

Course Syllabus: CAS CS 551 Streaming and Event-driven Systems

Instructor Name: Vasiliki Kalavri

Course Time & Location:

Office Location: MCS 227F

Contact Information: vkalavri@bu.edu

Instructor Office Hours: Monday 9am-10am EST, Friday 9am-11am EST

Teaching Fellow: TBA

TF Office Hours: TBA

1. Overview

1.1 Course Description

Modern data-driven applications increasingly require continuous, low-latency processing of large-scale, rapid data events such as clicks, search queries, online interactions, financial transactions, traffic records, and sensor measurements. Distributed stream processing has become highly relevant to industry and academia due to its capabilities to both improve established data processing tasks and to facilitate novel applications with real-time requirements. In this course, you will learn how to design, implement, and evaluate scalable and reliable stream processing and event-driven applications.

Specifically, we will cover the following topics:

- Publish/Subscribe systems
- Architecture of distributed stream processing systems
- Dataflow programming
- Fault-tolerance and processing guarantees
- Streaming state management
- Windowing semantics and optimizations
- Complex event processing
- Microservice architectures
- Serverless functions and their relationship to stream processing

1.2 Course Objectives

At the end of the course, successful students will have gained skills and hands-on experience on the following methods and technology:

- Design and implementation of dataflow stream processing applications
- Message queues, log-based message brokers, and publish/subscribe systems
- Ability to comprehensively compare the features, architecture, and processing guarantees of modern streaming systems
- Implementing, deploying, and evaluating event-based applications with Apache Flink and Apache Kafka.
- Operations for scalable and reliable stream processing, including logging, monitoring, debugging, and upgrading streaming applications
- A solid understanding of the challenges and trade-offs one needs to consider when designing and deploying streaming applications

Further, students will be exposed to recent developments in stream processing research through paper assignments and presentations. The collaborative semester long project will prepare them for the practical aspects of their future careers and expose them to project management tools and software engineering best practices.

1.3 Prerequisites

CAS CS 112 and CAS CS 210; CAS CS 451 and CAS CS 460 or consent of instructor.

To be successful in this course, students will need to have strong programming skills, a solid understanding of Computer Systems fundamentals (CS 210) and some prior experience with object-oriented programming / Java (CS 211). Familiarity with Distributed Systems (CS 451) and Database Systems (CS 460) is highly recommended.

2. Instructional Format, Course Pedagogy, and Approach to Learning

2.1 Courseware

- We will use the **course website** to maintain an up-to-date class schedule: <https://vasia.github.io/dspa21/>
- We will use **Piazza** for announcements, questions, discussions, and all other communication: <https://piazza.com/bu/spring2021/cs591k1>
- We will use **Gitlab** for the hands-on sessions, the semester projects, and homework submission: <https://cs591k1-gitlab.bu.edu/>
 - Sign up for an account using your BU email.
 - Once approved, you will be able to login and create projects.

2.2 Lectures

Lectures will be held during the assigned time slots. Section 4 of the syllabus provides the topic and assigned readings for each lecture. You are expected to complete the readings before the day of the lecture and actively participate in class discussions. Lecture slides will be made available on the class website prior to the lectures or shortly after.

2.3 Discussions

Students are expected to attend the weekly discussion section they have been assigned to. The Teaching Fellows will lead the discussion sessions. The objectives are: to present material on the required tools such as “Apache Flink”, “Apache Kafka”, and “Stateful Functions”, that reinforce the concepts covered in the lectures, and answer questions (or provide clarifications) regarding the assignments and projects. The Teaching Fellow will post information to Piazza as necessary. In addition to the discussions, the Teaching Fellows will hold weekly Office Hours.

2.3.1 Software requirements

During the discussions sessions, you will solve a set of programming exercises using Apache Flink and Apache Kafka. Please use your own laptop or desktop computer and make sure to setup your environment correctly as described below.

You can develop and execute Flink applications on Linux, macOS, and Windows. However, UNIX-based setups have complete tooling support and are generally preferred by Flink developers. **All assignments assume a UNIX-based setup.** If you are a Windows user, you are advised to use Windows subsystem for Linux (WSL), Cygwin, or a Linux virtual machine to run Flink in a UNIX environment.

