

# CHERI-SIMT: Implementing Capability Memory Protection in GPUs

Matthew Naylor

University of Cambridge  
Cambridge, United Kingdom

Alexandre Joannou

University of Cambridge  
Cambridge, United Kingdom

A. Theodore Markettos

University of Cambridge  
Cambridge, United Kingdom

Paul Metzger

University of Cambridge  
Cambridge, United Kingdom

Simon W. Moore

University of Cambridge  
Cambridge, United Kingdom

Timothy M. Jones

University of Cambridge  
Cambridge, United Kingdom

## Abstract

Governments are increasingly advising software manufacturers to employ memory-safe languages and technologies to combat adversarial attacks on modern computing infrastructure. This introduces pressures across the entire computing industry, including GPU vendors who provide implementations of unsafe C/C++-based languages, such as CUDA and OpenCL, for programming the devices they produce. One of the memory-safety technologies being recommended is *Capability Hardware Enhanced RISC Instructions* (CHERI). CHERI builds strong and efficient memory safety into underlying instruction-set architectures allowing continued, but memory-safe, use of C/C++-based languages on top.

In this paper, we evaluate the feasibility of incorporating CHERI into GPU architectures by extending a prototype, open-source, synthesisable, SIMT core and CUDA-like programming environment with support for CHERI. We present techniques to considerably ameliorate the costs of CHERI in SIMT designs, reducing register-file storage overheads from 103% to 7%, logic-area overheads by 44% to a cost comparable to one additional multiplier per vector lane, and execution-time overheads to 1.6%. With the proposed techniques, CHERI offers a viable path to strong and efficient GPU memory safety, while avoiding the need to replace established programming practices.

**CCS Concepts:** • Computer systems organization → Parallel architectures; • Security and privacy → Hardware security implementation; • Software and its engineering → Parallel programming languages.

**Keywords:** Memory Safety; CHERI; GPU; SIMT

## ACM Reference Format:

Matthew Naylor, Alexandre Joannou, A. Theodore Markettos, Paul Metzger, Simon W. Moore, and Timothy M. Jones. 2026. CHERI-SIMT: Implementing Capability Memory Protection in GPUs. In *Proceedings of the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1 (ASPLOS '26)*, March 22–26, 2026, Pittsburgh, PA, USA. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3760250.3762234>

## 1 Introduction

Recent studies report that memory-safety bugs account for around 70% of security vulnerabilities in major software projects [13, 50, 68]. These bugs typically arise from mistakes involving C/C++ pointers, such as out-of-bounds accesses (violating *spatial* memory safety) and use-after-free errors (violating *temporal* memory safety). In response, government agencies around the world are increasingly advising software manufacturers to employ memory-safe languages and technologies, as set out in recent joint whitepapers [14, 15]. There are also calls to standardise principles and practices for software memory safety, allowing future systems to comply with agreed criteria for preventing this whole class of vulnerabilities [61, 70].

While the above-mentioned studies focus on CPU codebases, GPU codebases are also commonly written in C/C++-based languages, such as CUDA and OpenCL, and have been found to suffer from many of the same problems. For example, Erb et al. [23] report 13 cases of buffer overflows in a set of 175 GPU applications taken from standard benchmark suites. Separately, researchers have shown that buffer overflows on GPUs can lead to data corruption on the stack and heap, control-flow hijacking, code injection, and arbitrary code execution [22, 30, 49, 60]. A simple example of a buffer overflow in CUDA is shown in Figure 1.

Although GPU code is generally less security critical than CPU code, the distinction between the two is becoming increasingly blurred. For example, NVIDIA's `nvc++` compiler allows GPU code to transparently access the memory of the CPU-side process that invoked it through a unified address space [59], and a single pragma around a CPU-side loop is all that's needed to offload it to the GPU. This brings great



This work is licensed under a Creative Commons Attribution 4.0 International License.

ASPLOS '26, Pittsburgh, PA, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2165-6/2026/03

<https://doi.org/10.1145/3760250.3762234>

```

__global__ void overread() {
    int data = 0xdata;
    int secret = 0xc0de;
    int *ptr = &data;
    printf("%x\n", ptr[1]);
}

```

**Figure 1.** Simple example of a buffer overread in a CUDA GPU kernel. In practice, this function prints the value of `secret`. Although `ptr` points to `data`, it is accessed out of bounds to obtain the value of a different variable.

flexibility to programmers, but equally great privilege to the GPU.

One of the memory-safety technologies being mentioned both in the joint government whitepapers and the memory-safety standardisation call is *Capability Hardware Enhanced RISC Instructions* (CHERI) [74]. CHERI enhances existing instruction-set architectures by replacing integer memory addresses with *capabilities*. Capabilities augment integer addresses with *metadata* including bounds and permissions, enforcing constrained access to bounded memory regions. In addition to providing deterministic spatial memory safety using bounds, they also provide *referential integrity* (they cannot be forged, corrupted, or confused with non-capabilities), which in turn supports temporal memory safety [25, 26] as well as efficient software compartmentalisation [3, 71], limiting what attackers can do if they gain access through other kinds of vulnerabilities. A key attraction of CHERI is that pointers in historically unsafe C/C++ can be automatically compiled down to capabilities while requiring only minor changes (if any) to existing codebases.

Over the past decade, CHERI has been studied heavily in the context of CPUs [65, 75, 76] and various challenges have been overcome to reduce hardware and performance costs to satisfactory levels [37, 72, 77]. This has led to industrial implementations from Arm [5], Microsoft [3], and CodaSip [16]. However, CHERI has not yet been applied to GPUs, which bring new challenges. Modern GPUs support hundreds of thousands of hardware threads leading to huge register files that account for a significant proportion of overall silicon area and power usage [27, 51]. CHERI increases both the pointer size and the architectural register size by a factor of two, potentially doubling the already-large register-file cost, which would be prohibitive. Indeed, prior work dismisses the use of CHERI on GPUs for precisely this reason [41]. Additionally, CHERI introduces a range of new instructions to manipulate capabilities and implementing these in every execution unit also incurs costs.

In this paper, we demonstrate techniques to considerably ameliorate the costs of CHERI in *Single Instruction, Multiple Thread* (SIMT) GPUs. We observe a large amount of redundancy in the metadata of capabilities between hardware threads, which can be readily exploited in SIMT architectures to reduce onchip storage overheads. We also observe that

several instructions introduced by CHERI are rarely found in the hot code paths of GPU kernels, allowing them to be implemented in a shared-function unit without impacting run-time performance. Exploiting these observations, we develop an efficient implementation of CHERI on top of an existing, prototype, open-source, synthesisable, SIMT GPU (SIMTIGHT) and CUDA-like C++ programming environment (NoCL) [52]. Our main results are as follows:

- We reduce the register-file storage overheads of CHERI from 103% to 14% by exploiting value regularity in capability metadata. With basic compiler support to limit the number of registers used to hold capabilities, we forecast that this overhead could be reduced to 7% without impacting run-time performance. Estimating that the register file accounts for less than half of the total onchip storage in a full GPU design, the overall storage overhead of CHERI would fall below 3.5%.
- We reduce the logic-area overheads of CHERI by 44% per streaming multiprocessor (SM) by moving non-critical logic into a shared-function unit. The absolute overhead is comparable to (but slightly larger than) the cost of an additional multiplier per vector lane.
- We find the run-time overhead of CHERI to be 1.6% on average across a range of GPU benchmarks — lower than reported CHERI overheads for CPU workloads.
- GPU benchmarks are simply recompiled to achieve full spatial memory safety and referential integrity. No source code changes to the benchmarks are required.

For comparison, we develop and evaluate an experimental Rust port of NoCL for SIMTIGHT, which shows an average 34% overhead for software bounds checking on the same GPU benchmarks ported ‘like-for-like’ from C++ to Rust (Section 4.7). We also compare qualitatively against existing approaches to GPU memory safety, covering both hardware support and the use of memory-safe languages (Section 5).

## 2 Background

### 2.1 Single Instruction, Multiple Threads (SIMT)

SIMT is a parallel execution model, popularised by NVIDIA and AMD GPUs, that combines the flexibility of a multi-thread programming model with the efficiency of SIMD hardware. The idea is to execute multiple hardware threads in lockstep with the aim of exploiting regularity between them. Known collectively as a *warp* (or *wavefront*), these threads can exhibit three main kinds of regularity [39]:

- *Control-flow regularity*, where threads in a warp follow the same path through the program. This allows the costs of fetching and decoding instructions to be amortised over multiple execution units.
- *Memory-access regularity*, where threads in a warp access neighbouring addresses in memory. This allows a

large number of narrow memory requests to be coalesced into a small number of wide requests, which can be more efficiently handled by the memory subsystem.

- *Value regularity*, where threads in a warp compute the same or similar intermediate values. This allows common data and computations to be shared between threads, reducing onchip storage costs and energy consumption.

While SIMT processors rely on inter-thread regularity to achieve optimal performance, they nevertheless permit general, independent, scalar computation in each thread. When threads in a warp *diverge* and take different paths through the program they can no longer execute in lockstep, leading to lower utilisation of the SIMD execution units. However, SIMT aims to *reconverge* these threads at the earliest opportunity to restore regularity and performance.

In addition to executing multiple threads per warp in parallel, SIMT processors also execute multiple warps concurrently, frequently context-switching between them (every cycle or every few cycles) to mask memory and compute latencies. A SIMT processor that executes multiple warps in this manner is sometimes referred to as *streaming multiprocessor*, or SM for short. NVIDIA’s GA100 GPU supports 128 SMs, with up to 64 warps per SM and 32 threads per warp, yielding over 250K hardware threads in total. Each of these hardware threads has access to its own set of private registers. The GPU register file can therefore require many megabytes of fast onchip storage, accounting for a major proportion of overall silicon area and power usage [27, 51].

It can be convenient to reuse SIMD terminology when working with SIMT. Although SIMT involves execution of scalar threads, it is possible to view each scalar value in a thread as an element of a *vector* value in a warp. Similarly, it is possible to view the scalar execution unit for each thread in a warp as a *vector lane*.

In modern GPUs, the number of execution units (or vector lanes) per SM typically equals the number of threads per warp, but this is not strictly necessary. Early NVIDIA GPUs shared 8 execution units across 32 threads in a warp, serialising each instruction over 4 consecutive cycles.

## 2.2 Value Regularity

Collange et al. first identified the prevalence of value regularity in SIMT workloads [20]. They use the term *uniform vector* to refer to a variable that has the same value in every thread in a warp, and *affine vector* to refer to a variable whose value is of the form  $base + t \times stride$  for each thread  $t$  in a warp with a fixed *base* and *stride*. A uniform vector is a special case of an affine vector with a zero stride. The authors report that, on average over a range of CUDA benchmarks running in simulation, 27% of vectors read from the register file, and 15% of vectors written, are uniform. These numbers rise to 44% and 28% respectively for affine vectors.

