Mon CHÉRI ❤️

Mitigating Uninitialized Memory Access with Conditional Capabilities

Merve Gülmez, Håkan Englund, Jan Tobias Mühlberg, Thomas Nyman

ULB UNIVERSITÉ LIBRE DE BRUXELLES

KU LEUVEN

CHERI-extension for **conditional capabilities** preventing **uninitialized memory accesses**

Memory-safe hardware, such as **CHERI**, is attractive for addressing **memory-safety** for code-bases in unsafe languages, e.g., C and C++

- CHERI addresses spatial and temporal safety
- Other aspects, e.g., **initialization safety**, is not addressed by CHERI

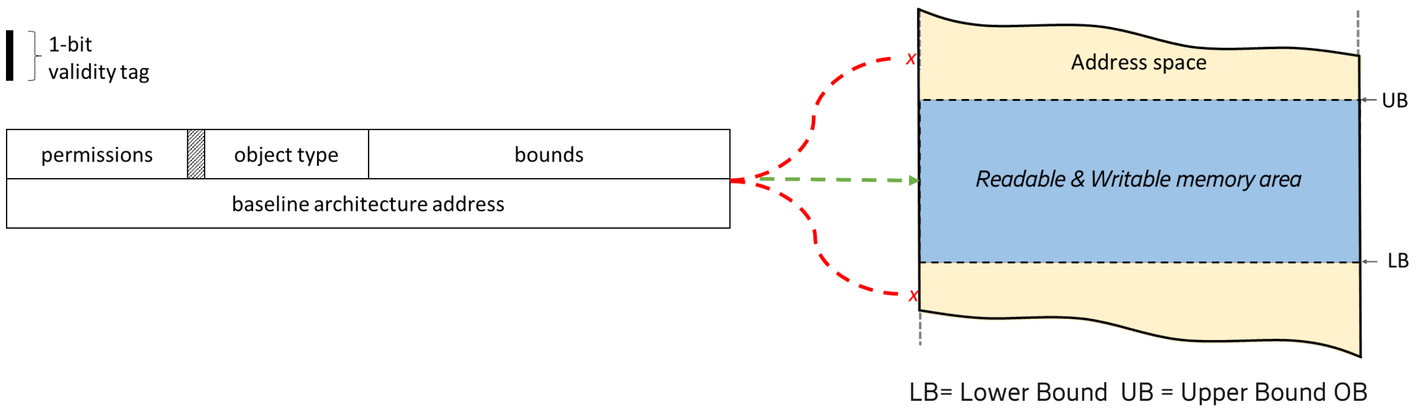

LB= Lower Bound UB = Upper Bound OB

Each CHERI capability is twice the size of a native pointer and includes memory address plus metadata: permissions, object type, and bounds

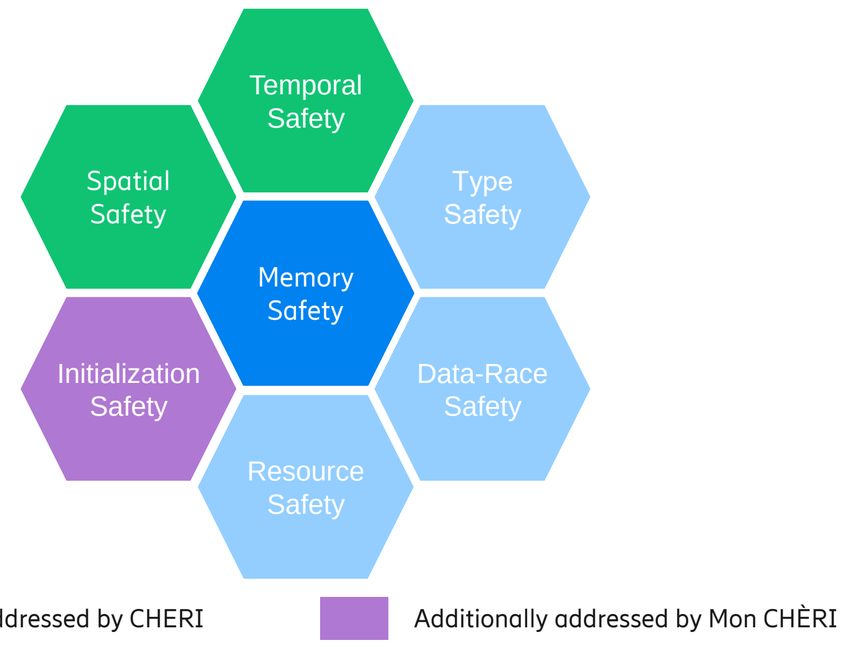

Use-before-initialized conditions the fourth largest class of memory-safety vulnerabilities (~10%)

