

# Crash Consistency Testing for Non-Volatile Memory Systems

Lukas Werling (Operating Systems Group)



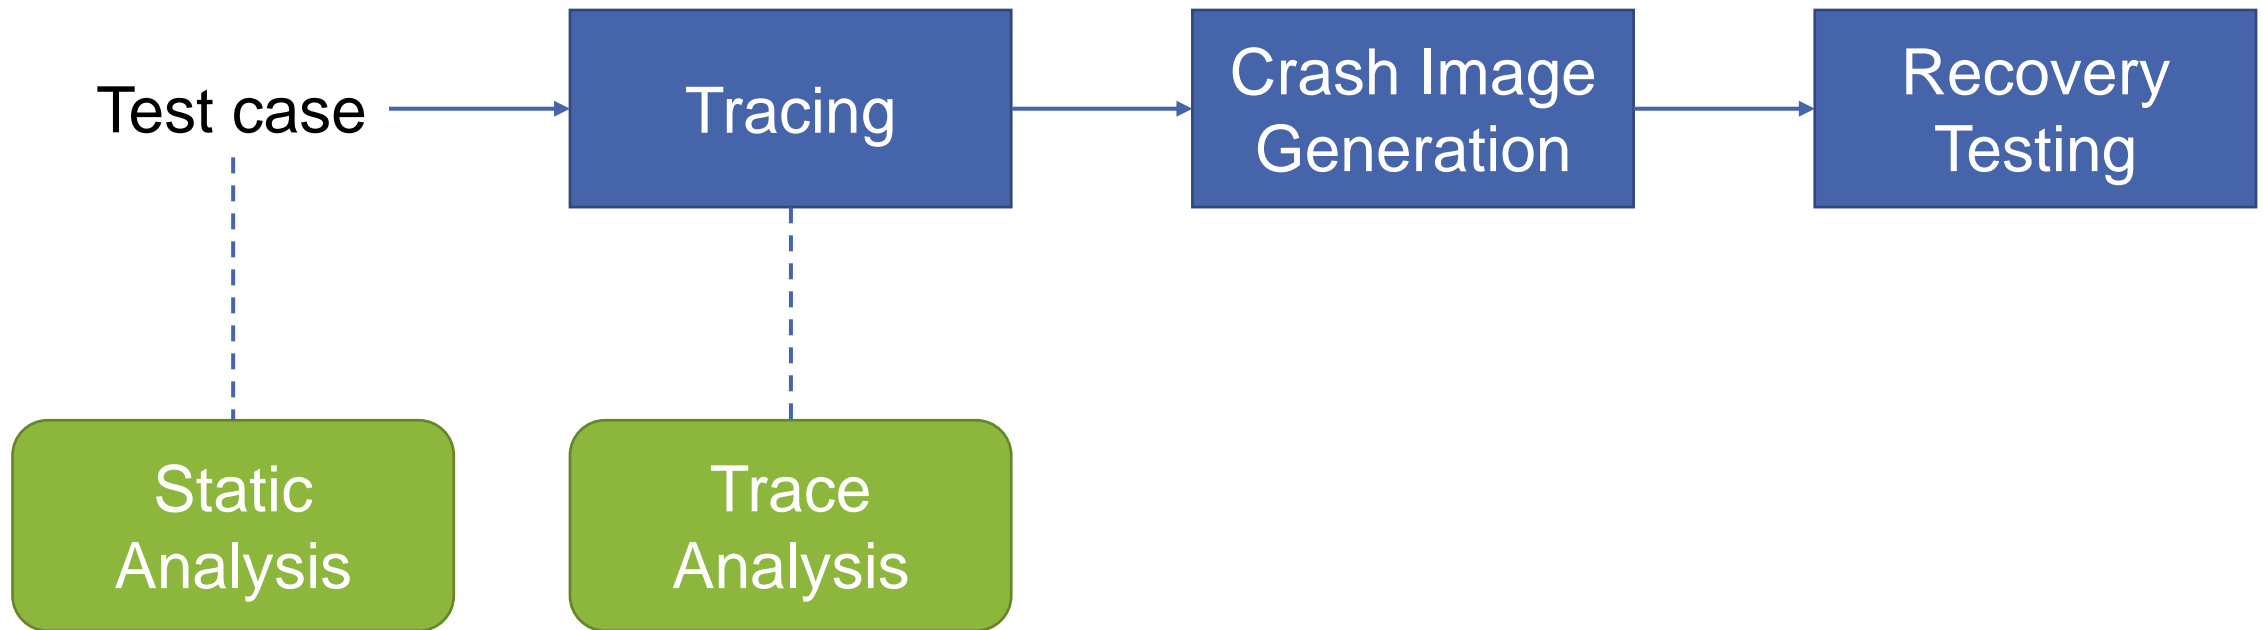
# Motivation

- Crash consistency is a common and important goal
  - Recover semantically correct state
- ... but rarely evaluated!
  - 1 of 11 papers advertising crash consistency published this year
- Vinter (ATC'22): Automatic testing of file systems
- More tools available for other use cases

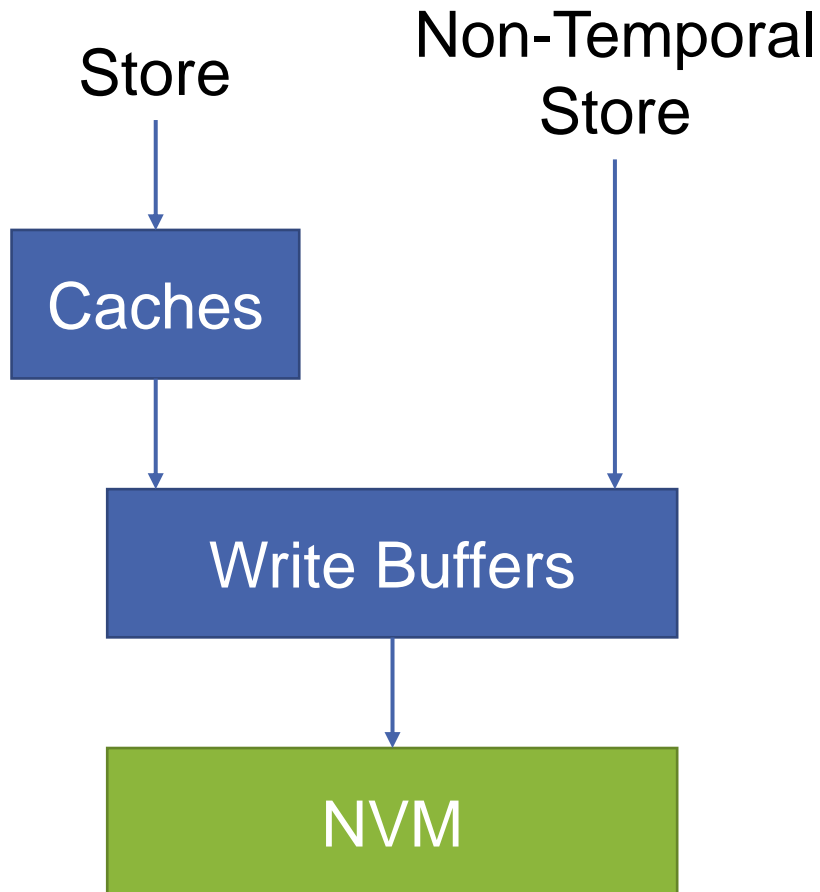


Consider testing crash consistency in **your** systems!

# Crash Consistency Testing Pipeline



# Tracing: ISA Semantics



Which store instructions?  
Weak ordering → fence

Volatile caches → cflush  
Intra cache line ordering

Volatile buffers → commit

# Tracing Methods

## Manual Annotation

- Easy to implement
- Fast
- Usually high level (e.g., library calls)
- Mistakes likely

## Binary Instrumentation

- Automatic (e.g., compiler plugin)
- Limited to user space applications
- Additional context from source code

## Virtualization

- Automatic, black box
- Captures full system (user and kernel)
- Expensive to trace arbitrary instructions (e.g., emulation)