The prevalence of value regularity in SIMT workloads can be understood by looking at the data-parallel programming frameworks in which they are implemented, such as CUDA and OpenCL. In CUDA, each thread typically decides which part of the input to process based on its thread index within a block, and its block index within a grid. Threads in the same warp always reside in the same block, hence calculations on the block index will involve uniform vectors. Similarly, threads in the same warp will have consecutive thread indices, hence calculations on the thread index will involve affine vectors with a unit stride. To compound this, the CUDA programming guide [58] recommends that memory addresses for loads and stores should be the same (uniform) or consecutive (affine) across a warp, where possible, to maximise performance.

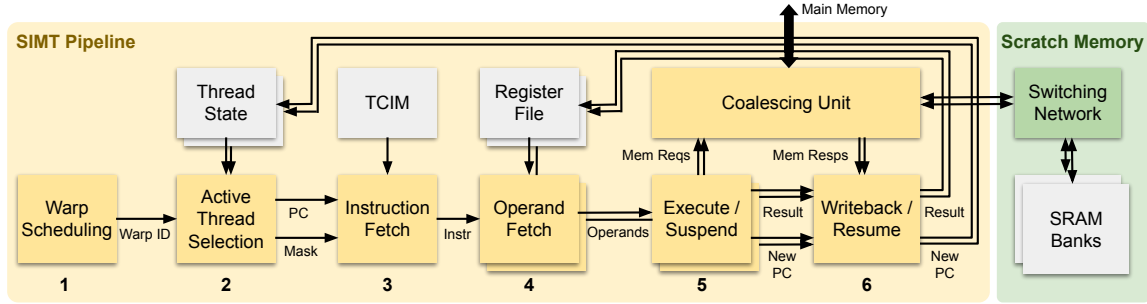
Value regularity can be exploited to reduce onchip storage requirements by storing uniform and affine vectors in a compact form. It can also be exploited to reduce energy requirements, or improve instruction throughput, by processing uniform and affine vectors in a specialised affine data path (with a single execution unit), separate from the general vector data path (with multiple execution units). These optimisations can be achieved using a technique known as *scalarisation*, which either operates *statically* in the compiler with instruction-set support [1, 10, 17, 36, 43, 69, 79] or *dynamically* in the microarchitecture [20, 28, 39, 42, 46, 52, 62, 78].

## 2.3 The SIMTIGHT GPU

Over the past decade, open-source GPUs have emerged as a promising platform for research and development. While early designs were based either on proprietary instruction sets [7], preventing free deployment of hardware, or custom instruction sets [2, 4, 9], with no access to a mature software stack or compiler, recent designs have started to employ the open RISC-V standard [19, 52, 63, 67], avoiding both problems. Researchers have shown that RISC-V is well-suited to SIMT execution, and how to extend RISC-V to support graphics. They have also demonstrated OpenGL, OpenCL, and CUDA-like software stacks on top of SIMT-style RISC-V GPUs. SIMTIGHT [52] is the latest in a line of such GPUs, which we use as a basis for the work presented in this paper.

SIMTIGHT implements RISC-V’s `rv32ima_zfinx` profile, i.e., a 32-bit machine with integer, multiply and divide, atomics, and single-precision floating-point support, with a merged integer and floating-point register file. The streaming multiprocessor (SM) component of SIMTIGHT is depicted in Figure 2 and has the following features:

- It is parameterised by the number of warps and the number of threads per warp, the latter of which is equivalent to the number of vector lanes as all threads in a warp can execute an instruction in parallel.
- It employs a 6-stage processor pipeline fed by a barrel scheduler that switches between warps on every cycle.



**Figure 2.** Diagram of the SIMTIGHT streaming multiprocessor (SM) [52], including the pipeline, tightly-coupled instruction memory (TCIM), coalescing unit, and scratchpad memory. Double boxes represent components containing logic or storage that is replicated per vector lane, and double lines represent per-lane wiring.

At most one instruction per warp is present in the pipeline at any time, avoiding data and control hazards.

- Each thread in a warp has its own program counter (PC), supporting control-flow divergence. The Active Thread Selection stage of the pipeline determines a subset of these threads which have the PC. The instruction at this common PC is fetched and decoded once, and then executed by all active threads in the warp, exploiting control-flow regularity. Thread convergence is achieved by prioritising the selection of threads residing at the deepest nesting level in the structured control-flow graph [18]. Further details about convergence can be found in the SIMTIGHT paper [52].
- Warps executing a multi-cycle instruction are suspended in the execute stage and resumed in the write-back stage without blocking the pipeline, tolerating high-latency operations such as memory loads and floating-point operations.
- Memory-access regularity is exploited using a coalescing unit that tries to pack memory requests from each vector lane into a smaller set of wider main-memory accesses using coalescing rules similar to those found in early NVIDIA Tesla devices [45].
- Value regularity is exploited using a set of microarchitectural techniques referred to as *advanced dynamic scalarisation* [52]. This includes *register-file compression* to reduce register-file storage requirements, a key feature of SIMTIGHT that we exploit in our implementation of CHERI, which is introduced in detail in Section 3.1.
- Efficient communication between hardware threads is facilitated through a scratchpad memory supporting parallel random access. This is implemented as a set of SRAM banks and a fast switching network. The scratchpad is critical for efficient implementation of `__shared__` memory in CUDA and `__local` memory in OpenCL.

```

0 struct Histogram : Kernel {
1   // Parameters
2   int len; unsigned char* in; int* out;
3   // Histogram bins in shared local memory
4   int* bins;
5
6   void init() {
7     declareShared(&bins, 256);
8   }
9
10  void kernel() {
11    // Initialise bins
12    for (int i=threadIdx.x; i<256; i+=blockDim.x)
13      bins[i] = 0;
14    __syncthreads();
15    // Update bins
16    for (int i=threadIdx.x; i<len; i+=blockDim.x)
17      atomicAdd(&bins[in[i]], 1);
18    __syncthreads();
19    // Write bins to global memory
20    for (int i=threadIdx.x; i<256; i+=blockDim.x)
21      out[i] = bins[i];
22  }
23 };

```

**Figure 3.** NoCL kernel to compute the 256-bin histogram of a given byte array using a single thread block. It is very similar to the CUDA version of the same kernel (lines 11–21 are identical) and demonstrates CUDA-style thread blocks, shared local memory, barrier synchronisation, and atomics.

The main limitation of SIMTIGHT is that it currently supports only a single SM. This is sufficient for the work described in this paper, where the majority of changes relate to components within an SM rather than the memory subsystem that connects SMs together. However, this does mean that the overheads we report are relative to a single SM rather than a full GPU design with multiple SMs and an advanced shared-memory subsystem.


The SIMTIGHT distribution ships with a programming API called NoCL that supports writing CUDA-like compute kernels in plain C++ (no special compute language is required).



It also includes a suite of benchmark compute kernels written in NoCL. A sample NoCL kernel to compute 256-bin histograms is shown in Figure 3. The authors report high IPC in many benchmarks as well as high performance density on FPGA compared to other open-source GPUs.

## 2.4 CHERI

CHERI extends conventional instruction-set architectures (MIPS, RISC-V, ARM, and x86) with *capabilities*. In the 32-bit RISC-V architecture, which we focus on in this paper, CHERI replaces 32-bit machine-word addresses with 64+1-bit capabilities. A capability captures an address together with bounds and permissions, and can be stored in registers or in memory. The bit representation of a 64+1-bit capability is as follows.



When using a capability to access memory (e.g., via load and store instructions), CHERI requires the address to lie within the bounds, throwing an exception otherwise. This is the basis for enforcing spatial memory safety. CHERI then provides instructions to manipulate capabilities, such as modifying the address (pointer arithmetic), reducing the permissions, and narrowing the bounds. When manipulating capabilities, the address is allowed to wander out of bounds to some extent. This is necessary to implement C/C++ pointers, which are allowed to point one byte beyond the end of an object and, in practice, often point further away than that [12].

A key property of CHERI is that capabilities are *unforgeable*: the only way to create a capability is to derive one from an existing capability using CHERI instructions, and doing so can never increase the bounds or permissions of the original. Consequently, the only memory that can be accessed by software running on a CHERI processor is that which is transitively reachable from capabilities stored in the register file. To achieve non-forgeability, CHERI stores a hidden *tag bit* for every register, and for every 64-bit word in memory, to distinguish valid capabilities from normal data. This makes it impossible to write arbitrary data to a register, or to memory, and subsequently interpret that data as a valid capability. It also makes pointers precisely distinguishable from data, supporting revocation and temporal memory safety [25, 26].

The tag bits of capabilities in memory are typically implemented by storing them in a reserved region of memory that is not architecturally addressable. A component called the *tag controller*, placed in front of main memory, ensures that each addressable 64-bit word and its corresponding tag bit are accessed atomically by processors. The tag controller includes a *tag cache* to optimise access to tag bits. It turns out that the miss rate of the tag cache, and hence the overhead of

<b>Get/clear tag bit</b>			<b>Load/store via capabilities</b>		
CGetTag	rd, cs1		CL[BHW][U]	rd, cs1, imm	
CClearTag	cd, cs1		CS[BHW]	rs2, cs1, imm	
<b>Get/reduce permissions</b>			<b>Load/store capabilities</b>		
CGetPerm	rd, cs1		CLC	cd, cs1, imm	
CAndPerm	cd, cs1, rs2		CSC	cs2, cs1, imm	
<b>Get/set bounds</b>			<b>Get/set/increment address</b>		
CGetBase	rd, cs1		CGetAddr	rd, cs1	
CGetLen	rd, cs1		CSetAddr	cd, cs1, rs2	
CSetBounds	cd, cs1, rs2		CIncOffset	cd, cs1, rs2	
CSetBoundsImm	cd, cs1, imm		CIncOffsetImm	cd, cs1, imm	
CSetBoundsExact	cd, cs1, rs2				
<b>Other</b>					
CGetType	rd, cs1		AUIPCC	cd, imm	
CGetSealed	rd, cs1		CJALR	cd, cs1, imm	
CGetFlags	rd, cs1		CJAL	cd, imm	
CSetFlags	cd, cs1, rs2		CSpecialRW	cd, cs1, imm	
CSealEntry	cd, cs1		CRRL	rd, rs1	
CMove	cd, cs1		CRAM	rd, rs1	

**Figure 4.** List of CHERI instructions implemented in SIMTIGHT (excluding atomics). CHERI extends each 32-bit general-purpose register with 33-bits of metadata. Operands *rd*, *rs1*, and *rs2* refer to the 32-bit general-purpose portion of a register while operands *cd*, *cs1*, and *cs2* refer to the full 65-bit contents. When an instruction writes to *rd*, the capability metadata for that register is set to a *null value* with the tag bit cleared.

accessing tag bits, can be reduced to almost zero in practice by exploiting the fact that blocks of memory (cache lines or pages) will often not hold any capabilities at all, allowing a highly compact representation in the tag cache [37].

