

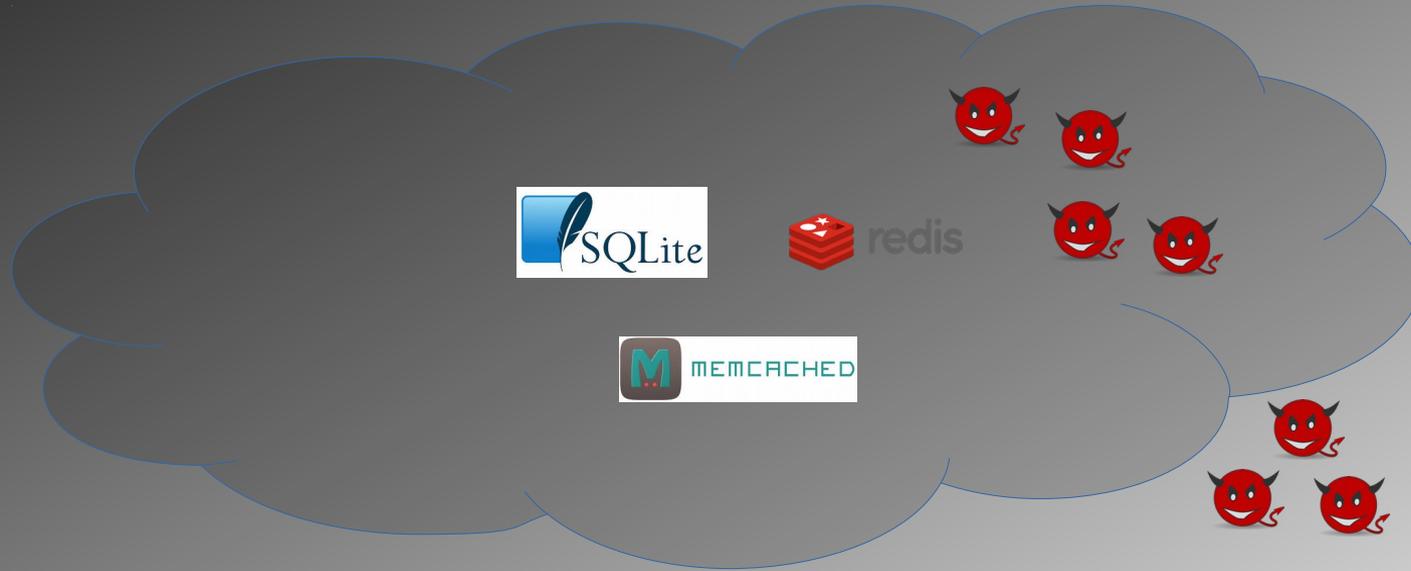
CoSMIX: A Compiler-based System for Secure Memory Instrumentation and Execution in Enclaves

Meni Orenbach
Yan Michalevsky
Christof Fetzer
Mark Silberstein

The ~~night~~ is dark and full of terrors

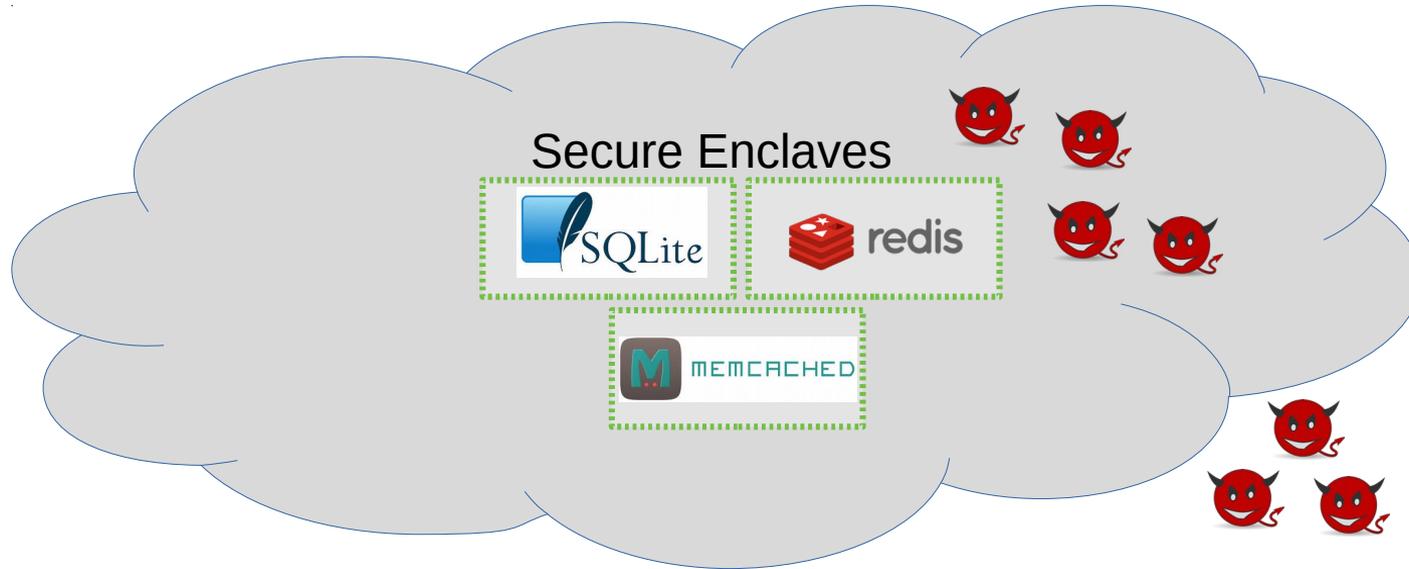
GAME OF THRONES

The cloud is dark and full of terrors



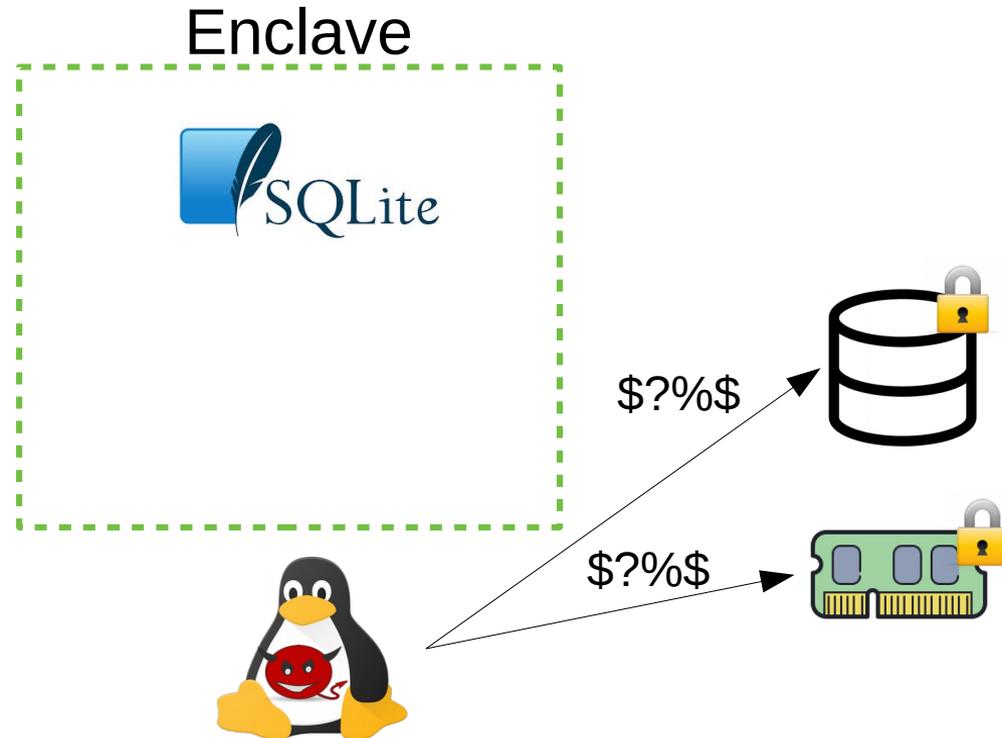
The cloud is dark and full of terrors

Yet, hardware enclaves can help

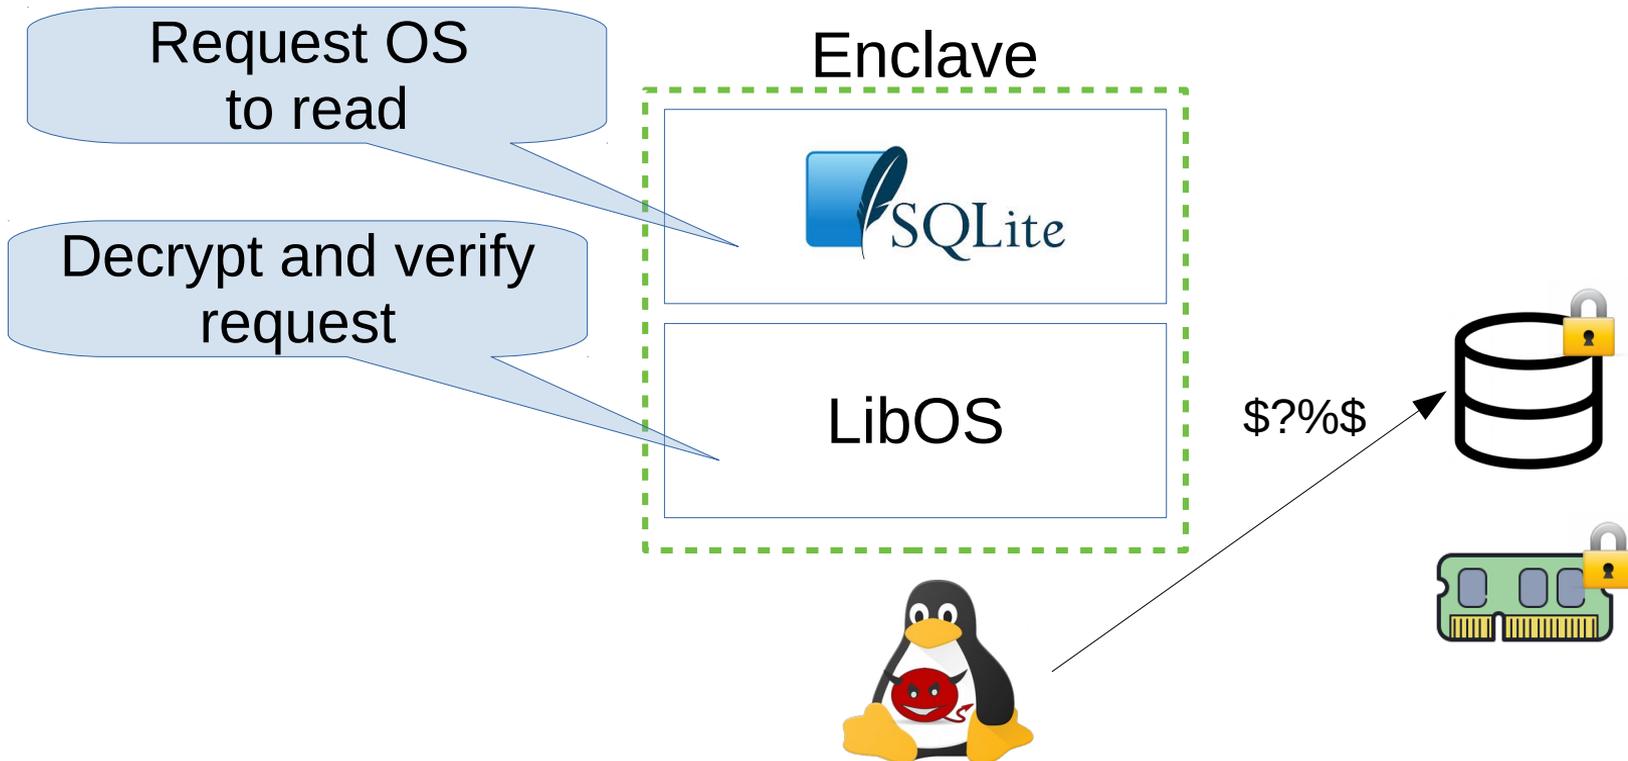


Enclaves shield applications from privileged adversaries

- ✓ Confidentiality
 - ✓ Integrity
- Untrusted OS



Run unmodified applications inside enclaves



Talk focus: Current SGX Enclaves

Integrated into Intel CPUs

Can we execute
any x86 application
inside **enclaves**?



Memory-mapped files in SGX?

