




A Formal CHERI-C Semantics for Verification

Seung Hoon Park , Rekha Pai , and Tom Melham 

Department of Computer Science, University of Oxford, Oxford, UK
{seunghoon.park, rekha.pai, tom.melham}@cs.ox.ac.uk

Abstract. CHERI-C extends the C programming language by adding *hardware capabilities*, ensuring a certain degree of memory safety while remaining efficient. Capabilities can also be employed for higher-level security measures, such as software compartmentalization, that have to be used correctly to achieve the desired security guarantees. As the extension changes the semantics of C, new theories and tooling are required to reason about CHERI-C code and verify correctness. In this work, we present a formal memory model that provides a memory semantics for CHERI-C programs. We present a generalised theory with rich properties suitable for verification and potentially other types of analyses. Our theory is backed by an Isabelle/HOL formalisation that also generates an OCaml executable instance of the memory model. The verified and extracted code is then used to instantiate the parametric *Gillian* program analysis framework, with which we can perform concrete execution of CHERI-C programs. The tool can run a CHERI-C test suite, demonstrating the correctness of our tool, and catch a good class of safety violations that the CHERI hardware might miss.

Keywords: CHERI-C · Hardware Capabilities · Memory Model · Semantics · Theorem Proving · Verification

1 Introduction

Despite having been developed more than 40 years ago, C remains a widely used programming language owing to its efficiency, portability, and suitability for low-level systems code. The language’s lack of inherent memory safety, however, has been the source of many serious issues [18]. While there have been significant efforts aimed at vulnerability mitigation, memory safety issues remain widespread, with a recent study stating that 70% of security vulnerabilities are caused by memory safety issues [31].

The Capability Hardware Enhanced RISC Instructions (CHERI) project offers an alternative model that provides better memory safety [44]. Its main features include a new machine representation of C pointers called *capabilities* and extensions to existing Instruction Set Architectures (ISA) that enable the secure manipulation of capabilities. Capabilities are in essence memory addresses bound to additional safety-related metadata, such as access permissions and bounds on the memory locations that can be accessed. As the hardware performs the safety checks on capabilities, legacy C programs compiled and run




Fig. 1: Simplified CHERI Capability Layouts

on CHERI architecture, i.e. CHERI-C code, acquire hardware-ensured spatial memory safety, while retaining efficiency. Porting code from one language to another generally requires significant efforts. But porting C codes to CHERI-C requires little, if any, changes to the original code to ensure the code runs on CHERI hardware [36, 39].

In 2019, the UK announced its *Digital Security by Design* programme with £190 million of funding distributed over more than 26 research projects and 5 industrial demonstrators [6] to ‘radically update the foundation of our insecure digital computing infrastructure, by demonstrating that mainstream processor technology ... can be updated to include new security technologies based on the CHERI Architecture’ [5]. A cornerstone of the programme is Morello [4], a CHERI-enabled prototype developed by Arm.

Over the several years that lead to the realisation of Morello, there were several design revisions made to the hardware; examples are depicted in Fig. 1. The refined designs used methods for compression of bounds that reduced cache footprints and improved overall performance while minimising incompatibility. Morello uses a very similar design to the compressed scheme for capabilities depicted in Fig. 1b, with the overall bit-representation of the layout differing slightly. Future capability designs may possibly incorporate a different bit-representation design, provided there are improvements in performance or compatibility. Due to the ever-changing design of capability bit-representations, it seems best to have an *abstract* representation of capabilities, so that CHERI-based verification tools can remain modular.

Checking for memory safety issues of legacy C code can, of course, be achieved using existing analysis tools for C, but there are new problems that arise when such code is run on CHERI hardware. Because the pointer and memory representations are fundamentally different in a CHERI architecture, there are non-trivial differences in the semantics between C and CHERI-C.

To illustrate this point, consider the C code in Listing 1.1. This code segment performs `memcpy` twice: once from `a` to `b`, where pointers/capabilities are stored misaligned in `b`, then from `b` to `c`, where pointers/capabilities are stored correctly again in `c`. In standard C, there are no problems accessing the pointer stored in `c`. But in CHERI-C, misaligned capabilities in memory are invalidated. That means the address and meta-data of the misaligned capabilities are accessible,

but such capabilities can no longer be dereferenced [41]. While `c` will contain the same capability value as that of `a`, the capability stored in `c` is invalidated. Thus, the last line will trigger an ‘invalid tag’ exception when the code is executed on ARM Morello and other CHERI-based machines.

```

1 #include <stdlib.h>
2 #include <string.h>
3 void main(void) {
4     int *n = calloc(sizeof(int), 1);
5     int **a = malloc(sizeof(int *));
6     *a = n;
7     int **b = malloc(sizeof(int *) * 2);
8     int **c = malloc(sizeof(int *));
9     memcpy((char *) b + 1, a, sizeof(int *));
10    memcpy(c, (char *) b + 1, sizeof(int *));
11    int x = **c;
12 }
```

Listing 1.1: C code example

Of course, existing C analysis tools cannot catch these cases, as such tools are not only unaware of the changes in the semantics that capabilities bring, but also the code is not problematic in conventional C. Moreover, while CHERI ensures spatial safety by the hardware, CHERI is still incapable of catching temporal safety violations, such as Use After Free (UAF) violations. There exists work that attempt to address temporal safety [11, 17, 42], but they are either a software-implemented solution [42], where overall performance is inevitably affected, or ongoing work [11]. There is, therefore, a need for program analysis tools that correctly integrate the semantics of CHERI-C.

To the best of our knowledge, there is no prior work on formalising a CHERI-C memory model. The Cerberus C work [30] is primarily designed to capture pointer provenance of C programs and uses CHERI-C as a reference for pointer provenance, but the tool lacks a formal CHERI-C memory model. ESBMC is a verification tool that supports CHERI-C code [15]. But support for tagged memory does not yet exist; ESBMC would not be able to catch the ‘invalid tag’ exception in the code in Listing 1.1. Furthermore, ESBMC’s memory model is not formally verified. Users of ESBMC must trust that the implementation of the memory model and its underlying theory are correct. SAIL formalisations for each CHERI architectures exist [3, 8, 9], but they only capture the low-level semantics of the architecture and not high-level C constructs such as allocation.

In this paper, we introduce a formal CHERI-C memory model that captures the memory semantics of the CHERI-C language. In Sect. 3, We formalise the memory and its operations and prove essential properties that provide correctness guarantees. We provide a rigorous logical formalisation of the CHERI-C memory model in Isabelle/HOL [32] (in Sect. 4.1) and use the code generation feature to generate a verified OCaml instance of the memory model [21]. We then show, in Sect. 4.2, the practical aspects of this work by providing the memory model to, and thereby instantiating, Gillian [20], a general, parametric verifica-

tion framework that supports concrete and symbolic execution and verification based on separation logic, backed by rich correctness properties. In Sect. 5, we demonstrate that the tool can capture the semantics of CHERI-C programs correctly. A discussion on the existing works can be found in Sect. 6 while Sect. 7 concludes this paper mentioning possible future directions. We first start with an introduction to the CHERI architecture.

