

P&S Modern SSDs

Understanding and Designing
Modern NAND Flash-Based Solid-State Drives

Dr. Jisung Park

Prof. Onur Mutlu

ETH Zürich

Spring 2022

18 March 2022

Course Info: Who Are We? (I)



■ Onur Mutlu

- ❑ Full Professor @ ETH Zurich ITET (INFK), since September 2015
- ❑ Strecker Professor @ Carnegie Mellon University ECE/CS, 2009-2016, 2016-...
- ❑ PhD from UT-Austin, worked at Google, VMware, Microsoft Research, Intel, AMD
- ❑ <https://people.inf.ethz.ch/omutlu/>
- ❑ omutlu@gmail.com
- ❑ <https://people.inf.ethz.ch/omutlu/projects.htm>

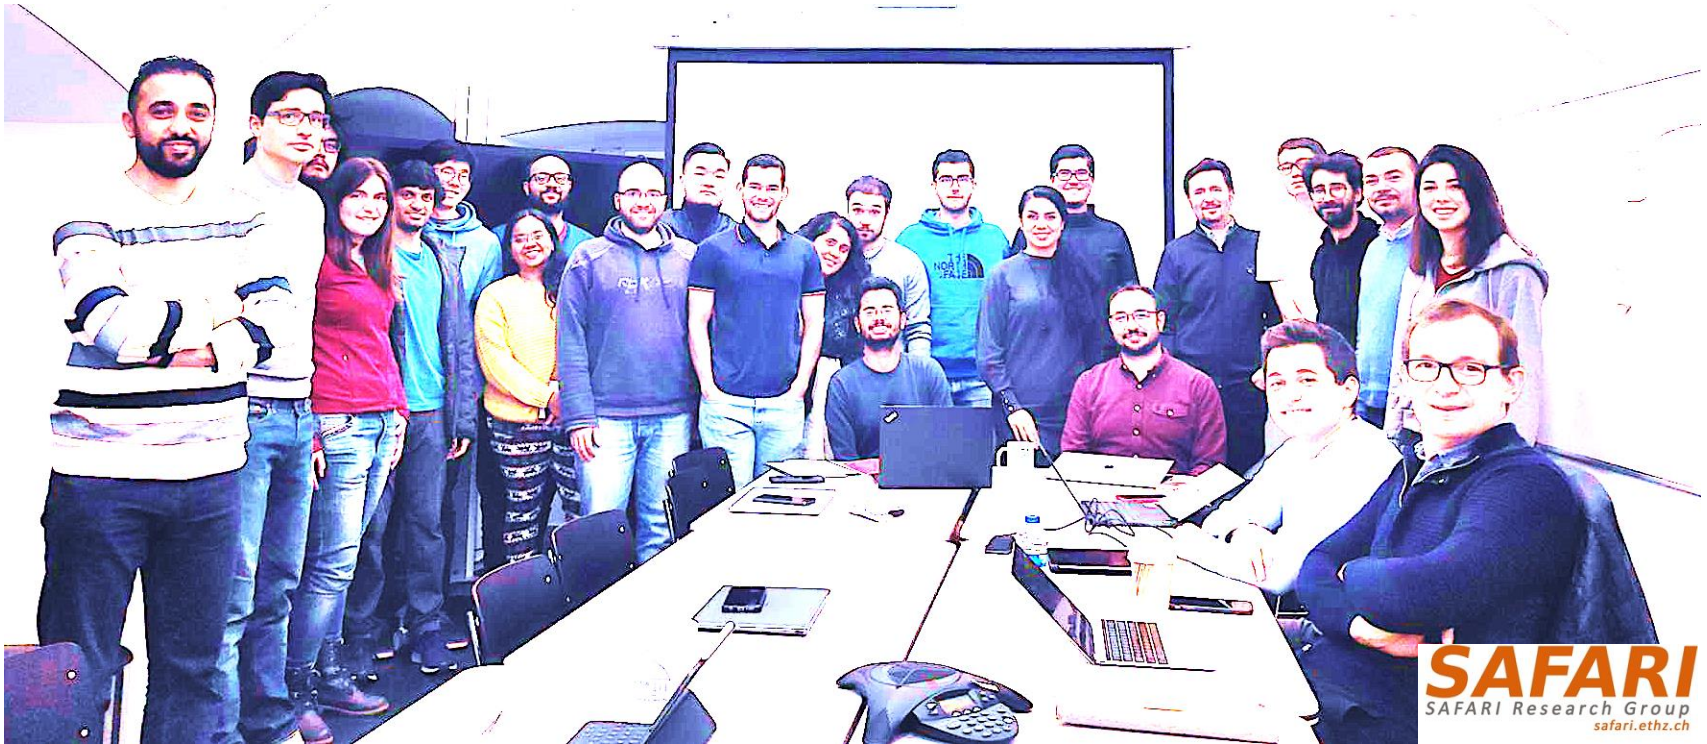
■ Research and Teaching in:

- ❑ Computer architecture, computer systems, hardware security, bioinformatics
- ❑ Memory and storage systems
- ❑ Hardware security, safety, predictability
- ❑ Fault tolerance
- ❑ Hardware/software cooperation
- ❑ Architectures for bioinformatics, health, medicine
- ❑ ...

Onur Mutlu's SAFARI Research Group

Computer architecture, HW/SW, systems, bioinformatics, security, memory

8 Postdoc, 13 PhD Students, 7 MS Students, 6 Affiliated Researchers



Think BIG, Aim HIGH!

<https://safari.ethz.ch>

SAFARI Newsletter: December 2021



Think Big, Aim High



ETH zürich

View in your browser

December 2021

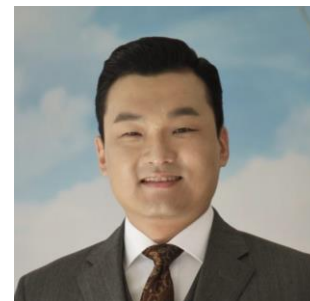


<https://safari.ethz.ch/safari-newsletter-december-2021>

Course Info: Who Are We? (II)

■ Dr. Jisung Park

- ❑ Senior researcher and lecturer
@ SAFARI research group since 2019
- ❑ PhD from Seoul National University
- ❑ Research Area: computer architecture, system software, memory/storage systems, system security, and HW/SW cooperation
- ❑ <http://jisungpark.kr/>
- ❑ jisung.park@safari.ethz.ch



Course Info: Who Are We? (III)

■ Dr. Mohammad Sadrosadati

- Senior researcher @ SAFARI research group since 2021
- Senior researcher @ IPM 2019–2021
- PhD from Sharif University of Technology
- Research Area: energy-efficient GPUs, solid-state drives, processing-in-memory, machine learning
- m.sadr89@gmail.com



■ Rakesh Nadig

- PhD Student @ SAFARI research group since 2021
- Senior staff engineer @ Samsung Electronics India 2014-2021
- MS from University of California Irvine
- Research Area: computer architecture, memory system design, multi-core architectures, near-memory processing, non-volatile memory
- rakesh.nadig@safari.ethz.ch



P&S: Modern SSDs (I)

227-0085-44L Projects & Seminars: Understanding and Designing Modern Solid-State Drives (SSDs)

Semester	Spring Semester 2022
Lecturers	J. Park
Periodicity	every semester recurring course
Language of instruction	English
Comment	Only for Electrical Engineering and Information Technology BSc. Course can only be registered for once. A repeatedly registration in a later semester is not chargeable.