Enclave



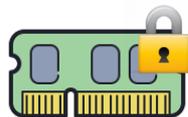
```
fast_read_db  
ptr = mmap(  
data = *ptr;
```

Handler customization is a powerful tool

- File mapping
- Disaggregated memory
- Compressed memory and more

Efficient and requires **no modifications**
Applications **rely** on this abstraction

Its Not Possible in SGX Today!



Memory-mapped files in SGX?

Enclave



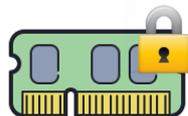
```
fast_read_db():  
ptr = mmap(...);  
data = *ptr;
```

Implementation options:

Use fault handler

Use SGX demand-paging instructions

Use in-enclave handler



Fault handler

Memory-mapped files in SGX?

Enclave



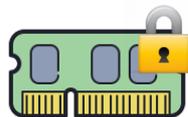
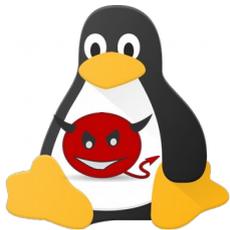
```
fast_read_db():  
ptr = mmap(...);  
data = *ptr;
```

Implementation options:

~~Use fault handler~~

Use SGX demand-paging instructions

Use in-enclave handler



Fault handler

Memory-mapped files in SGX?

Enclave



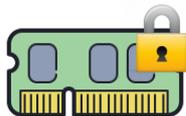
```
fast_read_db():  
ptr = mmap(...);  
data = *ptr;
```

Implementation options:

~~Use fault handler~~

~~Use SGX demand-paging instructions~~

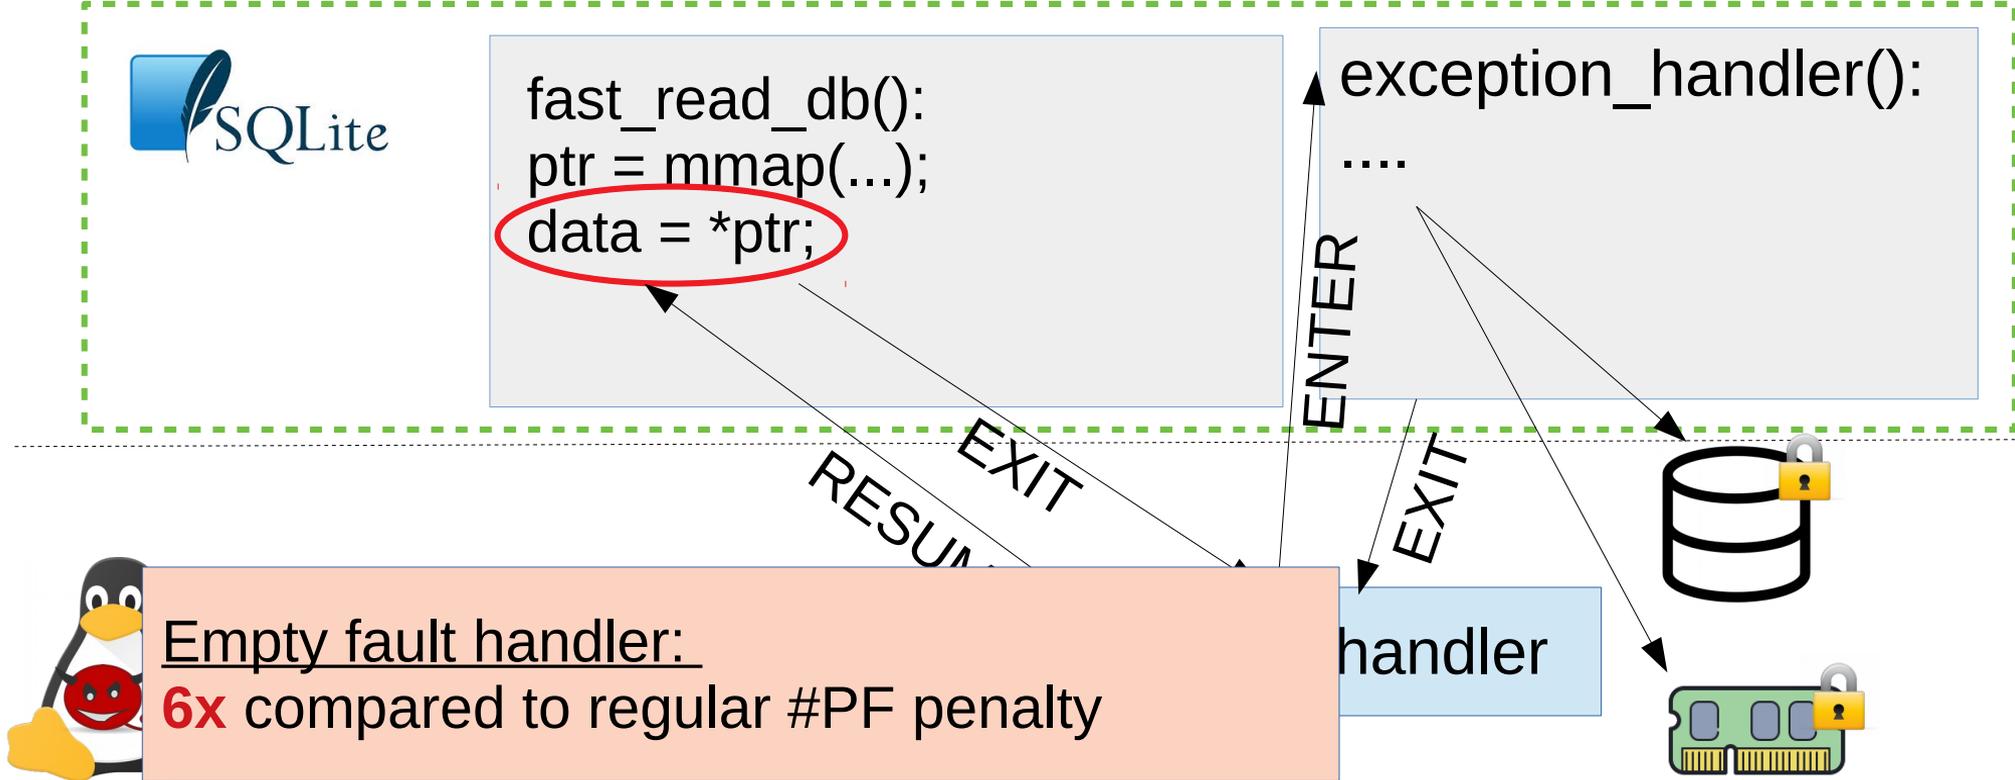
Use in-enclave handler



Fault handler

Inefficient: in-enclave handler

Enclave



Empty fault handler:
6x compared to regular #PF penalty

Insecure: in-enclave handler

Enclave



```
fast_read_db():  
ptr = mmap(...);  
data = *ptr;
```

```
exception_handler():  
....
```

RESUME
EXIT

Fault handler



Enclaves are missing
an OS abstraction!

Previous work sidesteps the lack of secure page faults

- In-enclave paging
- Oblivious page accesses
- Secure access to remote memory



Ad-hoc solutions

Requires applications modifications
What if we want to use more than one?

Efficient page fault customization **missing**

We cannot rely on **hardware**

We cannot **change applications**

So what can we do?

**Automatically change
applicataions with
memory instrumentation**



Agenda

- ~~Motivation~~
- CoSMIX
- Evaluation

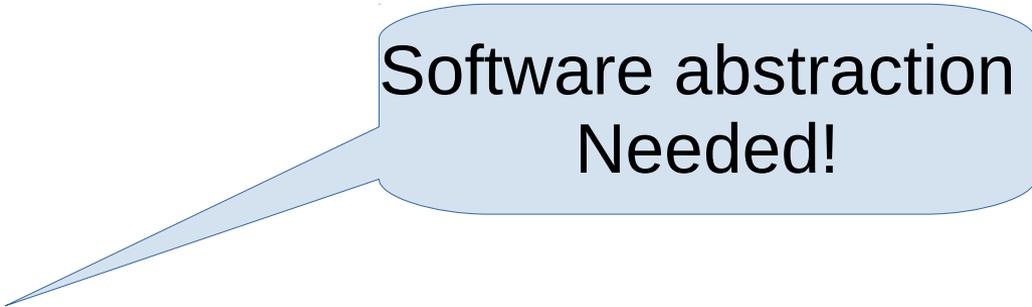
CoSMIX

Efficient **Instrumenting** compiler
and runtime system

What to instrument?

- Observation:
 - No need to instrument every memory access
- Instrumentation policy:
 - Annotations on memory allocations
 - Automatically inferring memory accesses to be instrumented

How to express memory behavior

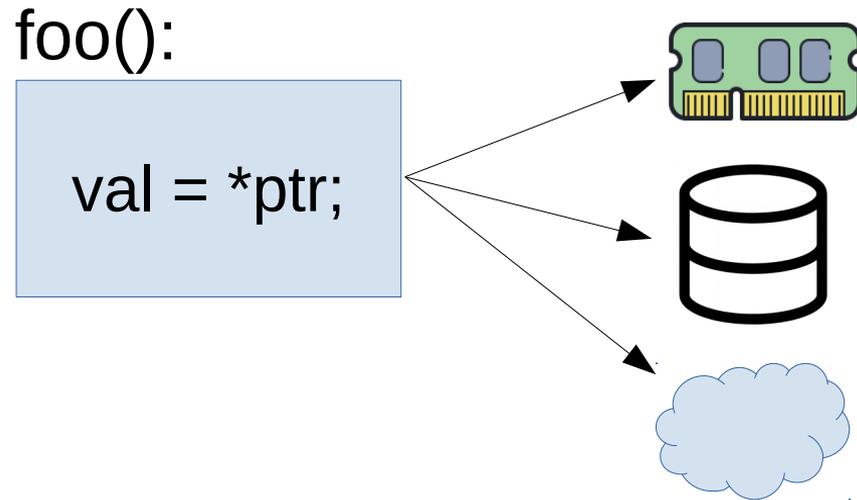


Software abstraction
Needed!

Interface for custom page fault handlers

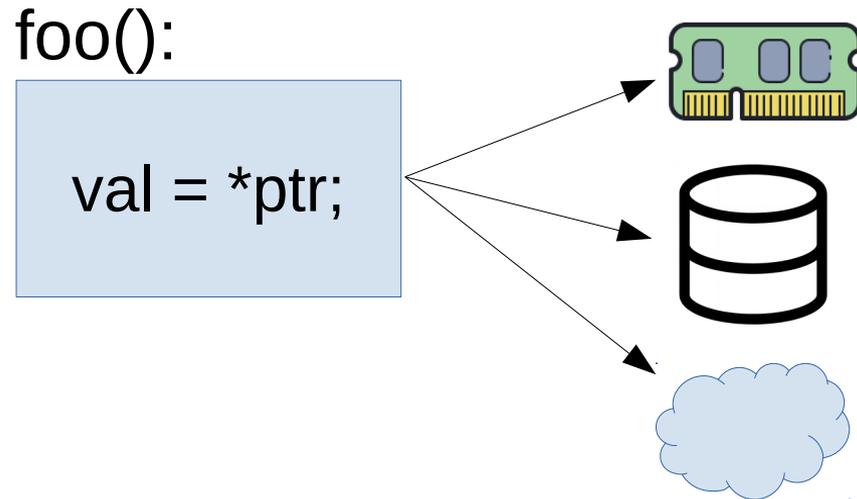
Memory Stores (Mstores)

Another layer of **virtual memory**
on top of an abstract **backing store**



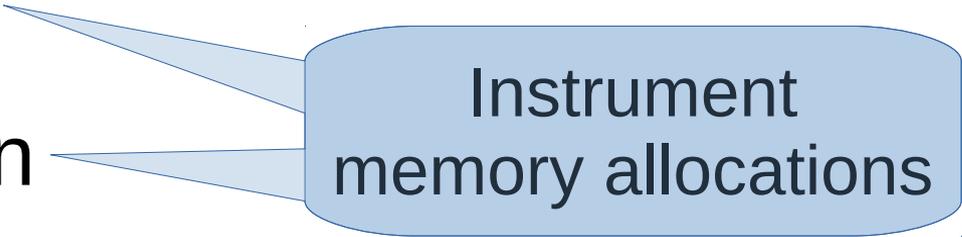
Memory Stores (Mstores)