2 CHERI

CHERI extends a conventional ISA by introducing *capabilities* which are essentially pointers that come along with metadata to restrict memory access. The ISA now has additional hardware instructions and exceptions that operate over capabilities. Register sets are extended to include capability registers, instructions are added that reference the capability registers, and custom hardware exceptions are added to block operations that would violate memory safety. Designs of CHERI capabilities have refined over the past several years and have been incorporated in several existing architectures, such as MIPS and RISC-V [40]. All CHERI-extended ISAs have been formally defined using the SAIL specification language, in which the logic of machine instructions and memory layout have been defined formally in a first-order language [13].

Regardless of the layout, CHERI capabilities include three important types of high-level information, in addition to a 64-bit address:

- **Permissions.** Permissions state what kind of operations a capability can perform. Loading from memory and storing to memory are examples of permissions a capability may possess.
- **Bounds.** Bounds stipulate the memory region that the address part of a capability can reference. The lower bound stipulates the lowest address that a capability may access, and the upper bound stipulates the highest address.
- **Tag.** Stored separately from the other components of a capability, the tag states the validity of the capability it is attached to. Capabilities with invalid tags can hold data but cannot be dereferenced. Attempts to forge capabilities out of thin air result in a tag-invalidated capability.

Fig. 1a show a 256-bit representation of a capability, which was one of the earlier designs. The lower and upper bounds are represented using the base and length fields. Here, the lower bound is the address stated by the base field, and the upper bound is the address in the base field plus the length field. Permissions and other metadata are stored in the remaining fields as a bit vector. The capability’s tag bit exists separately from the capability. Tag bits are, in practice, stored separately from the main memory where capabilities reside, so users cannot manipulate the tag bits of capabilities stored in memory. Furthermore, overwriting capabilities stored in memory with non-capability values invalidates their tag bits, which ensures capabilities cannot be forged out of thin air.

This representation, in theory, exercises a high level of compatibility with existing C code. But performance, particularly with regards to caching, is reduced

due to the size of the capability representation [43]. Refined designs ultimately resulted in a capability that utilises a floating-point-based lossy compression technique on the bounds [43], such as the one depicted in Fig. 1b. In many cases, the upper bits of the address fields are most likely to overlap with those of the lower and upper bounds. Knowing this, bounds can be compressed by having the upper bits of their fields depend on that of the address, which means only the lower bits need to be stored.

The lossy compression of bounds may result in some incompatibility. Bounds may no longer be represented exactly, and changes in the address field may result in an unintentional change in the bounds. Nonetheless, such representations give an acceptable level of compatibility, provided aggressive pointer arithmetic optimisations are avoided. The Morello processor incorporates a similar compression-based design in its architecture, though sizes of each field differ [12].

The added capability-aware instructions operate over capabilities. Conventional load and store operations are extended to first check that the tag, permissions, and bounds of the capability are all valid. Violations result in triggering a capability-related hardware exception. There are additional operations to access or change the tag, permissions, and bounds. To ensure spatial memory safety, these operations can, at most, make the conditions for execution more restrictive; they cannot grant that which was not previously available. For instance, one cannot lower the lower bound of a capability to access a region that was inaccessible before, or grant a store permission that was unset beforehand. Because of how tags work for capabilities stored in memory, one cannot grant capabilities larger bounds or more permissions by manipulating the memory—attempting this results in tag invalidation.

Library support for CHERI has grown over the past few years. In particular, a software stack for CHERI-C that utilises a custom Clang compiler now exists [41]. Users can compile their program either in ‘purecap’ mode, where all pointers in programs are replaced with capabilities, or in ‘hybrid’ mode, where both pointers and capabilities co-exist within the program. Because operations that change the fields of a capability does not generally exist in standard C, Clang incorporates additional CHERI libraries of operations that users may use to access or mutate capabilities.

3 CHERI-C Memory Model

Incorporating hardware-enabled spatial safety requires significant changes to the C memory model. Pointer designs must be extended to incorporate bounds, metadata, and the out-of-band tag bit. The memory, i.e. heap, must also be able to distinguish the main memory and the tagged memory. Operations with respect to the heap must also be defined such that tag preservation and invalidation are incorporated appropriately.

In this section, we provide a generalised theory for the CHERI-C memory model. We identify the type and value system used by the memory model. We then define the heap and the core memory operations. Finally, we state some

essential properties of the heap and the operations that (1) characterises the semantics and (2) states what types of verification or analyses could be supported. We make the assumption that we work on a ‘purecap’ environment, where all pointers have been replaced with capabilities.

3.1 Design

The CHERI-C memory model is inspired by that of CompCert [26]. The beauty of CompCert is that it is a verified C compiler. The internal components, which include the block-offset based memory model, are formalised in a theorem prover, with many of its essential properties verified. Using CompCert’s memory model as a basis, we design the CHERI-C memory model by providing extensions to ensure the modelling of correct semantics and the capture of safety violations:

- **Capability Values.** In addition to the standard primitive types, we incorporate abstract capabilities as values. We also incorporate capability fragments to provide semantics to higher-level memory actions like `memcpy`, which should preserve tags if copied correctly and invalidate otherwise [41].
- **Extended Operations.** Basic memory actions such as `load` and `store` now work on capabilities and will trigger the correct capability-related exception when required.
- **Tagged Memory.** Tags in memory are stored separately from the main heap, as could be seen by the formal CHERI-MIPS SAIL model [9]. So we provide a separate mapping for tagged memory for storing capability tags.
- **Freed Regions.** The standard CompCert memory model can mark which memory regions are valid but lacks the ability to distinguish which regions are marked as ‘Freed’. We incorporate freed regions as a means to catch temporal safety violations.