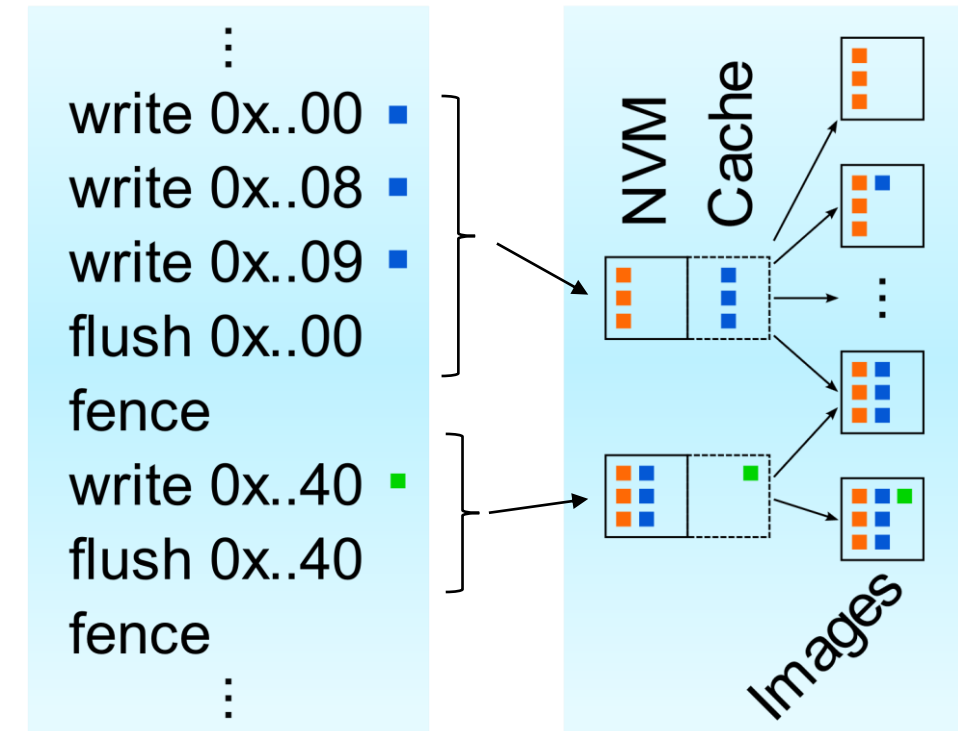
# Crash Image Generation

- Replay memory trace: NVM, cache
- Generate crash images at each fence
  - “Happy path”: every write persisted
  - Subsets of writes

Crash Image  
Generator

NOVA: Up to 512 modified cache lines

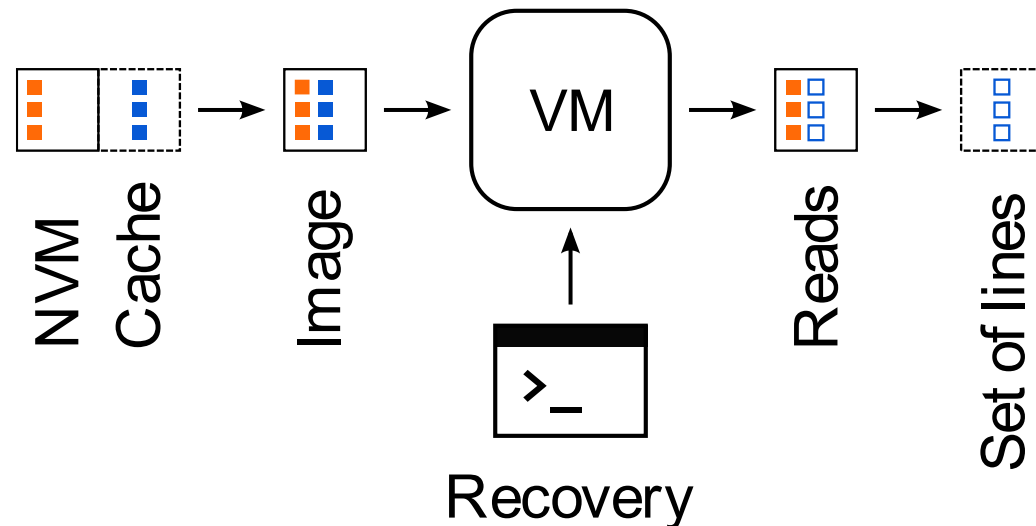
State explosion!