Another layer of **virtual memory**
on top of an abstract **backing store**

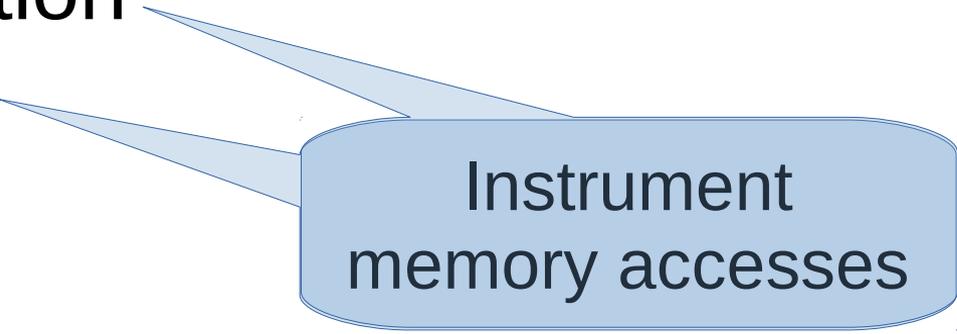


Memory store interface

- Allocation
- Deallocation
- Address translation
- Paging system



Instrument
memory allocations

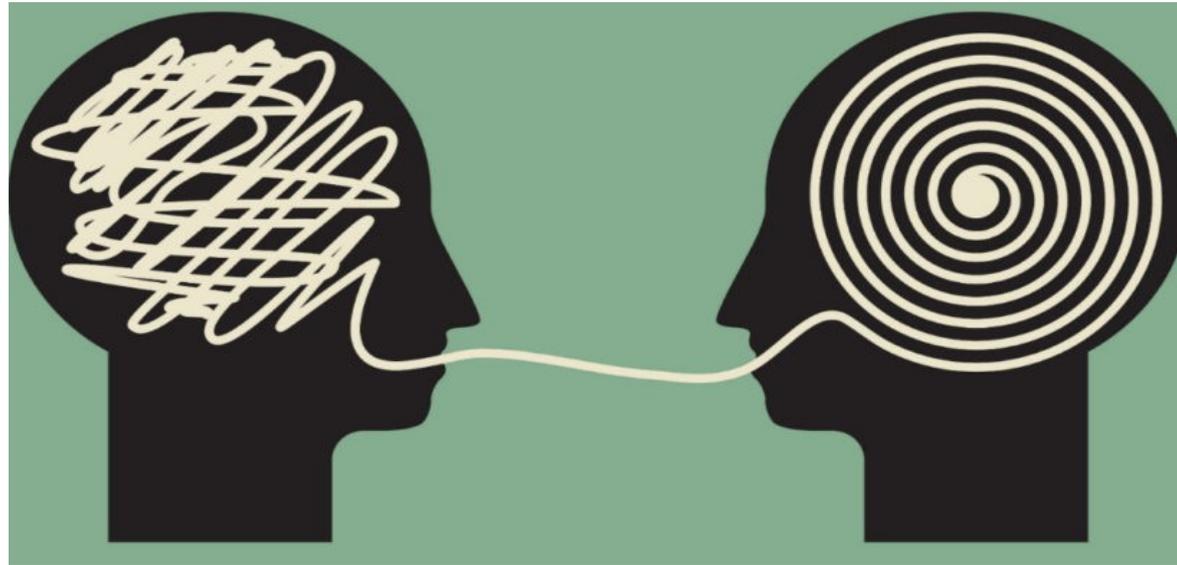


Instrument
memory accesses

Software address translation

Memory store address

Backing store address



More **flexible** than page fault handler customization

Memory store interface

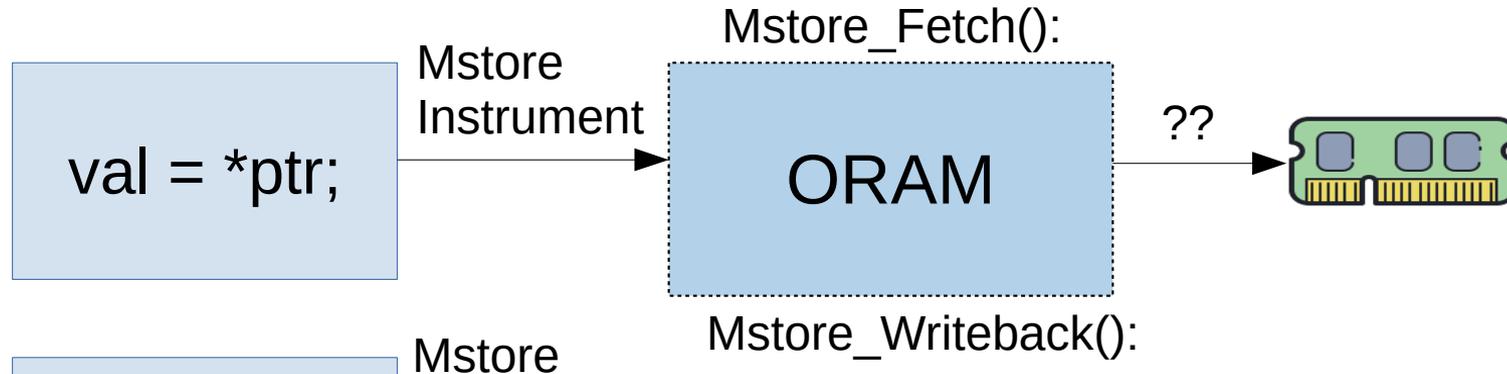
- Allocation
 - Deallocation
 - Address translation
 - Paging system
- 
- Direct-access
Mstore

Use case: Oblivious RAM (ORAM)

- Preserves input-output behavior
- Obfuscates distribution of memory accesses

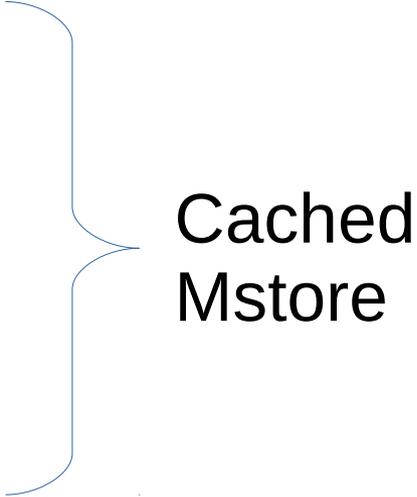


Direct-access memory store

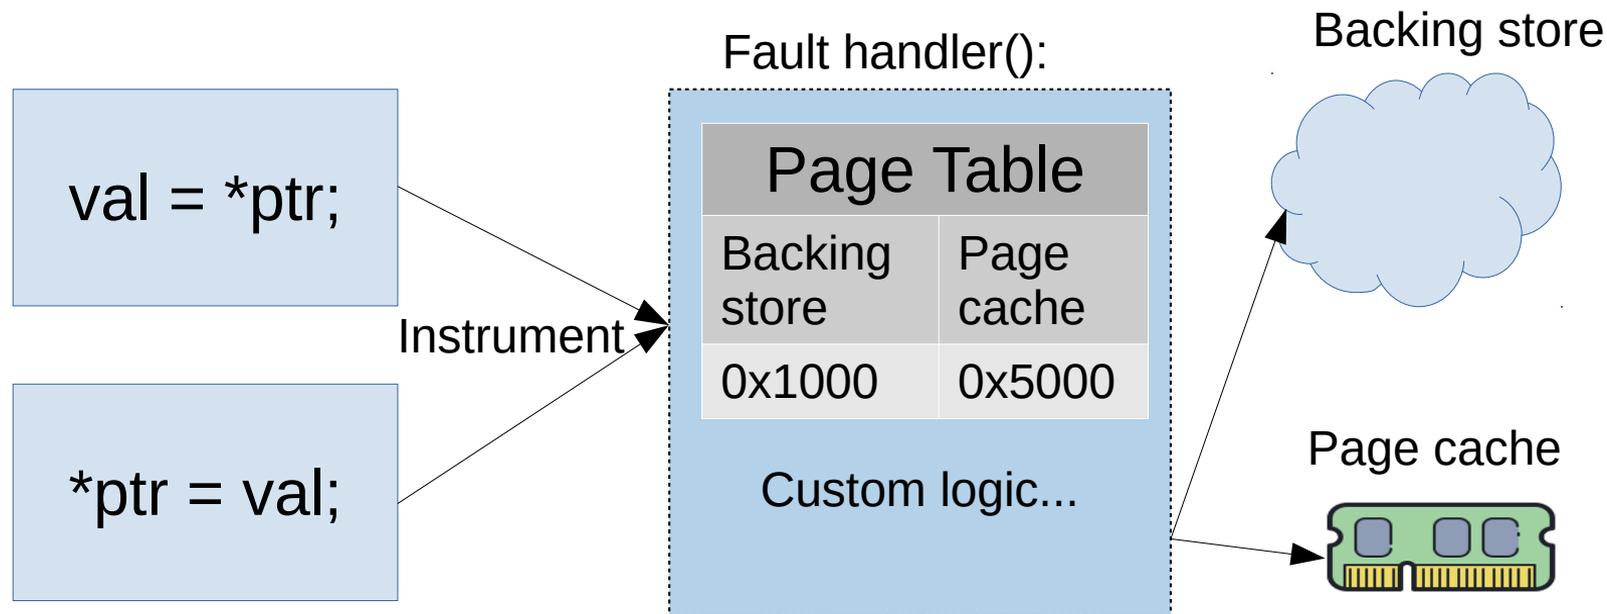


**Previous works required manual modifications
to use ORAM on every access**

Memory store interface

- Allocation
 - Deallocation
 - Address translation
 - Paging system
- 
- Cached
Mstore

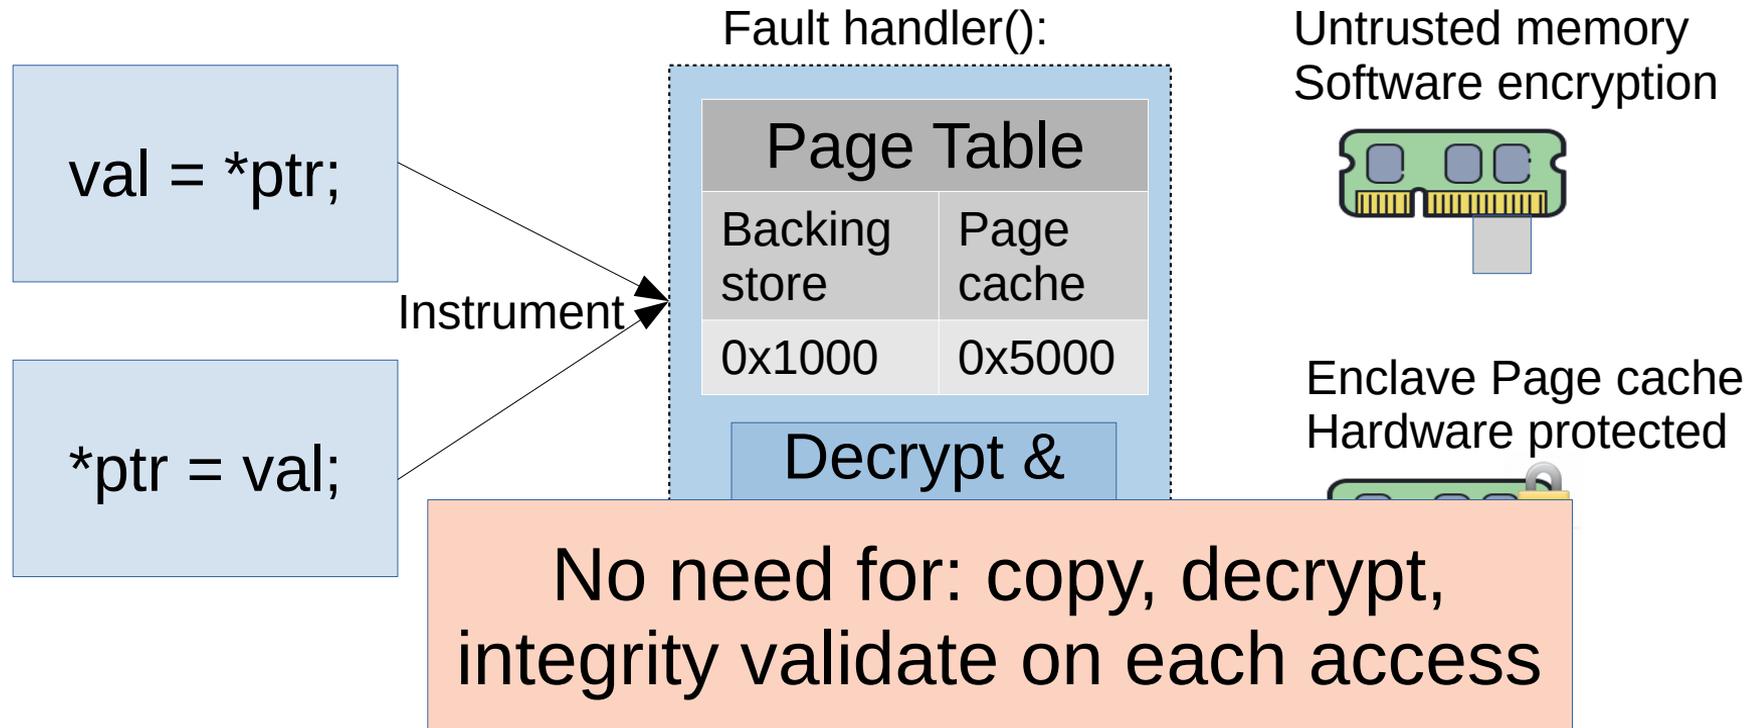
Cached memory stores



Use case: In-enclave demand-paging

- Maintains SGX enclaves security
- Removes costly enclave exits
 - Boost performance
- Previous work required **manual modifications**