3.2 Type and Value System

Figure 2 shows the formalisation of CHERI-C types and values. Types τ are analogous to chunks in CompCert terms. Types comprise primitive types (e.g. $U8_\tau$,

$$\begin{aligned}
 \tau &\triangleq U8_\tau \mid S8_\tau \mid \dots \mid U64_\tau \mid S64_\tau \mid Cap_\tau \\
 MCap &\triangleq \mathcal{B} \times \mathbb{Z} \times md \\
 Cap &\triangleq MCap \times \mathbb{B} \\
 \mathcal{V}_C &\triangleq U8_V \ :: \ 8 \text{ bits} \mid \dots \\
 &\quad \mid S64_V \ :: \ 64 \text{ sbits} \\
 &\quad \mid Cap_V \ :: \ Cap \\
 &\quad \mid CapF_V \ :: \ Cap \times \mathbb{N} \\
 &\quad \mid Undef \\
 \mathcal{V}_M &\triangleq Byte \ :: \ 8 \text{ bits} \\
 &\quad \mid MCapF \ :: \ MCap \times \mathbb{N}
 \end{aligned}$$

Fig. 2: CHERI-C Types and Values

$S64_\tau$, etc.) and a capability type Cap_τ . We define a function $|\cdot| : \tau \rightarrow \mathbb{N}$ that returns, in terms of bytes, the size of the type. For Cap_τ , the value is not fixed but requires that it must be divisible by 16. This requirement allows capabilities with 128- and 256-bit representations to have a valid size.

$MCap$ represents a *memory capability* value and is represented as a tuple (b, i, m) , which comprises the block identifier $b \in \mathcal{B}$, offset $i \in \mathbb{Z}$, and metadata $m \in md$, where md represents the bounds and permissions. Here, \mathcal{B} must be a countable set. Offsets are represented as integers, as CHERI allows out-of-bounds addresses, where the address may be lower than the lower bound. Because capabilities stored in memory have their tag bit stored elsewhere, we make the distinction between memory capabilities and *tagged capabilities*, Cap , which is a capability $((b, i, m), t)$ that contains the tag bit $t \in \mathbb{B}$.

Unlike those of CompCert, CHERI-C values \mathcal{V}_C are given type distinctions to ensure: (1) types can be inferred directly, and (2) they contain the correct values at all times. From a practical standpoint, this ensures that the proof of correctness of memory operations can be simplified, and bounded arithmetic operations can be implemented correctly. Capability values Cap_V and capability fragment values $CapF_V$ also exist as values. Provided some capability value $C \in Cap_V$, capability fragment values $C_n \in CapF_V$ correspond to the n -th byte of the capability C . For both cases, instead of fixing their representation concretely, we represent them abstractly using a tuple. This representation ensures that conversion to a compressed representation could be achieved when needed while avoiding the need to fix to one particular bit representation. Furthermore, this approach provides a reasonable way to correctly define `memcpy`, where capability tags must be preserved if possible. While capability fragments are extended structures of capabilities, operations that can be performed on capability fragments are limited. Finally, we have $Undef$, which represents invalid values. These values may appear when, for example, the user calls `malloc` and immediately tries to load the undefined contents. The idea behind incorporating capability fragments values is heavily inspired by the work from [25].

Because values are given a type distinction, identifying the types of values is straightforward. For capability fragments, we have two choices: they may either be a $U8_\tau$ or $S8_\tau$ type. Capability fragments are essentially bytes, so operations over capability fragments can be treated as if they were a $U8_\tau$ or $S8_\tau$ type. Since $Undef$ does not correspond to a valid value, it is not assigned a type.

$$\begin{aligned}
 \text{CapErr} &\triangleq \text{TagViolation} \mid \text{PermitLoadViolation} \mid \dots \\
 \text{LogicErr} &\triangleq \text{UseAfterFree} \mid \text{MissingResource} \mid \dots \\
 \text{Err} &\triangleq \text{CapErr} \mid \text{LogicErr} \\
 \mathcal{R} \rho &\triangleq \text{Succ } \rho \\
 &\quad \mid \text{Fail } \text{Err}
 \end{aligned}$$

Fig. 3: CHERI-C Errors

Memory operations, such as `load` and `store`, are defined so that, upon failure, the operation returns the type of error that lead to the failure. In general, partial functions, or function using the option type, can model function failure but cannot state what caused the failure. As such, the operations use the return type $\mathcal{R} \rho$, where ρ is a generic return type. For CHERI-C, we make the distinction between errors caused by capabilities, denoted by `CapErr`, and errors caused by the language, denoted by `LogicErr`. Figure 3 depicts the formalised Errors system used by the memory model.

3.3 Memory

We now formalise the memory. We use CompCert’s approach of using a union type $\mathcal{V}_{\mathcal{M}}$ that can represent either a byte or a byte fragment of a memory capability. Then it is possible to create a memory mapping $\mathbb{N} \rightarrow \mathcal{V}_{\mathcal{M}}$.¹ We also create a separate mapping of type $\mathbb{N} \rightarrow \mathbb{B}$ for tagged memory. When the user attempts to store a capability, it will be converted into a memory capability and then stored in the memory mapping. Separately, the tag bit will be stored in the tagged memory. When the tag bit is stored, adjustments are made to ensure tags are only stored in capability-size-aligned offsets.

To ensure we can catch temporal safety violations, we need to be able to make distinctions between blocks that are freed and blocks that are valid. One way to encode this is as follows: a block b may point to either a freed location (i.e. $b \mapsto \emptyset$), or point to the pair of maps we defined earlier. The idea is that if a block identifier points to a freed block, attempts to load such a block will trigger a ‘Use After Free’ violation and would otherwise point to a valid mapping pair. Ultimately, the heap has the following form:

$$\mathcal{H} : \mathcal{B} \rightarrow ((\mathbb{N} \rightarrow \mathcal{V}_{\mathcal{M}}) \times (\mathbb{N} \rightarrow \mathbb{B}))_{\emptyset}$$

3.4 Operations

We define the core memory operations, or *actions*, of the memory model. We use the same result type \mathcal{R} given in Fig. 3 instead of using a partial function to give the type of error, should the operation fail.

The memory actions $A_{\mathcal{C}} = \{\text{alloc}, \text{free}, \text{load}, \text{store}\}$ are given below with their respective signatures:

- `alloc` : $\mathcal{H} \rightarrow \mathbb{N} \rightarrow \mathcal{R} (\mathcal{H} \times \text{Cap})$
- `free` : $\mathcal{H} \rightarrow \text{Cap} \rightarrow \mathcal{R} (\mathcal{H} \times \text{Cap})$
- `load` : $\mathcal{H} \rightarrow \text{Cap} \rightarrow \tau \rightarrow \mathcal{R} (\mathcal{V}_{\mathcal{C}})$
- `store` : $\mathcal{H} \rightarrow \text{Cap} \rightarrow \mathcal{V}_{\mathcal{C}} \rightarrow \mathcal{R} (\mathcal{H})$

¹The notation \rightarrow denotes a partial map. Offsets in heaps are \mathbb{N} , whereas offsets stored in capabilities are \mathbb{Z} . Operations check whether the offsets are in bounds, which requires offsets to be non-negative. This means valid offset values can be converted from \mathbb{Z} to \mathbb{N} without issues.