To setup and run Flink, you additionally need:

A Java 8 or 11 installation. To develop Flink applications and use its DataStream API in Java or Scala you will need a Java JDK. A Java JRE is not sufficient!

Apache Maven 3.x. Flink provides Maven archetypes to bootstrap new projects.

An IDE for Java and/or Scala development. Common choices are IntelliJ IDEA, Eclipse, or Netbeans with appropriate plugins installed. We recommend using IntelliJ IDEA.

Even though Apache Flink is a distributed data processing system, you will typically develop and run initial tests on your local machine. This makes development easier and simplifies cluster deployment, as you can run the exact same code in a cluster environment without making any changes.

2.4 Classroom recordings (LfA only)

All class sessions will be recorded for the benefit of registered students who are unable to attend live sessions (either in person or remotely) due to time zone differences, illness or other special circumstances. Recorded sessions will be made available to registered students **ONLY** via their password-protected BU account. Students may not share such sessions with anyone not registered in the course and may certainly not repost them in a public platform. Students have the right to opt-out of being part of the class recording. Please consult the following site for further details: <https://digital.bu.edu/lfa-classroom-recordings>.

2.5 Course Materials

There is no required textbook for this class. Slides, lecture notes, and other publicly available resources will be published on the course website and on Piazza. A list of readings is provided in the course website: <https://vasia.github.io/dspa21/readings.html>. You should be able to access all readings when connected to the campus network. Please contact the instructor if any of the listed readings is unavailable or inaccessible.

3. Assignments and Grading Criteria

3.1 Semester Project

This class is highly collaborative and research-oriented. During the first week, you will be provided with a list of semester projects and you will be asked to select your top-3 preferences. You will then be assigned to a project team with 3-5 students. During the semester, the team will be working together to deliver:

- A **design document** outlining (1) the project goals, (2) an implementation and evaluation plan, (3) the task distribution among team members.
- A **midterm project demo**. Demos will be presented (live or pre-recorded) during class time on **March 9**.
- A **final demo and poster** to be presented during the last day of class.
- The project's **gitlab repository**, including code, tests, automation and plotting scripts, and documentation.

3.2 Paper assignment

Every team will be assigned one research paper to read, review, and present. Paper assignments consist of the following deliverables:

1. A written review, to be completed and submitted individually.
2. A **collaborative** paper presentation. Every team will present a research paper in class on their assigned date. Presentations will be 30'-40' long and can be delivered live or in pre-recorded format.

3.3 Grading Scheme

Your final grade will be determined as follows:

1. Participation & effort (10%):
 - In-class discussions.
 - Piazza contributions.
 - Office Hours participation.
 - Git activity.
2. Paper assignment (20%):
 - Written review 10%.
 - Oral presentation 10%.
3. Semester project (70%) (in teams of 3-5 students):
 - Design document (maximum 3 pages) 10%.
 - Midterm demo 20%.
 - Final demo and deliverables: 40%. The final deliverables include (1) the full code implementing the project tasks as defined in the project design document, (2) auxiliary code for data pre-processing, deployment, and testing, (3) complete supporting documentation.

Individual contributions to collaborative assignments will be assessed by taking into account the following:

- The quality of individual task deliverables outlined in the project design document.
- The individual's ability to answer questions about the project during demo presentations and office hours.
- The individual's performance during the paper and demo presentations.
- The individual's contribution to the project's gitlab repository (git history).

There is no final exam at the end of the course.

4. Class and University Policies

4.1 Homework submission

All assignments and the project deliverables will be submitted via the course Gitlab. All deliverables are **due by latest 11:59pm on the day of the respective deadline**.

4.2 Attendance

Students are expected to attend each class session unless they have a valid reason for being absent. Acceptable excused absences include observing religious holidays and illness. In such cases, students are advised to contact the instructor as soon as possible, so that reasonable accommodations can be provided. Please review the **BU attendance policy** and the **BU Policy on Religious Observance** for more information.

