

Carbide: A Safe Persistent Memory Multilingual Programming Framework

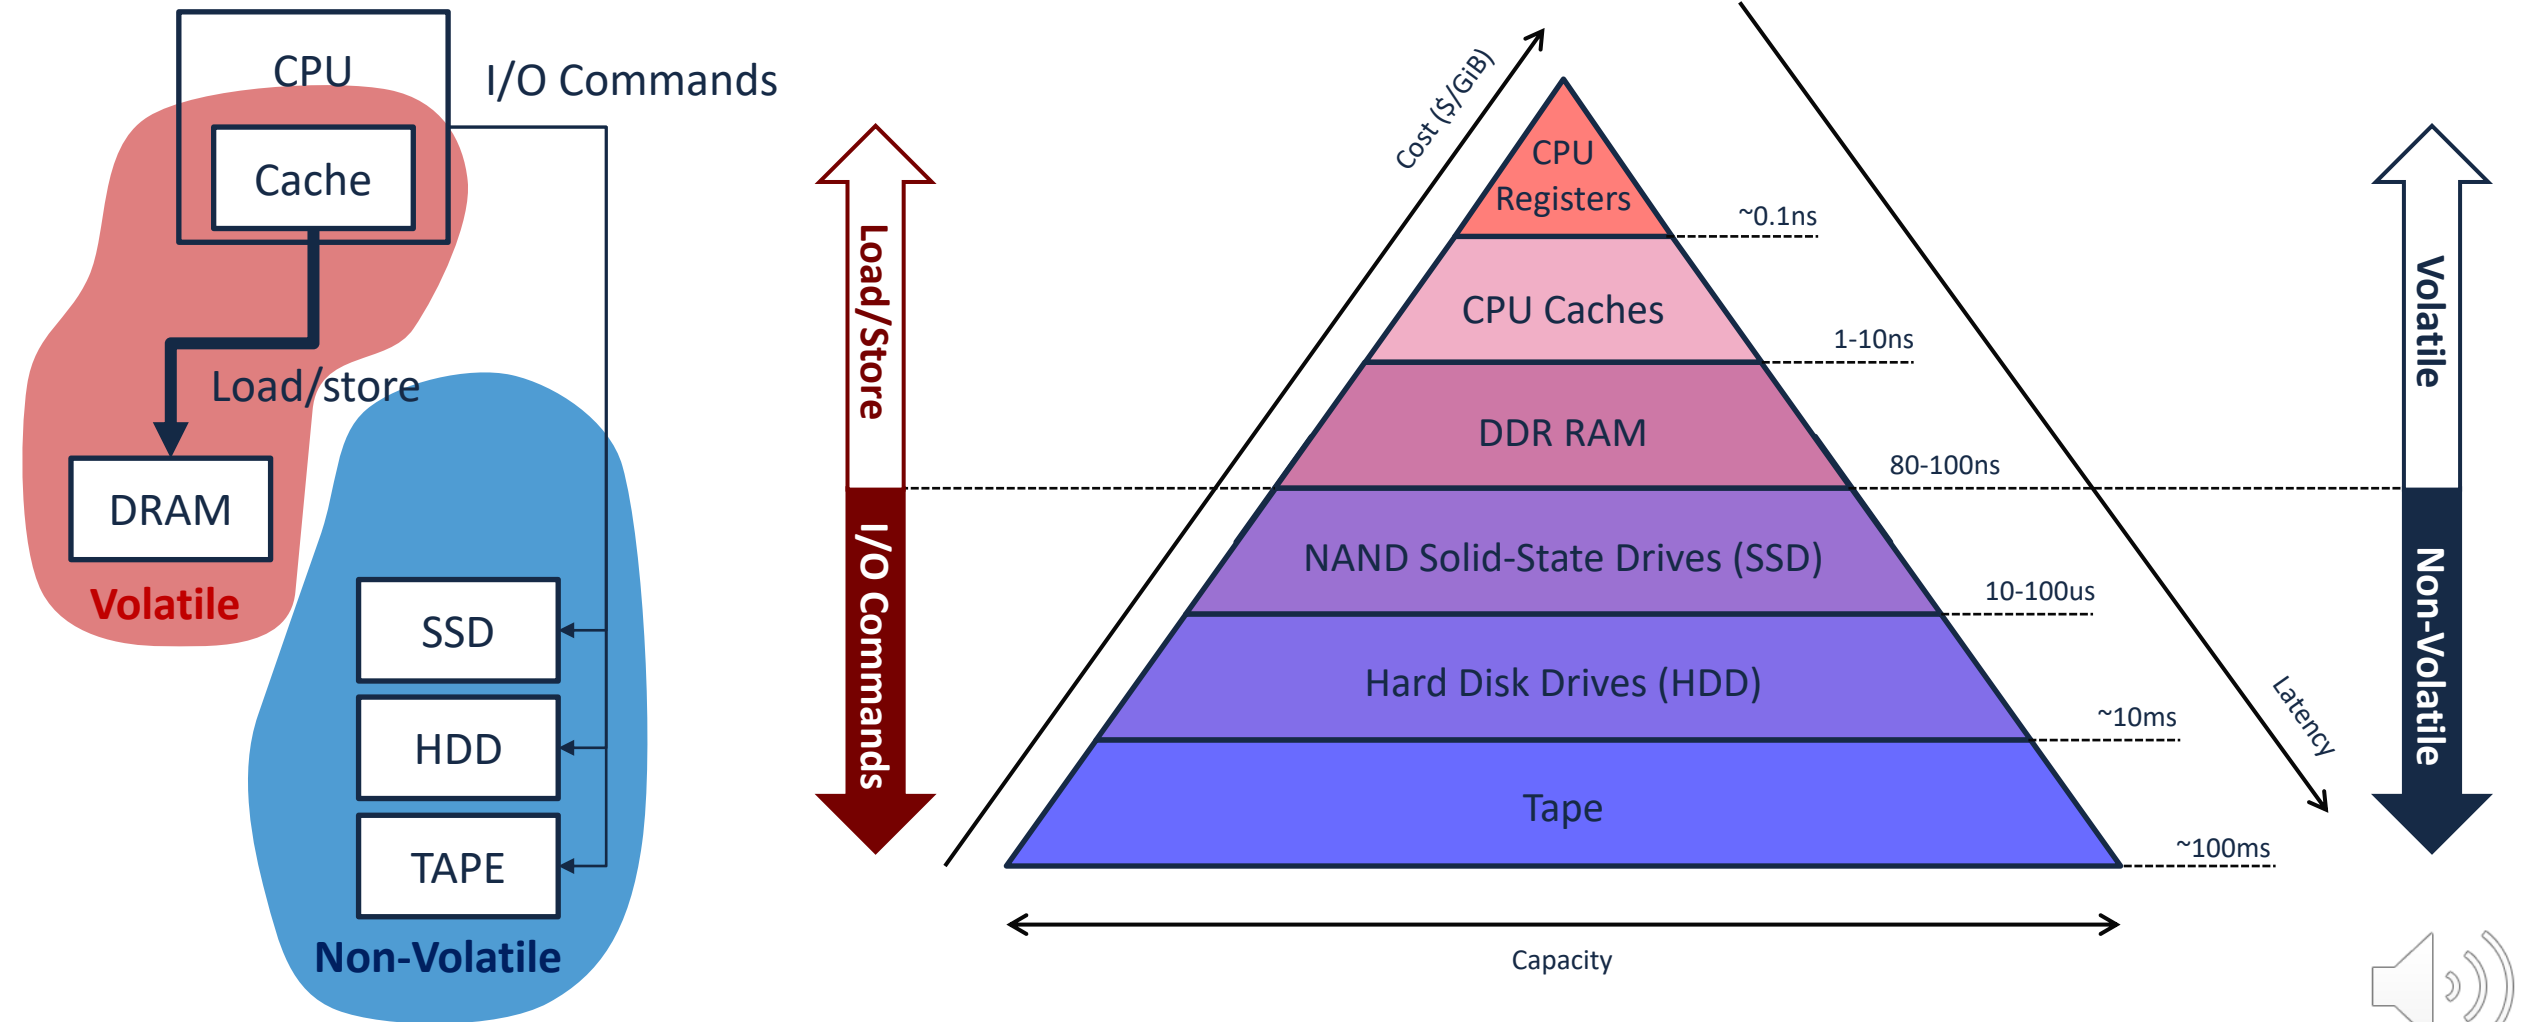
Morteza Hoseinzadeh* and Steven Swanson

*Non-Volatile Systems Laboratory
Department of Computer Science & Engineering
University of California, San Diego*