The function $\text{alloc } \mu n = \text{Succ } (\mu', c)$ takes a heap μ and size n input and produces a fresh capability c and the updated heap μ' as output. The bounds of c are determined by n . In the case of compressed capabilities, a sufficiently large n may result in the upper bound being larger than what was requested. The capability c is also given the appropriate permissions and a valid tag bit. Like that of CompCert, alloc is designed to never fail, provided that the countable set \mathcal{B} has infinite elements.

The function $\text{free } \mu c = \text{Succ } (\mu', c')$ takes a heap μ and capability $c = ((b, i, m), t)$ as input. Upon success, the operation will return the updated heap, where we now have $b \mapsto \emptyset$. The capability c' is also updated such that the tag bit of c is invalidated. This conforms to the CHERI-C design stated in [41]. We note that c should also be a valid capability, that is—at the very least—the tag bit should be set, and the offset should be within the capability bounds. The function free may fail if the block is invalid or already freed, even if the capability itself was valid. In such case, free returns a logical error.

The function $\text{load } \mu c t = \text{Succ } v$ takes a heap μ , capability c and type t as input, where t is the type the user wants to load. Upon success, the operation will return the value v from the memory, where v has the corresponding type t .² Before load attempts to access the block provided by c , it first checks that c has sufficient permissions to load. We use the CHERI-MIPS SAIL implementation of the CL[C] instruction [40] for the capability checks, implementing the extra checks provided that $t = \text{Cap}_\tau$. Once the capability checks are done, the operation attempts to access the blocks and the mappings, failing and returning the appropriate logical error if they do not exist.

When accessing both the main memory and tagged memory, there are a number of cases to consider. When loading primitive values, it is important that the region about to be loaded is all of *Byte* and not of *MCapF* type. Thus, before loading the values, we check whether the contiguous region in memory are all of *Byte* type. If this is not the case, load will return *Undef*. For capability fragments, the cell in memory has to be an *MCapF*. Finally for capabilities, not only do the contiguous cells have to be of *MCapF* type, but (1) they must have the same memory capability value, and (2) the fragment values must all be a sequence forming $\{0, 1, \dots, |\text{Cap}_\tau| - 1\}$. The idea is that even if the contiguous cells have the same memory capability values, they do not form a valid capability if the fragments are not stored in order. After all the checks, the tagged memory will be accessed, where the tag value is retrieved.³ The loaded memory capability and tag bit are then combined to form a tagged capability, which load returns.

The function $\text{store } \mu c v = \text{Succ } \mu'$ takes a heap μ , capability c , and value v . Upon success, the operations will return the updated heap μ' . Like load , store performs the necessary capability checks based on CHERI-MIPS' CS[C] instruction and attempts to access the blocks and mappings afterwards, returning the appropriate exception upon failure. For storing primitive values and capability

²For capability fragments, the corresponding type may be either $U8_\tau$ or $S8_\tau$.

³The tagged memory does not need to be accessed if c does not have a capability load permission. In such case, the loaded capability will have an invalidated tag.

fragment values, the main memory mapping will simply be updated to contain the values, and the associated tagged memories will be invalidated. For primitive values that are not bytes, the values will be converted into a sequence of bytes, where each byte in the list will be stored contiguously in memory. For a capability fragment value, it will be stored in the cell as an $MCapF$ type, where the tag value of the fragment will be stripped when storing in memory. Finally, for capability values, the value will be split into a list comprising $|Cap_\tau| - 1$ memory capability fragments, with the fragment value forming a sequence $\{0, 1, \dots, |Cap_\tau| - 1\}$, and a tag bit. The main memory will store the list of memory fragments contiguously, and the tagged memory will store the tag value in the corresponding capability-aligned tagged memory.

3.5 Properties

In the previous section, we have articulated a formal CHERI-C memory model, explaining how the heap is structured and how the operations are defined. It is essential that the formalisation we provided is correct and is also suitable for verification or other types of analyses. In this section, we first discuss the properties of the memory. We then discuss the properties of the operations themselves, primarily concerned with correctness.

When we observe the memory, it is important that we always work with a valid one, i.e. the memory is *well-formed*. In our formalisation, we require that all tags in the tagged memory are stored in a capability-aligned location. The well-formedness relation \mathcal{W}_f^C is defined as follows:

$$\mathcal{W}_f^C(\mu) \equiv \forall b \in \text{dom}(\mu). b \mapsto (c, t) \longrightarrow \forall x \in \text{dom}(t). x \bmod |Cap_\tau| = 0$$

The well-formedness property must hold when the heap is initialised and when memory operations mutate the heap. That is, provided μ_0 is the initialised heap where all mappings are empty, $\alpha \in A_C$ is a memory action, v are the arguments of the memory operation α and μ' is one of the return values denoting the updated heap, we have the following properties:

$$\mathcal{W}_f^C(\mu_0)$$

$$\mathcal{W}_f^C(\mu) \Longrightarrow \alpha \mu v = \text{Succ } \mu' \Longrightarrow \mathcal{W}_f^C(\mu')$$

The two properties above ensure that the heap is well-formed throughout the execution of the CHERI-C program.

For the correctness of the operations, we primarily consider soundness and completeness:

- If the inputs are valid for operation $\alpha \in A_C$ then the action should succeed.
- If the action α succeeds, the inputs provided to the operations are valid.
- If the inputs are invalid for the operation α , then the action should fail and return the correct error.

The first and second points are simple soundness and completeness properties. The third point is important in that the input may be problematic in many ways. For example, the NULL capability has an invalid tag bit, invalid bounds, and no permissions. The function `load` will fail if provided with the NULL capability, as it violates many of the checks. Because the SAIL specification states that tags are always checked first, the error must be a `TagViolation` type.

Next, we need to ensure successive operations yield the desired result. The primary properties to consider are the *good variable* laws [26]; examples of properties encoding this law include *load after allocation*, *load after free*, and *load after store*. It is worth mentioning there are some caveats. For example, the *load after store* case no longer guarantees that you will retrieve the same value you stored, unlike CompCert’s load after store property in [26], since the value that was stored and to be loaded again could have been either a capability or capability fragment. In such cases, the tag bit may become invalidated due to insufficient permissions on the capability, or because storing capability fragments resulted in the tagged memory being cleared. The solution is to divide the general property into a primitive value case and a capability-related value case. Ultimately, the idea is to prove that the loaded value is *correct* rather than exact, i.e. capability-related values when loaded with have the correct tag value.