Courses	Catalogue data	Performance assessment	Learning materials	Groups	Restrictions	Offered in	» Overview
-------------------------	--------------------------------	--	------------------------------------	------------------------	------------------------------	----------------------------	----------------------------

Abstract	The category of "Laboratory Courses, Projects, Seminars" includes courses and laboratories in various formats designed to impart practical knowledge and skills. Moreover, these classes encourage independent experimentation and design, allow for explorative learning and teach the methodology of project work.
----------	--

Objective	<p>NAND flash memory is the de facto standard in architecting a storage device in modern computing systems. As modern computing systems process a large amount of data at an unprecedented scale, a storage device needs to meet high requirements on storage capacity and I/O performance. A NAND flash-based SSD can provide an order(s) of magnitude higher I/O performance compared to traditional hard-disk drives (HDDs), with a much lower cost-per-bit value over any other SSDs based on emerging non-volatile memory (NVM) technologies.</p> <p>NAND flash memory has several unique characteristics, such as the erase-before write property (i.e., a flash cell needs to be first erased before programming it), limited lifetime (i.e., a cell can reliably store data for a certain number of program/erase cycles), and large operation units (e.g., a NAND flash chip reads/writes data in a page (e.g., 16 KiB) granularity). To achieve high performance and large capacity of the storage system while hiding the unique characteristics of NAND flash memory, it is critical to design efficient SSD firmware, commonly called Flash-Translation Layer (FTL). An FTL is responsible for many critical management tasks, such as address translation, garbage collection, wear-leveling, and I/O scheduling, that significantly affect the performance, reliability, and lifetime of the SSD.</p> <p>In this P&S, we will cover how a modern NAND flash-based SSD is organized and operates, from the basics of underlying NAND flash devices and various SSD-management tasks at the FTL-level. You will build a practical SSD simulator by refactoring MQSim, a state-of-the-art simulator for high-end SSDs, to support advanced features of modern NAND flash chips and essential SSD-management tasks. This will allow you to have the chance to obtain a comprehensive background of modern storage systems and research experience on system optimization with rigorous evaluation.</p> <p>Prerequisites of the course:</p> <ul style="list-style-type: none">• No prior knowledge in NAND flash-based storage systems is required.• Digital Design and Computer Architecture (or equivalent course)• Good knowledge in C/C++ programming language is required.• Interest in system optimizations
-----------	--

The course is conducted in English.

<http://www.vvz.ethz.ch/Vorlesungsverzeichnis/lerneinheit.view?lang=en&semkez=2022S&ansicht=KATALOGDATEN&lerneinheitId=157085&>

P&S: Modern SSDs (II)

NAND flash memory is the de facto standard in architecting a storage device in modern computing systems. As modern computing systems process a large amount of data at an unprecedented scale, a storage device needs to meet high requirements on storage capacity and I/O performance. A NAND flash-based SSD can provide an order(s) of magnitude higher I/O performance compared to traditional hard-disk drives (HDDs), with a much lower cost-per-bit value over any other SSDs based on emerging non-volatile memory (NVM) technologies.

NAND flash memory has several unique characteristics, such as the erase-before write property (i.e., a flash cell needs to be first erased before programming it), limited lifetime (i.e., a cell can reliably store data for a certain number of program/erase cycles), and large operation units (e.g., a NAND flash chip reads/writes data in a page (e.g., 16 KiB) granularity). To achieve high performance and large capacity of the storage system while hiding the unique characteristics of NAND flash memory, it is critical to design efficient SSD firmware, commonly called Flash-Translation Layer (FTL). An FTL is responsible for many critical management tasks, such as address translation, garbage collection, wear-leveling, and I/O scheduling, that significantly affect the performance, reliability, and lifetime of the SSD.

