Verified Compilation of C Programs with a Nominal Memory Model

YUTING WANG, Shanghai Jiao Tong University, China

LING ZHANG, Shanghai Jiao Tong University, China

ZHONG SHAO, Yale University, USA

JÉRÉMIE KOENIG, Yale University, USA

Memory models play an important role in verified compilation of imperative programming languages. A representative one is the block-based memory model of CompCert—the state-of-the-art verified C compiler. Despite its success, the abstraction over memory space provided by CompCert’s memory model is still primitive and inflexible. In essence, it uses a fixed representation for identifying memory blocks in a global memory space and uses a globally shared state for distinguishing between used and unused blocks. Therefore, any reasoning about memory must work uniformly for the global memory; it is impossible to individually reason about different sub-regions of memory (i.e., the stack and global definitions). This not only incurs unnecessary complexity in compiler verification, but also poses significant difficulty for supporting verified compilation of open or concurrent programs which need to work with contextual memory, as manifested in many previous extensions of CompCert.

To remove the above limitations, we propose an enhancement to the block-based memory model based on *nominal techniques*; we call it the *nominal memory model*. By adopting the key concepts of nominal techniques such as *atomic names* and *supports* to model the memory space, we are able to 1) generalize the representation of memory blocks to any types satisfying the properties of atomic names and 2) remove the global constraints for managing memory blocks, enabling flexible memory structures for open and concurrent programs.

To demonstrate the effectiveness of the nominal memory model, we develop a series of extensions of CompCert based on it. These extensions show that the nominal memory model 1) supports a general framework for verified compilation of C programs, 2) enables intuitive reasoning of compiler transformations on partial memory; and 3) enables modular reasoning about programs working with contextual memory. We also demonstrate that these extensions require limited changes to the original CompCert, making the verification techniques based on the nominal memory model easy to adopt.

CCS Concepts: • **Software and its engineering** → **Software verification**; **Compilers**; • **Theory of computation** → **Program verification**.

Additional Key Words and Phrases: Memory Models, Nominal Techniques, Verified Compilation

ACM Reference Format:

Yuting Wang, Ling Zhang, Zhong Shao, and Jérémie Koenig. 2022. Verified Compilation of C Programs with a Nominal Memory Model. *Proc. ACM Program. Lang.* 6, POPL, Article 25 (January 2022), 31 pages. <https://doi.org/10.1145/3498686>

Authors’ addresses: Yuting Wang, John Hopcroft Center for Computer Science, Shanghai Jiao Tong University, China, yuting.wang@sjtu.edu.cn; Ling Zhang, John Hopcroft Center for Computer Science, Shanghai Jiao Tong University, China, ling.zhang@sjtu.edu.cn; Zhong Shao, Yale University, USA, zhong.shao@yale.edu; Jérémie Koenig, Yale University, USA, jeremie.koenig@yale.edu.

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

© 2022 Copyright held by the owner/author(s).

2475-1421/2022/1-ART25

<https://doi.org/10.1145/3498686>

1 INTRODUCTION

Memory models are critical components for formalizing the semantics of imperative programming languages. They determine an abstraction over memory and provide necessary operations for manipulating memory states at the corresponding level of abstraction. In the setting of verified compilation of imperative programs, a memory model with appropriate abstraction not only reduces the complexity of verification but also enables rich forms of semantics preservation. Therefore, a lot of research has been done for developing memory models for verified compilation, ranging from highly abstract models that treat memory states as mappings from addresses to values to highly concrete ones that capture the linear layout of physical memory [Tuch et al. 2007].

Among the memory models developed so far, the most representative one is the *block-based memory model* [Leroy et al. 2012; Leroy and Blazy 2008] defined in CompCert—the state-of-the-art verified C compiler [Leroy 2021]. It provides a medium level of abstraction—neither too abstract nor too concrete—by modeling memory space as a collection of contiguous memory regions (also called *blocks*). Isolation between different memory blocks is simply captured by giving each of those memory blocks a unique identifier (called its *block id*). Furthermore, pointers are naturally defined as pairs of block ids and offsets to memory cells, and definitions of pointer operations follow in a straightforward manner. The block-based memory model has been highly successful. By uniformly applying it to all of CompCert’s languages, the developers of CompCert were able to verify over 20 compiler passes containing advanced optimizations. It has also been widely adopted in other verification projects, including various extensions of CompCert for supporting compositionality and concurrency (e.g., [Besson et al. 2017; Jiang et al. 2019; Kang et al. 2016; Sevcík et al. 2011, 2013; Song et al. 2020; Stewart et al. 2015; Wang et al. 2019, 2020]) and formalization of language semantics for program verification and program analysis (e.g., [Appel 2011; Gu et al. 2015, 2018]).

1.1 Deficiency of the Block-Based Memory Model

Despite its previous success, CompCert’s block-based memory model is still quite primitive and inflexible, making it difficult to support intuitive and flexible reasoning in verified compilation. This includes the following points:

- **Inflexibility 1: Fixed Representation of Block IDs.** In the block-based memory model, block ids are represented by a fixed type, i.e., positive numbers. With this uniform and fixed representation, it is impossible to distinguish between memory regions playing different roles (such as the stack, heap, and memory regions for global definitions). Therefore, it is difficult to reason about specific parts of memory in isolation.
- **Inflexibility 2: Sequential Numbering of Valid Blocks.** In any memory state, a special positive number named `nextblock` provides the next available block id. The global memory space is divided by `nextblock` into two parts: blocks with ids less than `nextblock` have already been allocated (called *valid blocks*), while the remaining blocks are waiting to be allocated (called *invalid blocks*). In any program semantics, the ids of valid blocks must be numbered sequentially by $[1, \dots, \text{nextblock} - 1]$. This creates unintuitive and unnecessary dependency between valid blocks playing different roles.
- **Inflexibility 3: Global Constraint for Allocation.** In any memory state, there is only a single `nextblock` for assigning new block ids upon allocation. In the setting where multiple open programs or threads work on the same memory state, every open program or thread must keep track of how other programs or threads modify `nextblock`. This global constraint imposes a complex dependency between open programs and threads.

Because of the complex dependencies created by the above inflexibilities, verification of any compiler transformation must treat the memory space *as a whole*. This incurs significant difficulties

both in verified compilation of whole programs and in that of open programs. Specifically, in the former setting, many compiler transformations only work on certain sub-regions of memory (e.g., transformations on stack frames). However, because of **Inflexibilities 1 and 2**, the reasoning must be lifted to the whole memory, therefore becoming significantly more involved and less intuitive. In the latter setting, one would like to get a compositional approach to verifying open or concurrent programs. However, such an approach must be compatible with the evolution of nextblock which is inherently non-compositional by **Inflexibilities 2 and 3**. This problem plagues many projects on extending CompCert to support compositionality or concurrency (e.g., CASCompCert [Jiang et al. 2019] and Thread-Safe CompCertX [Gu et al. 2018]). To circumscribe it, various ad hoc modifications to the block-based memory model were invented, leading to verification results that are overly technical and less reusable.

1.2 Nominal Memory Model and Verified Compilation

In this paper, we develop a systematic, clean and lightweight approach to eliminating the above inflexibilities of the block-based memory model. Our approach is based on *nominal techniques* [Gabbay and Pitts 2002; Pitts 2016]. Specifically, by adopting the very basic concepts of nominal techniques to formally model the block-based memory space, we obtain a natural generalization of the block-based memory model which we call the *nominal memory model*. Our key ideas include the following:

- **Generalizing Block IDs to Atomic Names.** In nominal techniques, the elements of countably infinite sets are used to represent available names; they are called *atomic names*. By generalizing the fixed type of block ids to a type parameter representing countably infinite sets, we allow block ids to be freely instantiated by any kinds of atomic names. For example, because the positive numbers are countably infinite, the original block ids become a special case after the generalization. More importantly, the type of block ids may be instantiated with rich data types, by which we can divide global memory into separate memory regions with clear roles. This eliminates **Inflexibility 1**.
- **Generalizing Valid Blocks to Supports.** In nominal techniques, dependency of objects on names is described via the concept of *supports*. In its most basic form, the *support* of an object is a finite set of atomic names that the object contains. By generalizing valid blocks to a type parameter that satisfies the basic requirements of supports, we allow new names to be freely generated as long as they are fresh w.r.t. the support. This eliminates **Inflexibility 2**.
- **Eliminating Global Constraints via Supports.** Like the type of atomic names, the support type can also be instantiated with sophisticated data types. This enables formalization of memory space with complex structures (e.g., multiple call stacks) and separate growth of memory regions for individual open modules or threads. This eliminates **Inflexibility 3**.