Finally, we have properties suitable for verification. We note that the memory \mathcal{H} can be instantiated as a separation algebra by providing the partial commutative monoid (PCM) $(\mathcal{H}, \uplus, \mu_0)$, where \uplus is the disjoint union of two heaps and μ_0 is the empty initialised heap. For tools that rely on using partial memories, it is also imperative to show that the well-formedness property is compatible with memory composition:

$$\mathcal{W}_f^c(\mu_1 \uplus \mu_2) \implies \mathcal{W}_f^c(\mu_1) \wedge \mathcal{W}_f^c(\mu_2)$$

We also note that the current heap design keeps track of *negative* resources [28], which may potentially be useful for incorrectness logic based verification [33].

4 Application

The overall memory model provided in Sect. 3 has been designed to be applicable for verification tools. In this section, we explain how we use the theory provided above to create a verified, executable instance of the memory model. We then explain how this executable model can be used to instantiate a tool called Gillian [20]. Using the instantiated tool, we demonstrate the concrete execution of CHERI-C programs with the desired behaviour.

4.1 Isabelle/HOL

Isabelle/HOL is an interactive theorem prover based on classical Higher Order Logic (HOL) [32]. We use Isabelle/HOL to formalise the entirety of the CHERI-C memory model discussed in Sect. 3. Types, values, heap structure,

etc. were implemented, memory operations were defined, and properties relating to the heap and the operations were proven. Memory capabilities, tagged capabilities, and capability fragments were represented using records, a form of tuple with named fields. For code generation, we instantiated the block type \mathcal{B} to be \mathbb{Z} . For showing that \mathcal{H} is an instance of a separation algebra, we use the `cancellative_sep_algebra` class [23] and prove that the heap model is an instance. This proof ultimately shows that \mathcal{H} forms a PCM. Proving that well-formedness is compatible with memory composition is stated slightly differently. The `cancellative_sep_algebra` class takes in a total operator \cdot_t instead of a partial one and requires a ‘separation disjunction’ binary operator $\#$, which states disjointedness. Ultimately, the compatibility property can be given as:

$$\mu_1 \# \mu_2 \implies \mathcal{W}_f^c(\mu_1 \cdot_t \mu_2) \implies \mathcal{W}_f^c(\mu_1) \wedge \mathcal{W}_f^c(\mu_2)$$

For partial mappings of the form $A \multimap B$, we use Isabelle/HOL’s finite mapping type `(’a, ’b)mapping` [22]. To ensure we obtain an OCaml executable instance of the memory model, we use the Containers framework [27], which generates a Red-Black Tree mapping provided the abstract mapping in Isabelle/HOL. All definitions in Isabelle were either defined to be code-generatable to begin with (i.e. definitions should not comprise quantifiers or non-constructive constants like the Hilbert choice operation *SOME*), or code equations were provided and proven to ensure a sound code generation [21]. For bounded machine words, which is required for formalising the primitive values, we use Isabelle/HOL’s word type `’a word`, where `’a` states the length of the word [14]. Types like `’a word`, `nat`, `int` and `string` were also transformed to use OCaml’s `Zarith` and native string library for efficiency [21].

4.2 Gillian

Gillian is a high-level analysis framework, theoretically capable of analysing a wide range of languages. The framework allows concrete and symbolic execution, verification based on Separation Logic, and bi-abduction [28]. The crux of the framework lies in its parametricity, where the tool can be instantiated by simply providing a compiler front end and OCaml-based memory models of the language. So far, CompCert C and JavaScript have both been instantiated for Gillian, giving birth to Gillian-C and Gillian-JS.

The underlying theoretical foundation of Gillian has its essential correctness properties like soundness and completeness already proven [20, 29]. Thus, users who instantiate the tool only need to prove the correctness of the implementation of their compiler and memory models to ensure the correctness of the entire tool. From the perspective of someone trying to instantiate Gillian with their compiler and memory models, it is essential to understand the underlying intermediate language GIL and the overall memory model interface used by Gillian.

GIL GIL is the GOTO-based Intermediate Language used by Gillian which is used for all types of analyses the tool supports. For concrete execution, GIL

supports basic GOTO constructs and assertions. For symbolic execution, the GIL grammar is extended to support path cutting, i.e. assumptions, and generation of symbolic variables. For separation logic based verification, the GIL grammar is further extended to support core predicates and user-defined predicates [28] that can be utilised to form separation logic based assertions. Furthermore, function specifications in the Hoare-triple form $\{P\}f(\bar{x})\{Q\}$ can be provided, where P and Q are separation logic based assertions.

Note that Gillian uses a value set \mathcal{V} which differs from that used in the CHERI-C memory model. As we are only interested in the values used in the CHERI-C memory model, it is possible to implement a thin conversion layer between the two value systems. We note that a list of GIL values also constitutes a GIL value, so arguments for functions can be expressed as a single GIL value. This is important when understanding the memory model layout of Gillian.

Memory Model Memory Models in Gillian have a specific definition and have properties that state what kind of analysis is supported. Proving that the provided memory models satisfy certain properties is essential in understanding what the instantiated tool supports.

Gillian differentiates between concrete and symbolic memory models, which are used for concrete and symbolic execution, respectively. As we are concerned with concrete execution, we will consider only concrete memory models here.

At the highest level, there are two kinds of memory model properties: *executional* and *compositional*. The *executional* memory model states properties a memory model must have for whole-program execution, and the *compositional* memory model states properties a memory model must have for separation logic based symbolic verification. Each paper in the Gillian literature states slightly different definitions for the memory models [20, 28, 29, 37]—in Definitions 1 and 2 below, we present unified, consistent definitions for each of the memory model properties. We ignore contexts, as there exists only one context in concrete memories, which is the GIL boolean value `true`.

Definition 1. (*Execution Memory Model*). Given the set of GIL values \mathcal{V} and an action set A , an execution memory model $M(\mathcal{V}, A) \triangleq (|M|, \mathcal{W}_f, \underline{ea})$ comprises:

1. a set of memories $|M| \ni \mu$
2. a well-formedness relation $\mathcal{W}_f \subseteq |M|$, with $\mathcal{W}_f(\mu)$ denoting μ is well-formed
3. the action execution function $\underline{ea} : A \rightarrow |M| \rightarrow \mathcal{V} \rightarrow \mathcal{R} (|M| \times \mathcal{V})$

Definition 2. (*Compositional Memory Model*). Given the set of GIL values \mathcal{V} and core predicate set Γ , a compositional memory model, $M(\mathcal{V}, A_\Gamma) \triangleq (|M|, \mathcal{W}_f, \underline{ea}_\Gamma)$ comprises:

1. a partial commutative monoid (PCM) $(|M|, \cdot, 0)$
2. A well-formedness relation $\mathcal{W}_f \subseteq |M|$ with the following property:

$$\mathcal{W}_f(\mu_1 \cdot \mu_2) \implies \mathcal{W}_f(\mu_1) \wedge \mathcal{W}_f(\mu_2)$$

3. *the predicate action execution function $\underline{ea}_\Gamma : A_\Gamma \rightarrow |M| \rightarrow V \rightarrow \mathcal{R} (|M| \times V)$*

First, we note that for concrete execution, Gillian also uses the return type \mathcal{R} in the action execution function \underline{ea} .⁴ For \mathcal{W}_f defined in Definition 1, the main properties that must be satisfied are Properties 3.1, 3.2, and 3.6 in [29].