Cached memory stores: In-enclave demand-paging



Memory stores are easy

- C/C++
- Small and simple interface: Function callbacks!
- Common building blocks provided

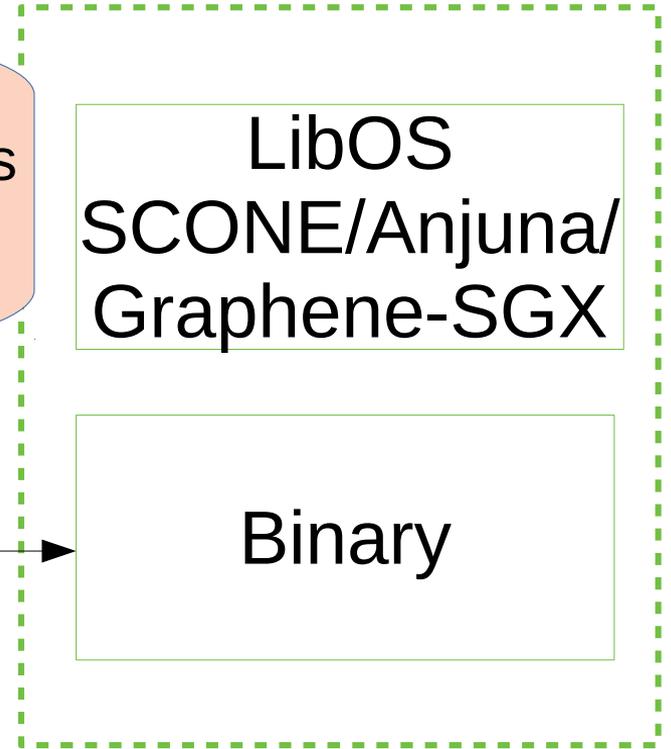
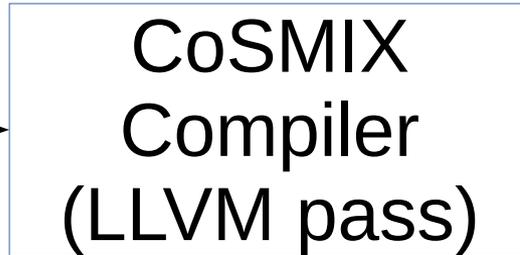
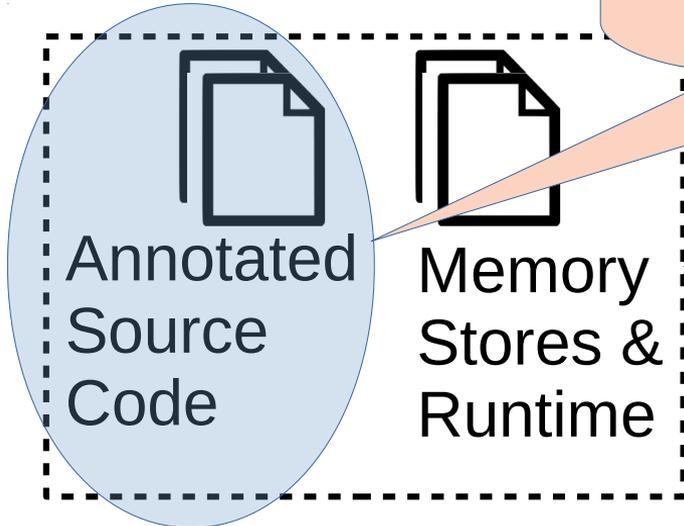