To demonstrate that the nominal memory model can indeed enable more flexible and intuitive reasoning about compiler transformations on both whole programs and open programs, we develop a series of extensions of CompCert based on it. First, we develop *Nominal CompCert*, a full extension of CompCert with a basic nominal memory model that is parameterized over the types of block ids and supports. Nominal CompCert provides a skeleton for developing richer extensions of CompCert that exploit the full power of the above generalization. Second, we instantiate the block ids and supports with rich data types for explicitly dividing global memory into separate regions based on their roles. With this specialized nominal memory space, we give precise definitions of transformations on partial memory as *concrete injection functions*, resulting in more intuitive proofs for the key passes of CompCert. Third, we exploit rich instantiation of the support type to verify compilation of open concurrent programs. Specifically, we develop *Multi-Stack CompCert*, a combination of Nominal CompCert and Stack-Aware CompCert [Wang et al. 2019] with additional

support for multiple call stacks. With the flexibility to independently grow individual stacks, Multi-Stack CompCert becomes the first compiler that supports verified compilation of C programs all the way down to multi-stack machines, which provides a new foundation for thread-safe compilation of Certified Concurrent Abstraction Layers [Gu et al. 2018].

1.3 Contributions

We summarize our contributions as follows:

- We introduce the nominal memory model, a natural extension of the block-based memory based on nominal techniques (Sec. 3.2). It enables flexible formalization of block identifiers and memory space and eliminates global assumptions on memory states in the original block-based memory model. The block-based memory model aligns well with such generalization as it becomes an instance of the nominal memory model.
- We develop Nominal CompCert, an extension to CompCert with its full compilation chain verified based on the nominal memory model (Sec. 3.3). Nominal CompCert is a general framework for verified compilation of C programs in the following sense: with its interface for supporting nominal names established, the compiler and its proofs work regardless what instances of block ids (names) and valid blocks (supports) are provided through this interface.
- We develop two extensions of Nominal CompCert that illustrate its power in verified compilation of whole programs and open programs. The first one enables intuitive reasoning of compiler transformations on partial memory (Sec. 4), while the second one enables modular reasoning about open programs working with contextual memory (Sec. 5).
- We demonstrate that the above developments require limited changes to CompCert, making verification techniques based on the nominal memory model lightweight and easy to adopt.

The developments presented in this paper are based on CompCert version 3.8 and fully formalized in the Coq proof assistant version 8.12.0. The complete artifact in Coq is located at <https://doi.org/10.5281/zenodo.5553752>.

1.4 Structure of the Paper

In Sec. 2, we first introduce the necessary background of the block-based memory model and elaborate on its problems. We then present the nominal memory model in its full detail and discuss how to extend CompCert to Nominal CompCert in Sec. 3. In the subsequent two sections, we successively introduce the key applications of the nominal memory model and Nominal CompCert, including verification of transformations on partial memory (in Sec. 4) and verification of open programs working with contextual memory (in Sec. 5). We provide an evaluation of our proof effort in Sec. 6. Finally, we give a comparison with existing work in Sec. 7 and conclude in Sec. 8.

2 BACKGROUND AND APPROACH

We first give a brief introduction to the block-based memory model and verified compilation of C programs based on it. We then elaborate on the deficiency of this memory model that prevents it from supporting intuitive proofs for verified compilation of whole programs and open programs.

2.1 The Block-Based Memory Model

In the block-based memory model, the memory space is represented as a countably infinite set of blocks where each block is given a unique identifier called its *block id* (denoted by b). CompCert uses positive numbers to represent block ids. The entire memory space is divided into two parts by a special block id called `nextblock`. A block with id less than `nextblock` has been allocated (called a *valid block*). Otherwise, it has not been allocated yet (called an *invalid block*). `nextblock`—initially

Definition $\text{block} : \text{Type} :=$ positive.	$\text{alloc} : \text{mem} \rightarrow \mathbb{Z} \rightarrow \mathbb{Z} \rightarrow \text{mem} \times \text{block}$
Record $\text{mem} : \text{Type} := \{$	$\text{free} : \text{mem} \rightarrow \text{block} \rightarrow \mathbb{Z} \rightarrow \mathbb{Z} \rightarrow [\text{mem}]$
$\text{mem_contents} : \text{block} \rightarrow \mathbb{Z} \rightarrow \text{memval};$	$\text{load} : \text{mem} \rightarrow \text{chunk} \rightarrow \text{block} \rightarrow \mathbb{Z} \rightarrow [\text{val}]$
$\text{nextblock} : \text{block}; \}$	$\text{store} : \text{mem} \rightarrow \text{chunk} \rightarrow \text{block} \rightarrow \mathbb{Z} \rightarrow \text{val} \rightarrow [\text{mem}]$
(a) Blocks and Memory States	(b) Memory Operations

Fig. 1. Definitions for the Block-Based Memory Model

with the value 1—denotes the id of the next block to be allocated and will be increased after each allocation. A valid memory block is a finite array of bytes with a lower and upper bound. A *memory state* (denoted by m) consists of a mapping from addresses to *in-memory values* in valid blocks and the value of `nextblock`. Its type `mem` is defined in Fig. 1a where `block` is the type of block ids, $\mathbb{Z}$ is the type of integers, `memval` is the type of in-memory values, and $m.\text{mem_contents } b \ o = v$ iff b is a valid block in m and v is the in-memory value at the o -th memory cell (byte) of b .

A pointer or a memory address is a pair (b, o) where b is the memory block it points to and o is an index to a memory cell in block b (also called an offset). Pointer arithmetic is represented by adjustments to offsets. For example, $(b, o) + n$ is defined as $(b, o + n)$. This simple block-based abstraction already provides basic support for memory isolation, in the sense that given a pointer to block b_1 we can never reach b_2 by performing pointer arithmetic if $b_1 \neq b_2$.

The main operations over memory are depicted in Fig. 1b where `val` is the type of *regular values* defined as follows:

Inductive $\text{val} := \text{Vundef} \mid \text{Vint } i_{32} \mid \text{Vlong } i_{64} \mid \text{Vsingle } f_{32} \mid \text{Vfloat } f_{64} \mid \text{Vptr } (b, o).$

Here, `Vundef` represents undefined values, `Vptr` (b, o) represents a pointer (b, o) , and the remaining forms denote 32- and 64-bit integer and floating point values. `chunk` is the type of *memory chunks* containing information necessary for performing conversion between in-memory values and regular values. Note that the option type `[]` is used to describe the possible failure of some operations. Given a memory state m , a lower bound l and a higher bound h , `alloc` $m \ l \ h$ returns a new block whose id is $m.\text{nextblock}$ and whose valid offsets are in the range of $[l, h)$. It also returns a new memory state where `nextblock` is increased by 1 and the newly allocated memory cells are initialized with undefined values. Note that `alloc` *always succeeds* because the memory space has infinitely many blocks. The `free` operation frees memory cells in a certain range. Given m , k b and o , `load` $m \ k \ b \ o$ loads a value starting from the address (b, o) such that the size and type of value are determined by k . Conversely, `store` stores a value into a certain location in memory. Note that we have omitted a discussion of permissions in the block-based memory as they are mostly orthogonal to this research.

