OPTR: Order-Preserving Translation and Recovery Design for SSDs with a Standard Block Device Interface

Yun-Sheng Chang and Ren-Shuo Liu

System and Storage Design Lab
Department of Electrical Engineering
National Tsing Hua University, Taiwan

國立清華大學
NATIONAL TSING HUA UNIVERSITY

USENIX ATC '19

Solid-State Drives (SSDs)

- Inherit the **interface** and a **weak guarantee** from HDDs
  - **Permit** persisting write requests in an **arbitrary order**
- Implication to FS and DBS
  - Need to frequently **flush** SSDs to ensure order
  - At the cost of performance degradation

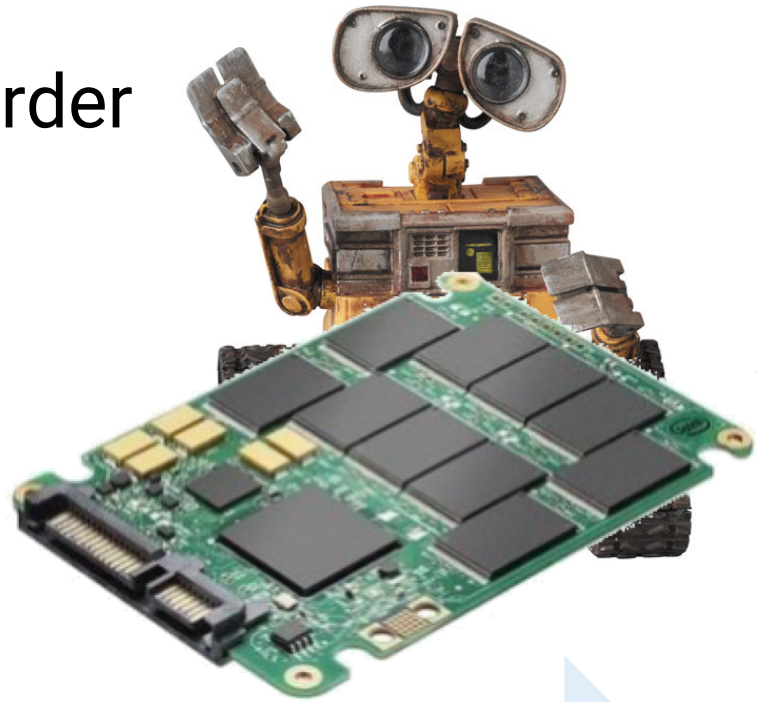

1989 1999 2009 2019

Order-Preserving SSDs (OP-SSDs)

- **Strong** request-level guarantees
  - Persist all write requests **in order**
  - Persist each write request **atomically** (a bonus)
- Invariants
  - **Identical** interface to existing software, i.e., read, write, and flush
  - **Comparable** performance to traditional SSDs

