

# SPP

## Safe Persistent Pointers for Memory Safety

**Dimitrios Stavrakakis**, Alexandra Panfill


MJin Nam, Pramod Bhatotia




THE UNIVERSITY  
*of* EDINBURGH

Memory safety violations : Illegal accesses to unintended memory regions

Spatial memory safety  
e.g., buffer overflow, stack overflow



Temporal memory safety  
e.g., dangling pointer, double free



# Memory safety in practice

Prevalent in almost all low-level unsafe C/C++ code



Chromium project <sup>1</sup>

- 70% of vulnerabilities are memory safety problems



Microsoft <sup>2</sup>

- 70% of vulnerabilities fixed in security patches are memory safety violations



Android <sup>3</sup>

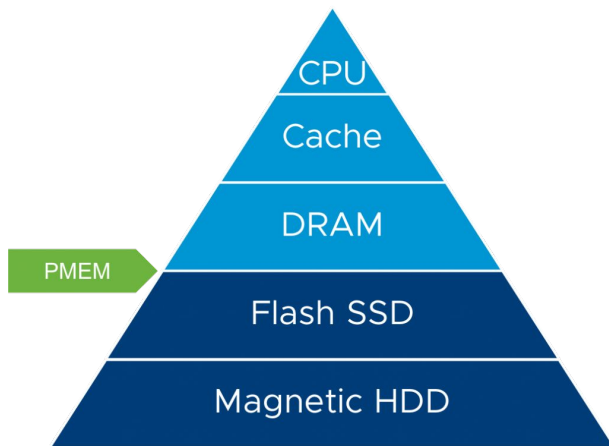
- 75% of vulnerabilities are memory safety issues

<sup>1</sup> Chromium project: <https://www.chromium.org/Home/chromium-security/memory-safety>

<sup>2</sup> Microsoft: <https://msrc-blog.microsoft.com/2019/07/16/a-proactive-approach-to-more-secure-code/>

<sup>3</sup> Android: <https://security.googleblog.com/2019/05/queue-hardening-enhancements.html>

Persistent memory management is susceptible to memory safety vulnerabilities



- Persistent memory programming model
- Durability & crash consistency
- Recovery code paths
- Performance & memory/storage overheads

Memory safety approaches for PM are **non-practical** for production deployment

How to design a **practical memory safety solution** for PM applications with minimal performance overheads?

Memory safety mechanism for PM-based applications

## System properties:

- Spatial memory safety
- Transparency
- High coverage
- Crash consistency



**Performance**

- ~~Motivation~~

- Design

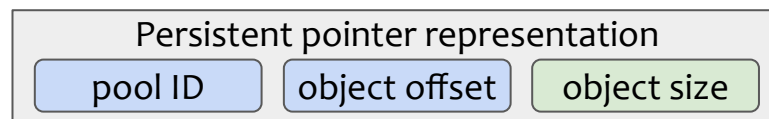
- Overview
- Challenges
- Example
- Persistent memory operations

- Implementation


- Evaluation

SPP enforces a **tagged pointer**-based approach for memory safety


- Enhanced PM pointer representation



- Native tagged pointer scheme



- Implicit runtime checks





# Design challenges

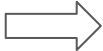


**#1** Performance & PM storage  
overheads

**#2** Transparent integration in existing  
toolchains

**#3** Crash consistency for PM safety  
metadata

**#4** Compatibility with existing  
applications and libraries


# Challenge #1: Performance & PM storage overheads

- PM pointer representation  64 bits of PM safety metadata per object
- Pointer tracking  omit checks for pointers to volatile memory
- Bound checks preemption  reduce the inserted runtime operations

Minimal PM safety metadata & optimized runtime instrumentation

## Challenge #2: Integration in existing toolchains


- Transparent creation of the tagged pointers via the PM API
- Wrappers for PM management operations
- Wrappers for memory intrinsic functions



SPP supports the PM API and memory intrinsic functions without modifications

# Challenge #3: Crash consistency for PM safety metadata

- PM guarantees atomicity only for aligned 8-byte stores
- Atomic operations & software transactions
- Include PM safety metadata in transaction logs