2.2 Memory Injections

A main task of compilers is to transform abstract data structures (e.g., stacks) for describing the semantics of the high-level source languages into concrete data structures for the low-level target languages. These transformations may drastically change the structure of memory.

CompCert introduces *injections* to capture such changes. An injection j is a partial function of type $\text{block} \rightarrow [\text{block} \times \mathbb{Z}]$, such that $j(b) = [(b', o)]$ iff a block b is inserted into block b' at offset o by the transformation. If $j(b) = \emptyset$ (we use $\emptyset$ to represent the `None` constructor), then b is removed from the memory after the transformation.

Given an injection function j , we can relate the source and target values via a binary relation $\hookrightarrow_j^v$ called a *value injection*. $v_1 \hookrightarrow_j^v v_2$ trivially holds iff v_1 is `Vundef`, indicating that undefined values can be relaxed to more concrete values by compiler transformations. If $v_1 = \text{Vptr } (b, o)$,

Fig. 2. Verified Compilation Chain of CompCert

then $v_1 \hookrightarrow_j^v v_2$ holds iff $j(b) = \lfloor (b', o') \rfloor$ and $v_2 = \text{Vptr}(b', o + o')$, i.e., the original address is shifted by injection j accordingly. If v_1 is neither undefined nor a pointer, then $v_1 \hookrightarrow_j^v v_2$ holds iff $v_1 = v_2$. A similar injection relation $\hookrightarrow_j^i$ is defined for in-memory values of which we elide a discussion. The in-memory value injection is lifted to an injection relation $\hookrightarrow_j^m$ between memory states (i.e., *memory injections*). Roughly speaking, $m_1 \hookrightarrow_j^m m_2$ holds if 1) given any two valid addresses of m_1 and m_2 related by j , the in-memory values at these addresses are related by $\hookrightarrow_j^i$, and 2) certain well-formedness conditions are satisfied, including that invalid blocks in m_1 must be mapped to $\emptyset$ and if $j(b) = \lfloor (b', o') \rfloor$ then b' must be valid in m_2 . A *memory extension* is a specialized memory injection with an identity injection function (i.e., $j(b) = \lfloor (b, 0) \rfloor$) and $m_1.\text{nextblock} = m_2.\text{nextblock}$. Memory extensions are used to describe transformations that do not change the structure of memory.

2.3 Correctness of Compilation in CompCert

CompCert takes C modules (parsed from .c files) as input, transforms them through a sequence of compiler passes (about 20 of them) to an architecture-independent language called Mach, from which it finally generates assembly programs on a pre-configured architecture, as depicted in Fig. 2.

Much of the success of CompCert comes from the block-based memory model used uniformly across the whole compilation and a uniform interface for describing the semantics of all of its languages. In any language $\mathcal{L}$ of CompCert, a program P is represented as a list of global definitions associated with its own identifiers and a main function as its entry point:

```
Record program (F V: Type) := { prog_defs: list (ident * globdef F V); prog_main: ident }.
```

These identifiers are also represented by using positive numbers and have a global scope. To describe the semantics of P , a global environment $\mathcal{G}(P)$ of type *genv* is generated which provides two pieces of information: a mapping from definition identifiers to memory blocks holding such definitions and a mapping from the ids of these blocks to actual definitions. By using $\mathcal{G}(P)$, the semantics of P is described as a small-step transition system labeled with traces of events, given by the following initialization and step relations where state is the type of abstract machine states:

```
initial_state: mem × state → Prop    step: (mem × state) → trace → (mem × state) → Prop.
```

The memory is initialized by allocating one block for every global definition. As the execution goes on, more blocks will be allocated (e.g., stack blocks created by function calls). We denote this semantics as $\llbracket P \rrbracket_{\mathcal{L}}$ and write $\llbracket P \rrbracket$ for $\llbracket P \rrbracket_{\mathcal{L}}$ when the language is known from the context.

Given a compiler pass C , its correctness property is formulated as follows where $\geq$ is a forward simulation relation between the semantics of the source and target programs:

$$\forall P P', C(P) = \lfloor P' \rfloor \implies \llbracket P' \rrbracket \geq \llbracket P \rrbracket.$$

That is, if P' is the result of compiling P by C , then any transitions performed by P can be simulated by corresponding ones by P' . To prove this theorem, one needs to establish an invariant between the source and target program states. A critical component of this invariant is a memory injection between the source and target memory states that should hold throughout the entire execution.

Fig. 3. Compiler Passes Transforming Stack Frames in CompCert

The complete CompCert compiler $C_{\text{compCert}}: \text{C.program} \rightarrow [\text{Asm.program}]$ is a transitive composition of its passes $C_1, \dots, C_n$. By composing the forward simulation proofs of C_i , we get

$$\forall P_s P_t, C_{\text{compCert}}(P_s) = \lfloor P_t \rfloor \implies \llbracket P_t \rrbracket \geq \llbracket P_s \rrbracket.$$

In the end, we would like to know if all the behaviors of executable target code are exhibited at the source level. For this, one can flip the forward simulation into a backward simulation by exploiting the facts that the target semantics of CompCert is *determinate* and the source one is *receptive* [Sevcík et al. 2013]. The final correctness theorem of CompCert is stated as follows:

THEOREM 2.1 (CORRECTNESS OF COMPCERT).

$$\forall P_s P_t, C_{\text{compCert}}(P_s) = \lfloor P_t \rfloor \implies \llbracket P_t \rrbracket \leq \llbracket P_s \rrbracket.$$

2.4 Problems

Having discussed the block-based memory model in detail, we can see that it indeed has the inflexibilities described in Sec. 1.1. Now, we elaborate on the problems they cause in verified compilation of whole programs and open programs.

2.4.1 Verifying Transformation on Partial Memory. The vanilla CompCert only supports simulation proofs for whole programs. Even in this setting, one compiler pass only operates on certain part of memory and has local effects. More specifically, excluding advanced optimizations such as tail-call elimination and inlining, only four of CompCert's passes change the memory structure, of which three change the stack frame by frame, as depicted in Fig. 3. Those are:

- **SimplLocals.** In a source language called Clight, a stack frame consists of a collection of blocks, one for each local variable in the associating function. For example, the left column of Fig. 3 shows a stack frame containing three blocks (b_1 , b_2 and b_3) for three local variables in some function. The SimplLocals pass transforms Clight programs by turning scalar local variables whose address are not taken into temporary variables, effectively removing their blocks from the stack frame. For example, in Fig. 3 the local variable associated with b_3 is turned into a temporary one. The remaining local variables are called *stack-allocated variables*.
- **Cminorgen.** This pass transforms programs in C#minor (a variant of Clight) into programs in an intermediate language called Cminor. The stack-allocated variables are merged into a single block called the *stack-allocated data* (sb in Fig. 3). From this point on, a stack frame only contains a single block.
- **Stacking.** This pass lays out the concrete stack frame containing function parameters, return addresses, spilled registers, etc. (the gray areas in Fig. 3). The stack-allocated data is inserted into the middle of the concrete frame.

Fig. 4. Transformation of Memory by Unusedglob

The remaining pass transforming memory is called Unusedglob. It drops *static* global definitions which are unused by the program. Fig. 4 shows an example where the blocks for the second and fourth global definitions (gb_2 and gb_4) are dropped by Unusedglob (where the numbers represent block ids). For this pass, the stack memory is completely untouched.

When proving the correctness of the above passes, one would expect that we could exploit the locality of memory transformations to construct intuitive proofs. For example, for SimplLocals, Cminorgen and Stacking, one would like to exploit the facts that 1) blocks for global definitions are unchanged, and 2) the changes to individual stack frames cannot interfere with each other. However, it is impossible to formalize these facts directly because we can neither distinguish blocks for global definitions from blocks in stack frames, nor tell the differences between blocks in different stack frames: they are all indexed by positive numbers. To bypass this problem, CompCert relies on extra invariants for classifying block ids into different ranges of positive numbers. Furthermore, it introduces general memory injections for relating source and target blocks, essentially lifting reasoning about local stack transformations to a global level. As we can see, the above problems are exactly caused by **Inflexibility 1** introduced in Sec. 1.1.