# Crash Image Generation: Heuristic

- Observation: Recovery ignores incomplete journal entries
- Idea: Trace NVM reads during recovery

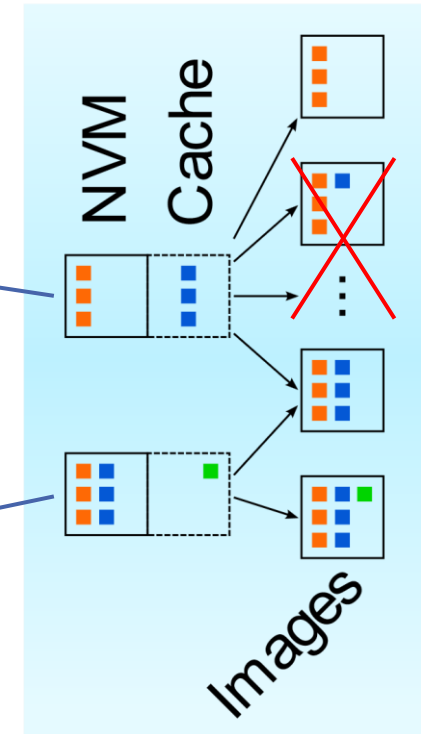
Crash Image  
Generator



## Journaling

write journal  
entry


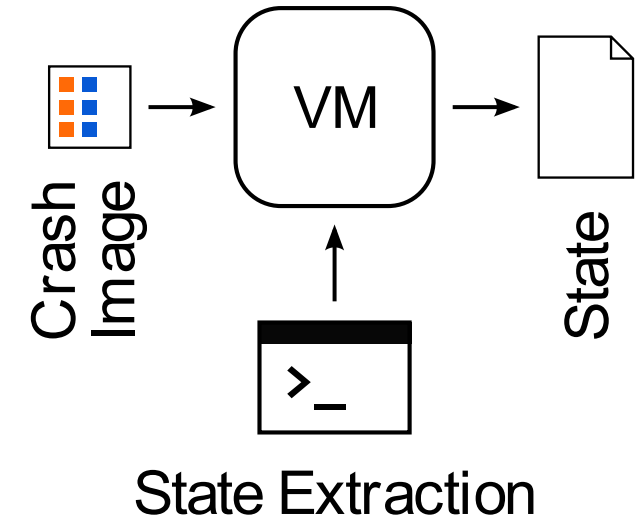
mark valid



Efficient generation of interesting crash images


# Recovery Testing

- Idea: Extract semantic state from crash images
  - File systems: serialized file listing
- Find unique states for analysis





# Pipeline Inputs



- Hand-written
- Existing test suite
- Fuzzing

- Recovery reads (Vinter)
- Data/control dependencies (Witcher<sup>1</sup>)
- Library calls (Chipmunk<sup>2</sup>)
- Ignore partial flushes

- Serialize all state
- Determine expected semantics

<sup>1</sup> SOSP'21

<sup>2</sup> EuroSys'23

# Crash Consistency Bugs in Practice

## ■ Vinter: Analysis of three NVM file systems

Missing flush in NOVA

	write	append	atime	[cm]time	chmod	chown	link	symlink
NOVA	🍎💔	✓*	✓	✓	✓	✓	⚡	🍎💔
NOVA-Fortis	🕒	🍎	✓	✓	✓	✓	🕒	🍎💔
PMFS	✓	✓	✓	✓	✓	✓	✓	✓
	mkdir rmdir	rename overwrite	rename directory	rename long name	touch	long name	unlink	update
NOVA	✓	⚡💔	⚡💔	⚡💔🕒	✓	🕒	✓	✓*
NOVA-Fortis	🕒	⚡💔🕒	⚡💔🕒	⚡💔🕒	🕒	🕒	🕒	💔
PMFS	(⚡)💣	(⚡)💣	✓	✓	✓	(⚡)	(⚡)💣	✓

🍎 data loss    💣 crash    ⚡ atomicity violation    🕒 read/write fails after recovery    🍎 multiple final states (SFS violation)

Figure 4: Crash consistency bugs discovered by VINTER.

# Bug Detail: Missing Flush in NOVA

Test command: `echo HelloWorld > /mnt/myfile`

Vinter report: 7 states, 4 final states

HelloWorld\n

HelloWor\0\0\0

HelloWorl\0\0

HelloWorld\0

```
NT-write 46784 + 0 "HelloWor"  
write    46784 + 8 'l'  
write    46784 + 9 'd'  
write    46784 + 10 '\n'
```