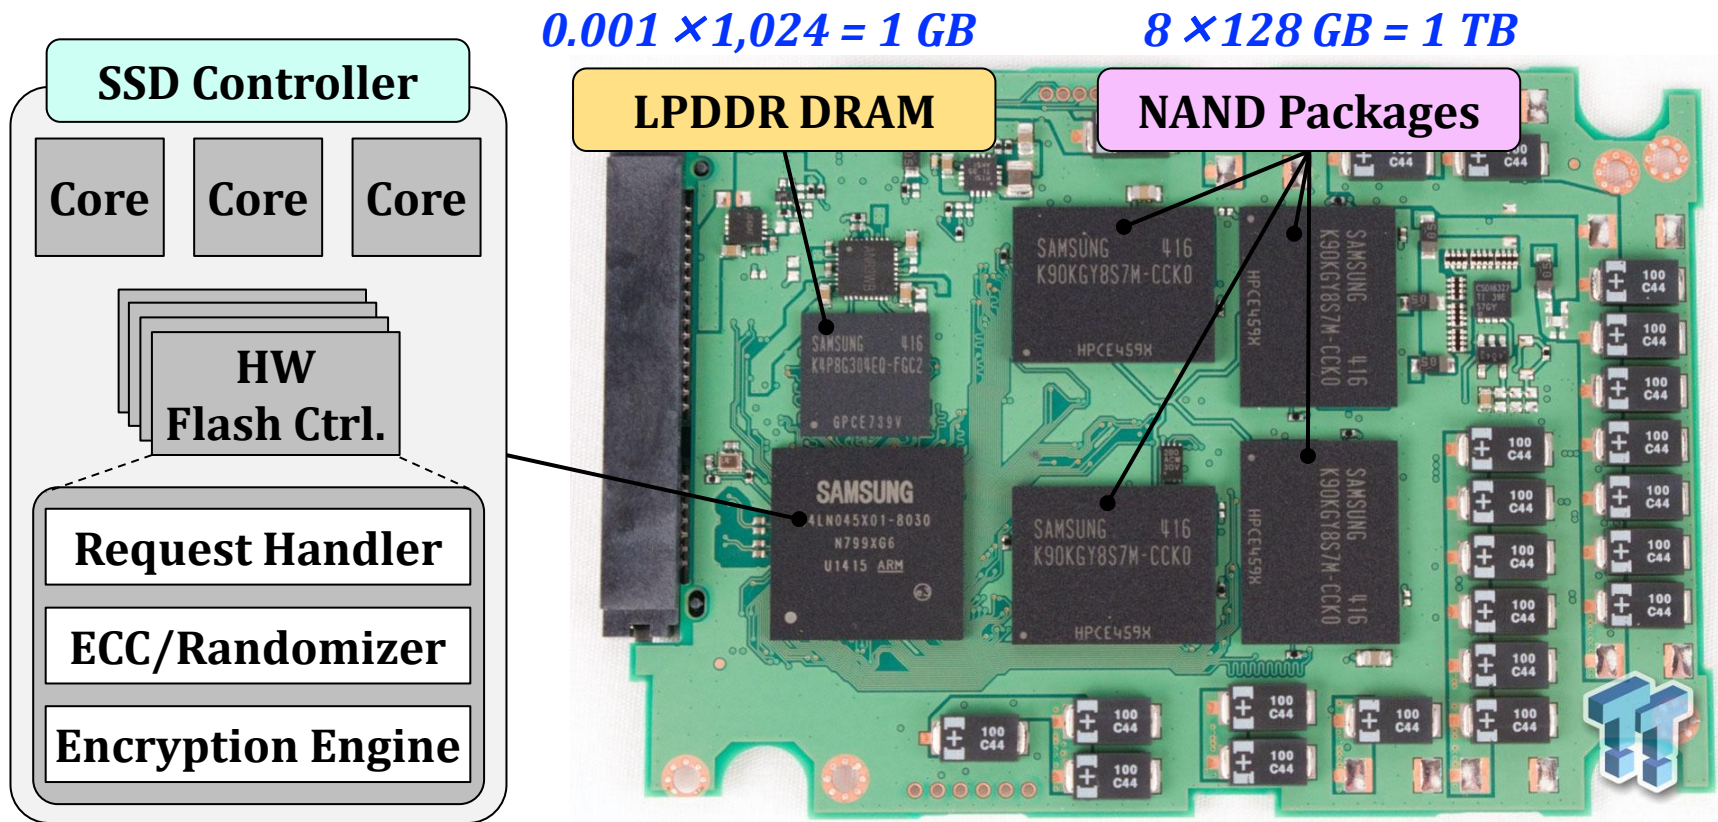
In this P&S, we will cover how a modern NAND flash-based SSD is organized and operates, from the basics of underlying NAND flash devices and various SSD-management tasks at the FTL-level. You will build a practical SSD simulator by refactoring MQSim, a state-of-the-art simulator for high-end SSDs, to support advanced features of modern NAND flash chips and essential SSD-management tasks. This will allow you to have the chance to obtain a comprehensive background of modern storage systems and research experience on system optimization with rigorous evaluation.