To keep the size of a capability to 64 bits, CHERI represents the bounds in a compressed format known as *CHERI Concentrate* [77]: a 32-bit lower bound and a 33-bit upper bound are together stored in just 15 bits by encoding them in a floating-point-like format relative to the address. The hardware costs of bounds compression are discussed in Section 3.3.

The properties provided by CHERI are *deterministic* and *enforced*, not subject to probabilities or bypassing.

## 3 Design and Implementation

In this section, we present the main changes needed to support CHERI efficiently in the SIMTIGHT GPU. Specifically, we implement a large subset of version 9 of the 32-bit CHERI instruction set [75], as shown in Figure 4.

### 3.1 Register File

Adapting a 32-bit RISC-V implementation to support CHERI requires extending every 32-bit general-purpose register to 65 bits. This is potentially a large cost in a streaming multiprocessor with thousands of hardware threads, each with their own set of private registers. However, our hypothesis is that there is likely to be a lot of value regularity in the capability metadata between threads executing in lockstep: threads accessing different elements of the same array at

the same time will likely involve the same bounds and permissions. Before trying to exploit this hypothesis, it is first useful to look at SIMTIGHT’s existing register-file compression mechanism in more detail.

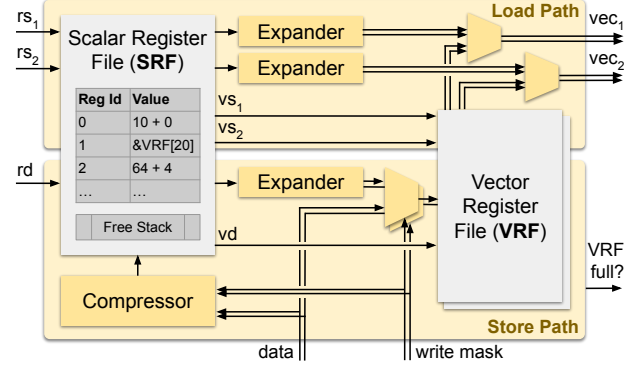
A single SIMTIGHT streaming multiprocessor contains  $32 \times \text{numWarps} \times \text{numThreadsPerWarp}$  architectural registers or, equivalently,  $32 \times \text{numWarps}$  architectural vector registers (each scalar register in a thread can be viewed as an element of vector register in a warp). SIMTIGHT’s compressed register file aims to reduce onchip storage requirements by exploiting the property that vector registers will often hold uniform or affine vectors that can be stored compactly. It detects these uniform and affine vectors at run time using an array of comparators in the register-file write path (cheaper inference-based mechanisms for detecting uniform and affine vectors are also possible [20, 39, 78] but are not yet supported in SIMTIGHT). Uniform and affine vectors are then stored in a scalar register file (SRF) while general (non-compressible) vectors are allocated on-demand in a larger, size-constrained vector register file (VRF). For every architectural vector register, the SRF either holds a compressed vector (a *base + stride* pair) or a pointer to a register in the VRF. The size of the VRF can be set arbitrarily. Currently in SIMTIGHT it is chosen statically at synthesis time based on experimental evaluation, but in principle it can be set dynamically, which is useful if, as in modern NVIDIA GPUs, the physical memory implementing the register file is shared with other storage structures such as the scratchpad and L1 cache [27]. VRF overflow is handled in hardware by dynamically spilling vector registers to main memory. The overall structure of the compressed register file is shown and explained in Figure 5.

The SIMTIGHT authors report that a compressed register file with a quarter-sized VRF (i.e., a VRF big enough to hold a quarter of all architectural vector registers) has a minimal impact on run-time performance while reducing register-file storage requirements by 68% per 2,048-thread SM.

### 3.2 Metadata Register File

To support CHERI in the register file, we instantiate two compressed register files: a 32-bit general-purpose register file and a new 33-bit capability-metadata register file. This has the advantage that integer addresses and capability metadata are compressed separately: if we have a vector of capabilities that all have the same metadata but different (non-affine) addresses, then the metadata can still be compressed even though the addresses cannot. We enable the detection of only uniform (not affine) vectors in the metadata register file as the notion of a stride does not really exist for capability metadata.


The main drawback of this simple approach is *fragmentation*: there are two VRFs, each capable of holding on only one kind of vector (data or metadata), and if one becomes full then spilling will occur even if there is space available



**Figure 5.** Overview of SIMTIGHT’s compressed register file. Registers  $rs_1$  and  $rs_2$  are looked up to produce vectors  $vec_1$  and  $vec_2$  respectively. The active elements of vector *data* (as specified by *write mask*) are written to register *rd*. If any of the registers  $rs_1$ ,  $rs_2$ , and  $rd$  are not held in the SRF then the SRF emits the locations  $vs_1$ ,  $vs_2$ , and  $vd$  of these registers respectively in the VRF. The *Compressor* attempts to transform a vector to a *base + stride* pair that can be stored in the SRF, while the *Expander* performs the inverse transformation. For vectors that cannot be compressed, the *Free Stack* tracks unused locations in the VRF where they can be stored. If this stack becomes near-empty, the *VRF full* flag is asserted, triggering the pipeline to spill registers from the VRF to main memory.

in the other. Furthermore, the SIMTIGHT baseline enforces a minimum VRF capacity of four vector registers per thread and capability metadata may be more compressible than that. We therefore implement a *shared VRF* between the general-purpose and capability-metadata register files. To avoid additional read ports on the shared VRF, we serialise data and metadata accesses, i.e., accessing a register that requires both an uncompressed data vector *and* an uncompressed metadata vector will result in a pipeline stall. Our hypothesis is that this will not happen very often due to inter-thread regularity.

As shown in Figure 5, the baseline SRF requires three read ports, two in the load path and one in the store path. SIMTIGHT implements this using two instances of a two-read-port SRAM primitive, with each instance holding identical data. This can be avoided in the capability-metadata SRF. As shown in Figure 4, the vast majority of CHERI instructions use only one capability source operand (i.e., *cs1* and not *cs2*). Only the CSC instruction to store a capability via a capability refers to both capability source operands. We therefore reduce the number of read ports on the capability-metadata SRF in exchange for taking an extra cycle to implement CSC. This means that the capability-metadata SRF essentially uses half the amount of storage as the baseline SRF. As shown in Figure 6, the execution frequency of the CSC instruction is quite low, around 2%.



**Figure 6.** Average execution frequency of CHERI instructions on GPU workloads relative to total instructions executed, obtained using the experimental setup introduced in Section 4. CHERI instructions that are not shown here are not executed at all in the analysed workloads.

The baseline register file restricts itself to *total* scalarisation: for a vector to be compressible, *all* elements must satisfy the uniform or affine requirements. One can envisage a generalisation of this where, for example, a vector can be partitioned into two different uniform or affine vectors and still held more compactly than an uncompressed vector. This would be useful in the presence of control-flow divergence, when a uniform or affine vector gets partially overwritten with a different uniform or affine vector. To implement this, each SRF entry would need to store an additional *base+stride* pair and a bit mask denoting which partition each vector element belongs to. Unfortunately, we have found that this SRF cost generally outweighs the associated savings in the VRF. However, the situation is slightly different for the capability-metadata SRF. First, it is half the size of the baseline SRF so increasing its size by a constant factor is not as expensive in absolute terms. Second, when a register holds an integer or floating-point value, rather than a capability, which is often the case, the metadata for that register is known to be a constant *null* value. This raises the possibility of extending the SRF with just a mask denoting which vector elements are null, supporting a form of partial scalarisation. We have implemented this as an optional feature in the metadata register file, which we refer to as the *null-value optimisation* (NVO). When a uniform vector is partially overwritten with a null vector, or vice-versa, that vector remains in the SRF. Furthermore, when a partially uniform vector is partially overwritten with a null vector or the same partially uniform vector, that vector also remains in the SRF.

### 3.3 Pipeline

Besides increasing the register size, another cost of adding CHERI to SIMTIGHT lies in the additional logic required to implement CHERI instructions in the pipeline and, in particular, to handle compressed bounds in capabilities. We use a standard library implementation of the CHERI Concentrate compressed bounds format [77] called CHERICAPLIB [64] whose main functions, along with their logic-area costs, are shown in Figure 7. While CHERICAPLIB functions for getting

- In-memory capability format including tag bit  
`type CapMem = Bit 65`
- In-pipeline (partially decompressed) capability format  
`type CapPipe = Bit 91`
- Memory access width:  $2^0$ ,  $2^1$ ,  $2^2$ , or  $2^3$  bytes  
`type AccessWidth = Bit 2`

– Convert from the in-memory format	ALMs
<code>fromMem :: CapMem → CapPipe</code>	46
– Convert to the in-memory format	
<code>toMem :: CapPipe → CapMem</code>	0
– Set the address, invalidating if too far out-of-bounds	
<code>setAddr :: (CapPipe, Bit 32) → CapPipe</code>	106
– Check that an access of given width is within bounds	
<code>isAccessInBounds :: (CapPipe, AccessWidth) → Bit 1</code>	25
– Return the base (lower bound) of the capability	
<code>getBase :: CapPipe → Bit 32</code>	50
– Return the length of the capability	
<code>getLength :: CapPipe → Bit 33</code>	20
– Return the top (upper bound) of the capability	
<code>getTop :: CapPipe → Bit 33</code>	78
– Narrow bounds of a capability to given base and length	
<code>setBounds :: (CapPipe, Bit 32) → CapPipe</code>	287

**Figure 7.** Key functions of the CHERICAPLIB library [64] to handle compressed bounds in 64+1-bit capabilities. The logic area requirement of each function in ALMs (Intel Stratix 10 Adaptive Logic Modules) is listed on the right-hand side. As a point of reference, a 32-bit multiplier requires 567 ALMs. As shown, the *isAccessInBounds* function can check against partially decompressed bounds much more cheaply than fully decompressing the bounds via *getBase* and *getTop* and then using two address-width comparators.

and setting bounds are quite expensive, the CHERI instruction histogram in Figure 6 shows that instructions to get and set bounds are not frequently executed in GPU workloads. This motivates the separation of CHERICAPLIB functions into those that are frequently needed (such as pointer arithmetic and bounds checking) and those that are infrequently needed (such as getting and setting bounds). The former can be instantiated per vector lane (fast path) while the latter can be instantiated per SM in a shared-function unit (slow path).

The shared-function unit (SFU) is a common feature of GPU designs and indeed SIMTIGHT includes one that is used to implement floating-point square root and division. SIMTIGHT’s SFU connects to every vector lane via a request serialiser and a response deserialiser. To support CHERI instructions in the SFU, we increase the size of SFU requests and responses to hold capability-sized operands and results respectively. This, in turn, increases the logic needed for serialisation and deserialisation but, overall, area is substantially reduced by moving logic out of the per-lane ALUs. We implement the CGetBase, CGetLen, CSetBounds[...], CRRL, and CRAM instructions in the SFU. All other CHERI instructions



are implemented per vector lane using four `CHERICAPLIB` function calls, as shown in Figure 8.

