

The slide features a background of a blue sky with white clouds. In the center, there is a green printed circuit board (PCB) with various electronic components. Overlaid on the PCB is the text "In-Memory" in a stylized font. Above the PCB, the text "CS 351" is displayed. Below the PCB, the text "Distributed Data Processing" is written in a large, bold, red font. Further down, the names of the professors and their advisors are listed.

CS 351

In-Memory

Distributed Data Processing

Prof. Panos K. Chrysanthis
Prof. Alexandros Labrinidis

Advisors:
Adam Lee, Bruce Childers, Jack Lange

Fall Term 2015 (16-1 or 2161)

Agenda for Today

- 5:30-6:00 Logistics & Intro to Stream (Panos)
- 6:00-6:45 CE-Storm - Confidential Elastic Processing of Data Streams (Alex)
- 6:45-7:15 Dinner break
- 7:15-7:45 Stream Processing on Emerging Memory Architectures (Santiago)
- 7:45-8:30 Discussion - Brainstorming

Logistics

- When : Wed 5:30 – 8:30pm (regular)
Mon 7:00 – 8:30 pm (make-up)
- Where: 6516 SENSQ
- What: *Quest Course*
 - Project-oriented investigation
 - Study state of the art-papers

No Classes & Make-up Classes

- Oct. 5 Make-up class
- Oct. 7 No class
- Oct. 26 Make-up class
- Nov. 4 No class
- Nov. 11 No class (?)
- Nov. 16 Make-up class



- Understand the state-of-the-art in in-memory data management on new memory architectures and distributed and cloud infrastructures
- Discover unsolved problems and challenges
- Learn (practice) how to give a good presentation
- Learn (practice) how to review papers
- Learn (practice) how to write a good technical paper
- Design **AstrapiDB** and produce a publishable paper

Course Requirements

- Participation: **20%**
- Presentations: **25%**
- Paper Reviews: **15%**
- Term Project & Report: **40%**

Administrative

- web page:
<http://db.cs.pitt.edu/courses/cs3551/fall2015>
– check often!
- use keyword **cs3551** in all emails to instructor (as part of the subject line)
- class mailing list: You would be signed up.

Paper Discussion Structure

1. One student will be the presenter
 - Select papers from the bibliography or come with alternative list of papers
 - Present the papers (1-2 talks)
2. A second student will lead the discussion on the Pros and provide a summary of the Pros
3. A third student lead the discussion on the Cons and provide a summary of the Cons

Project Structure

- Each group (group of 2-3 students is permitted) will select a topic to work on
 - Each student will act as a project leader for a month
 - The project leader will do a project presentation at the end of each month and the group at the end of the term
 - Select papers from the bibliography or come with alternative list of papers
 - Present the papers (1-2 talks)
 - Write a project report

Preparation of your Talk

- Reading: Read the papers but read others as well:
 - Cites and Cited, follow-ups by the same author, etc.
- Assume that the average reader has understood the easiest 2/3 of the paper.
- You, the expert on the papers, need to supply the rest.

Talk Outline

- Categorize issues and solutions in your topic
 - those that are unique to the new environments and systems
 - those that are shared with any distributed system
 - Broad-brush sketch of important results
 - give outline of talk in this context
 - Postpone discussion of things you are going to treat in detail later
 - Details of 2-3 chosen issues/solutions
 - Summary of solved problems, unsolved problems, non-problems
- A peek into your paper & project ideas

Topics & Projects

Topics	Metrics	Projects
	QoD, QoS	
?	Energy	?
	Time	
	Space	