For Unusedglob, because it only drops certain blocks for global definitions, one would expect that its memory invariant uses a *partial identity mapping*. However, in contrast to this intuition, the memory invariant for Unusedglob cannot be anything less than a general memory injection. To understand this, note that blocks for global definitions are allocated before any stack blocks. Because valid block ids must be consecutive integers starting from 1, dropping even a single global definition will result in shifting of all the block ids allocated after. This situation is illustrated by the example in Fig. 4 where the removal of gb_2 and gb_4 causes the ids of gb_3 and gb_5 to change from 3 to 2 and 5 to 3, respectively. Moreover, all the stack block ids are shifted by -2 . In the end, a non-trivial memory injection is required to connect the source and target blocks. With this injection the correctness proof of Unusedglob involves complicated reasoning over stack memory even when the stack is not touched at all. These problems are exactly caused by **Inflexibility 2** in Sec. 1.1.

2.4.2 Verifying Transformations on Open Programs. Verification techniques for open programs are compositional only if the semantics of open programs are compatible with each other. However, the existence of a global nextblock makes this compatibility very difficult to establish, even when different open programs operate on completely separate memory regions.

We take the compilation and linking of Certified Concurrent Abstraction Layers (CCAL) as an example to illustrate the above problem [Gu et al. 2018]. A concurrent certified abstraction layer L consists of shared and private memory states, abstract states, and a set of possibly shared primitive operations (like external functions). The semantics of a language (e.g., C or assembly) built on top of L forms an abstract machine in which concurrent programs can be formally described. A CCAL object is a formally verified open program built on top of some layer L . A multi-threaded program is developed by abstracting individual threads into CCAL objects. These objects are then compiled by an extension of CompCert called Thread-Safe CompCertX into assembly code. Finally, CCAL objects need to be linked together to form a complete program.

Fig. 5. Linking of Multiple Stacks into a Single Stack in CCAL

For the above linking to be possible, the views of memory of CCAL objects must be compatible with each other. Achieving this compatibility is a non-trivial task. To understand this, observe that a new stack block is allocated for every function call in CompCert’s assembly. To link threads together, it is necessary for each thread to have the same sequences of stack blocks and `nextblock`, meanwhile preventing one thread from accessing the private stack memory of the others.

To solve the above problem, the authors of Thread-Safe CompCertX modify the assembly semantics so that when a thread yields to its context, a sequence of dummy stack blocks is allocated to increase `nextblock` in accordance with the actual allocation of stack blocks by the context. Moreover, to avoid any interference between memory operations on the stacks in different threads, the dummy blocks do not carry any read or write permission—they are “dead” memory cells from the perspective of the focused thread. With those devices, it is possible to “thread” the private stack frames of each thread into a single stack. As an example, the linking of stacks for two threads is depicted in Fig. 5. Here, the call to the `yield` primitive from thread 1 to 2 in the function *g* allocates two dummy blocks (marked with diagonal lines) for the corresponding execution in thread 2. Accesses to these dummy blocks are invalid in thread 1.

The solution above has two problems. First, intuitively, each thread should have its own private stack: the context should not interfere with operations on this stack. Contrary to this assumption, the growth of dummy frames depends on how contextual threads change `nextblock`. This introduces unnecessary complication to verification of compilation. Second, in the linked program, stack frames for different threads are interleaved with each other. This makes the semantics of linked programs much more complex. It is also extremely difficult to further verify their compilation to actual machine code where each thread should have its own contiguous and private stack. As we can see, the above problems are exactly caused by **Inflexibility 2 and 3** in Sec. 1.1.

2.5 Our Approach

As we have discussed in the introduction, we shall solve the above problems by getting rid of the inflexibilities of the block-based memory model through a generalization based on nominal techniques. We shall present this nominal memory model and the development of Nominal CompCert in Sec. 3. After that, we demonstrate how the two kinds of problems discussed above can be solved by further extending Nominal CompCert in Sec. 4 and Sec. 5.

3 NOMINAL MEMORY MODEL AND VERIFIED COMPILATION

3.1 Key Ideas

Nominal techniques [Gabbay and Pitts 2002; Pitts 2016] provide a mathematical foundation for managing named resources. In this setting, any countably infinite set $\mathbb{A}$ can represent a set of available names. Elements in these sets are called *atomic names* (or simply *names* in short). Dependency

```

Module Type BLOCK.
  Parameter block : Type.
  Parameter eq_block :  $\forall x\ y : \text{block}, \{x = y\} + \{x \neq y\}$ .
End BLOCK.

Module Type SUP.
  Parameter sup : Type.
  (** Operations *)
  Parameter sup_empty : sup. (* Empty support *)
  Parameter fresh_block : sup  $\rightarrow$  block. (* Generation of fresh blocks *)
  Parameter sup_incr : sup  $\rightarrow$  sup. (* Increment of supports *)
  Parameter sup_in : block  $\rightarrow$  sup  $\rightarrow$  Prop. (* Check validity of block ids *)
  (** Properties *)
  Parameter sup_dec :  $\forall b\ s, \{\text{sup\_in } b\ s\} + \{\neg \text{sup\_in } b\ s\}$ .
  Parameter empty_in :  $\forall b, \neg \text{sup\_in } b\ \text{sup\_empty}$ .
  Parameter freshness :  $\forall s, \neg \text{sup\_in } (\text{fresh\_block } s)$ .
  Parameter sup_incr_in :  $\forall a\ s, \text{sup\_in } a\ (\text{sup\_incr } s) \leftrightarrow a = (\text{fresh\_block } s) \vee \text{sup\_in } a\ s$ .
End SUP.

```

Fig. 6. Interfaces of the Nominal Memory Model

of objects upon names is captured by the notions of *permutations* and *supports*. A permutation π is a bijection from $\mathbb{A}$ to itself that only renames a finite subset of names in $\mathbb{A}$. Given a set of objects X and some $x \in X$, $A \subseteq \mathbb{A}$ supports x (or A is a support of x) iff, for any permutation π on $\mathbb{A}$ that is an identity mapping for names in A , we have $\pi \cdot x = x$ where $_ \cdot _$ denotes the “application” (known as an *action*) of π to the object x . This effectively means that x is independent of any name outside of A . Only objects that can be supported by a finite set of names are of interest to us. A binary relation called *freshness* makes the independence relation concrete. A name $a \in \mathbb{A}$ is fresh w.r.t. x (written as $a \# x$) iff x is supported by some finite set $A \subseteq \mathbb{A}$ and $a \notin A$.

The above concepts can be used to characterize the block-based memory model. By taking memory states as objects containing names, we adopt the following analogies:

- Block ids represent names that memory states depend upon;
- Given a memory state m , the set of valid blocks in m represents a support of m ;
- Given a memory state m , its `nextblock` is fresh w.r.t. m .

Note that the set of valid blocks is always finite. To some extent, the existing memory operations in CompCert already exploit the properties of atomic names and supports. For example, `alloc` always succeeds because there is always an infinite amount of ids outside the set of valid blocks.

However, the block-based memory also makes rigid assumptions about names and supports:

- Block ids are fixed to positive numbers;
- For any memory state m whose `nextblock` is n , its support must be the fixed set of consecutive numbers $\{1, \dots, n - 1\}$;
- For any memory state m whose `nextblock` is n , its fresh block must be assigned with the id n .

As we have seen in Sec. 2.4, these rigid assumptions cause serious problems in compiler verification. We shall remove those assumptions by generalizing the block-based memory model to work with any valid atomic names, supports and freshness relations.