In ChERI, the program counter is also considered to be a capability. Accordingly, we extend `SIMTight`'s per-thread program counters (PCs) to be program-counter capabilities (PCCs). In the Active Thread Selection stage of the pipeline, we require that the chosen threads not only have the same PC but the same PCC. This means that only a single program-counter bounds check is required per SM. Nevertheless, these changes do increase logic area and, in software, we do not currently exploit the ability to change PC metadata dynamically in GPU kernels. We therefore provide an optional feature whereby PC metadata can be set once per warp at kernel-invocation time but never changed. This allows Active Thread Selection to disregard PC metadata. We refer to this feature as the *static PC metadata restriction*.

### 3.4 Memory Subsystem

`SIMTIGHT`'s memory subsystem supports 8-bit, 16-bit, and 32-bit accesses natively. To implement 64-bit (capability width) accesses, we use *multi-flit transactions* whereby a series of contiguous memory requests terminated by an *is-final* bit are treated atomically by the memory subsystem. A 64-bit access is achieved using two inseparable 32-bit accesses. This avoids increasing the data-path width in the memory subsystem in exchange for a two-cycle capability access time. As shown in Figure 6, the CLC and CSC instructions for loading and storing capabilities are executed fairly infrequently. The logic required to serialise/deserialise 64-bit requests/responses is placed between the pipeline and the coalescing unit.

CHERI requires a 1-bit tag to be maintained for every naturally aligned 64-bit value in memory, indicating whether or not that value holds a capability. As `SIMTight`'s memory subsystem is natively 32-bit, we opt to maintain a 1-bit tag for every 32-bit naturally aligned value for simplicity (not necessity). We therefore introduce the invariant that for a 64-bit capability to be valid, the tag bits of both its upper and lower halves must be set. To allow capabilities to be stored in scratch memory, we extend the data width of each onchip SRAM scratchpad bank from 32 to 33 bits. To allow capabilities to be stored in main memory, we follow the approach taken by existing ChERI-enabled CPUs: tag bits are stored in a reserved region and a tag controller is placed just in front of main memory providing the illusion of atomic access to each value and its tag bit [37, 65, 76].

## 4 Evaluation

### 4.1 Experimental Setup

We use the suite of 14 CUDA-like NoCL benchmark programs shipped as part of the standard `SIMTIGHT` distribution and listed in Table 1. To compile the benchmarks, we use the ChERI fork of Clang 13 (the latest compiler supporting ChERI-RISC-V at the time). For a fair comparison, we disable

Benchmark	Description	Origin
VecAdd	Vector addition	[56]
Histogram	256-bin histogram calculation	[55]
Reduce	Vector summation	[55]
Scan	Parallel prefix sum	[53]
Transpose	Matrix transpose	[55]
MatVecMul	Matrix $\times$ vector multiplication	[56]
MatMul	Matrix $\times$ matrix multiplication	[55]
BitonicSm	Bitonic sorter (small arrays)	[56]
BitonicLa	Bitonic sorter (large arrays)	[56]
SPMV	Sparse matrix $\times$ vector multiplication	[8]
BlkStencil	Block-based stencil computation	In house
StrStencil	Stripe-based stencil computation	In house
VecGCD	Vectorised greatest common divisor	In house
MotionEst	Motion estimation	In house

**Table 1.** NoCL benchmark suite.

*scalar evolution* (SCEV) of pointers, a compiler optimisation that is not yet supported when targeting ChERI-RISC-V but which is expected to be in future. To support ChERI, some minor changes to the NoCL library are needed, such as setting the bounds of the stack and dynamically allocated buffers, but the benchmarks themselves do not require any modifications at all.

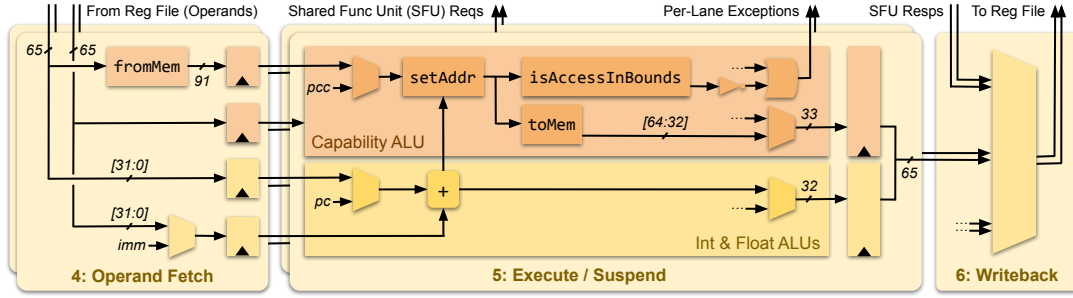
We use the standard optimisation level `-O2` but force the compiler to inline the `kernel()` method of every NoCL compute kernel. This avoids function-call overhead in the NoCL inner loop, which invokes this method for every thread in a block, and every block in a grid. Aggressive inlining is standard when compiling GPU code, e.g., all device functions in CUDA are inlined by default [58]. In Section 4.4, we discuss some of the advantages of inlining in more detail.

We obtain all results on a Terasic DE10-Pro development board with a Stratix-10 FPGA holding a single `SIMTIGHT` SM connected to a DDR4 DIMM and a ChERI-enabled host CPU, as depicted in Figure 9. Following modern NVIDIA devices, and prior work on `SIMTIGHT`, we use 64 warps per SM and 32 threads per warp providing 2,048 threads per SM in total. This number of warps is sufficient to mask the latency of DDR4 memory on FPGA, achieving good performance without caches. In Section 4.4, we explain that the addition of caches and multiple SMs are not expected to impact the reported ChERI overheads.


For the compressed register file, the `SIMTIGHT` authors report that a  $1/4$ -size VRF provides a 68% storage reduction with negligible run-time and memory-access overheads. However, they were using a modern GCC (version 12) while we are using a relatively old Clang (the latest version that supports ChERI at the time), which yields inferior results. We therefore opt for a  $3/8$ -size VRF in the baseline providing a 55% storage reduction, as shown in Table 2.

We consider three main configurations of `SIMTIGHT`:





**Figure 8.** SIMTight pipeline modifications to support CHERI. The output of the integer ALU’s adder either contains the address result for `CIncOffset[...]/CSetAddr`, the address to access for `CL[...]/CS[...]`, or the address to jump to for `CJAL[...]`. This address is fed into `setAddr` to yield a capability that is either written to a result register via `toMem`, fed to `isAccessInBounds` for bounds checking, or written to the PCC.



**Figure 9.** Diagram of the SIMTight evaluation SoC, with data-bus widths (excluding tag bits).

VRF Size (Registers)	Storage (Kb)	Compress Ratio	Cycle Overhead	Mem Access Overhead
1,024 (1/2)	1,202	1 : 0.57	0.8%	0.1%
768 (3/8)	937	1 : 0.45	0.9%	2.2%
512 (1/4)	672	1 : 0.32	4.3%	39.9%

**Table 2.** Results of register-file compression in the SIMTight baseline, using Clang 13, for a 1/2, 3/8, and 1/4-size VRF.

- **Baseline** Baseline configuration with a compressed general-purpose register file, but no CHERI. Benchmarks run with no memory safety.
- **CHERI** Extension of the baseline with CHERI. Value regularity in capability metadata is not detected or exploited. No CHERI instructions are implemented in the shared function unit. Benchmarks run with full spatial memory safety and referential integrity.
- **CHERI (Optimised)** Extension of the CHERI configuration with optimisations. Value regularity in capability metadata is detected and exploited. The capability-metadata register file is compressed. The shared VRF and null-value optimisations are both enabled. CHERI instructions for getting and setting bounds are implemented in the shared function unit. This configuration also enables the static PC metadata restriction.

## 4.2 Threat Model


We consider a simple threat model in which an attacker seeks to exploit an out-of-bounds memory access to exfiltrate or corrupt data, or hijack control flow. When CHERI is enabled, applications are compiled in pure capability mode with all C++ pointers (and architectural addresses, such as the stack pointer and return addresses) being implemented as capabilities, enabling deterministic prevention of all such attacks. This covers code running on the SIMTight SM and the host CPU, which can freely exchange capabilities via main memory. CHERI lays the foundation for protection against a much wider range of threats, including exploitation of use-after-free bugs, and interaction with untrusted software components. These threats are beyond the scope of this paper, but we refer the reader to security evaluations of CHERI on CPUs, which may be indicative [3, 26, 35, 38, 73].

## 4.3 Register-File Overhead


Figure 10 shows the proportion of register values that need be stored as uncompressed vectors in the shared VRF, both for values stored in the general-purpose register file and the capability-metadata register file. With the null-value optimisation, only the `BlkStencil` benchmark uses space in the VRF for capability metadata (explained below). The register-file storage overhead of CHERI is therefore almost entirely accounted for by the cost of the capability-metadata SRF, which is 14% of the total register-file storage of the baseline.

Figure 11 shows that no benchmark uses more than half of the available registers to hold capabilities. Therefore, with compiler support to limit the number of registers can hold capabilities, the size of the capability-metadata SRF could be halved without impacting run-time performance. This would reduce the register-file storage overhead to 7%.

In the literature, the register file is typically considered to account for around 256KB of onchip storage per SM, compared to 64KB per SM for scratchpad memory and 64KB per SM for the L1 cache [27]. Furthermore, the shared L2 cache typically accounts for around 128KB to 256KB of storage per



**Figure 10.** Proportion of registers stored as vectors in the VRF (lower is better) for the general-purpose register file, and for the capability-metadata register file with and without the null-value optimisation (NVO). Remaining registers are stored compactly as scalars in the SRF.



**Figure 11.** Number of registers per thread used to hold capabilities. Each thread has access to 32 registers in total. The remaining registers are never used to hold capabilities.

SM in modern GPUs [54]. We therefore estimate that the register file accounts for less than half of total onchip GPU storage. The total storage overhead of CHERI in a full GPU design would therefore likely fall below 3.5%.

Upon inspection of the BlkStencil benchmark, we see that the capability-metadata divergence arises from a compiler optimisation. A line of source code of the form

```
if (cond) {acc += *p1;} else {acc += *p2;}

```

effectively gets transformed to


```
if (cond) {tmp = p1;} else {tmp = p2}; acc += *tmp;

```


where  $p_1$  and  $p_2$  point to elements of different arrays (one stored in global memory and the other in shared local memory). The compiler has therefore transformed control-flow divergence into pointer-value divergence. In this work, we are using entirely pre-existing compiler toolchains but the above optimisation could potentially be disabled for SIMT targets, to preserve value regularity.