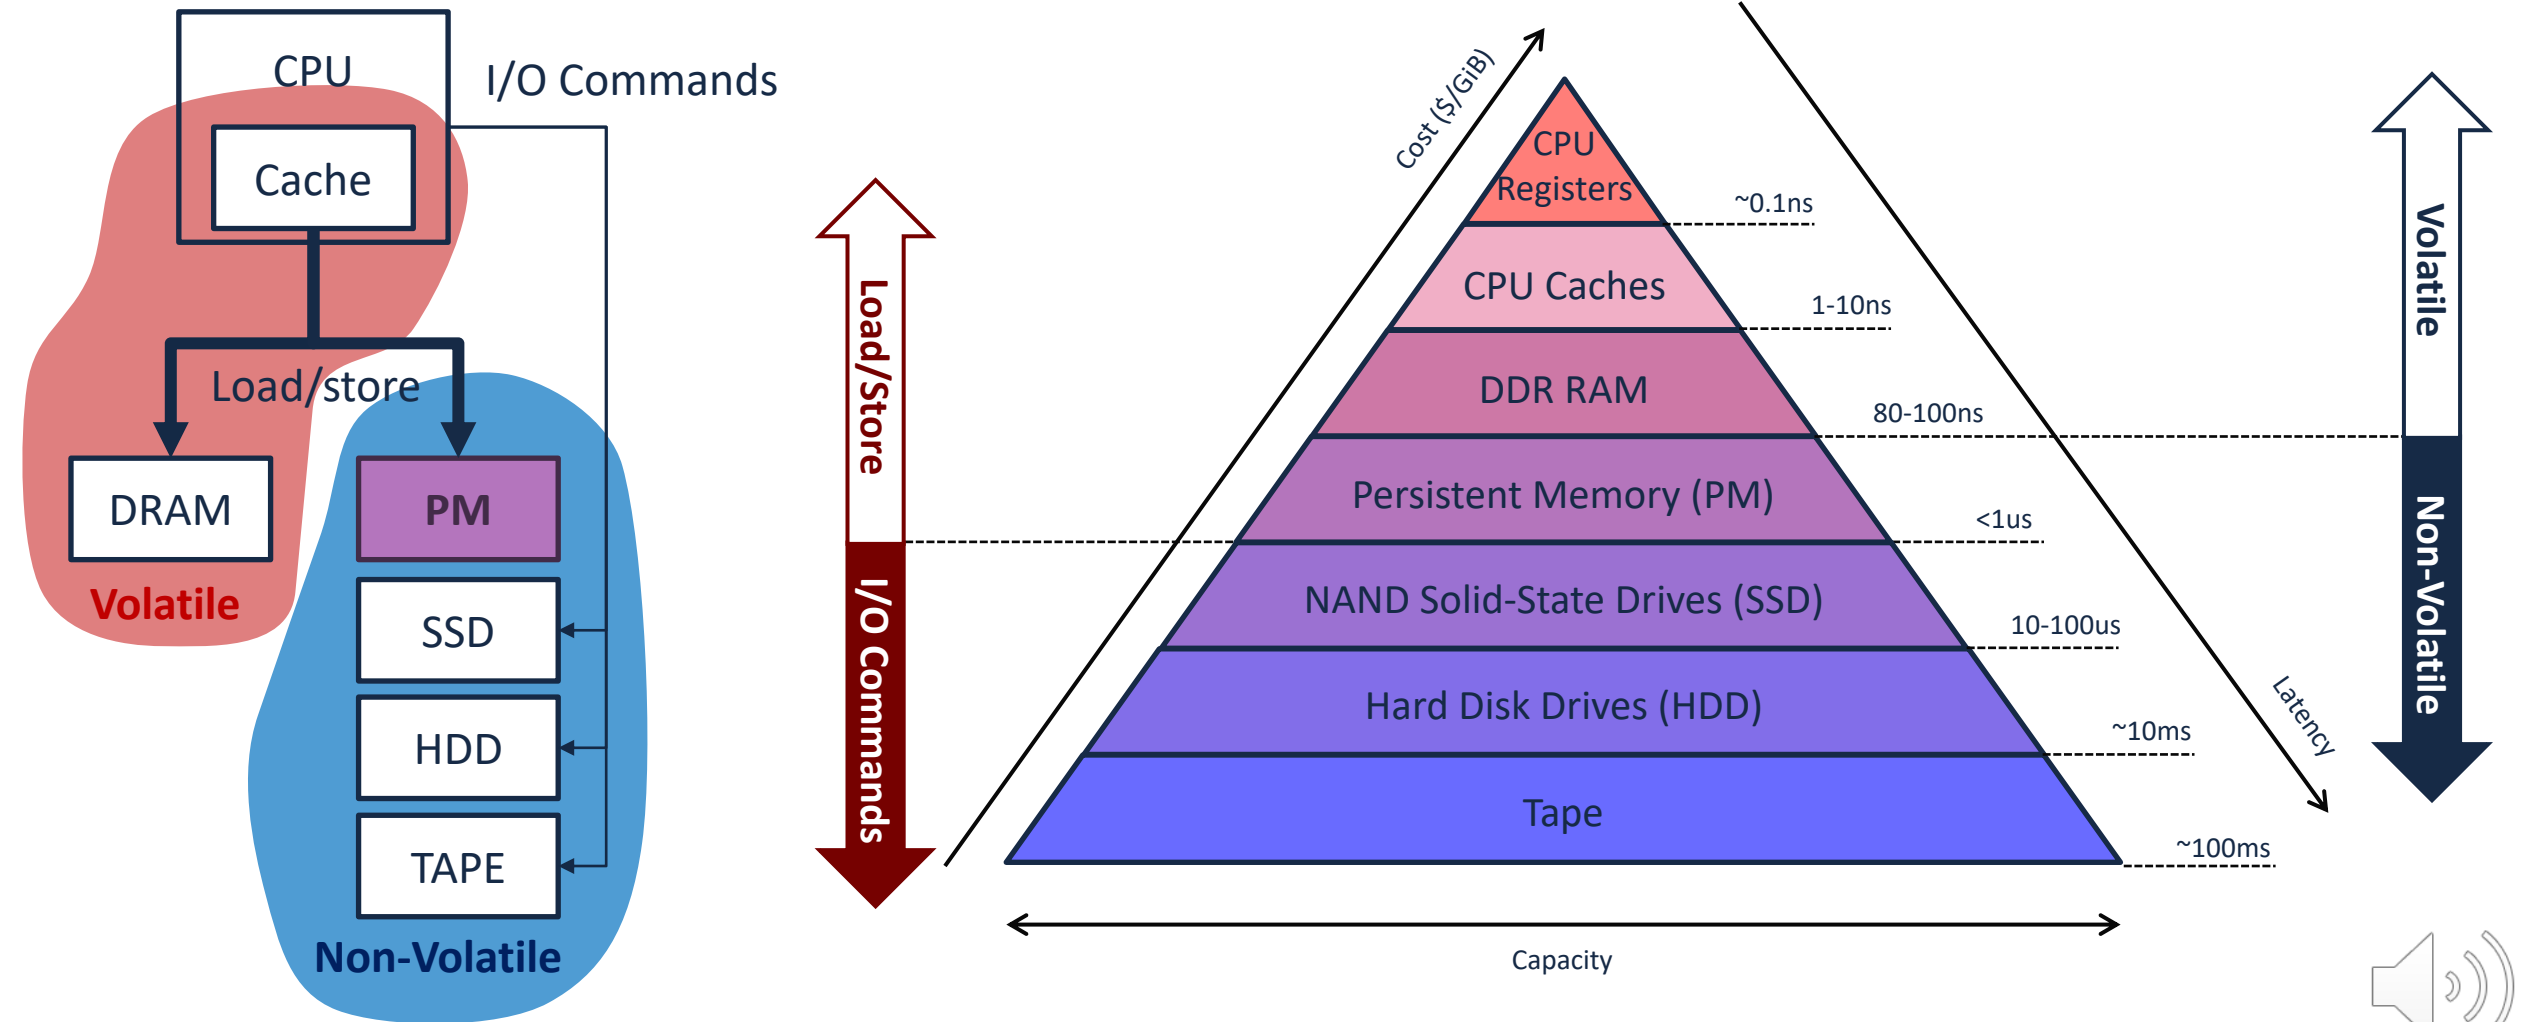
** The first author is currently employed by Oracle*

NVMW 2022

Persistent Memory (PM)



Persistent Memory (PM)