Crash-consistent (meta)data updates using atomic operations and transactions

# Challenge #4: Compatibility with applications/libraries

- Pre-compiled shared libraries/applications

```
void* tagged_pm_ptr = pmemobj_direct(obj_id); //get tagged ptr
...
internal_foo(tagged_pm_ptr); //internal function call
...
external_foo(tagged_pm_ptr); //external function call
external_foo(clean_pm_ptr); //external function call
...
```


SPP masks the tagged pointers passed to external shared libraries/applications

SPP sets and updates the native pointer **tag** on pointer operations

- `pm_ptr` : pointer to a 42-bytes PM object

- `pm_ptr += 21; // ptr is in bounds`

- `pm_ptr += 21; // ptr gets out of bounds`



# Design overview - PM layout

Virtual address space

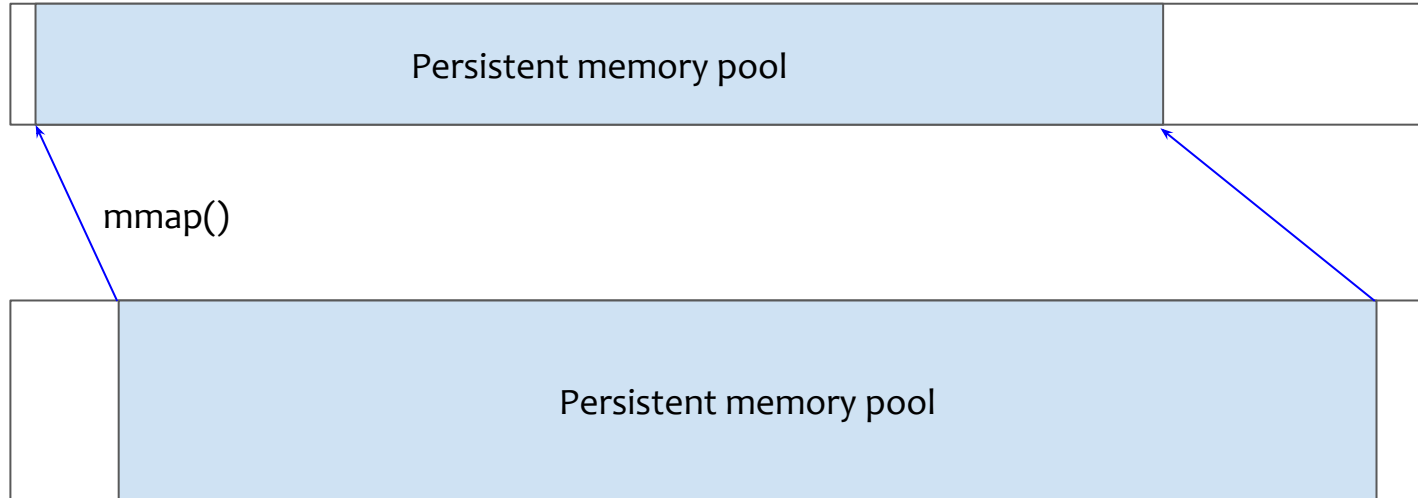


Persistent  
Memory



# Design overview - PM layout

Virtual address space




Persistent memory pools are directly mapped to the virtual address space of an application