4.3 Late work policy

Students who submit homework late will only be eligible for up to **50% of the original score**.

4.4 Academic conduct

Academic standards and the code of academic conduct are taken very seriously at our university. Please take the time to review the CAS Academic Conduct Code: <http://www.bu.edu/academics/resources/academic-conduct-code/> and the GRS Academic Conduct Code: <http://www.bu.edu/cas/students/graduate/grs-forms-policies-procedures/academic-discipline-procedures/>. Please review the sections regarding plagiarism and cheating carefully. Copies of the CAS Academic Conduct Code are also available in room CAS 105. A student suspected to violate this code will be reported to the Academic Conduct Committee, and if found culpable, the student will receive a grade of "F" for the course

All assignments must be completed individually, unless instructed otherwise. Discussion with fellow students via Piazza or in-person are encouraged, but presenting the work of another person as your own is expressly forbidden. This includes "borrowing", "stealing", copying programs/solutions or parts of them from others. Note that we may use an automated plagiarism checker. Cheating will not be tolerated under any circumstances.

Any resources, including material from other students (current or past), that are used, beyond the text or that provided by the TF or professor must be clearly acknowledged and attributed. Using such material may at the discretion of the TF or professor result in a lower grade. However, if such material is used and not acknowledged and attributed, it will automatically be considered as possible academic misconduct.

5. Accommodations

If you are a student with a disability or believe you might have a disability that requires accommodations, please contact the Office for Disability Services (ODS) at (617) 353-3658 or access@bu.edu to coordinate any reasonable accommodation requests. ODS is located at 25 Buick Street on the 3rd floor.

6. Detailed Schedule

The rest of the syllabus is tentative and might be updated during the semester. We will be keeping you informed of any changes made to the readings or assignment deadline via Piazza.

Make sure to become familiar with the **Official Semester Dates**.

Some of the critical Semester Dates are:

- The Last Day to DROP Classes (without a ‘W’ grade) March 1, 2021.
- The Last Day to DROP Classes (with a ‘W’ grade) April 2, 2021.

Date	Topic	Readings	Assignment/ Deliverable
01/25	Discussion #1: Introduction to Apache Flink	Flink quickstart guide https://ci.apache.org/projects/flink/flink-docs-release-1.12/try-flink/local_installation.html	
01/26	Introduction to stream processing	The 8 Requirements of Real-Time Stream Processing [1]	
01/28	Fundamentals of dataflow streaming systems	Beyond Analytics: The evolution of Stream Processing Systems [2]	
02/01	Discussion #2: Introduction to Apache Kafka	Kafka quickstart guide https://kafka.apache.org/quickstart	Project Selection
02/02	Stream ingestion and streaming sources	Helios: hyperscale indexing for the cloud & edge [3]	
02/04	Publish/subscribe systems	The Many Faces of Publish/Subscribe [4]	
02/08	Discussion #3: The DataStream API (part I)	https://ci.apache.org/projects/flink/flink-docs-release-1.12/dev/datastream_api.html	
02/09	Introduction to dataflow programming		
02/11	Notions of time in stream processing	Flexible Time Management in Data Stream Systems [5]	
02/16	Discussion #4: The DataStream API (part II)		
02/18	Progress tracking	Out-of-order processing: a new architecture for high-performance stream systems [6]	
02/19			Project deliverable #1: Design Document
02/22	Discussion #5: Event time		
02/23	Windowing semantics	SECRET: a model for analysis of the execution semantics of stream processing systems [7]	
02/25	Window evaluation and optimizations	Efficient Window Aggregation with General Stream Slicing [8]	
03/01	Discussion #6: Window operators	https://ci.apache.org/projects/flink/flink-docs-release-1.12/dev/stream/operators/windows.html	
03/02	Streaming state management	State Management in Apache Flink [9]	