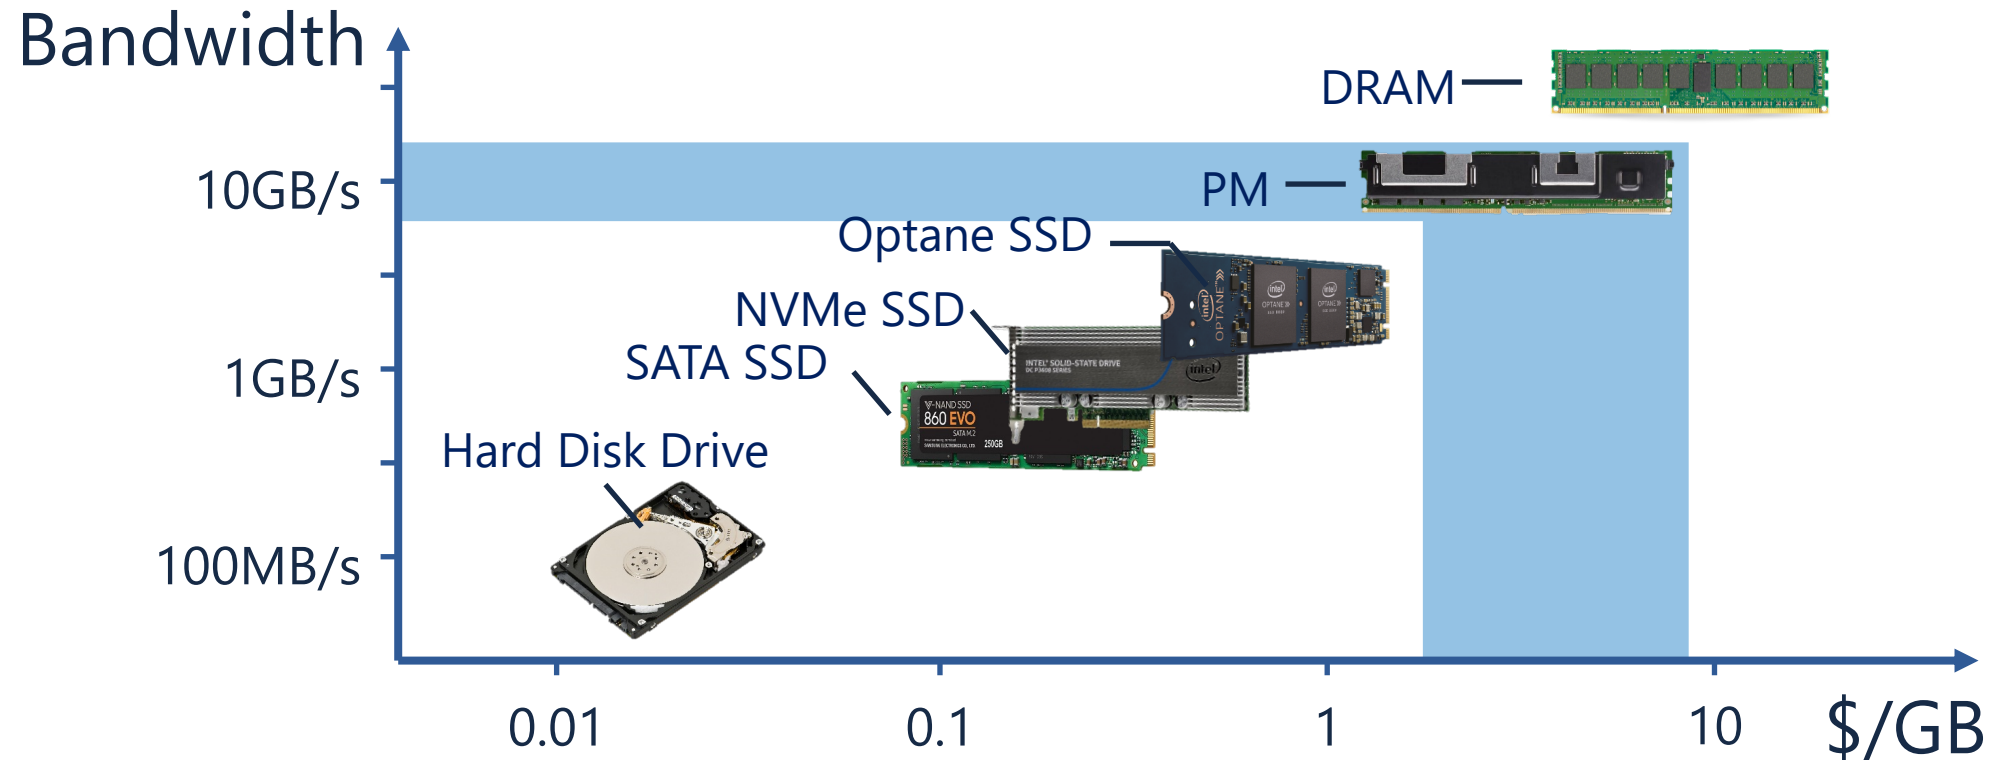


Challenges of Using PM in User-Level Applications

- Intensified current programming challenges (e.g., memory leaks)
- Persistent data consistency
 - Volatile CPU caches reorder the updates
 - No atomic compare-swap-persist instruction exists
 - Stores are not persistent until cache line is flushed
 - Non-temporal stores and cache-line flush instructions are expensive
- PM management burden is on user applications
- Handling hardware errors directly in the applications

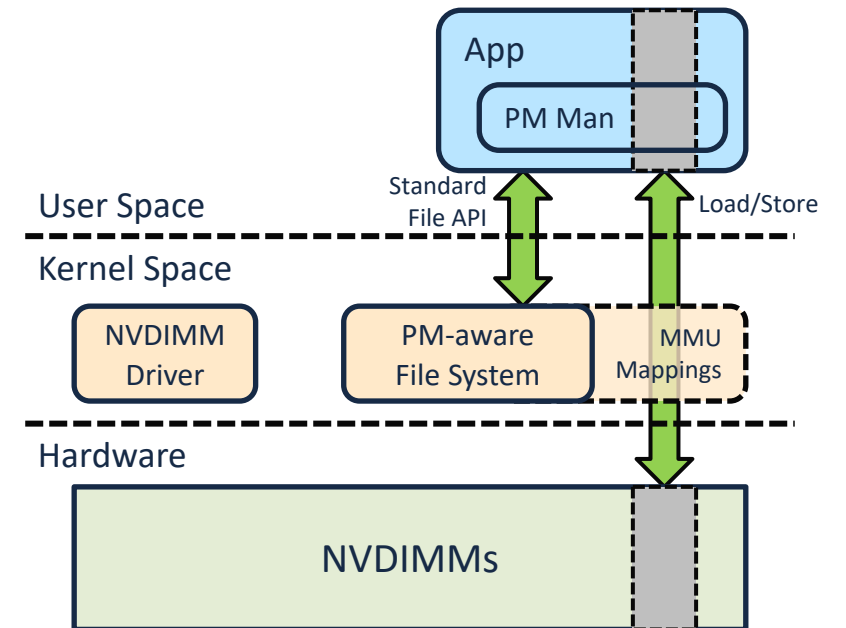
Challenges of Using PM in Storage Systems