#### 4.4 Memory-Bandwidth Overhead

Figure 12 shows that the introduction of CHERI does not significantly affect DRAM bandwidth usage in SIMTIGHT. Inlining of GPU functions plays a role here. Without inlining, CHERI incurs a slightly larger overhead due to increased function calls and associated loads/stores of double-sized pointers on the stack. Another common source of accessing pointers on the stack is compiler-inserted register



**Figure 12.** DRAM bandwidth usage with/without CHERI.



**Figure 13.** Execution-time overheads of the optimised CHERI configuration relative to the baseline configuration.

spills. SIMTIGHT provides a proof-of-concept *compressed stack cache* to reduce the cost of register spilling at low hardware cost by caching uniform/affine vectors in a compressed form. While this cache is particularly effective on capability metadata, it does not appear to have a noticeable impact on performance, at least in the existing benchmark suite.

Our experimental setup is currently limited to a single SM GPU without caches. However, as the memory subsystem and bandwidth usage of a single-SM GPU is largely unaffected by addition of CHERI, we project that a multi-SM memory subsystem (and its cache footprint) would be similarly unaffected.

#### 4.5 Execution-Time Overhead

Figure 13 shows the execution-time overhead of adding CHERI to SIMTIGHT. The geometric-mean cycle overhead is 1.6%. This is lower than CHERI overheads reported on CPUs [66], suggesting that CHERI is cheaper on GPUs than CPUs. Heavy function-call inlining and lack of pointer chasing in GPU workloads limit the number of accesses to double-sized pointers on the stack and heap respectively. The main outlier in the results is BlkStencil, which exhibits two uncommon but costly behaviours: capability-metadata divergence (Section 4.3) and execution of a relatively high number of CSC instructions, each of which can incur a 1-cycle performance penalty during operand fetch (Section 3.2).

#### 4.6 Synthesis Results

Table 3 shows the logic area and onchip storage overheads of adding CHERI to SIMTIGHT. Our optimisations reduce the area overhead by 44% to 708 ALMs per vector lane, comparable to (but slightly larger than) the cost of an additional

Configuration	Area (ALMs)	Area (DSPs)	Block RAM (Kilobits)	Fmax (MHz)
Baseline	126,753	0	2,156	180
CHERI	166,796	0	4,399	181
CHERI (Optimised)	149,356	0	2,394	180

**Table 3.** Synthesis results for a single SIMTIGHT SM on FPGA with/without CHERI. Use of DSP blocks has been disabled to obtain a single ALM count representing all logic used.

multiplier (567 ALMs) per vector lane. The onchip memory storage overhead, measured in bits, is largely eliminated.


Comparing against academic prototype implementations of CHERI on CPUs [65], our area overheads are lower; we have been able to amortise the cost of CHERI logic across multiple execution units. Arm’s prototype implementation of CHERI in their Neoverse N1 CPU exhibits much lower relative area overheads [72] than academic CPU prototypes as the baseline design is richer in terms of features and optimisations. For similar reasons, it is reasonable to expect that the relative area overheads of CHERI in a commercial-grade GPU would be lower than those for SIMTIGHT.

#### 4.7 Software Bounds Checking

For comparison against CHERI, we have developed an experimental Rust port of NoCL and its benchmark suite that runs on the SIMTIGHT GPU. Rust is a modern systems programming language providing similar levels of efficiency to C/C++ but with stronger correctness properties. In safe Rust, memory safety is enforced: all references point to valid live memory and cannot be accessed out of bounds.

Our Rust port is ‘like-for-like’ in the sense that C++ and Rust versions of each benchmark are defined very similarly. This gives confidence that different versions of the same benchmark are behaving in a similar way, allowing a fair performance comparison. On the other hand, it deviates from Rust’s thread model, which prohibits multiple threads from having write access to the same memory region at the same time, preventing data races. Such access is commonplace in CUDA and OpenCL, e.g., multiple threads writing different parts of a result array at the same time. Resolving this discrepancy is an open problem, and we discuss some recent research in this area in Section 5.1. In our experimental NoCL port, we do not attempt to solve this issue; the Rust compiler is simply unaware of the presence of multiple GPU threads.

Figure 14 compares the execution times of C++ and Rust versions of the benchmarks running on SIMTIGHT. The C++ and Rust compilers used are both based on v19.1.7 of LLVM. Overall, there is a geometric mean overhead of 46% from using Rust. Bounds checking alone accounts for a 34% overhead. This indicates the software bounds checking is expensive in low-level GPU code, at least without providing information about the relationships between the sizes of buffers, which the compiler could use to eliminate some of the checks.



**Figure 14.** Execution-time overheads of our Rust port of NoCL running on SIMTIGHT.

Prior work by Harris et al. [32] reports that the Rust compiler is very good at eliminating bounds checks, resulting in low overhead for software bounds checking in Rust code running on CPUs. However, care is needed when interpreting these results as it is unclear how many bounds checks are avoided due to the use of unsafe code blocks. For example, in the `matrixmultiply` benchmark they use, the key functions are declared as `unsafe` and operate directly on raw pointers. For performance-critical code, the developers in this case have bypassed Rust’s safety checks, an approach that is potentially quite common. In our experience, it is difficult for the Rust compiler to eliminate bounds checks in CUDA-style code because there is no general relationship between the bounds of an array and the thread id used to compute an index into that array.

## 5 Related Work

### 5.1 Safe Languages Targeting GPUs

Over the past decade, GPUs have become a target for several high-level data-parallel array programming languages [24, 29, 31, 34, 48, 57]. These languages replace low-level management of individual threads and memory accesses with high-level bulk parallel array operations such as `map`, `reduce`, `scan`, and so on. They have been shown to achieve useful levels of performance while often guaranteeing properties such as memory safety and data-race freedom. Interestingly, many bulk array operations are safe by construction and do not require bounds checks. However, some operations, such as `gather` and `scatter`, permit ad-hoc indexing and do require run-time safety checks. The authors of the high-level array language Futhark [34] report a 6% average performance overhead due to bounds checking on GPUs [33].

The flip side of abstracting away from low-level details is that the programmer loses control over these details. Köpcke et al. argue that such control is needed to extract the highest levels of performance from GPU hardware [40]. Inspired by Rust, they propose a safe low-level language for GPU programming called Descend. The concept of *views* is used to safely describe parallel accesses to shared memory in a way that can be statically checked for memory safety and data-race freedom. Essentially, this means that the programmer expresses both the individual memory accesses of each

thread and a form of proof that these accesses are safe. As a result, Descend programs look quite different to CUDA programs and are harder to write; there are likely to be major costs in terms of porting existing applications and re-educating users. As in Rust, there are also cases where it is difficult to express the desired behaviour in a way that can be checked by the compiler, and `unsafe` code blocks are provided as an escape hatch. Any programs containing such blocks are no longer guaranteed to be safe.

Compiling Rust itself down to GPUs has also been explored. The Rust-CUDA project [21] extends the Rust compiler with support for targeting NVIDIA GPUs. The project is described as being in the early stages and does not yet present any performance analysis. The developers have recently announced a new effort to pursue the project further [44].

## 5.2 Hardware Support for GPU Memory Safety

To our knowledge, the only prior hardware approach to GPU memory safety is GPUShield [41]. This extends NVIDIA and Intel GPU models with a *bounds table* and uses the top 16 bits of every 64-bit pointer to hold an index into the table. The bounds table is setup before a GPU kernel is launched, and remains unchanged for the duration of execution. Bounds on kernel arguments are provided by the caller, and the compiler is extended to emit bounds information for buffers declared within a kernel. The compiler is also modified to avoid bounds checking on accesses that are known to be safe by static analysis; this is achieved by marking pointers as ‘unprotected’ in the top 16 bits.

The authors argue that extending the size of registers is unacceptable on GPUs due to already-massive GPU register files. However, our results suggest that these costs can be largely eliminated by exploiting inter-thread redundancy in the metadata. Furthermore, extending the pointer size has several benefits, such as retaining access to the full address space, supporting 32-bit as well as 64-bit architectures, and increasing the bits available to encode metadata. This in turn avoids need for a limit on the number of buffers that can be protected, and the need for indirection via a bounds table, as bounds can be encoded directly within the metadata itself.

GPUShield has been developed with a strong emphasis on efficiency, and the authors report a low average execution-time overhead of 0.8%. However, there are limitations in terms of expressibility and security. From an expressibility perspective, bounds cannot be modified during kernel execution, meaning that dynamically allocated memory cannot be protected. From a security perspective, pointer metadata is unprotected, which means that the buffer id can be corrupted or forged. Furthermore, GPUShield’s ‘unprotected’ pointers, which allow the bounds table lookup to be bypassed for efficiency, makes it possible to forge a pointer to any address in memory. As GPUShield was developed in simulation, the area overhead is unclear, but it would probably be lower

Feature	GPUShield	CHERI
Supports spatial memory safety	✓	✓
Provides referential integrity	✗	✓
Supports 32-bit and 64-bit architectures	✗	✓
Permits use of entire address space	✗	✓
Supports an unlimited number of buffers	✗	✓
Supports dynamic allocation of buffers	✗	✓
Pointers can be distinguished from data	✗	✓
Applies to both CPUs and GPUs	✗	✓
Demonstrated in a synthesisable GPU	✗	✓
Performance overhead on GPUs	Low	Low
Silicon area overhead on GPUs	Low (likely)	Medium

**Figure 15.** Comparison of GPUShield [41] and CHERI.

than CHERI’s as data-path widths are largely unaffected. A feature comparison of GPUShield and CHERI is summarised in Figure 15.

## 5.3 Decoupled Capability Processors

For memory-accessing devices that require fine-grained memory protection but for which an ISA extension is not appropriate or too invasive, researchers have explored *decoupled* capability processors [6, 11, 47]. These take the form of a ‘bump in the wire’ between a device and main memory which holds a set of capabilities and checks that every memory access is permitted by one of these capabilities. In principle, GPU memory accesses could be protected in this way, with the GPU being treated as a black box. However, the protections provided would be quite limited. The main problem is how to associate each access made by the GPU with the capability that it should be checked against. Without this, each access would need to be checked against the union of the capabilities available, and fine-grained protection would be lost. Embedding metadata into the address to identify a capability would lead to a similar design to GPUShield, with the same drawbacks. Another problem with the decoupled approach is that it is not possible to protect memory that is internal to the GPU, such as scratchpad memory.

## 6 Conclusion

Despite huge numbers of registers in heavily threaded GPU cores, extending registers with metadata can, in fact, be feasible. Storage costs in SIMT designs are dependent on the amount of value regularity between threads executing in lockstep. Capability metadata exhibits substantial value regularity, which can be exploited to largely eliminate the storage overhead of CHERI’s double-size registers. Furthermore, the logic-area overhead of CHERI can be significantly reduced by amortising the cost of some CHERI instructions across multiple execution units. With these optimisations, CHERI offers a viable path to memory-safe C/C++-based languages running on GPUs, avoiding the need for widespread porting of applications to new memory-safe languages and the associated re-education of users. Compared to recent work



on hardware support for GPU memory safety, CHERI provides far stronger security properties with similar runtime overheads, but likely higher area overheads. In future, the much-broader threat model supported by CHERI would be interesting to explore in the context of GPUs.