**It took a week to implement
memory-mapped files mstore**

CoSMIX: end-to-end usage



Annotate memory allocations with memory stores

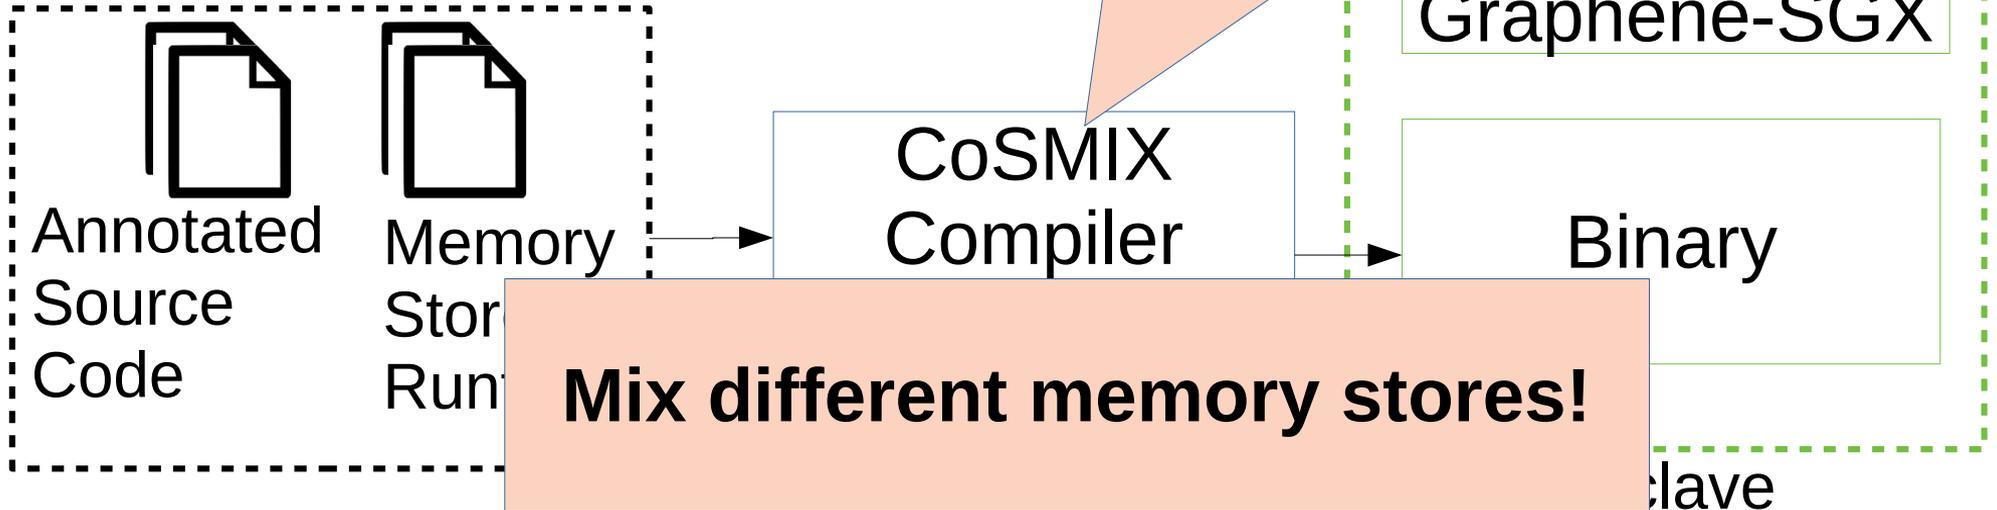


Enclave

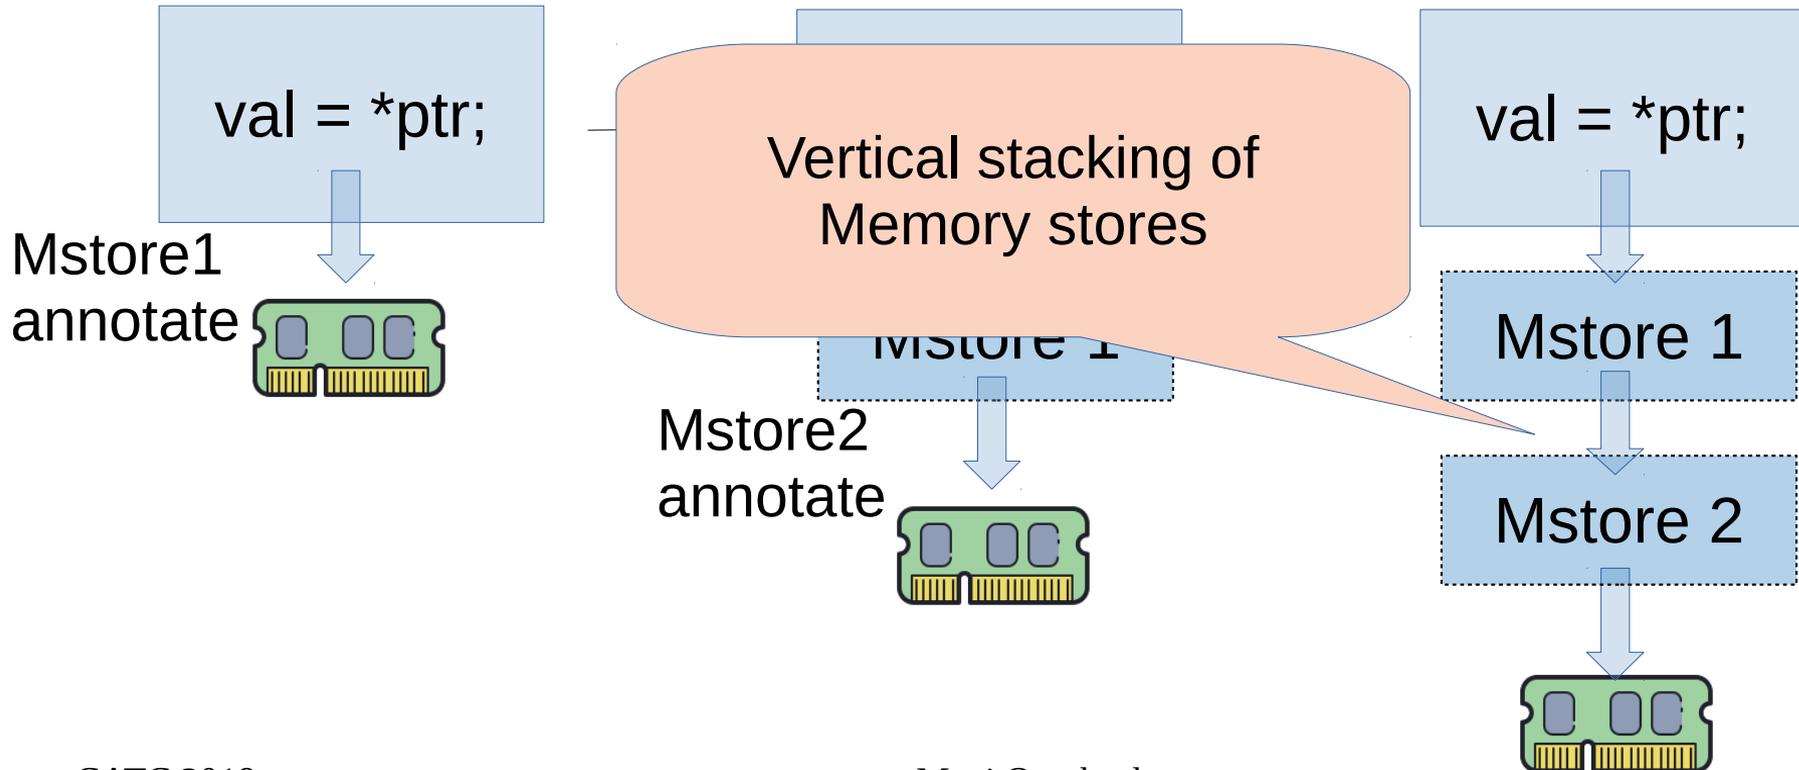
CoSMIX: end-to-end usage



Allocations instrumented to infer memory stores' callbacks for each access



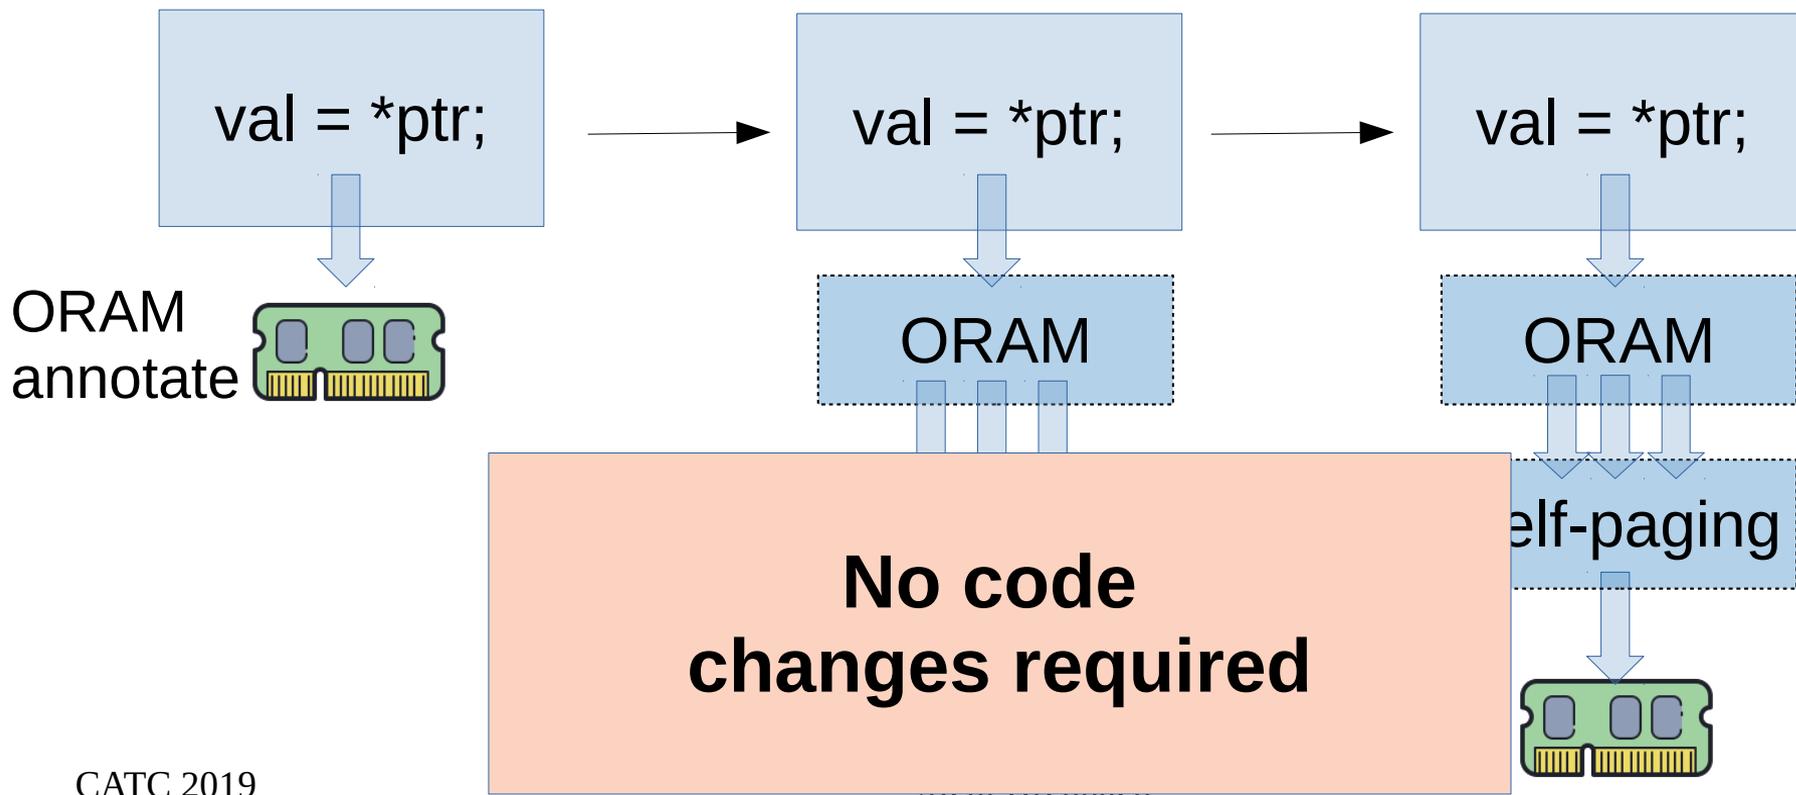
Memory stores are stackable



Stacking ORAM on in-enclave paging

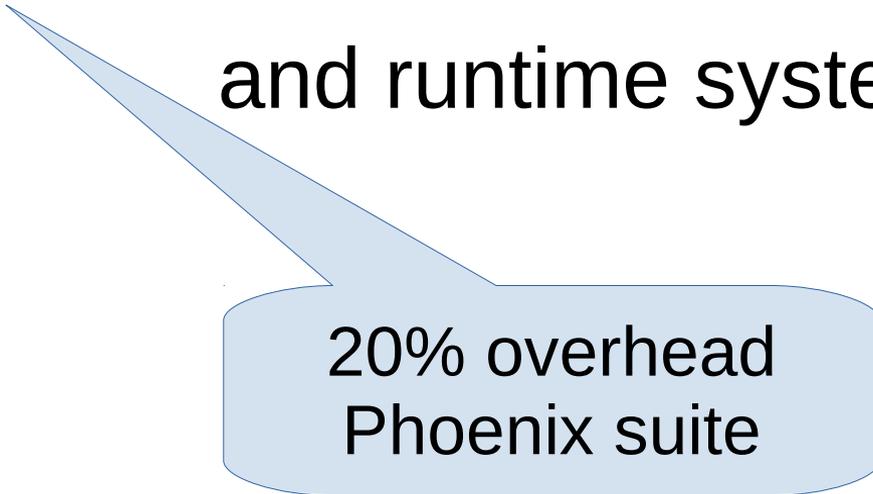
- ORAM performs multiple memory accesses to hide memory access patterns
- ORAM data structures increase its memory footprint
- ORAM mstore in SGX may cause **memory thrashing**

Stacking ORAM on in-enclave paging



CoSMIX

Efficient Instrumenting compiler
and runtime system



20% overhead
Phoenix suite

Efficient instrumentation



- **Selective** instrumentation with **pointer analysis**
- Temporal access locality
 - Software TLB: **Cache address translations**
 - Selective instrumentation: **TLB can be small**
- Intensive memory accesses in loops
 - **Hoist** instrumentation for loop-invariant accesses

Agenda

- ~~Motivation~~
- ~~CoSMIX~~
- Evaluation

Workloads

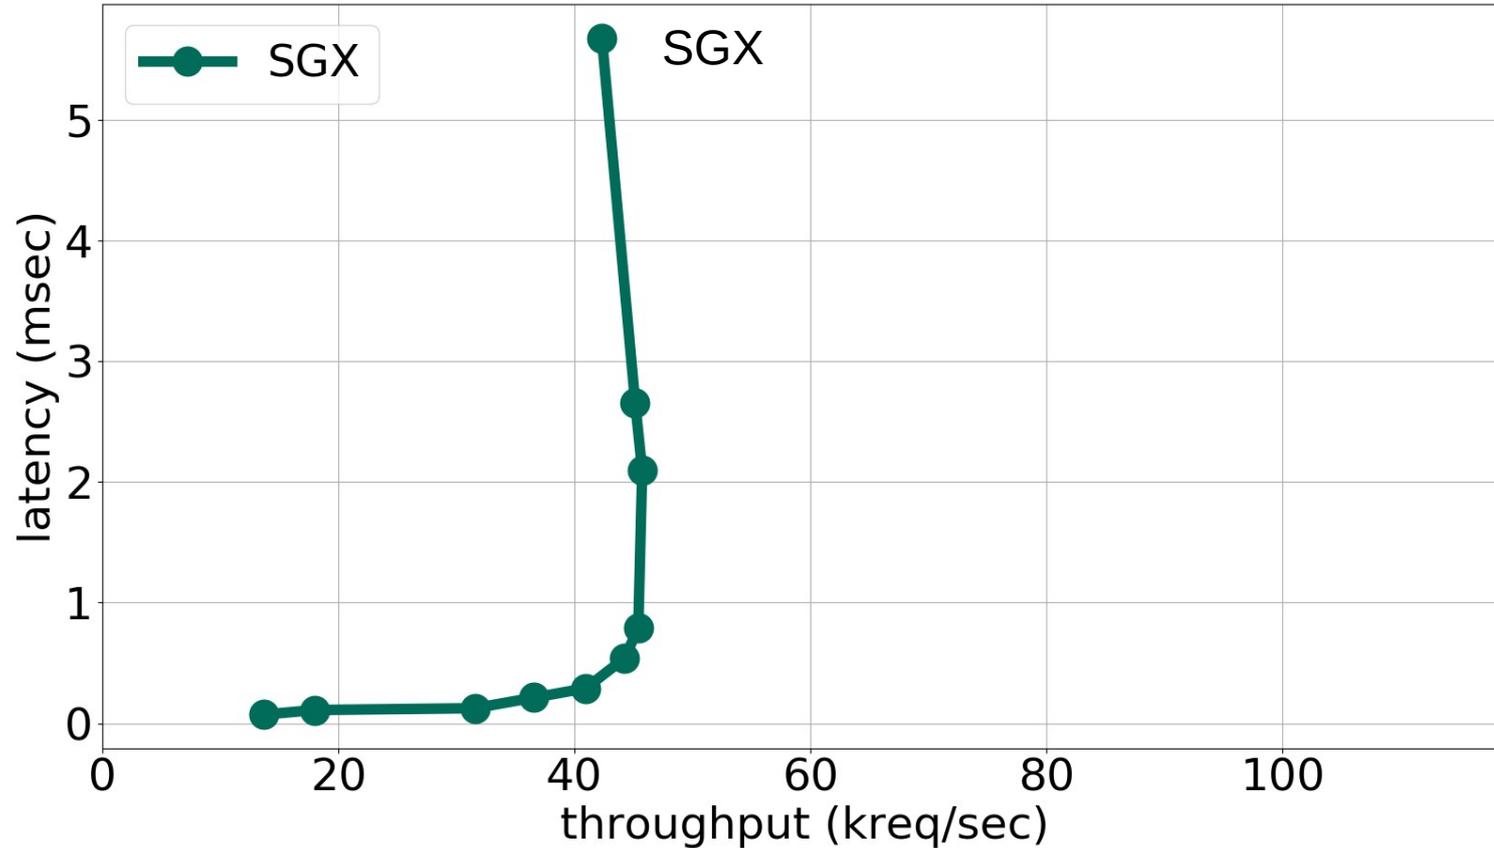
Workload	LOC	Added logic
Memcached	16k	In-enclave paging
REDIS	124k	In-enclave paging
SQLITE	135k	Secured file mapping
Face Verification	700	Oblivious page access (ORAM)

Workloads

Workload	LOC	Added logic
Memcached	16k	In-enclave paging
REDIS	124k	In-enclave paging
SQLITE	135k	Secured file mapping
Face Verification	700	Oblivious page access (ORAM)