- Inefficient usage of PM when used as a block device
- Limited scalability due to PM's expensive price



PM Programming Model

- A set of standards for enabling application development for persistent memory to address the PM programming challenges:
 - + DAX enabled file system on PM
 - + `mmap()` files (mem pools) to the virtual address
 - + User space memory management



PM Programming Frameworks

1. *Basic PM Programming Frameworks*
 - Provide interface to access PM
 - Make no safety guarantee on usage
 - Examples: PMDK, Atlas, go-pmem, Mnemosyne, and NV-Heaps
2. *Code Transformation Frameworks*
 - Statically analyze the code and inject PM operations
 - Limit the flexibility to make the program state machine smaller
 - Examples: AutoPersist, NVTraverse, Mirror, and Hippocrates
3. *Debugging/Bug-Fixing Tools*
 - Statically analyze the code and do symbolic execution to find the bugs
 - NP-Hard problem, and path explosion in large programs
 - Examples: NVL-C, Jaaru, and Agamoto

PM Programming Frameworks

4. *Testing Frameworks*

- Dynamically inject failures to test the program
- Completeness proof is not provided
- Examples: PMTest and XFDetector

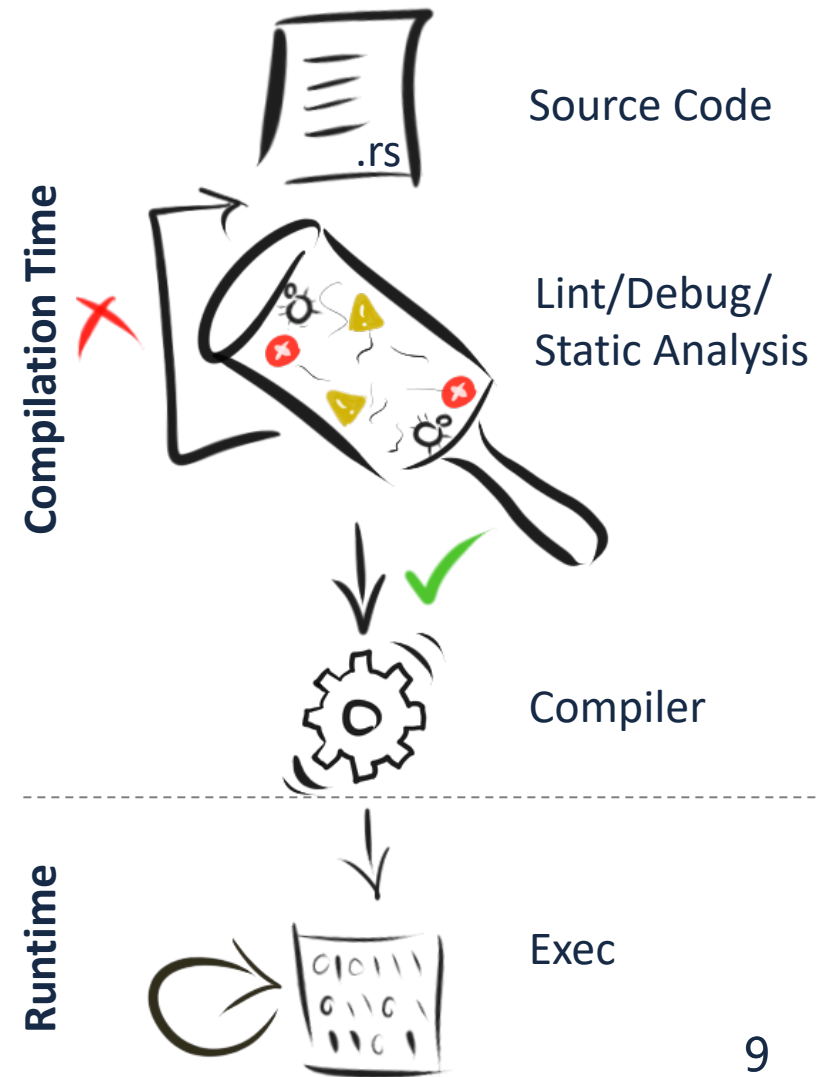
5. *Pre-Compilation Debugging Frameworks*

- Apply safety rules statically as it's being developed
- Limit the flexibility as they apply restrictive safety rules
- Example: Corundum

Corundum

- A PM programming library for Rust
- Enforces PM safety at compile time
- High performance due to static analysis
- Idiomatic approach to support PM
- It guarantees no PM-related bug

PM Safety \subset Rust's Type Safety



Corundum Challenges

- Too restrictive
- Risky optimizations are not possible
- Steep learning curve for non-Rust developers

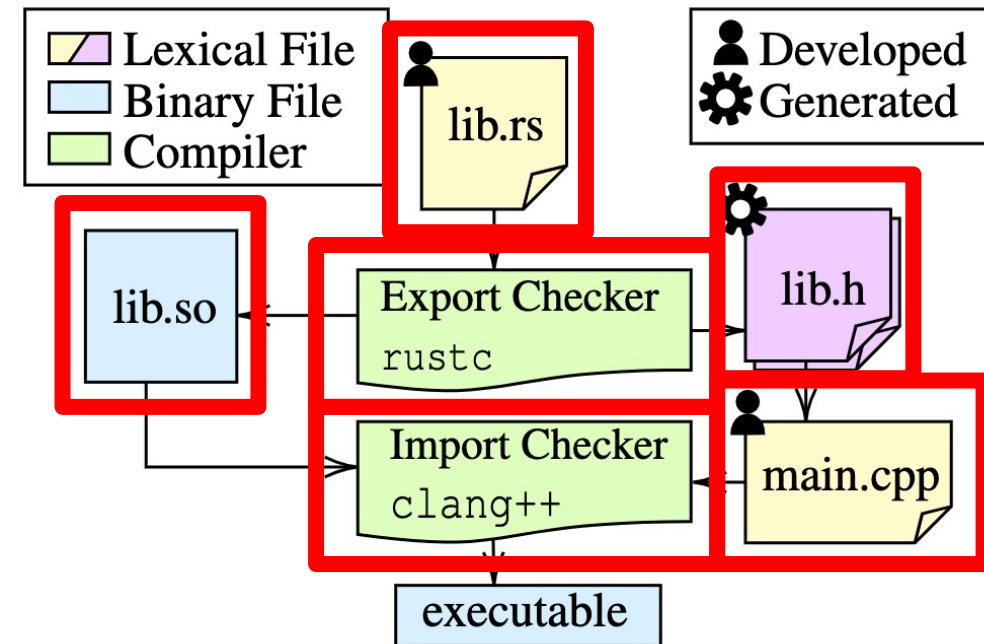
CARBIDE:

USE CORUNDUM FOR DEFINING PERSISTENT TYPES

USE C++ FOR DEVELOPING THE PROGRAM

Carbide

- Developing persistent data structure type separately using Corundum in Rust (*lib.rs*)
- Strict rules apply to persistent types only
- Data types are externally available through a dynamic library (*lib.so*) with an automatically generated **API** (*lib.h*)
- The Export Checker statically checks the container types for the capability of external usage
- The Import Checker statically checks the types being stored in PM



Example

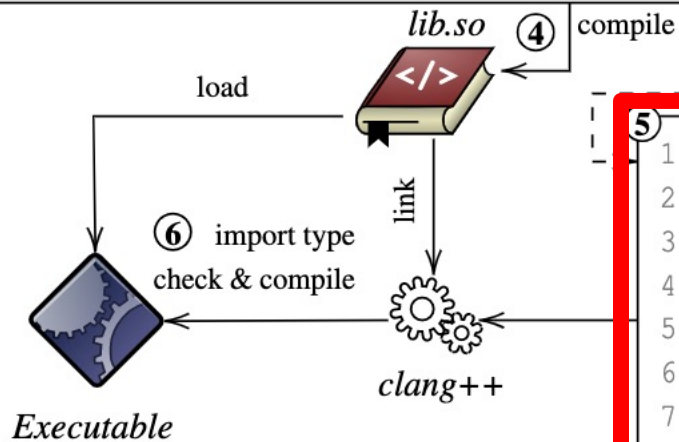
```
1 enum Elem<T: PSafe, P: MemPool> {  
2   Nil,  
3   Cons(T, Pbox<Elem<T, P>, P>),  
4 }
```

Carbide System (q is the pool type)

```
1 #[derive(Extern)]  
2 #[pools(q)]  
3 struct List<T,P:MemPool> {  
4   head: PRefCell<Elem<  
5     ↳ ByteArray<T,P>,P>,P>  
6 }  
7 #[extern]  
8 impl<T,P:MemPool> List {  
9   pub fn append(&self,  
10    ↳ v: Gen<T,P>) {  
11     /* elided for space */  
12 }  
list.rs
```

```
1 extern "C" {  
2   fn q_list_append(this: &List<q>,  
3     ↳ v: Gen<void,q>);  
}
```

```
1 template<class T, class P>  
2 class List: psafe_params {  
3   _List<P> *self;  
4 public:  
5   void append(const T& v) {  
6     list_traits<P>::append(  
7       ↳ self, v);  
8 }  
9 };  
10 template<>  
11 struct list_traits<q> {  
12   template <class T>  
13   void append(  
14     ↳ const _List<q> *self,  
15     ↳ const T& v) {  
16     q_list_append(self, v);  
17 };  
list.hpp
```



```
1 #include <list.hpp>  
2  
3 int main() {  
4   auto h = q::open( "foo");  
5   q::List<int> lst(h,"list1");  
6   lst.append(10);  
7 }  
main.cpp
```

Carbide's Design Goals

API Design

1. Preserve the same guarantees as Corundum's
2. Provide a seamless cross-language PM management system
3. Provide a safe C++ interface to interact with data as defined in Rust
4. Statically checked the external usage of the persistent type definition in Rust
5. Statically checked the usage of external persistent type declaration in C++
6. Specify a design pattern to make a C++ type persistent

API Design Challenges

- Type Interoperability
 - Rust and C++ layout memory differently
- Polymorphism
 - Polymorphic types are not available through dynamic libraries in Rust and C++
- Memory Leaks
 - C++ does not garbage collect when dynamic allocation is used
- Lifetime Conflict
 - The RAI model in C++ and Rust have distinct lifetime scopes

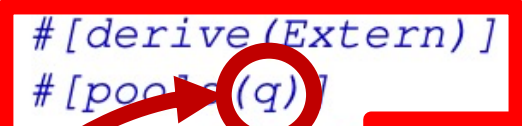


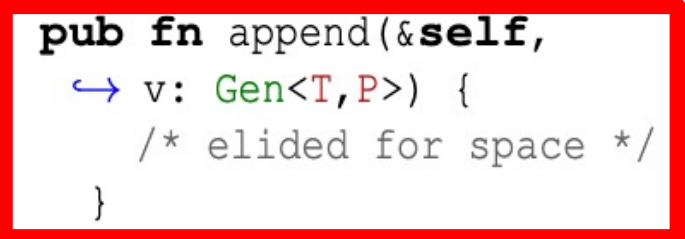
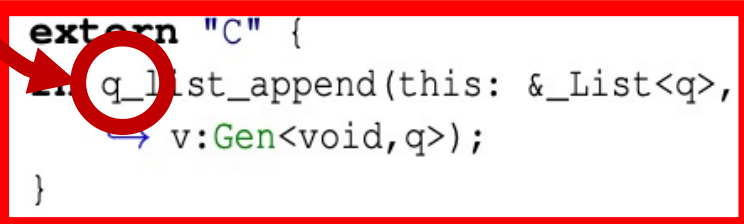

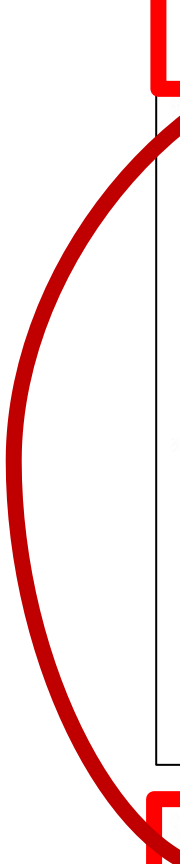


Type Interoperability


- Portable type:
 - Annotate external types for a specific set of pools
 - Corundum's rules apply (Rust's type system)
 - Exactly one pool type parameter
 - Other type parameters are used in form of byte arrays
 - External interface is FFI-compatible
 - Provide at least one transactional constructor
- Carbide exports the type's functionality by generating an FFI for every specified **pool**