## Acknowledgements

Thanks to Jianyi Cheng for improving the reproducibility of the results presented in this paper. Our work was supported by UK-EP SRC under the *CAPcelerate Project* (EP/V000381/1) and the *Chrompartments Project* (EP/X015963/1), both part of the Digital Security by Design (DSbD) Programme and the DSbDtech initiative. Additional data related to this publication is available at <https://doi.org/10.17863/CAM.120202>. CHERI support for SIMTIGHT has been merged into the main SIMTIGHT distribution, available at <https://github.com/CTSRD-CHERI/SIMTight>.

## References

- [1] Advanced Micro Devices (AMD). 2012. *Southern Islands Series Instruction Set Architecture 1.1*.
- [2] Muhammed Al Kadi, Benedikt Janssen, and Michael Huebner. 2016. FGPU: An SIMT-Architecture for FPGAs. In *ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA 2016)*. <https://doi.org/10.1145/2847263.2847273>
- [3] Saar Amar, David Chisnall, Tony Chen, Nathaniel Wesley Filardo, Ben Laurie, Kunyan Liu, Robert Norton, Simon W. Moore, Yucong Tao, Robert N. M. Watson, and Hongyan Xia. 2023. CHERIoT: Complete Memory Safety for Embedded Devices. In *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2023)*. <https://doi.org/10.1145/3613424.3614266>
- [4] Kevin Andryc, Murtaza Merchant, and Russell Tessier. 2013. FlexGrip: A soft GPGPU for FPGAs. In *International Conference on Field-Programmable Technology (FPT 2013)*. <https://doi.org/10.1109/FPT.2013.6718358>
- [5] Arm. 2024. Arm Morello Program. <https://www.arm.com/architecture/cpu/morello> (accessed 2024-01-30).
- [6] Leonid Azriel, Lukas Humbel, Reto Achermann, Alex Richardson, Moritz Hoffmann, Avi Mendelson, Timothy Roscoe, Robert N. M. Watson, Paolo Faraboschi, and Dejan Milojicic. 2019. Memory-Side Protection With a Capability Enforcement Co-Processor. *ACM Transactions on Architecture and Code Optimisation* 16, 1, Article 5 (2019), 26 pages. <https://doi.org/10.1145/3302257>
- [7] Raghuraman Balasubramanian, Vinay Gangadhar, Ziliang Guo, Chen-Han Ho, Cherin Joseph, Jaikrishnan Menon, Mario Paulo Drumond, Robin Paul, Sharath Prasad, Pradip Valathol, and Karthikeyan Sankaralingam. 2015. Enabling GPGPU Low-Level Hardware Explorations with MIAOW: An Open-Source RTL Implementation of a GPGPU. *ACM TACO* 12, 2, Article 21 (2015). <https://doi.org/10.1145/2764908>
- [8] Nathan Bell and Michael Garland. 2008. *Efficient Sparse Matrix-Vector Multiplication on CUDA*. Research Report. NVIDIA Corporation.
- [9] Jeff Bush, Mohammad A. Khasawneh, Khaled Z. Mahmoud, and Timothy N. Miller. 2016. NyuziRaster: Optimizing rasterizer performance and energy in the Nyuzi open source GPU. In *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS 2016)*. <https://doi.org/10.1109/ISPASS.2016.7482095>
- [10] Zhongliang Chen and David Kaeli. 2016. Balancing Scalar and Vector Execution on GPU Architectures. In *IEEE International Parallel and Distributed Processing Symposium (IPDPS 2016)*. <https://doi.org/10.1109/IPDPS.2016.74>
- [11] Jianyi Cheng, A. Theodore Markettos, Alexandre Joannou, Paul Metzger, Matthew Naylor, Peter Rugg, and Timothy M. Jones. 2025. Adaptive CHERI Compartmentalization for Heterogeneous Accelerators. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture (ISCA 2025)*. <https://doi.org/10.1145/3695053.3731062>
- [12] David Chisnall, Colin Rothwell, Robert N.M. Watson, Jonathan Woodruff, Munraj Vadera, Simon W. Moore, Michael Roe, Brooks Davis, and Peter G. Neumann. 2015. Beyond the PDP-11: Architectural Support for a Memory-Safe C Abstract Machine. In *20th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2015)*. <https://doi.org/10.1145/2694344.2694367>
- [13] Chromium Team. 2024. The Chromium Projects: Memory safety. <https://www.chromium.org/Home/chromium-security/memory-safety/> (accessed 2024-01-25).
- [14] Cybersecurity & Infrastructure Security Agency (CISA), NSA, FBI, Australian Cyber Security Centre (ACSC), Canadian Centre for Cyber Security (CCCS), UK's National Cyber Security Centre (NCSC-UK), Germany's Federal Office for Information Security (BSI), Netherlands' National Cyber Security Centre (NCSC-NL), Computer Emergency Response Team NZ (CERT NZ), and NZ's National Cyber Security Centre (NCSC-NZ). 2023. Shifting the Balance of Cybersecurity Risk: Principles and Approaches for Security-by-Design and -Default. [https://www.cisa.gov/sites/default/files/2023-04/principles\\_approaches\\_for\\_security-by-design-default\\_508\\_0.pdf](https://www.cisa.gov/sites/default/files/2023-04/principles_approaches_for_security-by-design-default_508_0.pdf) (accessed 2024-08-02).
- [15] Cybersecurity & Infrastructure Security Agency (CISA), NSA, FBI, Australian Cyber Security Centre (ASD's ACSC), Canadian Centre for Cyber Security (CCCS), UK's National Cyber Security Centre (NCSC-UK), Computer Emergency Response Team New Zealand (CERT NZ), and New Zealand's National Cyber Security Centre (NCSC-NZ). 2023. The Case for Memory Safe Roadmaps: Why Both C-Suite Executives and Technical Experts Need to Take Memory Safe Coding Seriously. <https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/02/Final-ONCD-Technical-Report.pdf> (accessed 2025-02-18).
- [16] CodaSIP. 2024. CodaSIP delivers processor security to actively prevent the most common cyberattacks. <https://codasip.com/press-release/2023/10/31/codasip-delivers-processor-security-to-actively-prevent-cyberattacks/> (accessed 2024-01-30).
- [17] Caroline Collange. 2011. *Identifying scalar behavior in CUDA kernels*. Research Report. ENS Lyon. <https://hal.science/hal-00555134>
- [18] Caroline Collange. 2011. *Stack-less SIMT reconvergence at low cost*. Research Report. ENS Lyon. <https://hal.science/hal-00622654>
- [19] Caroline Collange. 2017. Simty: generalized SIMT execution on RISC-V. In *1st Workshop on Computer Architecture Research with RISC-V (CARRV 2017)*.
- [20] Caroline Collange, David Defour, and Yao Zhang. 2010. Dynamic Detection of Uniform and Affine Vectors in GPGPU Computations. In *Euro-Par 2009 – Parallel Processing Workshops*.
- [21] Riccardo D'Ambrosio. 2021. The Rust CUDA Project. <https://github.com/Rust-GPU/Rust-CUDA> (commit 8a6cb73, accessed 2024-02-01).
- [22] Bang Di, Jianhua Sun, and Hao Chen. 2016. A Study of Overflow Vulnerabilities on GPUs. In *International Conference on Network and Parallel Computing*.
- [23] Christopher Erb, Mike Collins, and Joseph L. Greathouse. 2017. Dynamic buffer overflow detection for GPGPUs. In *IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. <https://doi.org/10.1109/CGO.2017.7863729>
- [24] August Ernstsson, Lu Li, and Christoph Kessler. 2018. SkePU 2: Flexible and Type-Safe Skeleton Programming for Heterogeneous Parallel Systems. *International Journal of Parallel Programming* 46, 1 (2018). <https://doi.org/10.1007/s10766-017-0490-5>
- [25] Nathaniel Wesley Filardo, Brett F. Gutstein, Jonathan Woodruff, Sam Ainsworth, Lucian Paul-Trifu, Brooks Davis, Hongyan Xia, Edward Tomasz Napierala, Alexander Richardson, John Baldwin,