# Design overview - PM allocation

Virtual address space

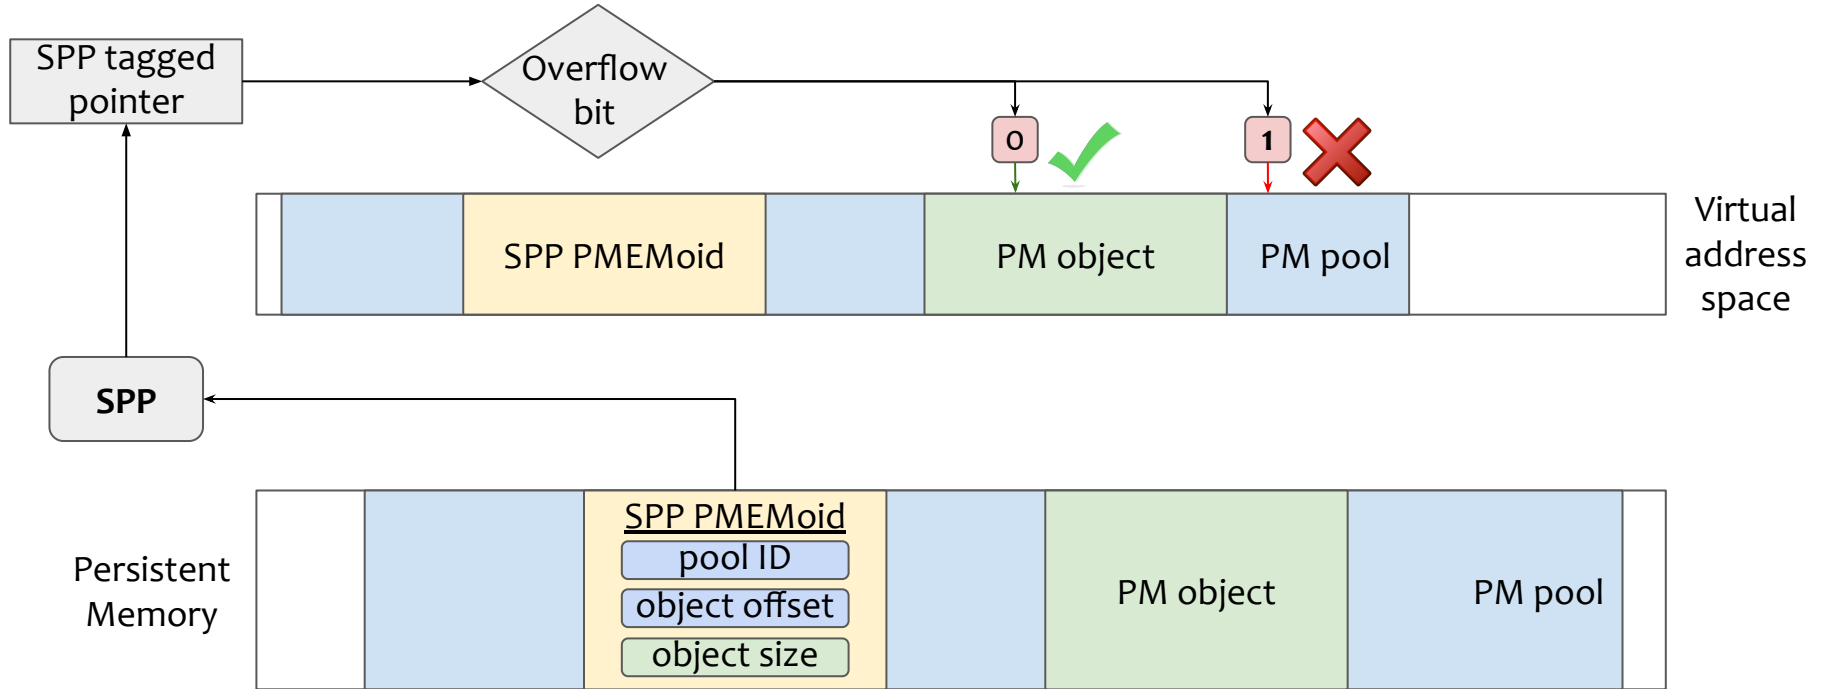


Persistent  
Memory



SPP allocates the object and atomically sets the **object size** field

# Design overview - PM access



On a memory access SPP preserves the **overflow bit** and performs an implicit bounds check