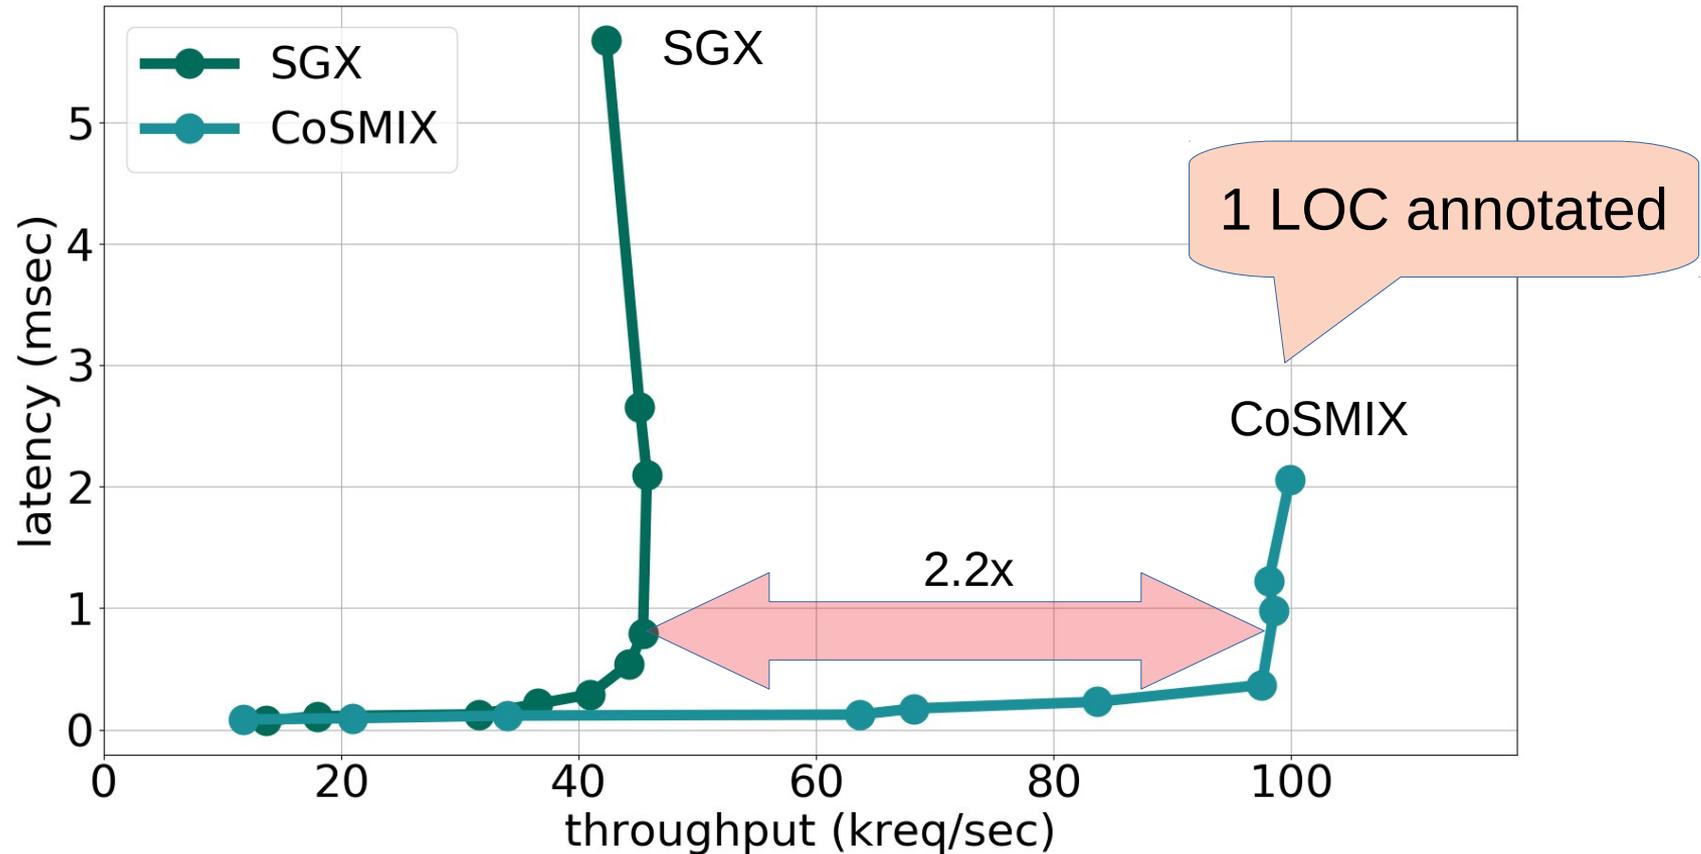
Memcached (600 MB)

Random access 1KB 10% SET



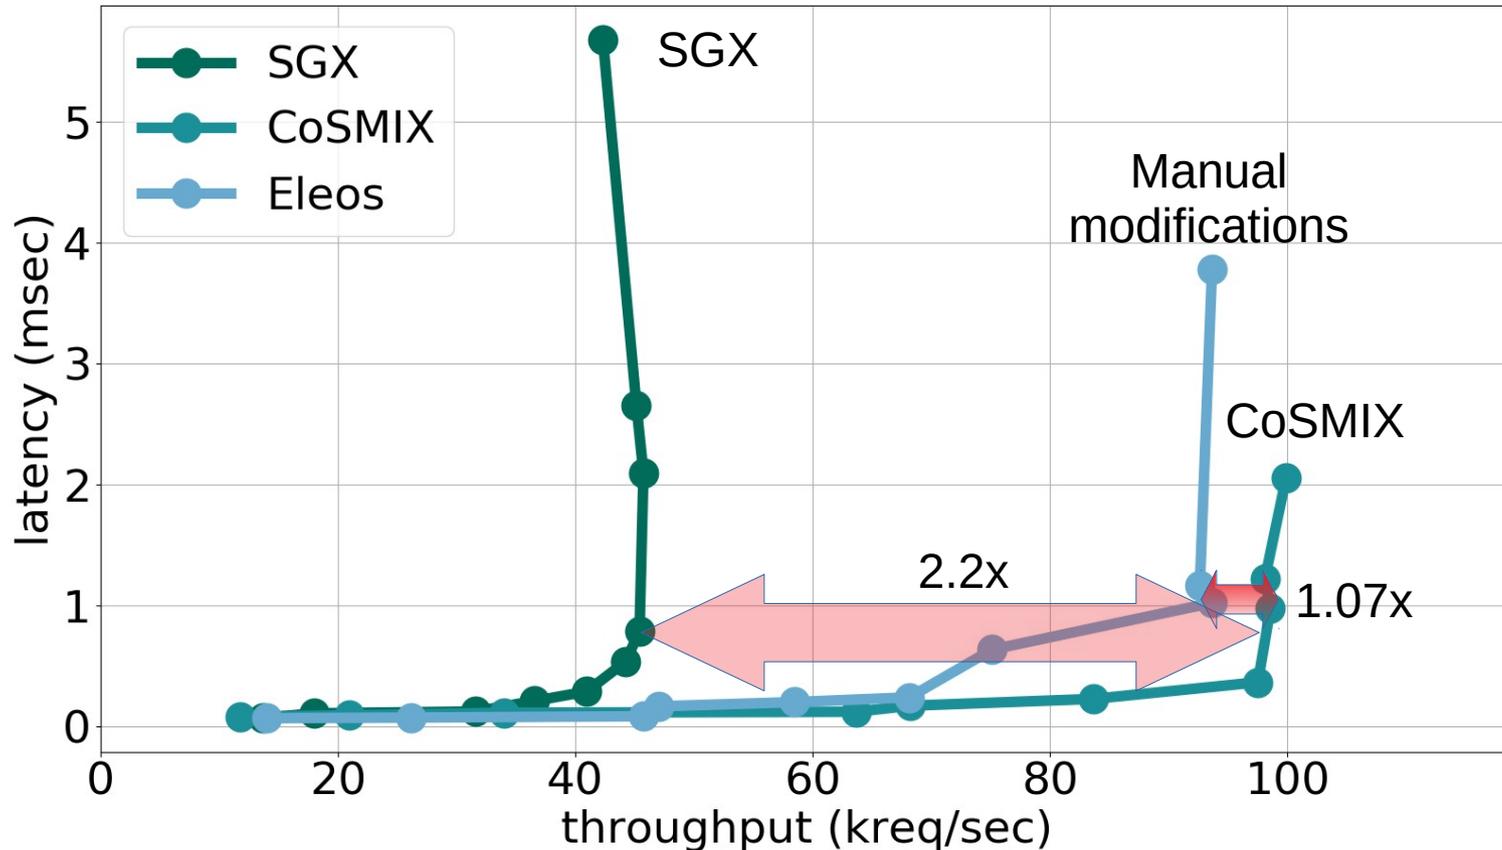
Memcached (600 MB)

Random access 1KB 10% SET



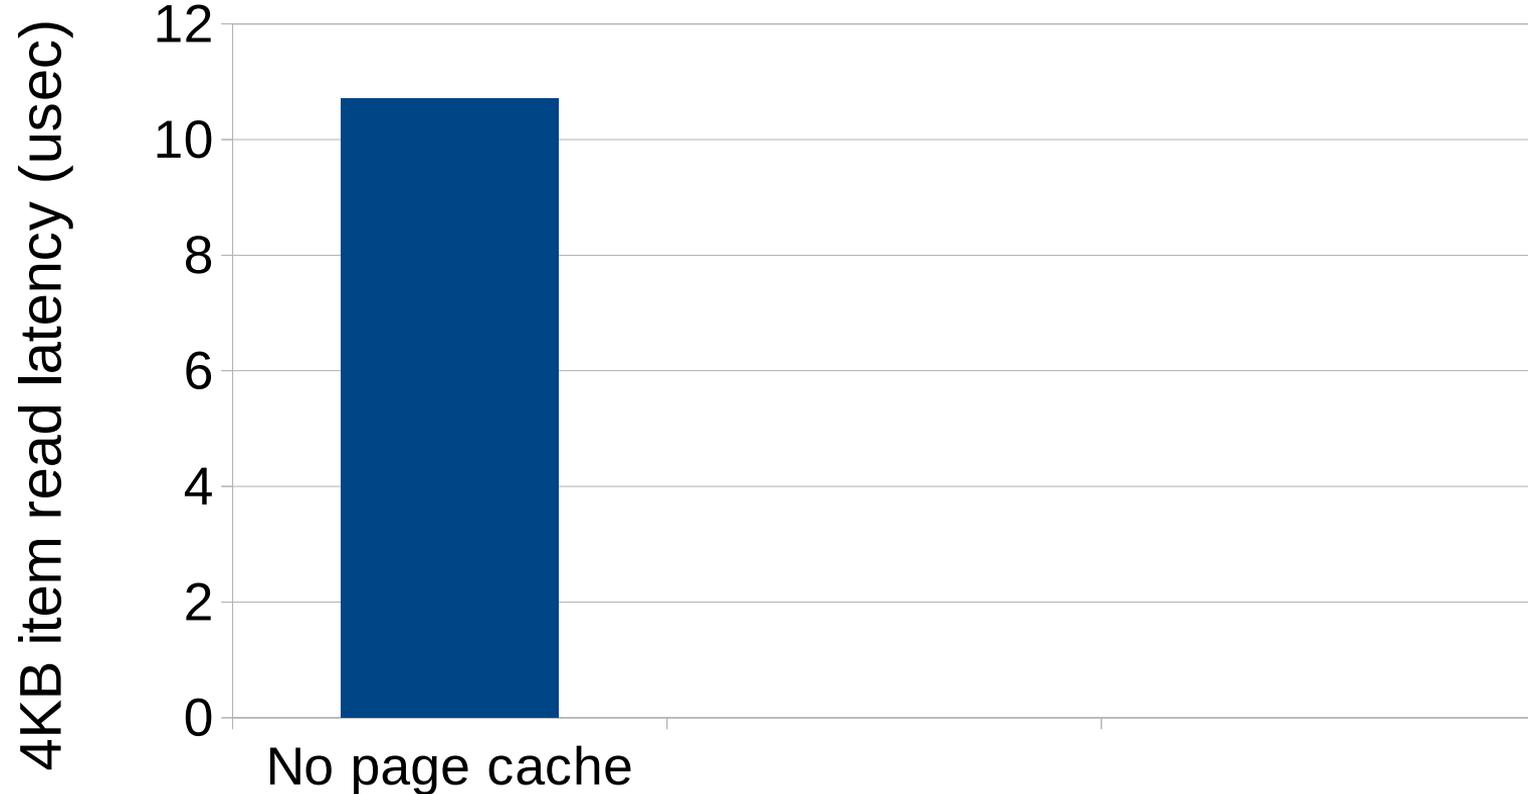
Memcached (600 MB)