P&S Modern SSDs: Contents

- We will introduce **how a modern NAND flash-based SSD is organized and operates** to provide high I/O performance while hiding unique characteristics of NAND flash memory
- You will learn **fundamentals and challenges** in designing modern SSDs
- You will review existing approaches that are widely adopted in modern SSDs and will **get familiar with new research proposals**
- You will **work hands-on**: refactoring MQSim, a state-of-the-art SSD simulator

Modern SSD Architecture

- A modern SSD is a complicated system that consists of multiple cores, HW controllers, DRAM, and NAND flash memory packages



Why So Complicated?

- To provide **backward compatibility** with traditional HDDs
 - ❑ Smaller sectors than file-system blocks: 512 Bytes vs. 4KiB
 - ❑ Support overwrites
- While **hiding unique characteristics** of NAND flash memory
 - ❑ Large operation units
 - ❑ Erase-before-write property
 - ❑ Asymmetry in operation units
 - ❑ Limited endurance
 - ❑ Various error sources
 - ❑ Asymmetry in operation latencies

Unique Characteristics of NAND Flash (I)

- Large operation units
 - Read/write granularity: **page** (4 – 16 KiB)
- Erase-before-write property
 - A page needs to be first **erased before programming**
- Operation-unit asymmetry
 - Erase granularity: **block** (hundreds or thousands of pages)

**In-place update (i.e., overwrite) is very inefficient
for NAND flash memory
→ Out-of-place write & garbage collection**

Unique Characteristics of NAND Flash (II)

■ Limited endurance

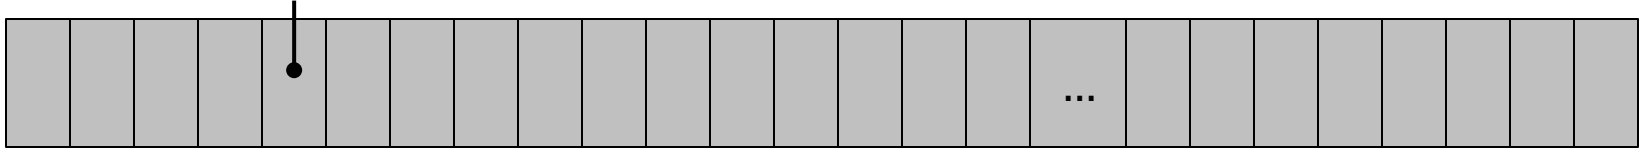
- ❑ A flash cell **cannot reliably store data** after experiencing a certain number of **program and erase (P/E) cycles**
- ❑ SLC (Single-Level Cell): > 100K P/E cycles
- ❑ MLC (Multi-Level Cell): ~ 10K P/E cycles
- ❑ TLC (Triple-Level Cell): < 3K P/E cycles
- ❑ QLC (Quad-Level Cell): < 1K P/E cycles

**Requires proper lifetime management techniques
(e.g., wear-leveling)**

Flash Translation Layer (FTL)