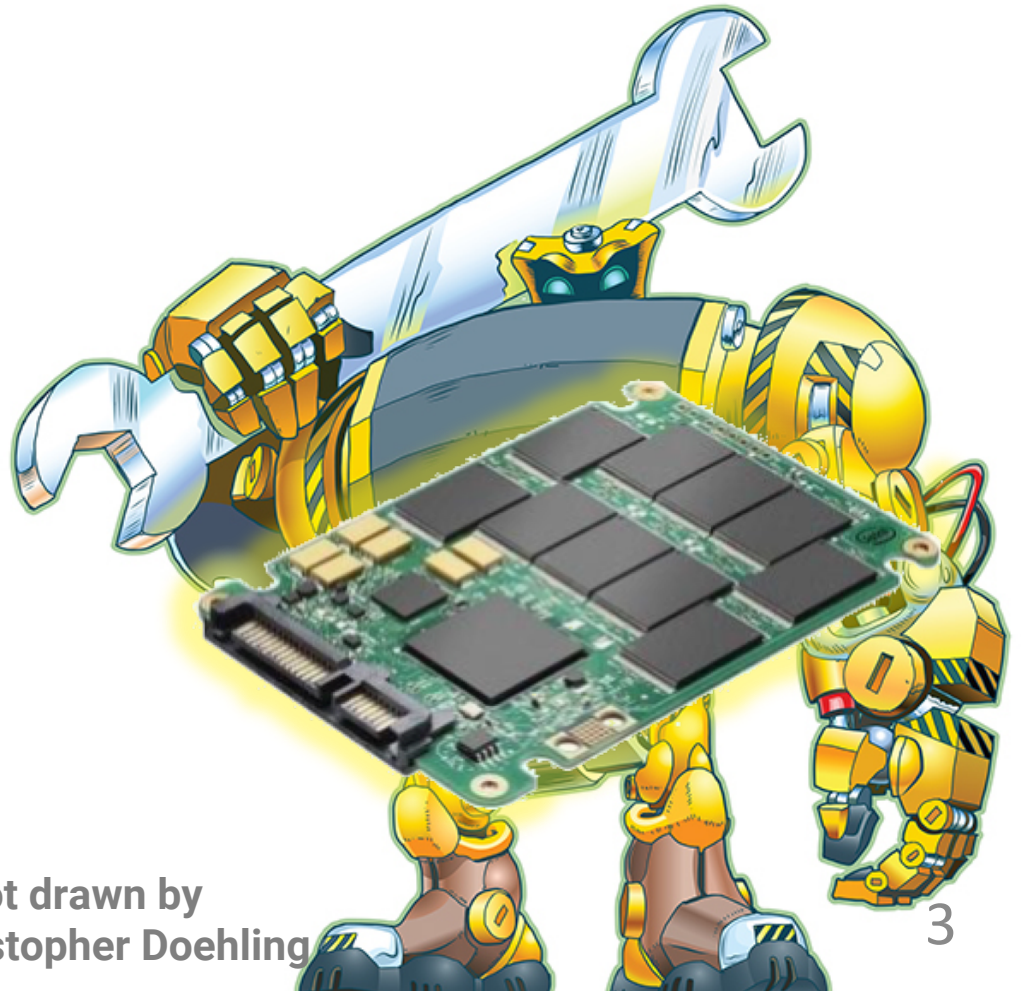

Robot drawn by Christopher Doehling

OP-SSDs in Computer Systems

- Optimize **existing FS and DBS**
  - Remove unnecessary flushes
  - Practical and manageable because OP-SSDs keep the interface intact
- Inspire **new FS and DBS**
  - Exploit the strong crash guarantees
- New **SSD firmware** research area
  - Flash-translation layers (FTLs)