Export Rules

```
#[derive(Extern)]
#[pool(q)]
struct List<T,P:MemPool> {
    head: PRefCell<Elem<
        ↪ ByteArray<T,P> P>,P>
}

#[extern]
impl<T,P:MemPool> List {
    pub fn append(&self,
        ↪ v: Gen<T,P>) {
        /* elided for space */
    }
}
```

 list.rs


```
extern "C" {
    ↪ q_list_append(this: &_List<q>,
        ↪ v:Gen<void,q>);
}
```

Polymorphism

- Type-parameter reduction and reparameterization
 - Specialize the data type parameters with **void**
 - Specialize the pool type parameters for every specified pools and generate the FFIs
 - Implement a C++ template class (vessel class) with the same parameters and functionality
 - Implement the type traits for the given pools in C++ to call the corresponding APIs

```
template<class T, class P>
class List: psafe_params {
    _List<P> *self;
public:
    void append(const T& v) {
        list_traits<P>::append(
            ↪ self, v);
    }
};
```

```
template<>
struct list_traits<q> {
    template <class T>
    void append(
        ↪ const _List<q> *self,
        ↪ const T& v) {
        q_list_append(self, v);
    }
};
```

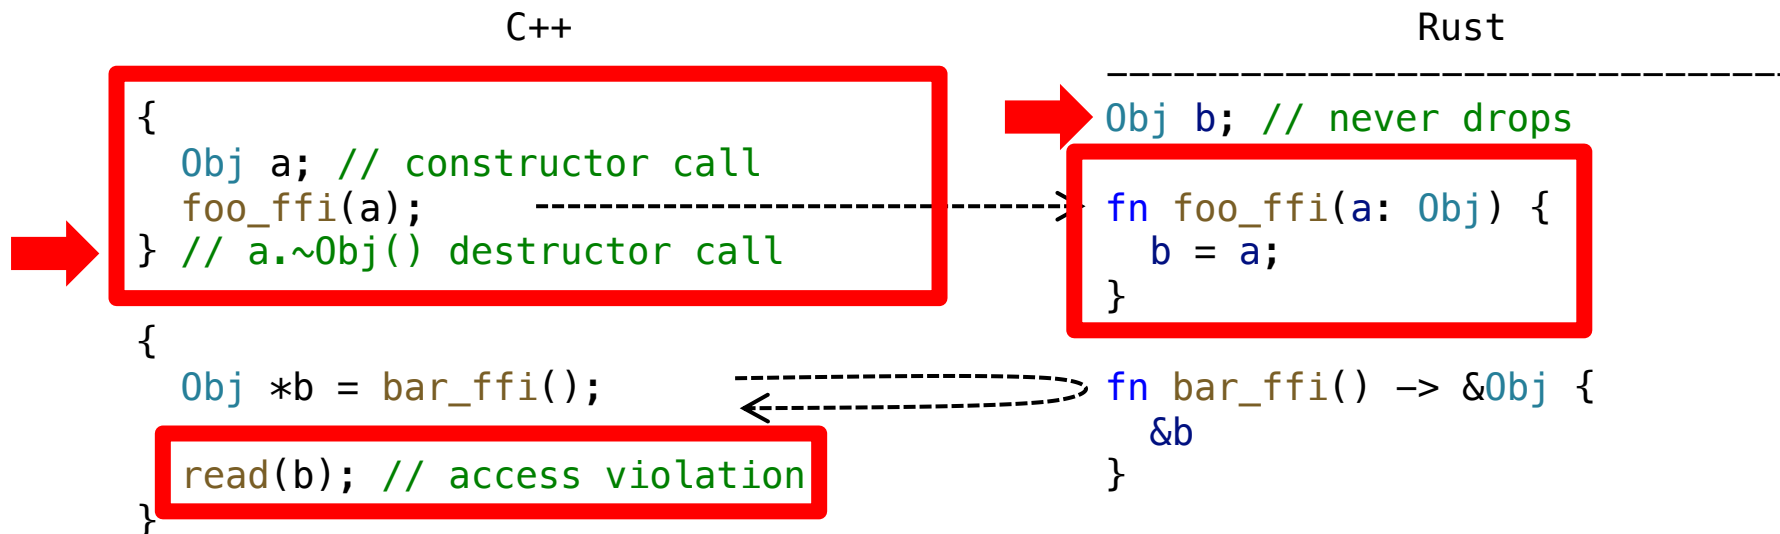
 *list.hpp*

Memory Leaks

- There is no PM dynamic allocation available in C++
- Only Carbide's internal types can manage PM
- Every allocation is owned by an object in Rust