Note that, since memory injection already provides a notion of permutation, it is natural to consider actions over memory states and permutation or swapping of names (block ids). We do not expand on this idea in the rest of this paper, because we will only consider programs with

```

Module Block < : BLOCK.
  Definition block := positive.      Definition eq_block := peq.
End Block.

Module Sup < : SUP.
  Definition sup := list block.      Definition sup_in (b: block) (s: sup) : Prop := b ∈ s.
  Definition sup_empty : sup := [].  Definition fresh_block (s: sup) := (find_max_pos s) + 1.
  Definition sup_incr (s: sup) := (fresh_block s)::s.    ...
End Sup.

```

Fig. 7. Block-based Memory Model as an Instance

fixed assignment of block ids. Permutations will be needed to avoid clashing of names for verifying compositional compilation in general, which we will investigate in the future (more on this in Sec. 8).

3.2 The Nominal Memory Model

Following the above ideas, we generalize the block-based memory model into the nominal memory model by introducing an abstraction of block ids and supports as module types, as depicted in Fig. 6. By the definition of BLOCK, block ids are names with decidable equality. By the definition of SUP, a support type must be accompanied by four kinds of operations: checking the membership of blocks in supports (`sup_in`), creating an empty support (`sup_empty`), generating fresh blocks (`fresh_block`) and increasing supports with new blocks (`sup_incr`). Furthermore, those operations must satisfy some well-formedness properties. Note that after the above abstraction, the rigid assumptions for block-based memory model are removed.

We also note that the above generalization does not exactly match with the standard definitions in nominal techniques. For example, BLOCK does not enforce that block ids are from a countably infinite set. Instead, the freshness property guarantees that any block fresh w.r.t a support s must not be already in s . Together with the totality of `fresh_block`, it implies the existence of an infinite number of block ids. We also define supports to be any type that has the interface of SUP, instead of a finite set of block ids. This generalization allows us to instantiate `sup` with complex data structures for formalizing memory space in fine-grained styles. We shall exploit this feature extensively in Sec. 4 and Sec. 5.

To make use of the above interfaces, we instantiate block ids and supports as follows:

```

Module Block < : BLOCK. ... End Block.      Module Sup < : SUP. ... End Sup.

```

Then the original block type is instantiated with `Block.block`. Moreover, the memory state carries a support instead of `nextblock`:

```

Record mem: Type := { mem_contents: block → Z → memval; support: Sup.sup;}.

```

The memory operations are adapted accordingly. For example, checking of valid blocks is done by using `sup_in` instead of comparing with `nextblock`:

```

Definition valid_block (m:mem) (b:block) := b < m.(nextblock). (* Old *)
Definition valid_block (m:mem) (b:block) := sup_in b m.(support). (* New *)

```

For another example, `alloc` now invokes `fresh_block` to allocate a new block instead of consulting `nextblock`. The properties of all memory operations can be easily reestablished because they are already ignorant of particular instantiations of block ids and supports.

Finally, the block-based memory model becomes a special case of the nominal memory model, as depicted in Fig. 7 where `find_max_pos` finds the maximal positive number in a list. We shall see more complex memory models that exploit the full power of the above abstractions later.

3.3 Nominal CompCert

We apply the nominal memory model to the full compilation chain of CompCert to get Nominal CompCert. First, we need to update the semantics of every language of CompCert. This is automatically achieved by replacing the old memory operations with the new ones. We use $\llbracket P \rrbracket^N$ to denote the new semantics of a program P . Second, we need to update the simulation proofs. Because the generalization of block ids and supports is mostly orthogonal to the simulation proofs in CompCert, the changes are minimal. The only slightly complicated changes are the results of losing the ability to identify valid blocks via comparison of positive numbers. For example, in the proof of the inlining optimization, the invariant `valid_block m sp` asserts that the stack block sp is valid. It also implies any block in m with id less than sp is also valid. After the generalization, we must make this fact explicit. This is achieved by breaking the above invariant into two: $sp = \text{fresh_block } sps$ and $\{sp\} \cup sps \subseteq m.\text{support}$. Here, sps denotes the blocks allocated before sp and the second new invariant asserts their validity.

By composing the updated simulation proofs, we get the final correctness theorem of Nominal CompCert, which is almost identical to Theorem. 2.1 except that the semantics of languages are based on the nominal memory model:

THEOREM 3.1 (CORRECTNESS OF NOMINAL COMPCERT).

$$\forall P_s P_t, C_{\text{compert}}(P_s) = \lfloor P_t \rfloor \implies \llbracket P_t \rrbracket^N \leq \llbracket P_s \rrbracket^N.$$

We note that Nominal CompCert is a lightweight extension to CompCert. The changes made to it amount to about 1% of the code (about 1.4k lines) of CompCert 3.8. A majority of these changes are trivial substitutions of `nextblock` with `supports`, except for about 200 lines of code for handling the complication resulting from the inability to compare block ids mentioned before. We also introduce about 200 lines of new code for implementing the nominal memory model as described in Sec. 3.2. A more detailed evaluation of Nominal CompCert can be found in Sec. 6.1.

3.4 Instantiation of Nominal CompCert

Nominal CompCert provides a skeleton for developing further extensions to CompCert that exploit the full power of nominal memory model. In these extensions, block ids and supports will be instantiated with more sophisticated structures. However, there is *zero immediate overhead* to introduce such instantiations because the entire proof of Nominal CompCert is ignorant of and holds for any instances of block ids and supports that meet their interfaces. Only after the user introduces mechanisms for exploiting the internal structures of such instances will any actual overhead be incurred. Even in those cases, the changes are built on well-defined interfaces and confined to local definitions and proofs.

As a side effect of CompCert's uniform implementation of memory model, the current realization of Nominal CompCert only allows *one* global instance of block ids and supports shared by the whole compiler. Therefore, when instantiating Nominal CompCert, we must design the instances of block ids and supports carefully so that they are rich enough for supporting the desired verification techniques while remain compatible with the abstract interfaces of Nominal CompCert.

With the above benefits and technical limitation in mind, we shall present a series of instantiations of Nominal CompCert in Sec. 4 and Sec. 5 that solve the problems in Sec. 2.4. These instantiations are developed *incrementally* such that the latter ones contain instances of block ids and supports with richer structures that subsume the former. In this sense, those instantiations are compatible with each other, with the final one (i.e., Multi-Stack CompCert) being most rich in features. Moreover, because they are fully compatible with the interfaces of Nominal CompCert, we can freely decide whether to leave the proofs for individual passes unchanged or to improve them by exploiting the

internal structures of block ids and supports. In other words, the benefits of instantiating block ids and supports can be exploited for different compiler passes *on an as-needed basis*. We shall witness this in detail in Sec. 4 and Sec. 5.

4 VERIFICATION OF TRANSFORMATIONS ON PARTIAL MEMORY

In this section, we discuss the application of our nominal memory model to verified compilation of whole programs. We shall first introduce an enriched instantiation of this model where the whole memory space is divided into memory regions with well-defined structures and clear roles. Based on this enriched memory model, we get a complete extension of Nominal CompCert. With this extension, we develop an effective solution to the first problem in Sec. 2.4, i.e., how to construct intuitive correctness proofs for compiler passes that operate on partial memory.

4.1 Key Ideas

To understand how intuitive correctness proofs for memory transformations may be obtained, we observe that transformations on memory in CompCert (excluding optimization passes) all operate on isolated local memory regions (see Sec. 2.4.1). Among them, `Unusedglob` only changes the memory for global definitions, leaving the stack untouched. On the other hand, `SimplLocals`, `Cminorgen` and `Stacking` change the stack frame by frame, leaving the global memory untouched. Therefore, if we organize the individual local memory regions into certain high-level structure, we can decompose a memory transformation into two parts: a transformation on local memory regions and another one on the high-level structure. Furthermore, if the high-level structure is stable under compilation, then only the local memory transformation needs to be explicitly captured in the proof invariant for memory states (i.e., memory injections). Because injections are central to the correctness proofs of memory transformations, this may make proofs much easier to understand.