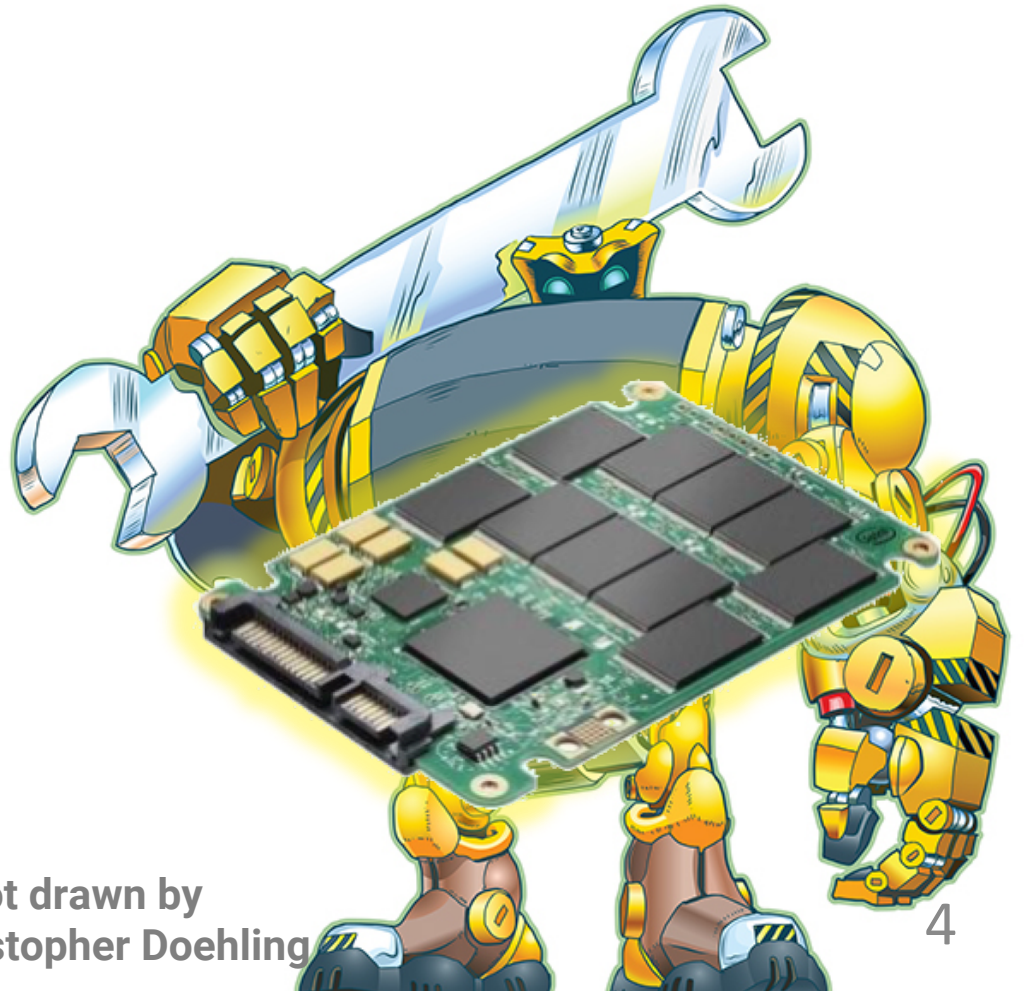

Robot drawn by Christopher Doehling

Contributions of this Work

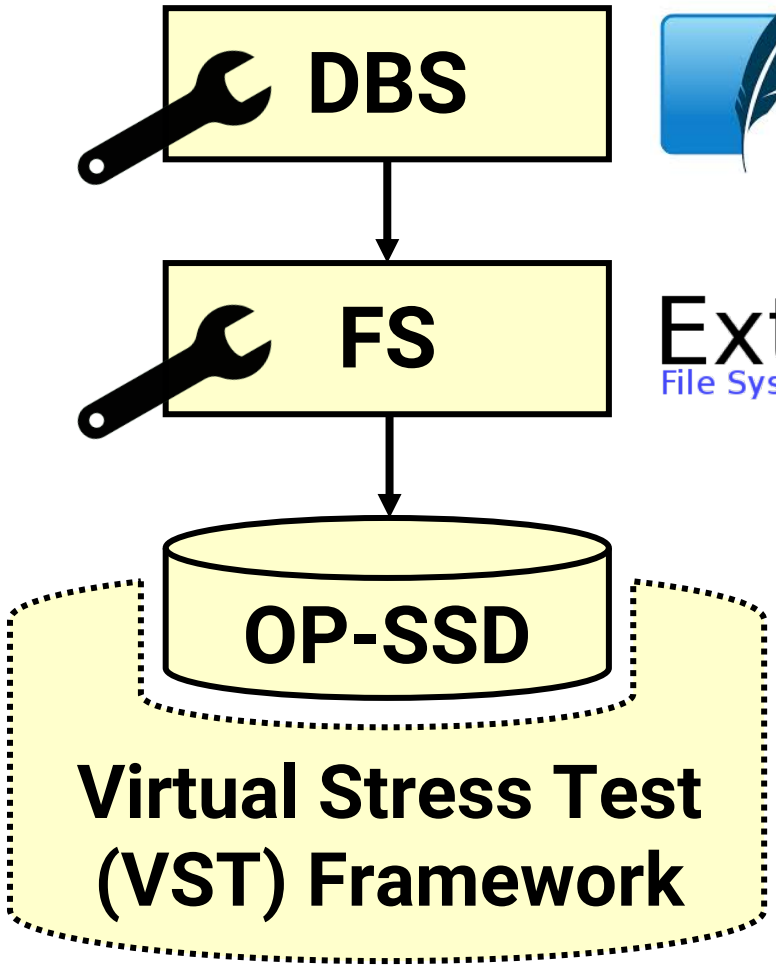

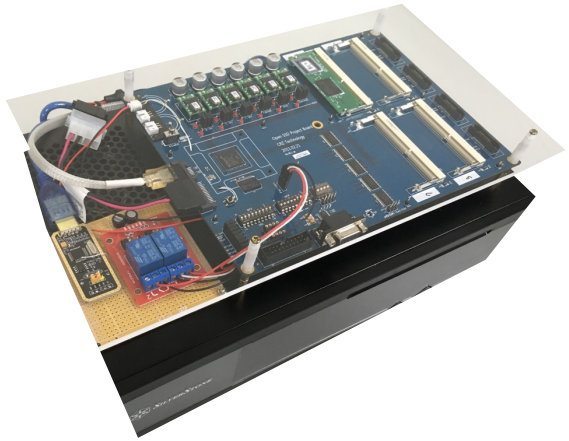

- Identify three **usage modes**

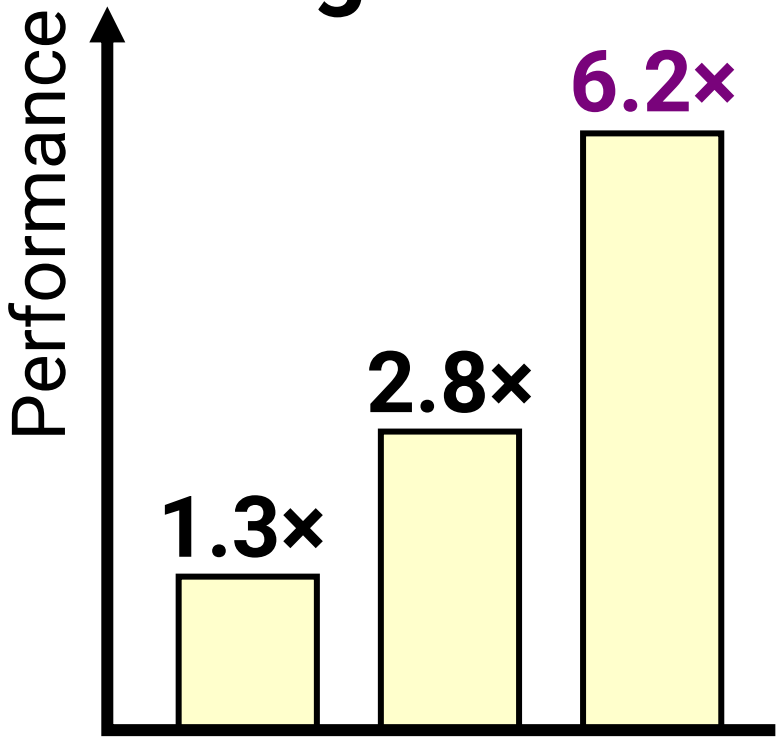