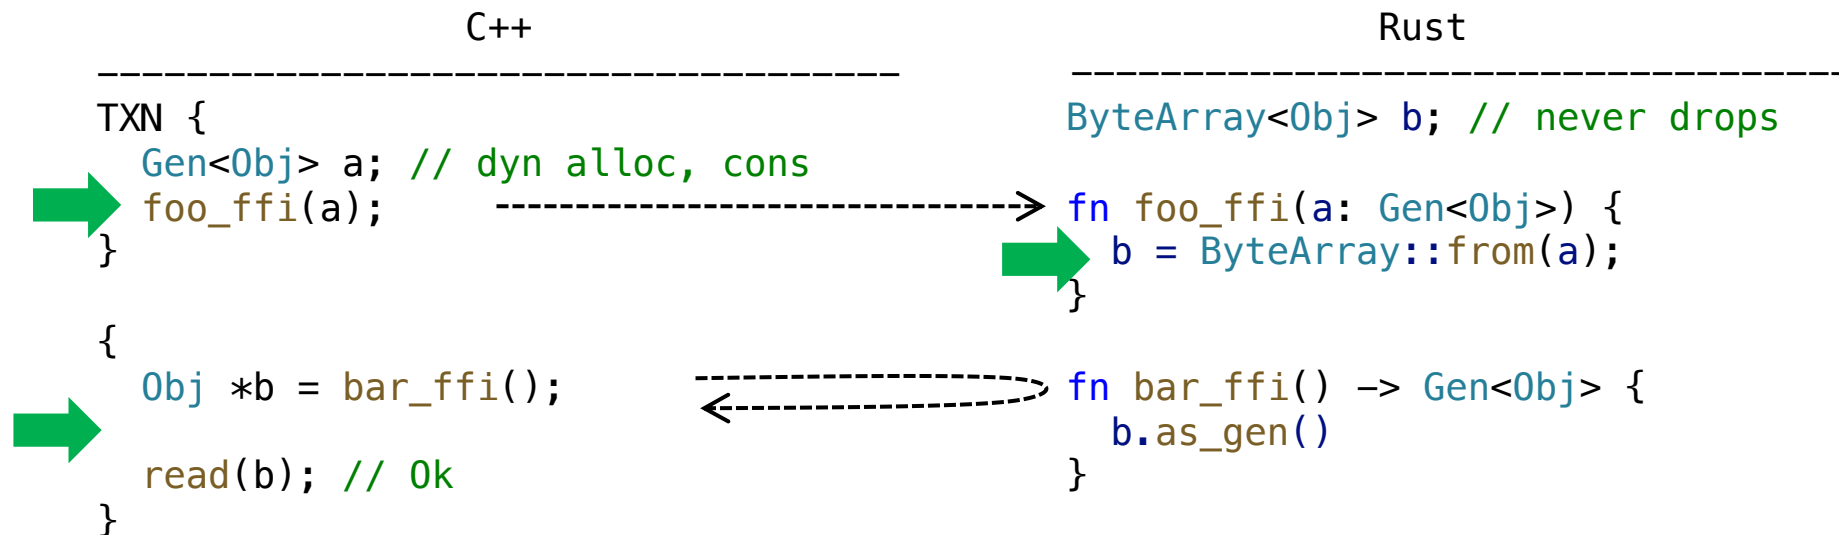
Lifetime Conflict

- Lifetime of a C++ object is unknown when passed to a foreign function
- The object's resources are released at the end of the scope




Extended RAI

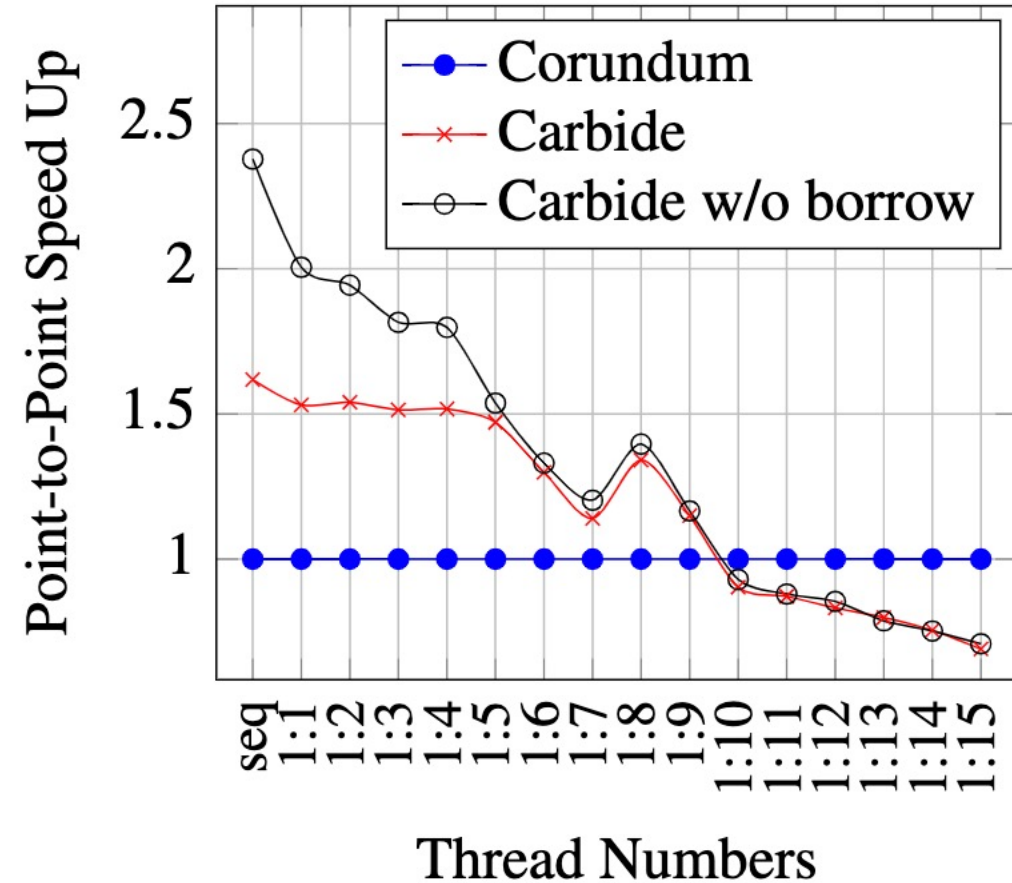

- A hyper scope is a scope extending from C++ to Rust
- **Gen<T, P>** as a cross-language reference type lives in a hyperscope
 - Defined in both Rust and C++
 - Contains a **relative pointer to the destructor** function to call from Rust
 - Dynamically allocates and construct the object when instantiated in C++
 - Does not immediately release resources in the destructor
 - Can merely move the resource to a **ByteArray<T, P>**



Performance Results



Optimization Impact and Scalability



Conclusion

- PM is an advanced memory technology that offers both high-performance and non-volatility
- PM programmers face a set of safety challenges, as well as higher price per GB compared to other NVM block devices
- Current PM programming frameworks exclusively offer safety or programming flexibility
- We presented Carbide, a PM programming framework that allows using Corundum data structures in C++ to guarantee safety as well as programming flexibility

Thank you!