Based on the above ideas, we divide the memory space into *global memory* for global variables and functions and *stack memory* for data allocated on the stack. We then organize each memory region into a high-level structure that is stable under compilation. This is easy for global memory: we only need to assign block ids to global definitions that are independent of compilation (e.g., their global names). For stack memory, we need to find a structure that will not be modified by any of the transformations mentioned previously, with which we can uniquely identify every stack memory block. Because in program execution the stack is managed by function calls and returns which are usually organized into a tree structure known as the *call tree* or *call graph*, a natural representation of the stack structure would be a *stack tree* that mirrors the call tree, except that its nodes are stack frames allocated by corresponding function calls. Because the above memory transformations do not change the order of function calls and returns at all, the stack tree remains stable under compilation. Finally, all the memory blocks allocated on the stack are classified and grouped into corresponding stack frames. With the stack tree, we will be able to decompose a transformation on the entire stack into local transformations on individual stack frames.

As we shall see below, we can implement the above structures of memory space by instantiating blocks and supports with rich data structures. With these instantiations, we are able to explicitly define the memory injections as *concrete functions* that precisely capture all the details of memory transformations. In contrast to the original proofs in CompCert, where memory injections are “black boxes” given by propositions asserting the *existence of some memory injections*, our concrete definitions of memory injections provide transparent formalization of memory transformations. This in turn leads to intuitive proofs of their correctness, as we shall see shortly in detail.

Note that we have not given a structure to the heap in this work. Instead of explicitly modeling heap allocation and deallocation (i.e., `malloc` and `free` functions) in the memory model by using the `alloc` and `free` operations in CompCert, our work and many others based on CompCert

```

Definition fid : Type := [ident].           Definition path : Type := list nat.
Module Block < : BLOCK.
  Inductive block :=
    | Stack : fid → path → positive → block
    | Global : ident → block.
  ...
End Block.

```

Fig. 8. Definition of Block IDs for Structured Memory Space

(including Stack-Aware CompCert [Wang et al. 2019], Verified Software Toolchain [Appel 2011], and CertiKOS [Gu et al. 2015, 2018]) treat `malloc` and `free` as library functions operating over a finite heap space represented by a global variable `block`. This is because `alloc` (which always succeeds) in CompCert’s memory model does not capture the possibility of allocation failures resulting from the finiteness of heap space. An important side effect is that, because all the heap variables reside in a finite and linear space in our model, the semantics of `malloc` becomes different in certain aspects. For example, the following permutation of two `malloc`s:

`p = malloc(2); q = malloc(4);` $\implies$ `q = malloc(4); p = malloc(2);`

is valid in CompCert because `malloc` never fails and `p` and `q` always reside in isolated blocks. In our new model, if the finite heap only have 4 bytes free, then for the left code fragment, `p = malloc(2)` would succeed but `q = malloc(4)` would fail, but after the permutation, `q = malloc(4)` would succeed and `p = malloc(2)` would fail. Also, in our new model, `p` and `q` are in the same block and `q < p` has defined semantics which is not true in CompCert’s heap model.

Another side effect of our treatment is that, after the initialization of global definitions, no new memory block will be ever allocated except for blocks on the stack. This will actually be useful when we discuss the implementation of supports later.

Finally, note that the above ideas are also applicable to advanced optimizations passes such as inlining and tailcall recognition, which are the only two other passes in CompCert that transform memory structures. However, because they change the structures of stack trees, more sophisticated instances for blocks and supports will be needed for realizing these ideas. Therefore, we have left the proofs for inlining and tailcall recognition as they are. This demonstrates that the benefits of nominal memory model can indeed be exploited *as needed*, as we have commented in Sec. 3.4. We will discuss how to improve the proofs for inlining and tailcall recognition in the future in Sec. 8.

4.2 Nominal Memory Model with Structured Memory Space

In this section, we introduce the instantiation of nominal memory model for formalizing the structured memory space proposed in Sec. 4.1. We first introduce the formal definitions of block ids and supports, which are driven by the division of memory space into global memory and a stack tree. In particular, a stack tree data structure is designed to implement the supports. With these formal definitions, we then define the operations for this new memory model and prove necessary properties of these operations.

4.2.1 Formalization of Block IDs. We define the type of block ids in Fig. 8. As expected, this definition captures the fact that a block is either in the global memory or on the stack. A block for the global definition with the unique name *id* is denoted by `Global id`. As discussed in Sec. 4.1, a block on the stack must reside in some stack frame which is in turn on the stack tree. We note that this stack frame is allocated by certain function call and is located at a unique node on the stack tree. If every stack block in the same stack frame has a distinct “local” block id, we can uniquely identify a stack block by using the position of its associate stack frame on the stack tree together


```

extern void g(int*, int*);
void k() { ... }
void h(int* w) { ... }
void f(int* e, int* f, int* g) {
  int x=*e, y=*f, z=*g;
  ...
  g(&x,&y);
  h(&z);
}
int main() {
  int a,b,c;
  ...
  k();
  f(&a,&b,&c); ...
}

```

(a) The Program

Local block ids: $a = 3, b = 7, c = 9, x = 2, y = 6, z = 9, \dots$

(b) The Stack Tree

Fig. 9. An Example of Stack Trees

with its local block id. Here, the unique position is represented as a *path* from the root of stack tree (corresponding to the entry function) to the stack frame of interest. In particular, a path p is a list $[i_1, i_2, \dots, i_n]$ denoting a sequence of indices to child nodes with which the stack frame can be reached from the root, i.e., the frame with path p can be reached by first visiting the i_1 -th child of the root frame, then the i_2 -th child of it, and so on until p is traversed.

The Stack constructor formalize the above ideas. For any stack block, if the path to its associating frame is p and its local id is b , it is denoted by $\text{Stack } fid \ p \ b$. The extra argument fid identifies whether the stack block is allocated by an internal function call or an external one. This is necessary for defining memory injections for *contextual memory* which we shall explain later. More specifically, when $fid = [id]$, the block is allocated by a call to the internal function id . Otherwise, it is allocated by some external function call.

As an example, Fig. 9 displays a program and a stack tree generated by its execution when the program counter is in h . Here, the rectangles represent stack frames, besides which the alphabetic names denote the internal functions that are invoked to create the frames; the rectangles with dashed borders represent frames allocated by external calls; the squares with rounded corners represent stack blocks, in which the numbers denote their local blocks ids; finally, the numbers besides arrows denote the indices to child nodes. With this figure we give some examples of formalized ids for stack blocks: the last block in f denotes its local variable z and has the id $\text{Stack } [f] \ [2] \ 9$; the first child of f is allocated externally by a call g and its second block has the id $\text{Stack } \emptyset \ [2, 0] \ 4$; finally, the only block in h has the id $\text{Stack } [h] \ [2, 1] \ 1$.

4.2.2 Formalization of Stack Trees. A stack tree provides a tree structure to stack memory; it also provides a stack structure for the allocation and deallocation of new stack blocks. The latter is needed for the implementation of supports (e.g., in the implementation of `fresh_block`). Therefore, the formal definition of the data structure for stack trees should look like a tree and a stack at the same time. This is achieved by separating the list of active frames which are on the right-most path of the stack tree from inactive ones (or “dead” frames). The active frames form a stack while the inactive frames form sub-trees that are children of active frames. This data structure is named *stree* and formally defined as follows:

Inductive *stree* : **Type** :=

| Node : $fid \rightarrow \text{list positive} \rightarrow \text{list stree} \rightarrow [\text{stree}] \rightarrow \text{stree}$.

A term of the form $\text{Node } fid \ bs \ dt \ at$ represents a node in the stack tree (i.e., a stack frame) that is either allocated by the function named id when $fid = [id]$ or allocated by external functions when $fid = \emptyset$. The list bs contains stack blocks allocated by the function that created the stack frame. The child nodes of this stack frame are in dt and at , where dt denotes a list of dead child trees and at denotes the only child frame that is still active when it is not $\emptyset$. By the definition of *stree*, we can traverse the stack of active frames by starting from the root node and repeatedly visit the argument at until $at = \emptyset$, at which point we have reached the top of the stack. The remaining frames in this tree are all dead. An empty tree (*empty_stree*) is defined as $\text{Node } \emptyset \ [] \ [] \ \emptyset$.