- David Chisnall, Jessica Clarke, Khilan Gudka, Alexandre Joannou, A. Theodore Markettos, Alfredo Mazzinghi, Robert M. Norton, Michael Roe, Peter Sewell, Stacey Son, Timothy M. Jones, Simon W. Moore, Peter G. Neumann, and Robert N. M. Watson. 2020. Cornucopia: Temporal Safety for CHERI Heaps. In *IEEE Symposium on Security and Privacy (SP)*. <https://doi.org/10.1109/SP40000.2020.00098>
- [26] Nathaniel Wesley Filardo, Brett F. Gutstein, Jonathan Woodruff, Jessica Clarke, Peter Rugg, Brooks Davis, Mark Johnston, Robert Norton, David Chisnall, Simon W. Moore, Peter G. Neumann, and Robert N. M. Watson. 2024. Cornucopia Reloaded: Load Barriers for CHERI Heap Temporal Safety. In *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Volume 2*. 18. <https://doi.org/10.1145/3620665.3640416>
- [27] Mark Gebhart, Stephen W. Keckler, Bruce Khailany, Ronny Krashinsky, and William J. Dally. 2012. Unifying Primary Cache, Scratch, and Register File Memories in a Throughput Processor. In *45th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. <https://doi.org/10.1109/MICRO.2012.18>
- [28] Syed Zohaib Gilani, Nam Sung Kim, and Michael Schulte. 2012. Power-efficient computing for compute-intensive GPGPU applications. In *21st International Conference on Parallel Architectures and Compilation Techniques (PACT 2012)*.
- [29] Jing Guo, Jeyarajan Thiagalingam, and Sven-Bodo Scholz. 2011. Breaking the GPU programming barrier with the auto-parallelising SAC compiler. In *Sixth Workshop on Declarative Aspects of Multicore Programming (DAMP)*. 10. <https://doi.org/10.1145/1926354.1926359>
- [30] Yanan Guo, Zhenkai Zhang, and Jun Yang. 2024. GPU memory exploitation for fun and profit. In *33rd USENIX Conference on Security*.
- [31] Bastian Hagedorn, Johannes Lenfers, Thomas Köhler, Xueying Qin, Sergei Gorlatch, and Michel Steuwer. 2020. Achieving High-Performance the Functional Way: A Functional Pearl on Expressing High-Performance Optimizations as Rewrite Strategies. *Proceedings of the ACM on Programming Languages* 4, ICFP (2020). <https://doi.org/10.1145/3408974>
- [32] Sarah Harris, Simon Cooksey, Michael Vollmer, and Mark Batty. 2023. Rust for Morello: Always-On Memory Safety, Even in Unsafe Code. In *37th European Conference on Object-Oriented Programming (ECOOP 2023)*. <https://doi.org/10.4230/LIPLcs.ECOOP.2023.39>
- [33] Troels Henriksen. 2021. Bounds Checking on GPU. *International Journal of Parallel Programming* 49, 6 (2021), 15. <https://doi.org/10.1007/s10766-021-00703-4>
- [34] Troels Henriksen, Niels G. W. Serup, Martin Elsmann, Fritz Henglein, and Cosmin E. Oancea. 2017. Futhark: purely functional GPU-programming with nested parallelism and in-place array updates. In *38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/3062341.3062354>
- [35] Graeme Jenkinson, Alfredo Mazzinghi, and Robert N. M. Watson. 2024. CHERI-based memory protection and compartmentalisation for web services on Morello. Technical Report, Capabilities Limited, [https://www.capabilitieslimited.co.uk/\\_files/ugd/893621\\_985a92a599bf41208e4c5710abcf3a68.pdf](https://www.capabilitieslimited.co.uk/_files/ugd/893621_985a92a599bf41208e4c5710abcf3a68.pdf) (accessed 2025-03-12).
- [36] Zhe Jia, Marco Maggioni, Jeffrey Smith, and Daniele Paolo Scarpazza. 2019. Dissecting the NVIDIA Turing T4 GPU via Microbenchmarking. arXiv:1903.07486 arXiv 1903.07486.
- [37] Alexandre Joannou, Jonathan Woodruff, Robert Kovacsics, Simon W. Moore, Alex Bradbury, Hongyan Xia, Robert N.M. Watson, David Chisnall, Michael Roe, Brooks Davis, Edward Napierala, John Baldwin, Khilan Gudka, Peter G. Neumann, Alfredo Mazzinghi, Alex Richardson, Stacey Son, and A. Theodore Markettos. 2017. Efficient Tagged Memory. In *International Conference on Computer Design (ICCD)*. <https://doi.org/10.1109/ICCD.2017.112>
- [38] Nicolas Joly, Saif ElSherei, and Saar Amar. 2020. Security Analysis of CHERI ISA. Technical Report, Microsoft Security Response Center, <https://github.com/microsoft/MSRC-Security-Research/blob/master/papers/2020/Security%20analysis%20of%20CHERI%20ISA.pdf> (commit 1372d4f, accessed 2025-03-12).
- [39] Ji Kim, Christopher Tornig, Shreesha Srinath, Derek Lockhart, and Christopher Batten. 2013. Microarchitectural mechanisms to exploit value structure in SIMT architectures. In *40th Annual International Symposium on Computer Architecture (ISCA 2013)*. <https://doi.org/10.1145/2485922.2485934>
- [40] Bastian Köpcke, Sergei Gorlatch, and Michel Steuwer. 2024. Descend: A Safe GPU Systems Programming Language. *Proceedings of the ACM on Programming Languages* 8, PLDI (2024). <https://doi.org/10.1145/3656411>
- [41] Jaewon Lee, Yonghae Kim, Jiashen Cao, Euna Kim, Jaekyu Lee, and Hyesoon Kim. 2022. Securing GPU via region-based bounds checking. In *49th International Symposium on Computer Architecture (ISCA)*. <https://doi.org/10.1145/3470496.3527420>
- [42] Sangpil Lee, Keunsoo Kim, Gunjae Koo, Hyeran Jeon, Won Woo Ro, and Murali Annavaram. 2015. Warped-Compression: Enabling Power Efficient GPUs through Register Compression. In *42nd Annual International Symposium on Computer Architecture (ISCA 2015)*. <https://doi.org/10.1145/2749469.2750417>
- [43] Yunsup Lee, Ronny Krashinsky, Vinod Grover, Stephen W. Keckler, and Krste Asanović. 2013. Convergence and scalarization for data-parallel architectures. In *IEEE/ACM International Symposium on Code Generation and Optimization (CGO 2013)*. <https://doi.org/10.1109/CGO.2013.6494995>
- [44] Christian Legnitto. 2025. Rebooting the Rust CUDA project. <https://rust-gpu.github.io/blog/2025/01/27/rust-cuda-reboot/> (accessed 2025-02-07).
- [45] Erik Lindholm, John Nickolls, Stuart Oberman, and John Montrym. 2008. NVIDIA Tesla: A Unified Graphics and Computing Architecture. *IEEE Micro* 28, 2 (2008). <https://doi.org/10.1109/MM.2008.31>
- [46] Zhenhong Liu, Syed Gilani, Murali Annavaram, and Nam Sung Kim. 2017. G-Scalar: Cost-Effective Generalized Scalar Execution Architecture for Power-Efficient GPUs. In *IEEE International Symposium on High Performance Computer Architecture (HPCA 2017)*. <https://doi.org/10.1109/HPCA.2017.51>
- [47] A. Theodore Markettos, John Baldwin, Ruslan Bukin, Peter G. Neumann, Simon W. Moore, and Robert N. M. Watson. 2021. Position Paper: Defending Direct Memory Access with CHERI Capabilities. In *Proceedings of the 9th International Workshop on Hardware and Architectural Support for Security and Privacy (HASP 2020)*. <https://doi.org/10.1145/3458903.3458910>
- [48] Trevor L. McDonnell, Manuel M.T. Chakravarty, Gabriele Keller, and Ben Lippmeier. 2013. Optimising purely functional GPU programs. In *18th ACM SIGPLAN International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/2500365.2500595>
- [49] Andrea Miele. 2016. Buffer overflow vulnerabilities in CUDA: a preliminary analysis. *Journal of Computer Virology and Hacking Techniques* 12, 2 (2016).
- [50] Matt Miller. 2019. Trends, challenges, and strategic shifts in the software vulnerability mitigation landscape. [https://github.com/Microsoft/MSRC-Security-Research/tree/master/presentations/2019\\_02\\_BlueHatIL](https://github.com/Microsoft/MSRC-Security-Research/tree/master/presentations/2019_02_BlueHatIL) (accessed 2024-01-25).
- [51] Sparsh Mittal. 2017. A Survey of Techniques for Architecting and Managing GPU Register File. *IEEE Transactions on Parallel and Distributed Systems* 28, 1 (2017). <https://doi.org/10.1109/TPDS.2016.2546249>
- [52] Matthew Naylor, Alexandre Joannou, Paul Metzger, A. Theodore Markettos, Simon W. Moore, and Timothy M. Jones. 2024. Advanced Dynamic Scalarisation for RISC-V GPGPUs. In *42nd IEEE International Conference on Computer Design (ICCD)*. <https://doi.org/10.1109/ICCD63220.2024.00047>
- [53] Hubert Nguyen. 2007. *GPU Gems 3* (first ed.). Addison-Wesley Professional.