Random access 1KB 10% SET



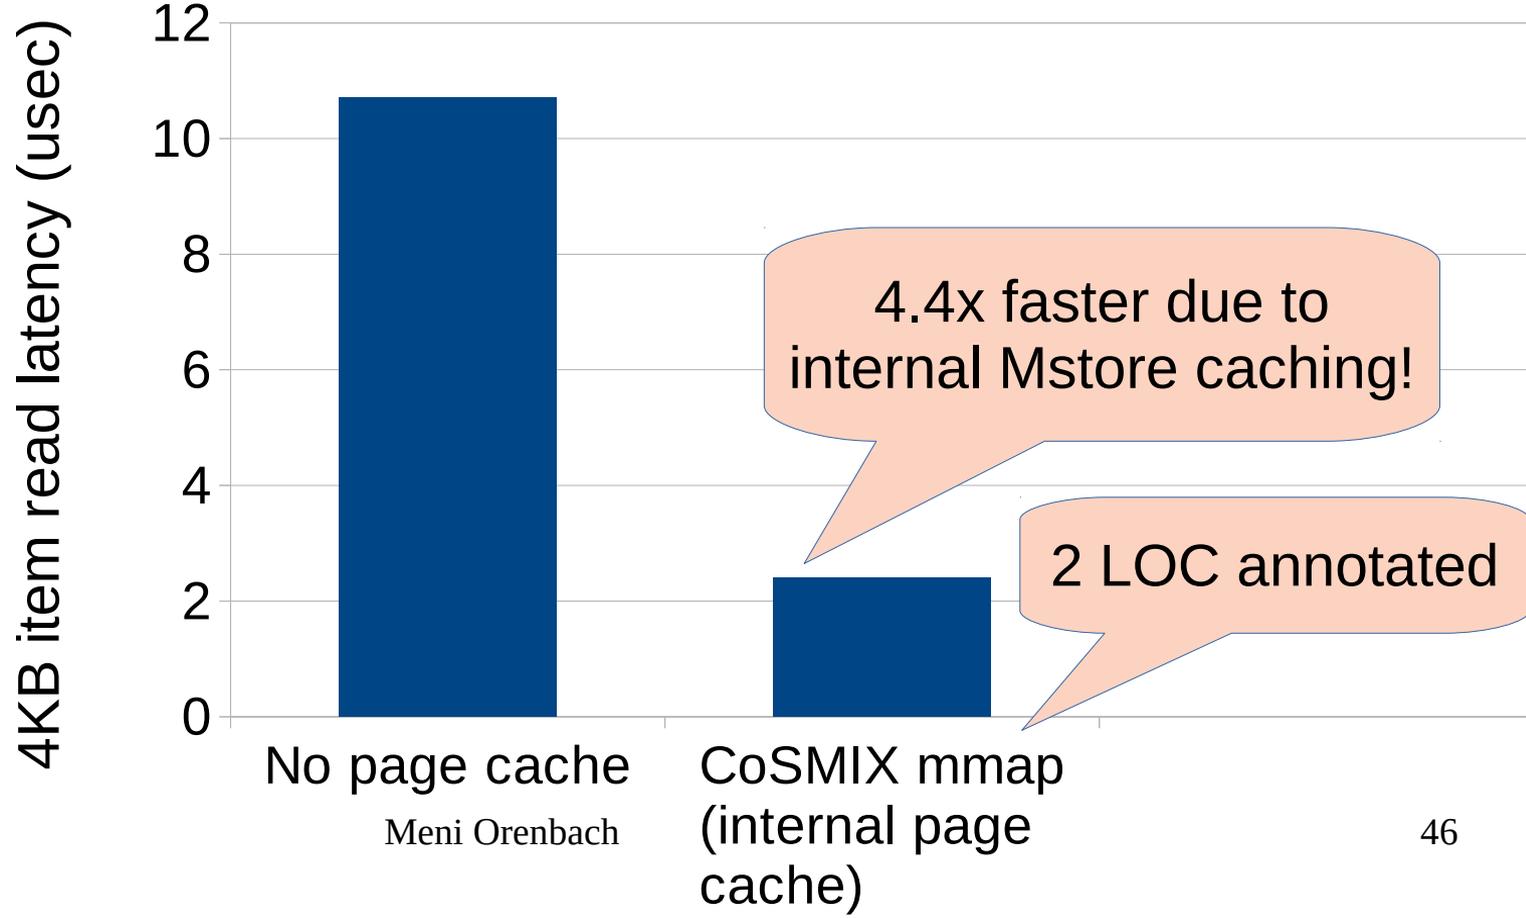
SQLite Encrypted DB file kvtest Random access

Lower is better!



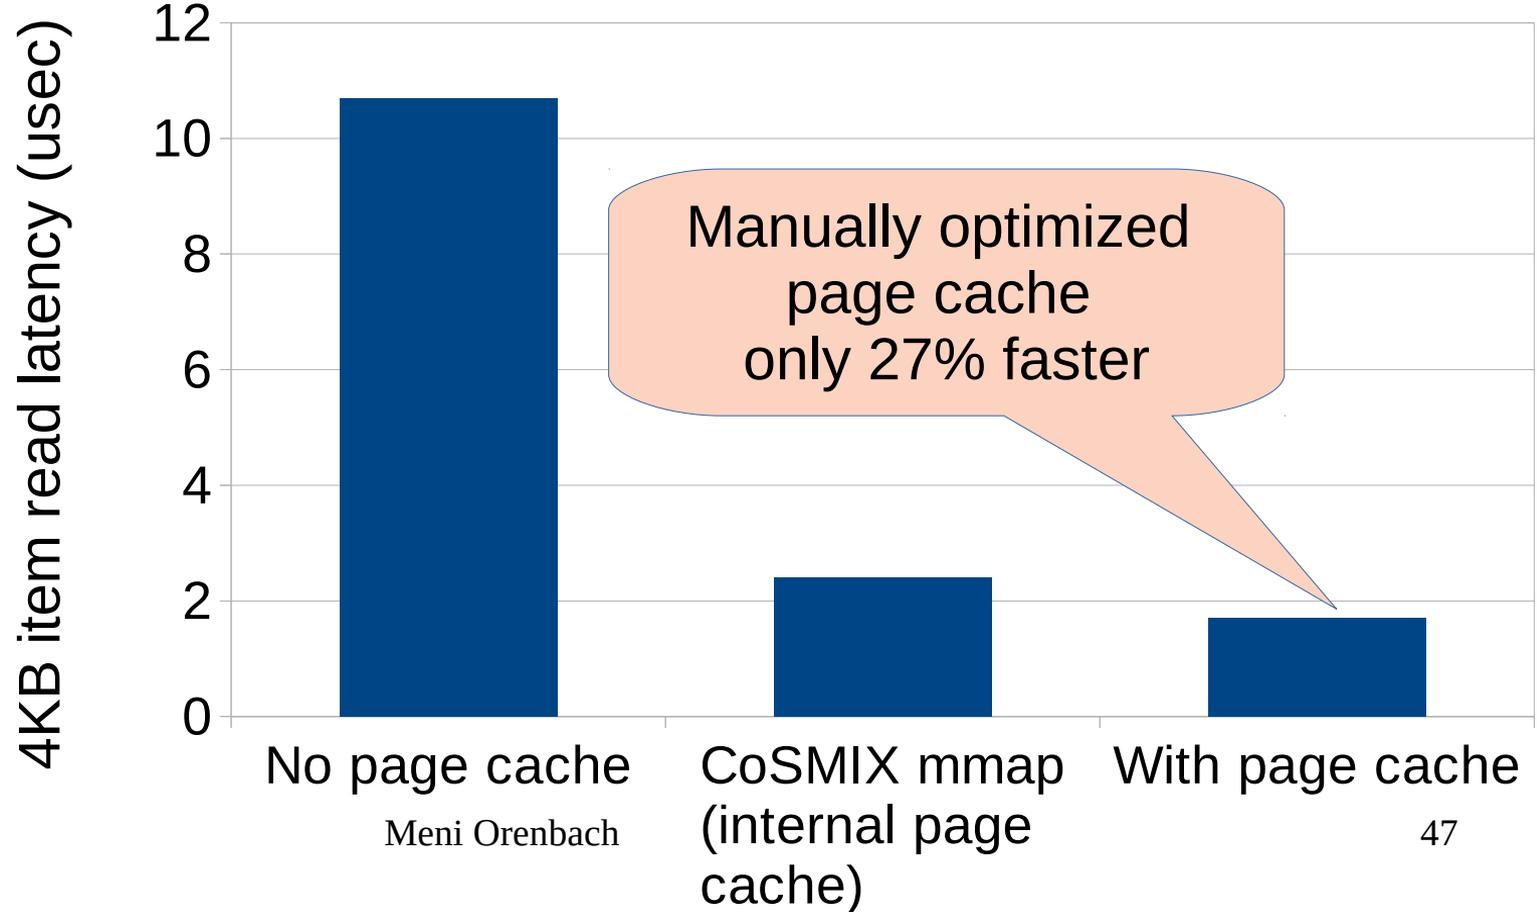
SQLite Encrypted DB file kvtest Random access

Lower is better!



SQLite Encrypted DB file kvtest Random access

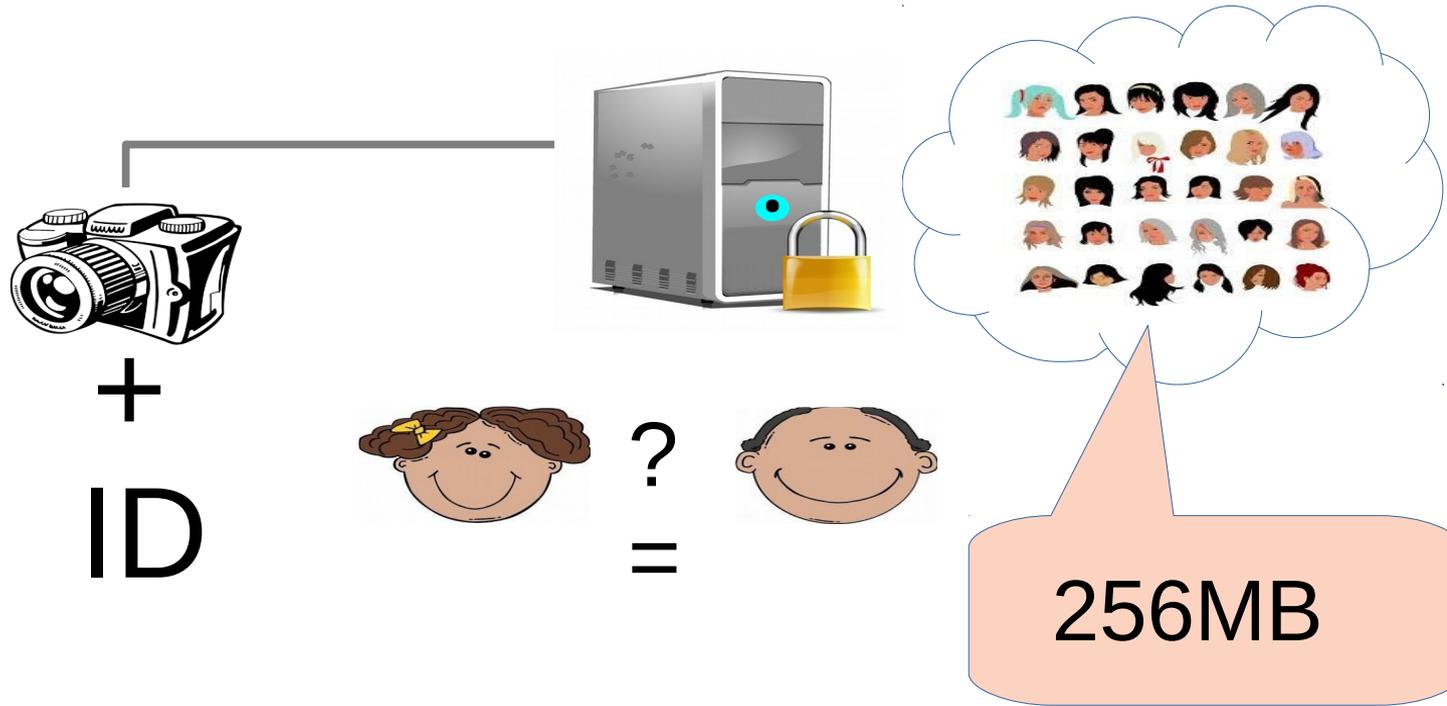
Lower is better!