**No flush!**

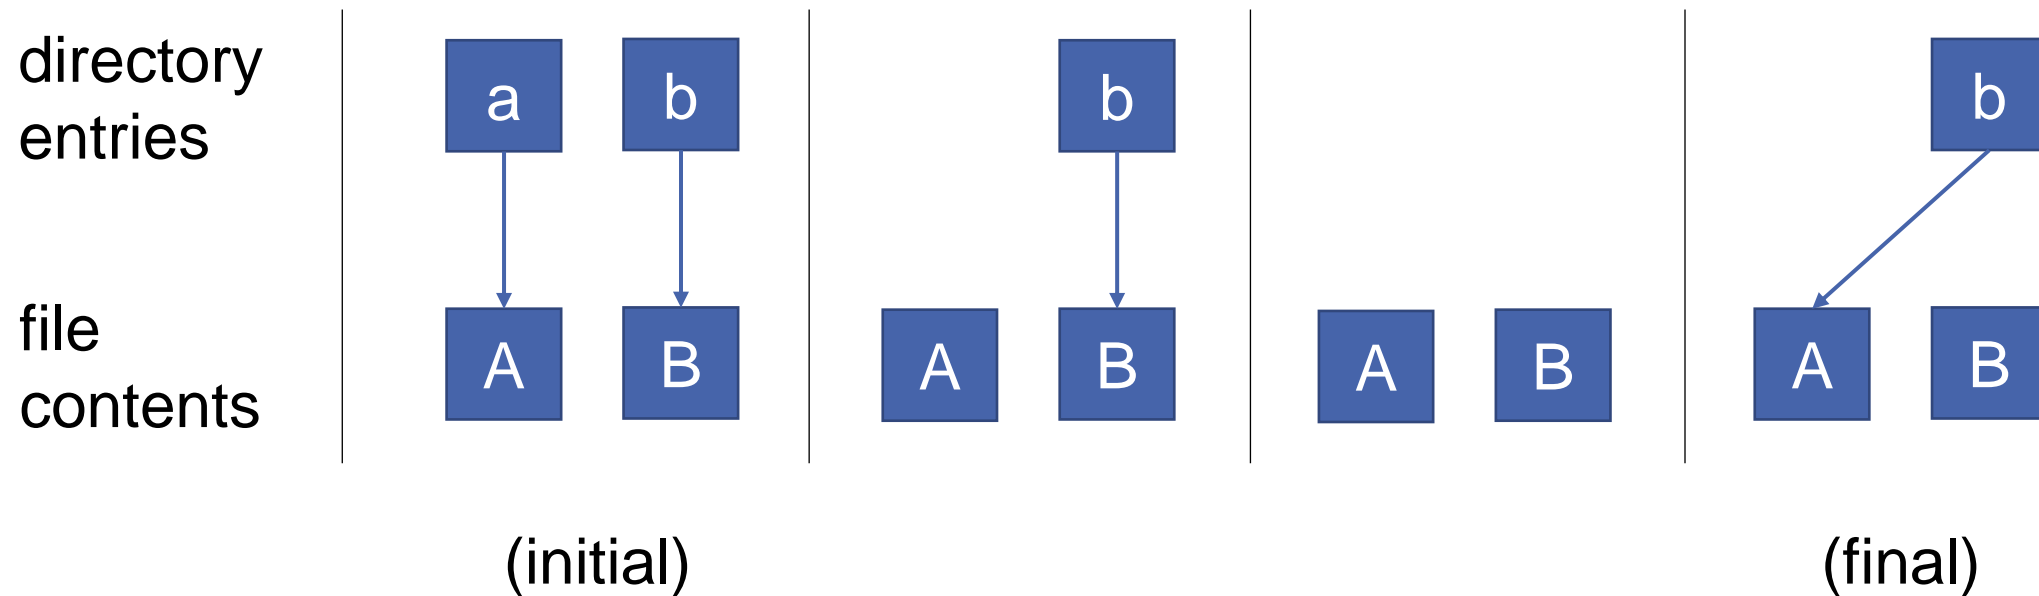
Stack trace:

```
__copy_user_nocache  
do_nova_inplace_file_write  
...  
vfs_write  
...
```

# Bug Detail: NOVA Rename Atomicity

Test command: `mv /mnt/a /mnt/b`

Vinter report: 4 states, 1 final state → not atomic!



# Observations

- Instruction-level tracing important for NVM primitives
  - ASM implementations!
- Testing of simple operation sufficient for good coverage
- Logic bugs are still prevalent

# Conclusion

- Crash consistency is important, but rarely evaluated
- Crash consistency testing pipeline
  - Tracing NVM events
  - Crash image generation with heuristic
  - Recovery testing
- Crash consistency bugs in practice

Consider testing crash consistency in **your** systems!

<https://github.com/KIT-OSGroup/vinter>