- [54] NVIDIA Corporation. 2023. NVIDIA Ada GPU Architecture (v2.02). <https://images.nvidia.com/aem-dam/Solutions/geforce/ada/nvidia-ada-gpu-architecture.pdf> (accessed 2025-01-07).
- [55] NVIDIA Corporation. 2024. CUDA Code Samples. <https://developer.nvidia.com/cuda-code-samples> (accessed 2024-01-23).
- [56] NVIDIA Corporation. 2024. NVIDIA OpenCL SDK Code Samples. <https://developer.nvidia.com/opencl> (accessed 2024-01-23).
- [57] NVIDIA Corporation. 2024. Thrust, the CUDA C++ template library, v2.1.0. <https://developer.nvidia.com/thrust> (accessed 2024-12-18).
- [58] NVIDIA Corporation. 2025. CUDA C++ Programming Guide (Release 12.8). [https://docs.nvidia.com/cuda/pdf/CUDA\\_C\\_Programming\\_Guide.pdf](https://docs.nvidia.com/cuda/pdf/CUDA_C_Programming_Guide.pdf) (accessed 2025-02-13).
- [59] NVIDIA Corporation. 2025. NVIDIA HPC Compilers User's Guide (v25.5). <https://docs.nvidia.com/hpc-sdk/compilers/hpc-compilers-user-guide/index.html> (accessed 2025-07-21).
- [60] Sang-Ok Park, Ohmin Kwon, Yonggon Kim, Sang Kil Cha, and Hyunsoo Yoon. 2021. Mind control attack: Undermining deep learning with GPU memory exploitation. *Computers & Security* 102, C (2021). <https://doi.org/10.1016/j.cose.2020.102115>
- [61] Alex Rebert, Ben Laurie, Murali Vijayaraghavan, and Alex Richardson. 2025. Google Security Blog: Securing tomorrow's software: the need for memory safety standards. <https://security.googleblog.com/2025/02/securing-tomorrows-software-need-for.html> (accessed 2025-02-26).
- [62] Sangpil Lee, Keunsoo Kim, Gunjae Koo, Hyeran Jeon, Murali Annavaram, and Woo Woo Ro. 2017. Improving Energy Efficiency of GPUs through Data Compression and Compressed Execution. *IEEE Trans. Comput.* 66, 05 (2017). <https://doi.org/10.1109/TC.2016.2619348>
- [63] Blaise Tine, Varun Saxena, Santosh Srivatsan, Joshua R. Simpson, Fadi Alzamar, Liam Cooper, and Hyesoon Kim. 2023. Skybox: Open-Source Graphic Rendering on Programmable RISC-V GPUs. In *28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2023, Volume 3)*. <https://doi.org/10.1145/3582016.3582024>
- [64] Peter Rugg, Alexandre Joannou, Jonathan Woodruff, and Ivan Ribeiro. 2024. CheriCapLib Library. <https://github.com/CTSRD-CHERI/cheri-cap-lib> (commit 354a673, accessed 2024-10-30).
- [65] Peter Rugg, Jonathan Woodruff, Alexandre Joannou, and Simon W. Moore. 2024. A Suite of Processors to Explore CHERI-RISC-V Microarchitecture. In *27th Euromicro Digital System Design Conference*.
- [66] Peter David Rugg. 2023. *Efficient spatial and temporal safety for microcontrollers and application-class processors*. Technical Report UCAM-CL-TR-984. University of Cambridge, Computer Laboratory. <https://doi.org/10.48456/tr-984>
- [67] Blaise Tine, Krishna Praveen Yalamarthi, Fares Elsabbagh, and Kim Hyesoon. 2021. Vortex: Extending the RISC-V ISA for GPGPU and 3D-Graphics. In *54th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 2021)*. <https://doi.org/10.1145/3466752.3480128>
- [68] Jeff Vander Stoep, Android Security & Privacy Team, Chong Zhang, and Android Media Team. 2019. Google Security Blog: Queue the Hardening Enhancements. <https://security.googleblog.com/2019/05/queue-hardening-enhancements.html> (accessed 2024-01-25).
- [69] Kai Wang and Calvin Lin. 2017. Decoupled affine computation for SIMT GPUs. In *ACM/IEEE 44th International Symposium on Computer Architecture (ISCA 2017)*. <https://doi.org/10.1145/3079856.3080205>
- [70] Robert N.M. Watson, John Baldwin, David Chisnall, Tony Chen, Jessica Clarke, Brooks Davis, Nathaniel Filardo, Brett Gutstein, Graeme Jenkinson, Ben Laurie, Alfredo Mazinghi, Simon Moore, Peter G. Neumann, Hamed Okhravi, Alex Richardson, Alex Rebert, Peter Sewell, Laurence Tratt, Murali Vijayaraghavan, Hugo Vincent, and Konrad Witaszczyk. 2025. It Is Time to Standardize Principles and Practices for Software Memory Safety. *Commun. ACM* 68, 2 (2025). <https://doi.org/10.1145/3708553>
- [71] Robert N.M. Watson, Jonathan Woodruff, Peter G. Neumann, Simon W. Moore, Jonathan Anderson, David Chisnall, Nirav Dave, Brooks Davis, Khilan Gudka, Ben Laurie, Steven J. Murdoch, Robert Norton, Michael Roe, Stacey Son, and Munraj Vadera. 2015. CHERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization. In *2015 IEEE Symposium on Security and Privacy*. <https://doi.org/10.1109/SP.2015.9>
- [72] Robert N. M. Watson, Jessica Clarke, Peter Sewell, Jonathan Woodruff, Simon W. Moore, Graeme Barnes, Richard Grisenthwaite, Kathryn Stacer, Silviu Baranga, and Alexander Richardson. 2023. *Early performance results from the prototype Morello microarchitecture*. Technical Report UCAM-CL-TR-986. University of Cambridge, Computer Lab.
- [73] Robert N. M. Watson, Ben Laurie, and Alex Richardson. 2021. Assessing the Viability of an OpenSource CHERI Desktop Software Ecosystem. Technical Report, Capabilities Limited, [https://www.capabilitieslimited.co.uk/\\_files/ugd/f4d681\\_e0f23245dace466297f20a0dbd22d371.pdf](https://www.capabilitieslimited.co.uk/_files/ugd/f4d681_e0f23245dace466297f20a0dbd22d371.pdf) (accessed 2025-03-12).
- [74] Robert N. M. Watson, Simon W. Moore, Peter Sewell, and Peter Neumann. 2019. An Introduction to CHERI. University of Cambridge Technical Report, UCAM-CL-TR-941.
- [75] Robert N. M. Watson, Peter G. Neumann, Jonathan Woodruff, Michael Roe, Hesham Almatary, Jonathan Anderson, John Baldwin, Graeme Barnes, David Chisnall, Jessica Clarke, Brooks Davis, Lee Eisen, Nathaniel Wesley Filardo, Franz A. Fuchs, Richard Grisenthwaite, Alexandre Joannou, Ben Laurie, A. Theodore Markettos, Simon W. Moore, Steven J. Murdoch, Kyndylan Nienhuis, Robert Norton, Alexander Richardson, Peter Rugg, Peter Sewell, Stacey Son, and Hongyan Xia. 2023. *Capability Hardware Enhanced RISC Instructions: CHERI Instruction-Set Architecture (Version 9)*. Technical Report UCAM-CL-TR-987. University of Cambridge. <https://doi.org/10.48456/tr-987>
- [76] Jonathan Woodruff. 2014. *CHERI: a RISC capability machine for practical memory safety*. PhD Thesis, UCAM-CL-TR-858. University of Cambridge.
- [77] Jonathan Woodruff, Alexandre Joannou, Hongyan Xia, Anthony Fox, Robert M. Norton, David Chisnall, Brooks Davis, Khilan Gudka, Nathaniel W. Filardo, A. Theodore Markettos, Michael Roe, Peter G. Neumann, Robert N. M. Watson, and Simon W. Moore. 2019. CHERI Concentrate: Practical Compressed Capabilities. *IEEE Trans. Comput.* 68, 10 (2019). <https://doi.org/10.1109/TC.2019.2914037>
- [78] Ping Xiang, Yi Yang, Mike Mantor, Norm Rubin, Lisa R. Hsu, and Huiyang Zhou. 2013. Exploiting Uniform Vector Instructions for GPGPU Performance, Energy Efficiency, and Opportunistic Reliability Enhancement. In *27th ACM International Conference on Supercomputing (ICS 2013)*. <https://doi.org/10.1145/2464996.2465022>
- [79] Ayse Yilmazer, Zhongliang Chen, and David Kaeli. 2014. Scalar Waving: Improving the Efficiency of SIMD Execution on GPUs. In *IEEE 28th International Parallel and Distributed Processing Symposium (IPDPS 2014)*. <https://doi.org/10.1109/IPDPS.2014.22>

## A Artifact Appendix

### A.1 Abstract

This appendix contains instructions for running the NoCL benchmark suite on the SIMTIGHT evaluation SoC both in simulation and on FPGA. All software dependencies for running benchmarks in simulation can be met using the provided Dockerfile. However, the Dockerfile does not cover dependencies needed to synthesise and run the SoC on FPGA, due to the need for a commercial license (Quartus 22.1 Pro) and specialised hardware (a Terasic DE10-Pro revD development board).



## A.2 Artifact Check List (Meta Information)

- **Program:** The SIMTIGHT hardware description, NoCL library, and NoCL benchmark suite
- **Compilation:** GHC 9.2.1, Quartus 22.1 pro, Verilator 4.0.38, and GCC 11.4.0 for the hardware description, and CHERI-LLVM for the benchmarks (build script for CHERI-LLVM included)
- **Data set:** Included (part of the NoCL benchmark suite)
- **Run-time environment:** Tested with Ubuntu 20.04 and Docker 26.1.3
- **Hardware:** An x86 machine (64GB of RAM recommended when synthesising with Quartus, and 256GB of RAM when using Design Space Explorer) and a Terasic DE10-Pro revD FPGA development board
- **Output:** Benchmark self tests, SIMTIGHT hardware performance counters, and Quartus synthesis results
- **Experiments:** Mostly scripted with some manual steps
- **How much time is needed to prepare workflow (approximately)?:** 10 mins to build the Dockerfile and 30 mins to build CHERI-LLVM
- **How much time is needed to complete experiments (approximately)?:** 90 mins to run the test suite in simulation for a single SIMTIGHT configuration, 60 mins for a single FPGA synthesis, one minute to program the FPGA and run the benchmarks on FPGA, and a day to do a 16-seed synthesis sweep for all three SIMTIGHT configurations using Quartus Design Space Explorer
- **How much disk space required (approximately)?:** 8GB for the Docker container (with CHERI-LLVM), 32GB for the FPGA synthesis sweep, and 75GB for the Quartus installation
- **Publicly available?:** Yes
- **Code licenses (if publicly available)?:** Apache License 2.0
- **Archived:** <https://doi.org/10.17863/CAM.120202>

## A.3 Description

**A.3.1 How to Access.** Our artifact is publicly available at <https://github.com/CTSRD-CHERI/SIMTight>.

**A.3.2 Software Dependencies.** Our flow has been tested with Ubuntu 20.04 and Docker 26.1.3. When targeting FPGA (not simulation), the environment variable `QUARTUS_ROOTDIR` should point to a working installation of Quartus 22.1 Pro and `LM_LICENSE_FILE` should point to a valid license.

**A.3.3 Hardware Dependencies.** When targeting FPGA (not simulation), we require a Terasic DE10-Pro FPGA development board (revision D) connected to the host Linux machine via USB and visible as the sole device when running the `jtagconfig` command.

## A.4 Installation

**Step 1.** Recursively clone the `asplos26` branch of the repository:

```
git clone -b asplos26 --recursive \
  https://github.com/CTSRD-CHERI/SIMTight
```

**Step 2.** To satisfy the dependencies of the project, build and enter the Docker container:

```
cd SIMTight && make shell
```

This takes around 10 minutes in the first instance, but only a matter of seconds thereafter.

**Step 3.** Download and build the exact version of the CHERI compiler used in this paper, and add it to the environment:

```
cd cheri-tools
./build-cheri.sh
source ./add-cheri-tools-to-path.sh
cd ..
```

This takes around 30 minutes to complete. Note that if you exit the Docker shell, and later re-enter it, you will not need to rebuild the CHERI tools but you will need to re-add them to the environment.

## A.5 Experiment Workflow

**Step 4.** To enable CHERI in SIMTIGHT, modify `inc/Config.h` to contain:

```
#define EnableTaggedMem 1
#define EnableCHERI 1
#define UseClang 1
```

The SIMTIGHT evaluation SoC can then be built and tested in simulation with the command:

```
cd test && ./test.sh && cd ..
```

This takes around 90 minutes to complete. If the final line of the output is `All tests passed` then the artifact is functioning as expected. In simulation, the `test.sh` script runs all benchmarks on small datasets to ensure timely completion. To obtain performance counters for each benchmark, pass the `--stats` option to the script (though this may not be particularly meaningful in simulation due to small datasets).

**Step 4.** The three main configurations of SIMTIGHT used in this paper – **Baseline**, **CHERI**, and **CHERI (Optimised)** – are defined in `scripts/sweep.py` and can be tested as follows.

```
cd scripts && ./sweep.py test && cd ..
```

This will take around 5 hours to complete. Test results for the three configurations are written to the file `test/test.log`. The final line of each test run should be `All tests passed`.

**Step 5.** Assuming `QUARTUS_ROOTDIR` points to a working installation of Quartus 22.1 Pro and `LM_LICENSE_FILE` points to a valid license, an FPGA bitfile for the SIMTIGHT evaluation SoC can be built as follows.

```
cd de10-pro && make && cd ..
```

Assuming the FPGA is connected via USB and visible as the sole device when running the `jtagconfig` command, the FPGA can be programmed with the bitfile using the command:

```
cd de10-pro && make download-sof && cd ..
```

This takes around 30 seconds to complete. To run all benchmarks on FPGA and obtain all hardware performance counters:

```
cd test && ./test.sh --stats --fpga-d --apps-only && cd ..
```



This takes around 30 seconds to complete, assuming that the FPGA image has already been built (the script will build and download the FPGA bitfile if it has not already been done, but a stepwise approach is more advisable when exercising the flow for the first time). The final line of the output should be `All tests passed`.

**Step 6.** To reproduce the results for the three main configurations of SIMTIGHT used in the paper:

```
cd scripts && ./sweep.py bench && cd ..
```

This takes around three hours to complete, as each configuration must be synthesised from scratch. The output is three

`.bench` files in the `test` directory, one for each configuration. The final line of each `.bench` file should be `All tests passed`.

**Step 7.** For FPGA synthesis results, we use Quartus Design Space Explorer to synthesise each configuration of SIMTIGHT across 16 different seeds, selecting the design with the highest Fmax. This long process can be initiated with the command:

```
cd scripts && ./sweep.py synth && cd ..
```

This should be run on a modern server with at least 256GB of RAM, where it takes around a day to complete. A summary of the synthesis results is written to `de10-pro/synth.log`.