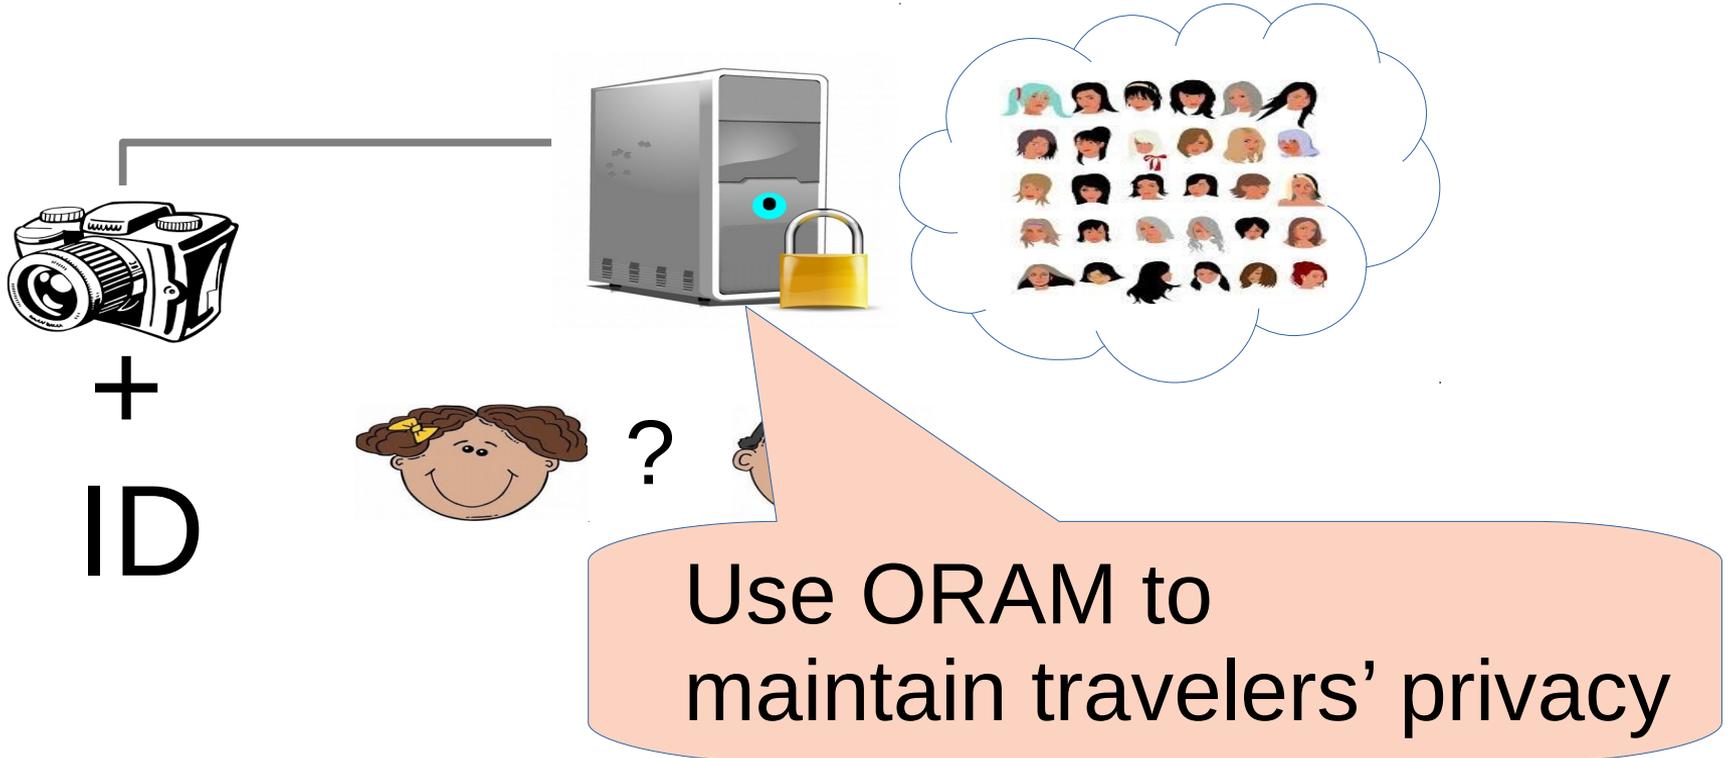
Biometric Identity checking server



Biometric Identity checking server



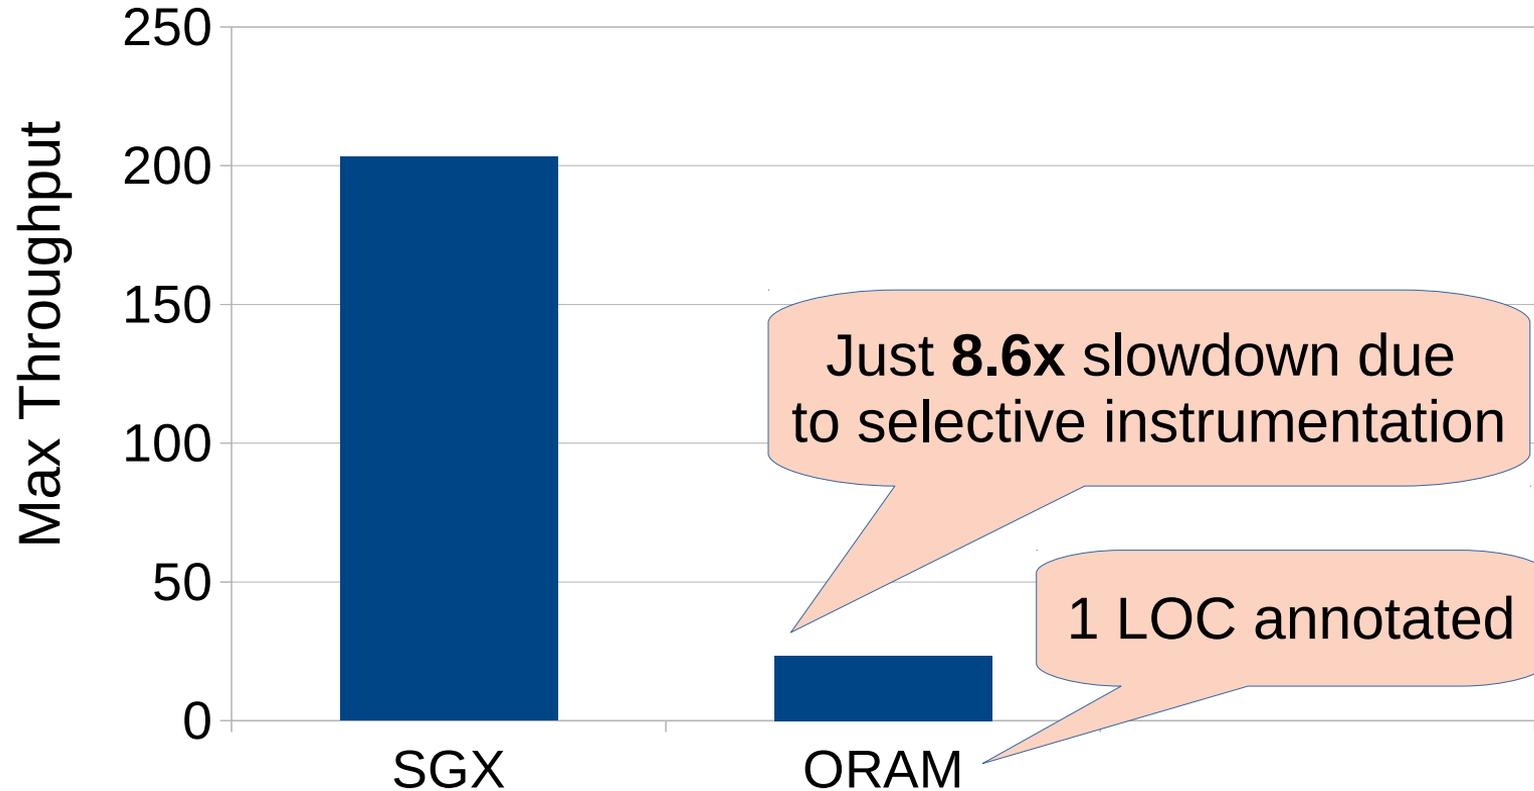
Biometric Identity checking server



ORAM notoriously inefficient

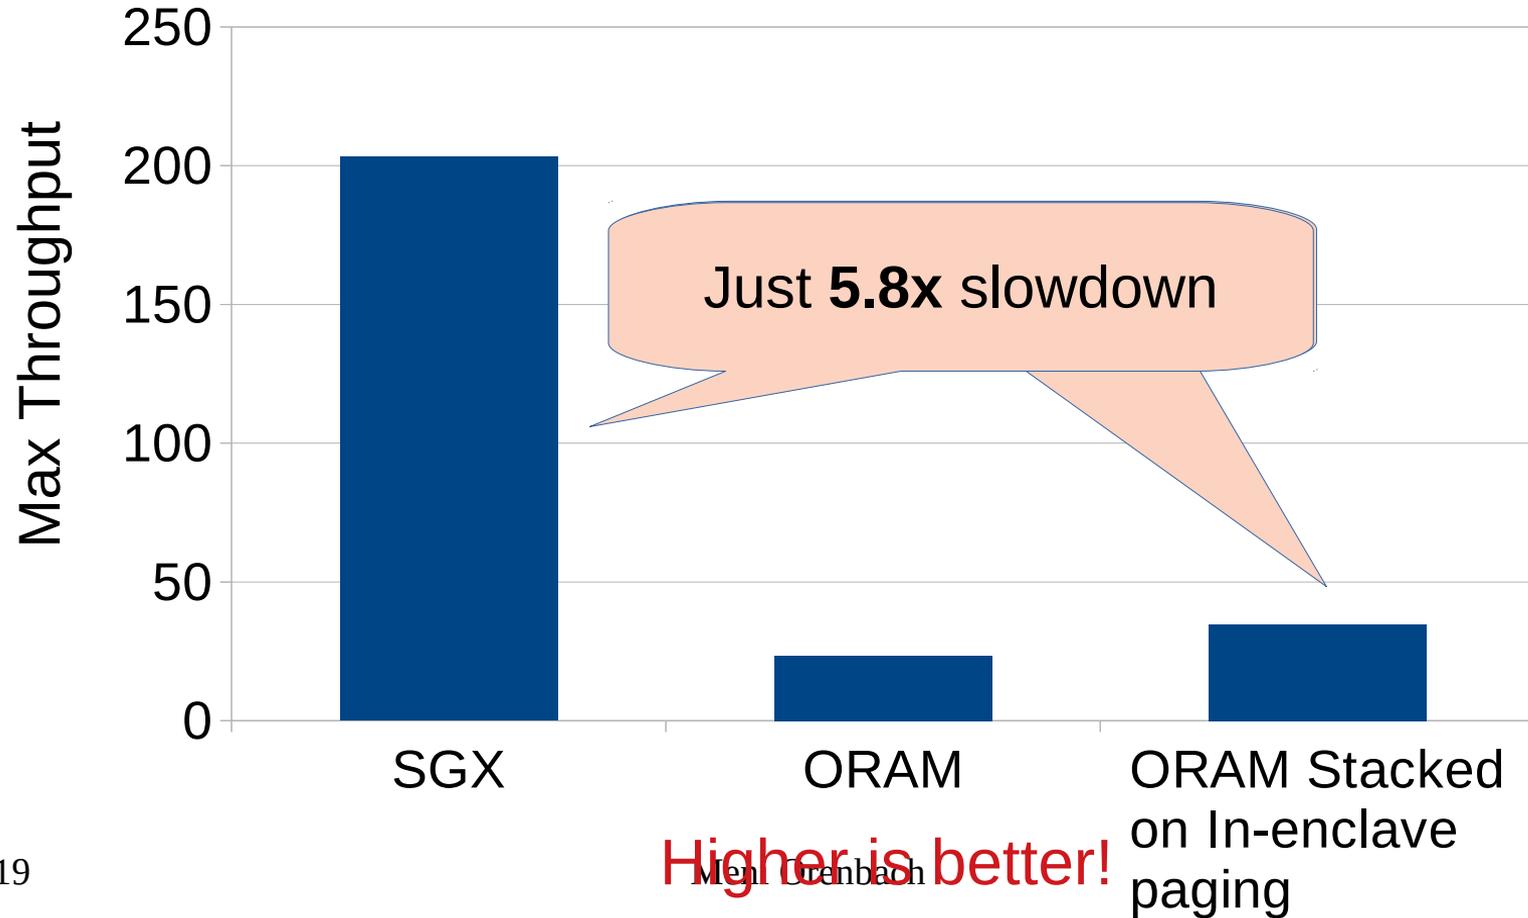
61x slowdown for memory accesses

Face verification



Higher is better!

Face verification



Conclusion

- Hardware enclaves improve the security of applications in public cloud environments
- Hardware **limits the use** of a trusted page fault handler
- **CoSMIX** provides trusted page fault handling today
 - Efficient instrumentation makes this a feasible approach

Thank You!



<https://github.com/acsl-technion/cosmix>