Mon CHÉRI is a further hardware extension to CHERI capabilities that enables memory access control to take previous operations on memory into account

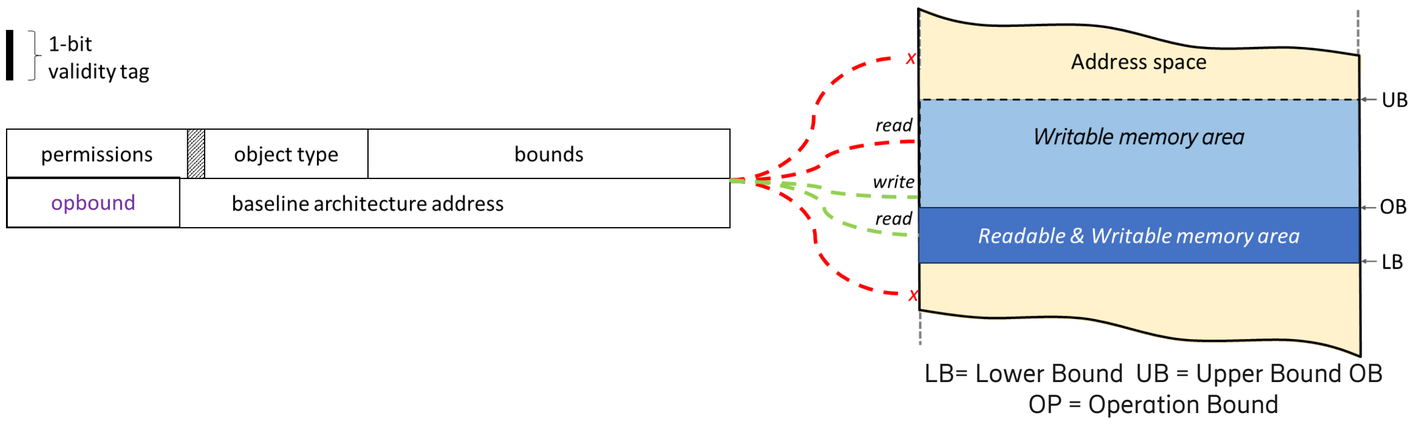

LB= Lower Bound UB = Upper Bound OB
OP = Operation Bound

Mon CHÉRI enables novel policies on memory such as:

- **Write-before-Read** – memory must be written to at least once before reading – initialization safety
- **Write-before-Read-Only** – `const` enforced by hardware
- **Write-before-Execute-Only** – emulation of XOM (eXecute Only Memory) using capabilities