# Outline

● ~~Motivation~~

● ~~Design~~

● Implementation

● Evaluation

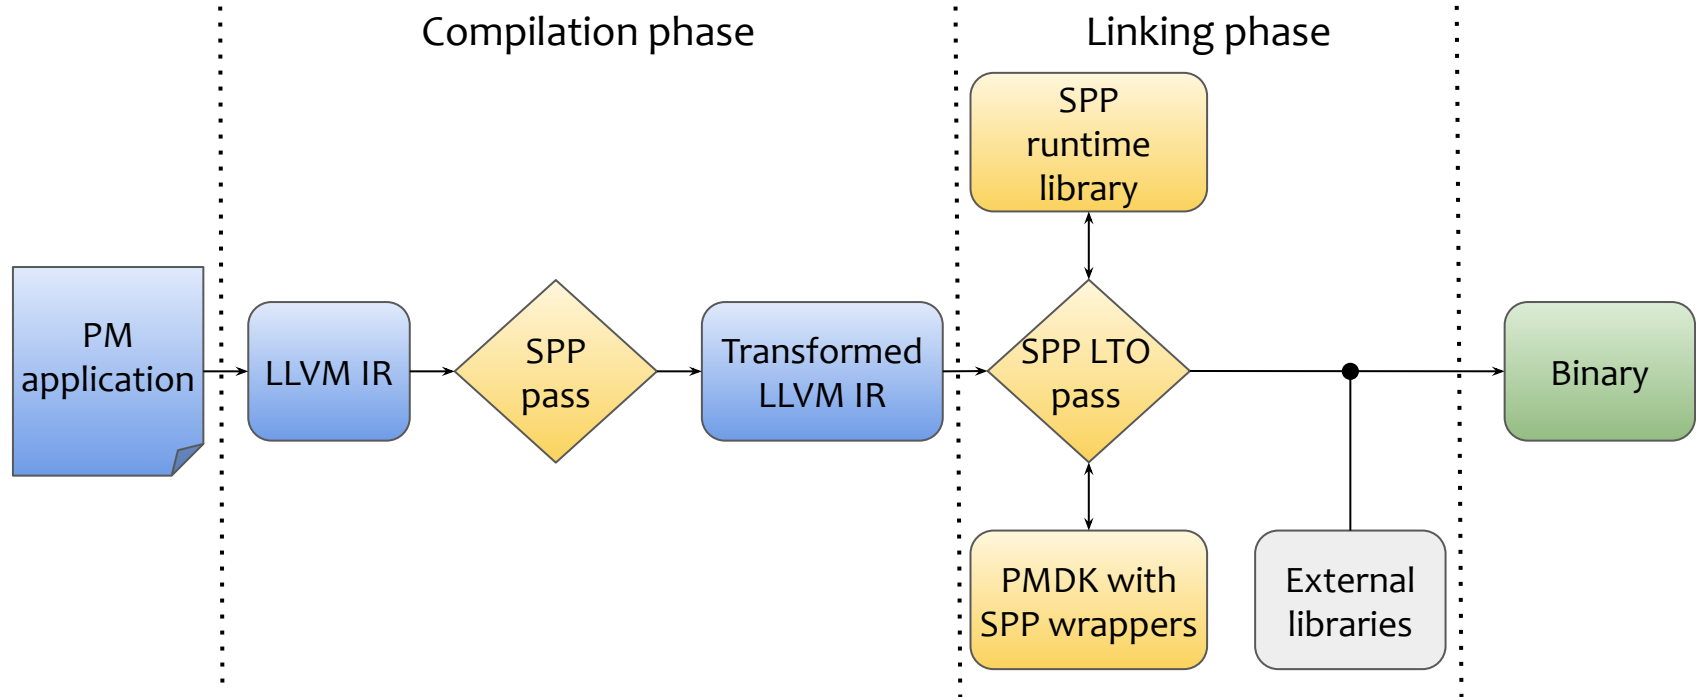
SPP is built on **PMDK**<sup>1</sup> and **LLVM**<sup>2</sup>

- SPP pointer representation – minimization of space overheads & runtime checks!
- Transformation & LTO compiler passes – performance & compatibility!
- PMDK programming model – transparent support!
- Crash consistency via PMDK transactions & atomic operations

<sup>1</sup>Persistent memory development kit (PMDK): <https://github.com/pmem/pmdk>

<sup>2</sup>LLVM: <https://github.com/llvm/llvm-project>

# SPP hardening workflow



The application is finally linked with SPP runtime library, PMDK and external libraries

# Outline

- ~~Motivation~~
- ~~Design~~
- ~~Implementation~~
- Evaluation


- What is the performance overhead of SPP?
  - Persistent memory KV store (pmemkv), PM phoenix benchmark suite
- How much PM space overhead does SPP introduce?
  - Persistent indices (ctree, rtree, rbtree, hashmap)
- How robust is SPP in detecting memory safety vulnerabilities?
  - RIPE benchmark framework