The PCM requirement is required to show that the heap forms a separation algebra [16]. \mathcal{W}_f is extended to state that memory composition must also be well-formed. Finally, the predicate action execution function \underline{ea}_Γ provides a way to frame on and off parts of the memory, though they are not required for concrete execution as they are not part of the GIL concrete execution grammar.

Using the CHERI-C memory model we defined earlier, we can show that our model conforms to both Definitions 1 and 2. Let A_C be the set of memory actions, \mathcal{H} be the memory, \underline{ea}_C be the action execution function of the CHERI-C memory model, and \mathcal{W}_f^C be the well-formedness relation. Then we observe that $(\mathcal{H}, \mathcal{W}_f^C, \underline{ea}_C)$ forms an execution memory model. We note that Properties 3.1 and 3.2 in [29] are satisfied, and Property 3.6 is trivial in that operations that return errors do not return an updated heap. We also note that the memory model also conforms to a compositional memory model, as we have the PCM $(\mathcal{H}, \uplus, \mu_0)$ along with the well-formedness property being composition-compatible. The predicate action execution function is not required to be given, as the concrete execution of Gillian does not utilise this feature.

4.3 Compiler

We implemented a CHERI-C to GIL compiler by utilising ESBMC’s GOTO language. The idea is that ESBMC uses its own intermediate representation for bounded model checking, which is the GOTO language. CHERI-enabled ESBMC uses Clang as a front end to generate the GOTO language. In our case we can build a GOTO to GIL compiler instead of building a CHERI-C compiler from scratch. The GOTO language is very similar to GIL in that they are both goto-based languages and uses single static assignment. For most parts, the compilation process is straightforward. As ESBMC’s GOTO language is typed while the CHERI-C memory model is untyped—untyped in the sense that the memory model does not support user-defined types like `structs`—we make sure that capability arithmetic and casts are applied correctly by inferring the sizes of the user-defined types.

5 Experimental Results

In Sect. 4, we have provided a way to instantiate the Gillian tool, where we obtain a concrete CHERI-C model using Isabelle/HOL and a CHERI-C to GIL

⁴In the Gillian literature, it is stated that \mathcal{R} can return both a return value and an error. The OCaml implementation of Gillian slightly differs from this and is more similar to \mathcal{R} used for the CHERI-C memory model.

compiler that utilises ESBMC’s GOTO language. Our framework can demonstrate that higher-level memory actions—such as `memcpy()`, which preserves tags when applicable—can be implemented. Furthermore, we can run concrete instances of programs that use `memcpy()` to show they emit the expected behaviour. This also means the tool can catch the `TagViolation` exception that is triggered in Listing 1.1. Our tool also allows capability-related functions defined in `cheriintrin.h` and `cheri.h`, to be usable, i.e. it is possible to call operations such as `cheri_tag_get()` and `cheri_tag_clear()`.

Filename	GC	GCC	AM	BMC
<code>buffer_overflow.c</code>	✓	✓	✓	✓
<code>dangling_ptr.c</code>	✓	✓	×	✓
<code>double_free.c</code>	✓	✓	×	✓
<code>invalid_free.c</code>	× ⁵	✓	✓	✓
<code>misaligned_ptr.c</code>	✓	✓	✓	×
<code>listing_1.c</code>	×	✓	✓	×

Table 1: Violation detection

Filename	Time(s)
<code>libc_malloc.c</code>	8.585
<code>libc_memcpy.c</code>	1.698
<code>libc_memmove.c</code>	0.318
<code>libc_string.c</code>	0.315

Table 2: GCC runtime performance

Table 1 shows a list of safety violations that Gillian-C, our tool, the ARM Morello hardware, and CHERI-ESBMC—labelled as GC, GCC, AM, and BMC, respectively—all catch. We observe that Morello fails to catch temporal safety violations such as dangling pointers and double frees. For the invalid free case, where we attempt to free a pointer not produced by `malloc`, we discovered a bug in the Gillian-C tool that fails to catch this violation.⁵ Gillian-C does not return any errors for the program in Listing 1.1, which is to be expected, as this is not problematic for conventional C. Finally, we observe that CHERI-ESBMC fails to catch the last two violations that relating to tag invalidation.

Table 2 shows the runtime performance of running the CHERI-C library test suites, based on the Clang CHERI-C test suite [1]. Tests were conducted on a machine running Fedora 34 on an 11th Gen Intel Core i7-1185G7 CPU with 31.1 GB RAM, with trace logging enabled. We note that when the test cases were executed on Morello without any modifications to the code, all of the tests terminated instantaneously without any issues. In the `libc_malloc.c` test case, we reduced the scope of the test⁶ to ensure the tool terminates within a reasonable time, though the performance can be drastically improved by turning logging off, e.g. the `libc_malloc.c` case would only take 0.686 seconds. For the remaining tests, we made modifications to the code to ensure the compiler can correctly produce the GIL code, and we made sure to preserve all the edge cases covered by the original tests. For example, in `libc_memcpy.c` we made sure to test all cases where both `src` and `dst` capabilities were aligned and misaligned in the beginning and the end, which affected tag preservation. We observed that no assertions were violated, and we also observed that the same

⁵The bug has since been fixed after a discussion with the developers [7].

⁶In particular, we reduced `max` from the `libc_malloc.c` case in [1] from 20 to 9.

code when run in Morello also resulted in no assertion violations, demonstrating a faithful implementation of CHERI-C semantics.

6 Related Work

The CompCert C memory model [26], CH₂O memory model [24], and Tuch’s C memory model [38] are C memory models formalised in a theorem prover, each focusing on different aspects of verification. Our model mostly draws inspiration from these models, extending such work to support CHERI-C programs.

VCC, which internally uses the typed C memory model [19], and CHERI-ESBMC [15] are designed with automated verification of C programs via symbolic execution in mind—in particular, CHERI-ESBMC supports hybrid settings and compressed capabilities in addition to purecap settings and uncompressed capabilities. Both tools rely on a memory model that is not formally verified, so the tools have components that must be trusted.

