD operations | 04 |
| 2 | Design API based server with GraphQL | 02 |
| 3 | Implement a producer-consumer problem using a message queue (RabbitMQ, Kafka) | 02 |
| 4 | Implement horizontal sharding for a basic database application | 04 |
| 5 | Set up master-slave replication between multiple databases | 04 |
| 6 | Design and implement a data pipeline to collect and preprocess raw data from multiple sources (CSV files, APIs, or databases). | 04 |
| 7 | Build a simple LRU (Least Recently Used) cache server | 02 |
| 8 | Implement a simple round-robin load balancer demonstrating how traffic is distributed across multiple servers | 02 |
| 9 | Build a fault-tolerant system by introducing retries and circuit breakers in an API based application. | 02 |
| 10 | Create a REST API to deploy a machine learning model using Flask or FastAPI. | 04 |