Conditional capabilities enforce conditional permissions

Operation bound tracks area for which condition is fulfilled

- Operation bound compressed into top 16 bits of address

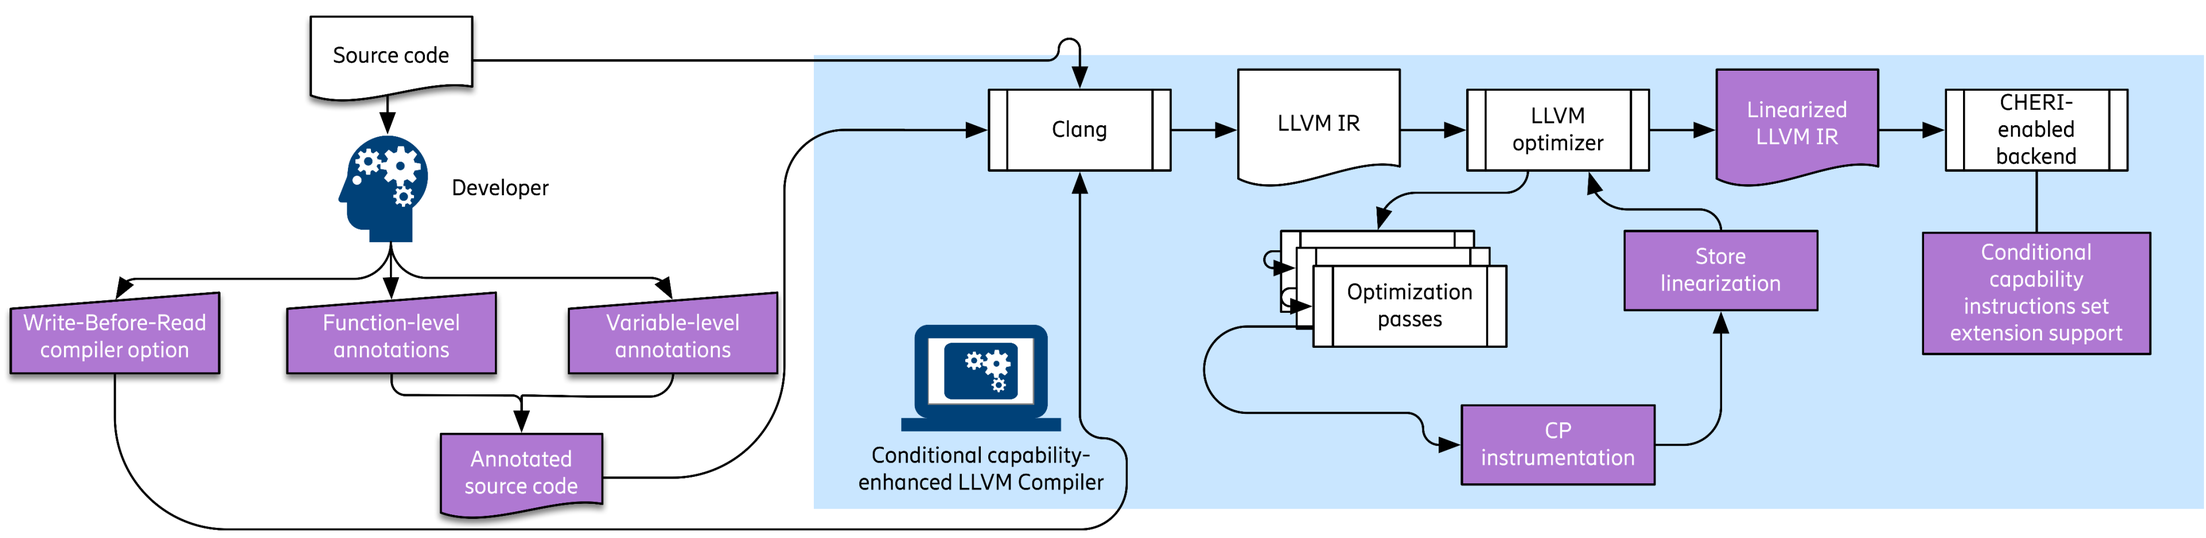

High-level overview of the conditional capability-enhanced LLVM compiler

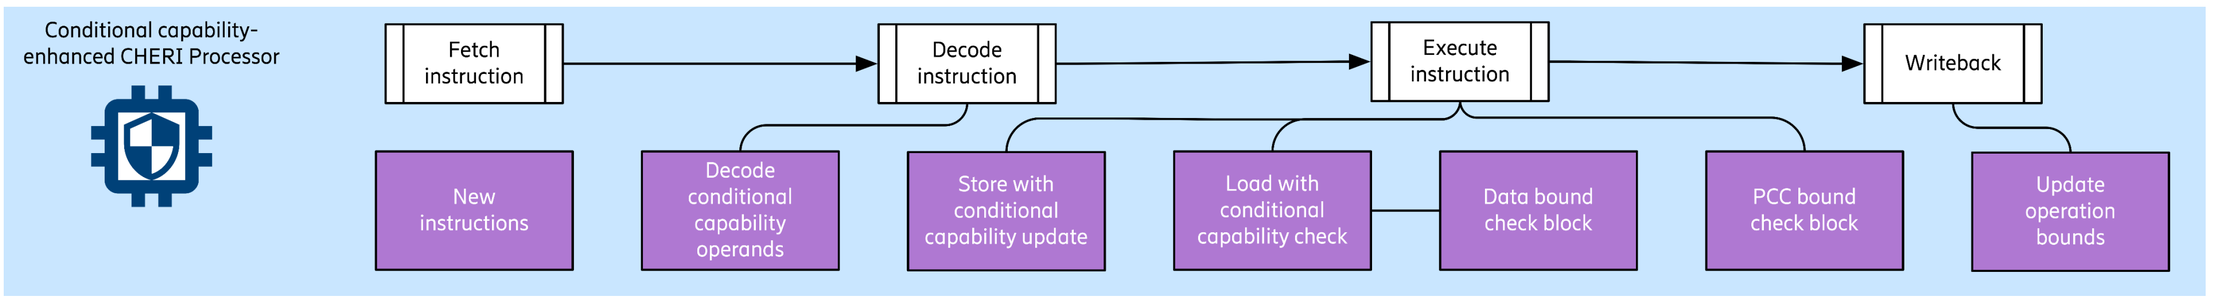

Conditional capability-enhanced CHERI processor

Evaluated using **MonCHÉRI-Flute FPGA softcore** and **QEMU-based full-system simulation** for performance impact and accuracy:

- Coremark performance: **3.5%** overhead over CHERI pure-capability mode
- Juliet Test Suite accuracy: **100%** detection rate with ≈ 1% false positives

	Positive	Negative
Bad	560 (100%)	0 (0%)
Good	6 (1%)*	554 (99%)

*) False negative cases exhibit uninitialized behavior, but such behavior that is benign from a security point of view