7 Conclusion and Future Work

We have provided a formal CHERI-C memory model and demonstrated its utility for verification. We formalised the entire theory in Isabelle/HOL and generated an executable instance of the memory model, which was then used to instantiate a CHERI-C tool. The result led to a concrete execution tool that is robust in terms of the properties that are guaranteed both by the tool and by the memory model. We demonstrated its practicality by running CHERI-C based test suites, capturing memory safety violations, and comparing the results with actual CHERI hardware—namely the physical Morello processor.

Currently there are a number of limitations provided by the memory model. Capability arithmetic is limited only to addition and subtraction, but the heap can be extended to incorporate mappings from blocks to physical addresses and vice versa. This provides a way to extend capability arithmetic. While the theory incorporates abstract capabilities, compression is still under work. We believe, however, that the abstract design itself does not need to change. It may be possible to utilise the compression/decompression work to convert between the two forms [2] when needed whilst retaining our design for the operations.

This theory serves as a starting point for much potential future work. A compositional symbolic memory model can be built from this design to enable symbolic execution and verification in Gillian. As we have already proven the core properties, proving the remaining properties for the extended model will allow automated separation logic based verification of CHERI-C programs.

Acknowledgements We are very grateful to the Gillian team, in particular, Sacha-Élie Ayoun, for providing assistance with instantiating the Gillian tool. We also thank Fedor Shmarov and Franz Brauße for providing assistance with building and modifying the ESBMC tool. This work was funded by the UKRI programme on Digital Security by Design (Ref. EP/V000225/1, SCorCH [10]).

Data-Availability Statement The Isabelle/HOL formalisation of the CHERI-C memory model described in Sect. 4.1 is available in the Isabelle Archive of Formal Proofs [34]. The artefact of the evaluation provided in Sect. 5, which includes Gillian-CHERI-C itself, CHERI-ESBMC, and other tools, is archived in the Zenodo open-access repository [35].

References

1. CHERI C Tests. <https://github.com/CTSRD-CHERI/cheri-c-tests>
2. cheri-compressed-cap. <https://github.com/CTSRD-CHERI/cheri-compressed-cap>
3. CHERI RISC-V Sail model. <https://github.com/CTSRD-CHERI/sail-cheri-riscv>
4. CHERI: The Arm Morello Board, <https://www.cl.cam.ac.uk/research/security/ctsr/cheri/cheri-morello.html>
5. CHERI: The Digital Security by Design (DSbD) Initiative, <https://www.cl.cam.ac.uk/research/security/ctsr/cheri/dsbd.html>
6. Digital Security by Design Challenge – UKRI, <https://www.ukri.org/our-work/our-main-funds/industrial-strategy-challenge-fund/artificial-intelligence-and-data-economy/digital-security-by-design-challenge/>
7. fix the behaviour of free, <https://github.com/GillianPlatform/Gillian/commit/6fa87b046f8d8f328c20b89cbdf1a00944da3fe>, GillianPlatform/Gillian@6fa87b0
8. Morello Sail specification. <https://github.com/CTSRD-CHERI/sail-morello>
9. Sail model of CHERI-MIPS ISA. <https://github.com/CTSRD-CHERI/sail-cheri-mips>
10. SCorCH: Secure Code for Capability Hardware, <https://scorch-project.github.io>
11. Armv8.5-A Memory Tagging Extension. Tech. rep. (Jun 2021), <https://documentation-service.arm.com/static/624ea580caabfd7b3c13e23f?token=>
12. ARM Ltd.: Arm Architecture Reference Manual Supplement Morello for A-Profile Architecture (2022), <https://documentation-service.arm.com/static/61e577e1b691546d37bd38a0?token=>
13. Armstrong, A., Bauereiss, T., Campbell, B., Reid, A., Gray, K.E., Norton, R.M., Mundkur, P., Wassell, M., French, J., Pulte, C., Flur, S., Stark, I., Krishnaswami, N., Sewell, P.: ISA Semantics for ARMv8-a, RISC-v, and CHERI-MIPS. Proc. ACM Program. Lang. **3**(POPL) (Jan 2019)
14. Beeren, J., Fernandez, M., Gao, X., Klein, G., Kolanski, R., Lim, J., Lewis, C., Matichuk, D., Sewell, T.: Finite Machine Word Library. Archive of Formal Proofs (Jun 2016), https://isa-afp.org/entries/Word_Lib.html, Formal proof development
15. Brauße, F., Shmarov, F., Menezes, R., Gadelha, M.R., Korovin, K., Reger, G., Cordeiro, L.C.: ESBMC-CHERI: Towards Verification of C Programs for CHERI Platforms with ESBMC. In: Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis. p. 773–776. ISSTA 2022, Association for Computing Machinery, New York, NY, USA (2022)

16. Calcagno, C., O’Hearn, P.W., Yang, H.: Local Action and Abstract Separation Logic. In: 22nd Annual IEEE Symposium on Logic in Computer Science (LICS 2007). pp. 366–378 (2007)
17. Chisnall, D.: Towards a Safe, High-Performance Heap Allocator (Sep 2022), <https://soft-dev.org/events/cheritech22/slides/Chisnall.pdf>, presented at CHERI Technical Workshop 2022
18. Chisnall, D., Rothwell, C., Watson, R.N., Woodruff, J., Vadera, M., Moore, S.W., Roe, M., Davis, B., Neumann, P.G.: Beyond the PDP-11: Architectural Support for a Memory-Safe C Abstract Machine. SIGPLAN Not. **50**(4), 117–130 (Mar 2015)
19. Cohen, E., Moskal, M., Tobies, S., Schulte, W.: A Precise Yet Efficient Memory Model For C. Electronic Notes in Theoretical Computer Science **254**, 85–103 (2009). <https://doi.org/https://doi.org/10.1016/j.entcs.2009.09.061>, proceedings of the 4th International Workshop on Systems Software Verification (SSV 2009)
20. Fragoso Santos, J., Maksimović, P., Ayoun, S.E., Gardner, P.: Gillian, Part i: A Multi-Language Platform for Symbolic Execution. In: Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation. p. 927–942. PLDI 2020, Association for Computing Machinery, New York, NY, USA (2020)
21. Haftmann, F.: Code generation from Isabelle/HOL theories (Dec 2021), <https://isabelle.in.tum.de/doc/codegen.pdf>
22. Haftmann, F., Krauss, A., Kunčar, O., Nipkow, T.: Data Refinement in Isabelle/HOL. In: Blazy, S., Paulin-Mohring, C., Pichardie, D. (eds.) Interactive Theorem Proving. pp. 100–115. Springer Berlin Heidelberg, Berlin, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39634-2_10
23. Klein, G., Kolanski, R., Boyton, A.: Mechanised Separation Algebra. In: Beringer, L., Felty, A. (eds.) Interactive Theorem Proving. pp. 332–337. Springer Berlin Heidelberg, Berlin, Heidelberg (2012). https://doi.org/10.1007/978-3-642-32347-8_22
24. Krebbers, R.: A Formal C Memory Model for Separation Logic. Journal of Automated Reasoning **57**(4), 319–387 (Dec 2016). <https://doi.org/10.1007/s10817-016-9369-1>
25. Krebbers, R., Leroy, X., Wiedijk, F.: Formal C Semantics: CompCert and the C Standard. In: Klein, G., Gamboa, R. (eds.) Interactive Theorem Proving. pp. 543–548. Springer International Publishing, Cham (2014)
26. Leroy, X., Appel, A.W., Blazy, S., Stewart, G.: The CompCert Memory Model, Version 2. Research Report RR-7987, INRIA (Jun 2012)
27. Lochbihler, A.: Light-Weight Containers for Isabelle: Efficient, Extensible, Nestable. In: Blazy, S., Paulin-Mohring, C., Pichardie, D. (eds.) Interactive Theorem Proving. pp. 116–132. Springer Berlin Heidelberg, Berlin, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39634-2_11
28. Maksimovic, P., Ayoun, S.E., Santos, J.F., Gardner, P.: Gillian, part II: real-world verification for javascript and C. In: Silva, A., Leino, K.R.M. (eds.) Proceedings of the 33rd Computer Aided Verification International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Part II. Lecture Notes in Computer Science, vol. 12760, pp. 827–850. Springer (2021). https://doi.org/10.1007/978-3-030-81688-9_38
29. Maksimovic, P., Santos, J.F., Ayoun, S.E., Gardner, P.: Gillian: A Multi-Language Platform for Unified Symbolic Analysis (2021). <https://doi.org/10.48550/ARXIV.2105.14769>, <https://arxiv.org/abs/2105.14769>