As an example, consider the tree in Fig. 9 again. Its rightmost path contains the list active frames allocated by *main*, *f* and *h*, and the remaining frames are all dead ones. It is formalized as t below where tj and tk denotes its first and second children:

$$\begin{aligned} tj &:= \dots & tk &:= \dots & tg &:= \text{Node } \emptyset \ [1, 4] \ [] \ \emptyset & th &:= \text{Node } [h] \ [1] \ [] \ \emptyset \\ tf &:= \text{Node } [f] \ [2, 6, 9] \ [tg] \ [th] & t &:= \text{Node } [main] \ [3, 7, 9] \ [tj, tk] \ [tf] \end{aligned}$$

Note that, although we have only shown an example of stack trees with externally allocated frames as leaves nodes, the non-leaf occurrences of these frames are allowed by the definition of *stree*. Because we do not deal with call back functions in this work, we elide a discussion of such examples.

We then formalize the operations over *stree* that are necessary for defining supports, including:

- $\text{next_stree} : \text{stree} \rightarrow \text{ident} \rightarrow \text{stree} * \text{path}$. It is for allocating a new frame. Given a stack tree t and the name f of the allocating function, $\text{next_stree } t f$ pushes a new active frame onto t and returns the updated tree together with the path to the new frame.
- $\text{next_block_stree} : \text{stree} \rightarrow \text{fid} * \text{positive} * \text{path} * \text{stree}$. It is for allocating a new block in the newest active frame F . Given a stack tree t , it returns the function name corresponding to F , a new block id that is *locally fresh* in F , the path to F , and the updated tree.
- $\text{return_stree} : \text{stree} \rightarrow [\text{stree} * \text{path}]$. It is for deallocating the newest active frame F , meaning that F is moved to the list of dead child trees of its parent. It returns $\emptyset$ if F does not have a parent. Otherwise, it returns the path to F and the updated tree.
- $\text{stree_in} : \text{fid} \rightarrow \text{path} \rightarrow \text{positive} \rightarrow \text{stree} \rightarrow \text{Prop}$. It is for checking the existence of blocks in a stack tree. Given a tree t , an optional name fid , a path p and a block b , $\text{stree_in } fid \ p \ b \ t$ holds when the frame on the path p in t is associated with the name fid and contains b .

Note that, by definition next_block_stree *always succeeds*, implying the existence of an infinite supply of stack block ids. This matches the nominal property of block ids in our memory model.

4.2.3 Formalization of Supports. The type of supports and its operations are defined in Fig. 10. By definition, a support consists of a stack tree together with a list of names for valid global definitions. Because of the separation of memory into stack and global regions, we introduce four functions for managing supports: $\text{sup_incr_glob} : \text{ident} \rightarrow \text{sup} \rightarrow \text{sup}$ for managing global definitions, and $\text{sup_incr_frame} : \text{sup} \rightarrow \text{id} \rightarrow \text{sup}$, $\text{sup_return_frame} : \text{sup} \rightarrow [\text{sup}]$, and $\text{sup_incr_block} : \text{sup} \rightarrow \text{sup}$ for managing the stack. The last three operations are implemented by using the corresponding operations over *stree*. Specifically, fresh_block and sup_incr_block make use of next_block_stree to generate a fresh block in the newest active frame and to update the stack tree, respectively, sup_incr_frame makes use of next_stree to allocate new frames, and sup_return_frame makes use of return_stree to convert active frames into dead ones. Note that sup_incr_block also implements sup_incr because it provides the only way to allocate new blocks after initialization.

4.2.4 Formalization of Memory Operations. With the new definition of supports, we update the memory operations accordingly. The key change is to introduce four new operations that operate on finer-grained memory regions. They are listed below, where the definition of memory states *mem* stays the same as described in Sec. 3.2:

```

Module Sup < : SUP.
  Record sup : Type := mk_sup { stack : stree; global : list ident; }.

  Definition sup_empty : sup :=
    mk_sup empty_stree [].

  Definition sup_incr_glob (i: ident) (s: sup) :=
    mk_sup (stack s) (i :: s.global).

  Definition sup_in (b: block) (s: sup): Prop :=
    match b with
    | Stack fid p i =>
      stree_in fid p i s.(stack)
    | Global id => id ∈ s.global
    end.

  Definition fresh_block (s: sup): block :=
    match next_block_stree s.(stack) with
    | (f, i, p, _) => Stack f p i
    end.

  ....

End Sup.

  Definition sup_incr_frame (s: sup) (id: ident): sup :=
    let (t', p) := next_stree s.(stack) id in
    mk_sup t' s.(global).

  Definition sup_return_frame (s: sup): [sup] :=
    match return_stree s.(stack) with
    | [(t', p)] => [mk_sup t' s.(global)]
    | [] => []
    end.

  Definition sup_incr_block (s: sup) :=
    let (_, t') := next_block_stree s.(stack) in
    mk_sup t' s.(global).

  Definition sup_incr := sup_incr_block.

```

Fig. 10. Definition of Supports for Structured Memory Space

- `alloc_glob`: $\text{ident} \rightarrow \text{mem} \rightarrow \mathbb{Z} \rightarrow \mathbb{Z} \rightarrow \text{mem} * \text{block}$. `alloc_glob id m l h` allocates a new block with range $[l, h)$ in m for the global definition `id` by invoking `sup_incr_glob`. It returns the updated memory and the block `Global id`.
- `alloc_frame`: $\text{mem} \rightarrow \text{ident} \rightarrow \text{mem} * \text{path}$. `alloc_frame m id` allocates a new active frame in m by invoking `sup_incr_frame` and returns the path to this new frame.
- `return_frame`: $\text{mem} \rightarrow [\text{mem}]$. `return_frame m` deallocates the newest active frame by invoking `sup_return_frame` and returns the updated state.
- `alloc_block`: $\text{mem} \rightarrow \mathbb{Z} \rightarrow \mathbb{Z} \rightarrow \text{mem} * \text{block}$. `alloc_block m l h` allocates a fresh block with range $[l, h)$ on the newest active frame of m by invoking `sup_incr_block`. This block's id is obtained by calling `fresh_block` on m .(support). The updated memory together with this id are returned.

The properties about the new and updated memory operations are proved by following the conventional pattern; we elide a discussion of these changes.

4.3 Nominal CompCert with Structured Memory Space

Even with the previous changes, Theorem. 3.1 still holds as its proof is compatible with any particular instances of BLOCK and SUP (as we have pointed out in Sec. 3.4). However, we would like to further exploit the enriched internal structure of the new memory model to improve the proofs for partial memory transformations. For this, we first need to update the semantics of Nominal CompCert's languages. This is achieved as follows:

- Make the global environments directly use identifiers of global definitions instead of block ids, because they are the same in the new memory model. Also, make use of `alloc_glob` for initialization of global memory instead of `alloc`.
- At every function call and return, invoke `alloc_frame` and `return_frame`, respectively.
- Replace all the invocations of `alloc` for allocating stack blocks with `alloc_block`.

<p>Variable ge: $genv$.</p> <p>Definition $struct_meminj$ ($s:sup$) ($b:block$): $\text{if } check_block\ s\ b$ $\text{then } [(b, 0)] \text{ else } \emptyset.$</p>	<p>Definition $check_block$ ($s:sup$) ($b:block$): $bool :=$ $\text{match } b \text{ with}$ $\text{ Global } id \Rightarrow \text{match } (find_symbol\ ge\ id) \text{ with}$ $\quad \emptyset \Rightarrow \text{false} \mid _ \Rightarrow \text{true end}$ $\text{ Stack } _ _ \Rightarrow \text{sup_dec } b\ s \text{ end.}$</p>
---	---

Fig. 11. Structural Injection Function for Unusedglob

With the above changes, we reprove the full compilation chain of Nominal CompCert correct by establishing simulations between the newly introduced operations under compilation. This requires a few predictable and local modifications to the correctness proof for each pass, including the advanced optimizations such as inlining and tailcall recognition. This indicates that our extension is still very compatible with the abstraction provided by Nominal CompCert. In the end, we got a correctness theorem similar to Theorem. 3.1.