Date	Topic	Readings	Assignment/ Deliverable
03/04	High-availability and recovery semantics	High-availability algorithms for distributed stream processing [10]	
03/08	Discussion #7: State management	https://ci.apache.org/projects/flink/flink-docs-release-1.12/dev/stream/state/state.html	
03/09			Project deliverable #2: Midterm demos
03/11	Distributed snapshots	Lightweight Asynchronous Snapshots for Distributed Dataflows [11]	
03/15	Discussion #8: Checkpoints and savepoints	https://ci.apache.org/projects/flink/flink-docs-release-1.12/dev/stream/state/checkpointing.html	
03/16	Exactly-once fault-tolerance	MillWheel: fault-tolerant stream processing at internet scale [12]	
03/18	Flow control and back pressure mechanisms	A Survey on the Evolution of Stream Processing Systems [13], Section 6.2	
03/16	Load shedding and approximate stream processing	Staying FIT: efficient load shedding techniques for distributed stream processing [14]	
03/18	Wellness Day		
03/22	Discussion #9: Reconfiguration and upgrading	https://ci.apache.org/projects/flink/flink-docs-release-1.12/ops/upgrading.html	
03/23	Elasticity and state migration	Three steps is all you need: fast, accurate, automatic scaling decisions for distributed streaming dataflows [15]	
03/25	Streaming query optimization	A catalog of stream processing optimizations [16]	
03/29	Discussion #10: Metrics and monitoring	https://ci.apache.org/projects/flink/flink-docs-release-1.12/ops/metrics.html	Paper reviews and presentations
03/30	Skew mitigation and streaming partitioning	The power of both choices: Practical load balancing for distributed stream processing engines [17]	
04/01	Stateful functions & event-driven applications	Cloudburst: Stateful Functions-as-a-Service [18]	
04/05	Discussion #11: Intro to Stateful Functions	https://ci.apache.org/projects/flink/flink-statefun-docs-release-2.2/getting-started/java_walkthrough.html	
04/06	Microservice architectures	FIRM: An Intelligent Fine-grained Resource Management Framework for SLO-Oriented Microservices [19]	
04/08	Transactional stateful functions	Fault-tolerant and transactional stateful serverless workflows [20]	
04/12	Discussion #12: Group project mentoring		

Date	Topic	Readings	Assignment/ Deliverable
04/13	Complex event processing	Processing flows of information: From data stream to complex event processing [21]	
04/15	Team paper presentation		
04/20	Team paper presentation		
04/21	Discussion #13: Group project mentoring		
04/22	Team paper presentation		
04/26	Discussion #14: Group project mentoring		
04/27	Team paper presentation		
04/29	Final project demos		Project deliverable #3: Demos and posters

- [1]: <http://cs.brown.edu/research/db/publications/8rulesSigRec.pdf>
[2]: <https://dl.acm.org/doi/abs/10.1145/3318464.3383131>
[3]: <https://dl.acm.org/doi/abs/10.14778/3415478.3415547>
[4]: <https://dl.acm.org/doi/pdf/10.1145/857076.857078>
[5]: <https://dl.acm.org/doi/10.1145/1055558.1055596>
[6]: <https://dl.acm.org/doi/10.14778/1453856.1453890>
[7]: <https://dl.acm.org/doi/10.14778/1920841.1920874>
[8]: https://openproceedings.org/2019/conf/edbt/EDBT19_paper_171.pdf
[9]: <https://dl.acm.org/doi/10.14778/3137765.3137777>
[10]: <https://ieeexplore.ieee.org/document/1410192>
[11]: <https://arxiv.org/pdf/1506.08603.pdf>
[12]: <https://dl.acm.org/doi/10.14778/2536222.2536229>
[13]: <https://arxiv.org/pdf/2008.00842.pdf>
[14]: <https://dl.acm.org/doi/10.5555/1325851.1325873>
[15]: <https://www.usenix.org/system/files/osdi18-kalavri.pdf>
[16]: <https://dl.acm.org/doi/10.1145/2528412>
[17]: <https://ieeexplore.ieee.org/document/7113279>
[18]: <http://www.vldb.org/pvldb/vol13/p2438-sreekanti.pdf>
[19]: <https://www.usenix.org/system/files/osdi20-qiu.pdf>
[20]: https://www.usenix.org/system/files/osdi20-zhang_haoran.pdf
[21]: <https://dl.acm.org/doi/10.1145/2187671.2187677>