30. Memarian, K., Gomes, V.B.F., Davis, B., Kell, S., Richardson, A., Watson, R.N.M., Sewell, P.: Exploring C Semantics and Pointer Provenance. *Proc. ACM Program. Lang.* **3**(POPL) (Jan 2019)
31. Miller, M.: Trends, challenges, and strategic shifts in the software vulnerability mitigation landscape (Feb 2019), <https://msrnd-cdn-stor.azureedge.net/bluehat/bluehatil/2019/assets/doc/Trends%20Challenges%20and%20Strategic%20Shifts%20in%20the%20Software%20Vulnerability%20Mitigation%20Landscape.pdf>, presented at BlueHat IL
32. Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL - A Proof Assistant for Higher-Order Logic. [ecture Notes in Computer Science, Springer (2002). <https://doi.org/10.1007/3-540-45949-9>
33. O’Hearn, P.W.: Incorrectness logic. *Proc. ACM Program. Lang.* **4**(POPL) (Dec 2019). <https://doi.org/10.1145/3371078>, <https://doi.org/10.1145/3371078>
34. Park, S.H.: A Formal CHERI-C Memory Model. *Archive of Formal Proofs* (Nov 2022), https://isa-afp.org/entries/CHERI-C_Memory_Model.html, Formal proof development
35. Park, S.H., Pai, R., Melham, T.: Artifact for Paper A formal CHERI-C Semantics for Verification (Jan 2023). <https://doi.org/10.5281/zenodo.7504675>, <https://doi.org/10.5281/zenodo.7504675>
36. Richardson, A.: Porting C/C++ software to Morello (Sep 2022), <https://soft-dev.org/events/cheritech22/slides/Richardson.pdf>, presented at CHERI Technical Workshop 2022
37. Santos, J.F., Maksimovic, P., Ayoun, S.E., Gardner, P.: Gillian: Compositional Symbolic Execution for All. *CoRR* **abs/2001.05059** (2020), <https://arxiv.org/abs/2001.05059>
38. Tuch, H.: Formal Verification of C Systems Code. *Journal of Automated Reasoning* **42**(2), 125–187 (Apr 2009). <https://doi.org/10.1007/s10817-009-9120-2>
39. Watson, R., Laurie, B., Richardson, A.: Assessing the Viability of an Open- Source CHERI Desktop Software Ecosystem. *Tech. rep., Capabilities Limited* (Sep 2021), <https://www.capabilitieslimited.co.uk/pdfs/20210917-capltd-cheri-desktop-report-version1-FINAL.pdf>
40. Watson, R.N.M., Neumann, P.G., Woodruff, J., Roe, M., Almatary, H., Anderson, J., Baldwin, J., Barnes, G., Chisnall, D., Clarke, J., et al.: Capability Hardware Enhanced RISC Instructions: CHERI Instruction-Set Architecture (Version 8). *Tech. rep., University of Cambridge, Cambridge, England* (Oct 2020), <https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-951.pdf>
41. Watson, R.N.M., Richardson, A., Davis, B., Baldwin, J., Chisnall, D., Clarke, J., Filardo, N., Moore, S.M., Napierala, E., Sewell, P., Neumann, P.G.: CHERI C/C++ Programming Guide. *Tech. rep., University of Cambridge, Cambridge, England* (Jun 2020), <https://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-947.pdf>
42. Wesley Filardo, N., Gutstein, B.F., Woodruff, J., Ainsworth, S., Paul-Trifu, L., Davis, B., Xia, H., Tomasz Napierala, E., Richardson, A., Baldwin, J., Chisnall, D., Clarke, J., Gudka, K., Joannou, A., Theodore Marketos, A., Mazzinghi, A., Norton, R.M., Roe, M., Sewell, P., Son, S., Jones, T.M., Moore, S.W., Neumann, P.G., Watson, R.N.M.: Cornucopia: Temporal Safety for CHERI Heaps. In: 2020 IEEE Symposium on Security and Privacy (SP). pp. 608–625 (2020). <https://doi.org/10.1109/SP40000.2020.00098>

43. Woodruff, J., Joannou, A., Xia, H., Fox, A., Norton, R.M., Chisnall, D., Davis, B., Gudka, K., Filardo, N.W., Marketos, A.T., Roe, M., Neumann, P.G., Watson, R.N.M., Moore, S.W.: CHERI Concentrate: Practical Compressed Capabilities. *IEEE Transactions on Computers* **68**(10), 1455–1469 (2019). <https://doi.org/10.1109/TC.2019.2914037>
44. Woodruff, J., Watson, R.N.M., Chisnall, D., Moore, S.W., Anderson, J., Davis, B., Laurie, B., Neumann, P.G., Norton, R., Roe, M.: The CHERI Capability Model: Revisiting RISC in an Age of Risk. In: 2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA). pp. 457–468. IEEE (Jun 2014)