■ Sophisticated SSD firmware

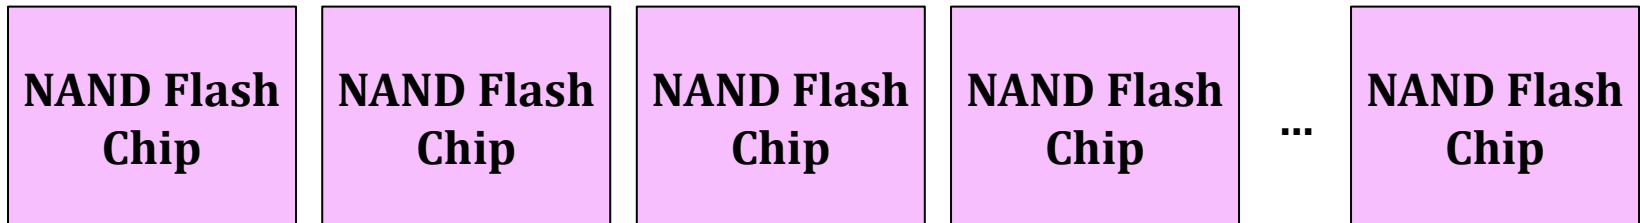
Logical Block (4 KiB, overwritable)



*Storage-device view at operating systems:
A series of blocks*

Flash Translation Layer (FTL)

- *Address translation (out-of-place write)*
- *Garbage collection*
- *Lifetime management (Wear-leveling)*
- *Reliability management (ECC, data refresh)*
- *I/O scheduling*



SSD Optimization

- Requires **comprehensive understandings** of
 - ❑ Microarchitecture of underlying NAND flash chips (HW)
 - ❑ Various internal management tasks (HW and SW)
 - ❑ OS & workload characteristics (SW)
- Optimization at one level may affect and/or be affected by the efficiency of designs at other levels.

We need an SSD simulator that accurately models various functionalities and components

What You Will Do

- Phase 1: Learning **fundamentals of NAND flash-based SSDs**
 - 3–4 background lectures on
 - Organization of NAND flash memory and NAND flash-based SSDs
 - NAND flash operations
 - SSD management tasks
 - Literature review

- Phase 2: Refactoring MQSim
 - Improving coding convention
 - Removing complicated features
 - For better readability and extendibility

Key Takeaways

- This P&S aims at improving your
 - **Knowledge** in Computer Architecture with a focus on modern storage systems
 - **Technical skills** required for good research
 - **Critical thinking and analysis**

Prerequisites of the Course

- Digital Design and Computer Architecture (or equivalent course)
- Familiarity with C++ programming (or high motivation to acquire it)
- Interest in
 - Computer architecture and systems
 - Discovering why things do or do not work
 - Solving problems
 - Designing an efficient and practical system

Course Info: How About You?

- Let us know your background, interests
- Why did you join this P&S?
- **HW0 – Student Information (Due: 24 March 2022)**

Course Requirements and Expectations

- Attendance required for all meetings
- Study the learning materials
- Each student will work on a hands-on project
- Participation
 - Ask questions, contribute thoughts/ideas

Course Website

- https://safari.ethz.ch/projects_and_seminars/spring2022/doku.php?id=modern_ssds
- Useful information about the course
- Check your email frequently for announcements

Meeting 1

■ Required Materials

- ❑ Arash Tavakkol, Juan Gomez-Luna, Mohammad Sadrosadati, Saugata Ghose, and Onur Mutlu, "[MQSim: A Framework for Enabling Realistic Studies of Modern Multi-Queue SSD Devices](#)," In USENIX FAST, 2018.
- ❑ MQSim GitHub Repository: <https://github.com/CMU-SAFARI/MQSim>

■ Recommended Materials

- ❑ Computer Architecture Fall 2020 – Lecture 26: Flash Memory and Solid-State Drives
 - <https://www.youtube.com/watch?v=rninK6KWBeM>
 - [PDF](#) and [PPT](#)
- ❑ Computer Architecture Fall 2020 – Lecture 14: Simulation (with a Focus on Memory)
 - <https://www.youtube.com/watch?v=3cI4zOoDk9Q>
 - [PDF](#) and [PPT](#)

Next Meetings

- We will meet weekly: Friday 10 AM
- Provide background lectures
- Discuss what each of you has done in the previous week
- Q&A for any difficulties in the previous week and directions for next weeks
- Presentation of your work

P&S Modern SSDs

Understanding and Designing
Modern NAND Flash-Based Solid-State Drives

Dr. Jisung Park

Prof. Onur Mutlu

ETH Zürich

Spring 2022

18 March 2022