Order-Preserving SSDs

Friday 11:50 AM

Last talk, last session

Track 1 Storage Failure & Recovery

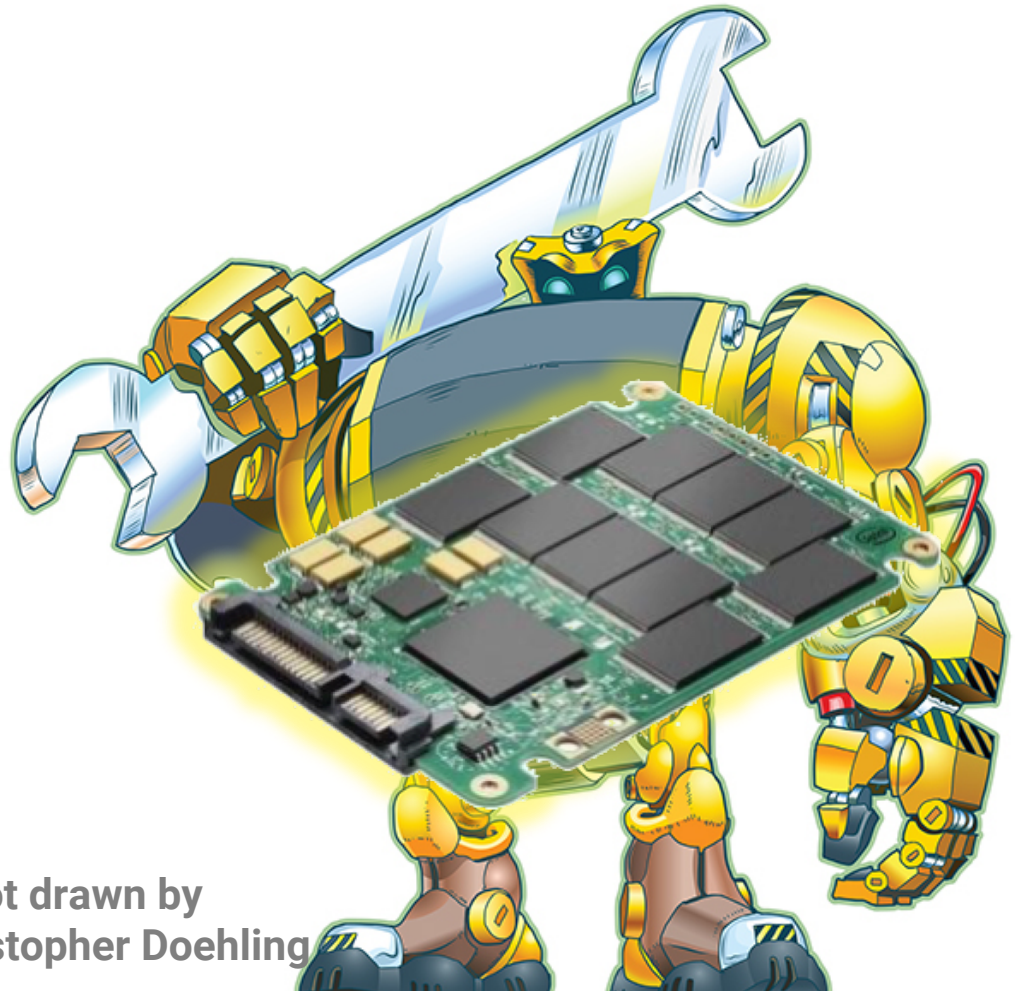

Robot drawn by
Christopher Doehling

國立清華大學
NATIONAL TSING HUA UNIVERSITY

USENIX
ATC '19