- Experimental setup:
  - Dual socket Intel Xeon Gold 6326 CPU (16 cores)
  - 64 GB DRAM / socket
  - 1 TB Intel Optane DC DIMMs / socket
  - PM configured in *App-Direct* mode
- Variants:
  - *PMDK*→No memory safety
  - *SafePM*→Application hardened with SafePM
  - *SPP*→ Application hardened with SPP



# Performance overhead

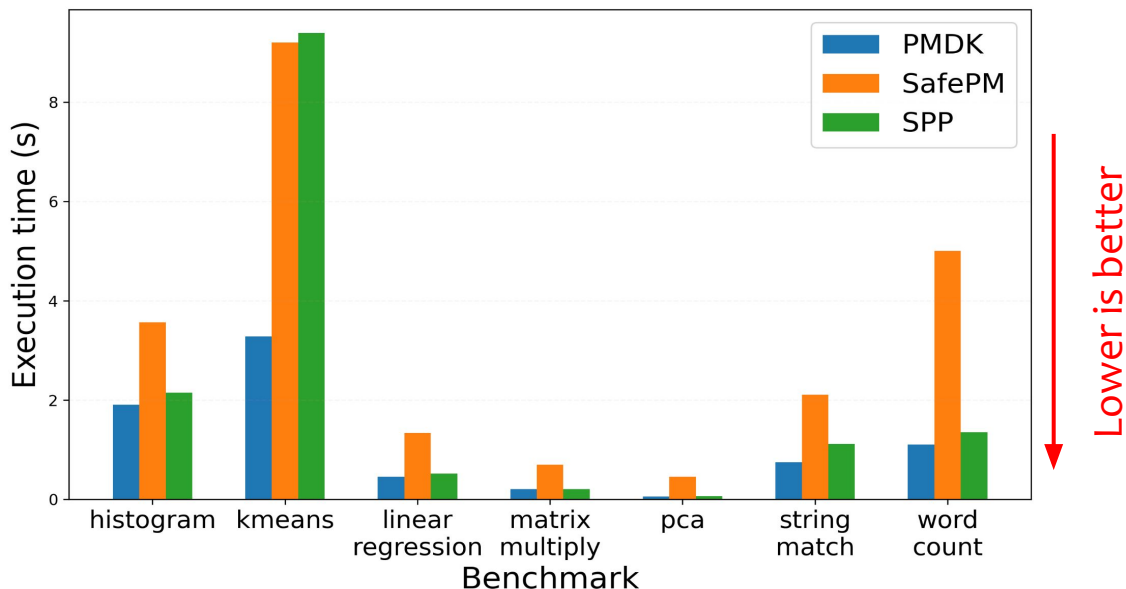
Persistent KV-store benchmark, **10M** ops, **50%** reads / **50%** writes



SPP incurs notably lower performance overheads compared to SafePM

# Performance overhead


PM Phoenix benchmark, 8 threads, largest dataset for each benchmark



SPP introduces minimal performance overheads even in CPU intensive scenarios

# Space overhead

Persistent indices, insert/get workloads, relative to PMDK



SPP does not introduce significant PM space overhead on average

RIPE benchmark, **223** buffer overflow exploits

Variant	Exploitable PM buffer overflows
PMDK	83
memcheck	20
SafePM	6
SPP	4

SPP is an efficient memory safety solution for PM with low performance overheads

**Current PM memory safety approaches are designed for **debugging** purposes**

- high performance overheads
- considerable PM storage overheads

## **Safe Persistent Pointers (SPP):**

- comprehensive spatial memory safety
- low performance & PM storage overheads
- no source code modifications
- crash consistency & durability



**Paper**



**Code**

**Try it out!**

<https://github.com/dimstav23/SPP>



Code  
Reproducible



Dataset  
Reproducible

[1] PM hierarchy image,

<https://www.starwindsoftware.com/blog/persistent-memory-in-vmware-vsphere-6-7-what-is-it-how-fast-is-it>