4.4 Intuitive Proofs for Partial Memory Transformations

With the instantiated Nominal CompCert in place, we are ready to show how to construct intuitive correctness proofs for the transformations on partial memory mentioned in Sec. 2.4.1. We shall take Unusedglob as the main example to illustrate the key ideas. The proof construction for the other passes follows a similar pattern, which we shall briefly discuss at the end of this section.

4.4.1 Structural Injection Functions. Previously, simulation proofs relying on general memory injections assume the *existence of some memory injection* that captures the transformation on memory. With our new memory model, it is now possible to explicitly *define the injection function* that precisely describes the transformation on memory. This definition is composed of three pieces of information:

- the shape of memory transformation that is *statically* known and *independent* of programs being transformed;
- the shape of memory transformation that is *statically* known but *depends* on programs being transformed;
- the *dynamic* information depending on *particular execution* of programs, e.g., the support of memory states.

We call it a *structural injection function*.¹ Unlike the existential injection functions used previously, structural injections explicitly capture the essence of memory transformations, hence enable intuitive characterizations of these transformations. As we shall see, concrete injection functions with the above structure can be derived for all the passes mentioned in Sec. 2.4.1.

We take Unusedglob as an example. Recall that, given a module M , it removes statically defined global variables or functions that are never used in M . Intuitively, its memory injection should map removed global blocks to $\emptyset$ and any other valid blocks to themselves. The structural injection function $struct_meminj$ exactly captures this intuition. It is defined in Fig. 11, where ge denotes the global environment of the target program and $check_block\ s\ b$ checks if the source block b should be mapped to a target block given the support s of the source program.

We explain how the definition in Fig. 11 is composed of the three pieces of information mentioned above. First, we know the injection function is a partial identity function because Unusedglob only drops memory blocks. This information is statically known and independent of programs being transformed. Second, to determine whether a global block id should be mapped to the target memory, we check if id is in the target environment ge . This information is statically known, but

¹It is different from and should not be confused with *structured injection* in Compositional CompCert [Stewart et al. 2015].

Fig. 12. Comparison of Injections for Unusedglob

```

Variable  $ge$ :  $\text{genv}$ .
Definition struct_meminj ( $s$ : $\text{sup}$ ) ( $b$ : $\text{block}$ ) :=
  if sup_dec  $b$   $s$ 
  then unchecked_meminj  $b$ 
  else  $\emptyset$ .

Definition unchecked_meminj ( $b$ : $\text{block}$ ) :=
  match  $b$  with
  | Stack [ $id$ ]  $p$   $i$   $\Rightarrow$ 
    do  $o \leftarrow \text{find\_frame\_offset } ge \ id \ i$ ;
    OK  $\lfloor (\text{Stack } [id] \ p \ 1, o) \rfloor$ 
  | _  $\Rightarrow \lfloor (b, 0) \rfloor$  end.

```

Fig. 13. Structural Injection Function for Cminorgen

depends on programs resulting from the transformation. Third, to determine if a stack block b should be mapped to a target block, we check if it is in the source support s . If not, it is an invalid block and should be mapped to $\emptyset$. This information depends on the dynamic execution of programs.

We illustrate the intuitiveness of the above structural injection by comparing it with the original injection. A concrete example is shown in Fig. 12 where Fig. 12a depicts CompCert’s injection for the example in Fig. 4 and Fig. 12b depicts the corresponding structural injection. As we can see, the original injection needs to describe shifting of positive block ids because of deletion of global definitions, while the new injection is simply a partial identity mapping (we have omitted the concrete ids for stack blocks for simplicity).

4.4.2 Verification Based on Intuitive Proof Invariants. With structural memory injections, we are able to improve the proofs of partial memory transformations so that they match with our intuition.

We continue with the Unusedglob example. The main complexity of the existing proof of Unusedglob lies in reasoning about shifted global definitions and stack frames. More specifically, the existing proof is based on an invariant of matching stack frames where each frame is associated with a source and target “bound” for describing the ranges of matching block ids in the source and target stack frames. The proof establishes an initial memory injection and shows it respects these bounds. It then shows that this invariant holds as the injection evolves with the allocation and deallocation of new stack blocks. However, since Unusedglob does not touch stack at all, we would have not expected such complicated reasoning about stack frames in the first place.

With structural memory injections, the excessive reasoning about stack frames is no longer needed. The invariant about stack now assumes the existence of identical stack frames w.r.t. struct_meminj. An intuitive simulation proof for Unusedglob follows from this simplified invariant.

4.4.3 Verification of Other Transformations. The above ideas are applicable to other transformations in a similar fashion. We take Cminorgen as an example whose structural injection function is defined in Fig. 13. By definition, Cminorgen injects stack blocks in each frame into a single block of stack-allocated data for the frame. Therefore, we know that block ids for global definitions are unchanged.

Furthermore, because the structure of the stack tree is unchanged, the paths to all stack frames remain the same after the transformation. Then, a source stack block b is injected into a block with the same path p , but with an index 1 because there is only one block in each stack frame after the transformation. Finally, the offset o at which the source stack block is inserted into stack allocated data is obtained by querying the global environment ge of the source program using the function `find_frame_offset`. With this structural memory injection, we are able to convert the reasoning about the stack as a whole into that about individual stack frames. This makes the simulation proof of `Cminorgen` significantly easier to understand than before. Similar observations can be made for the remaining two transformations on the stack (i.e., `SimplLocals` and `Stacking`).

5 VERIFIED COMPILATION OF PROGRAMS WITH CONTEXTUAL MEMORY

In this section, we discuss the application of our nominal memory model to verified compilation of programs that work with contextual memory. We achieve this by enriching the nominal memory model with even more fine-grained structures to represent contextual memory. These developments demonstrate an effective solution to the second problem in Sec. 2.4.

5.1 Key Ideas

The idea of *contextual compilation* is widely adopted in the research of verified compilation of open programs (e.g., [Gu et al. 2018; Song et al. 2020; Wang et al. 2019]). It is based on the assumption that, when compiling multiple open modules or threads, only one of them is compiled while the others (the context) are fixed. The individually compiled modules or threads are then linked together at the target level to form the whole program. We have already seen such an example in Sec. 2.4.2.

To verify contextual compilation, we need to not only keep track of how memory is transformed by internal functions, but also do that for external functions. Since the context is fixed, the transformation on contextual memory should always be described as *an identity mapping* from source to target memory blocks, regardless of how complicated the memory transformation is for internal executions. However, this seemingly simple task is extremely difficult to complete in the original `CompCert` because it lacks the ability to distinguish memory blocks allocated by internal functions from those by external ones.

We argue that the fine-grained representation of block ids together with structural memory injections provide an elegant solution to contextual compilation of open programs and threads. Let us first take a look at contextual compilation of open programs worked on by *a single thread*. Recall that we have assumed that the whole memory space consists of memory for global definitions and the stack. Because the former is fixed after initialization, contextual programs can only modify the stack space by allocating and deallocating new frames. Then, the stack blocks allocated by external calls should always be mapped to themselves. This is indeed the case in our structural memory injections: the external stack blocks has the form `Stack 0 p b` by the definition of block ids in Sec. 4.2.1 and `Stack 0 p b` is always mapped to itself at offset 0 by the definitions of structural injections (e.g., those in Sec. 4.4.1 and Sec. 4.4.3). Based on this observation, we have already proved that, given any external calls whose semantics simulate themselves under the identity mapping of contextual memory, they are compatible with the simulation proofs for internal executions. This indicates that the extension in Sec. 4 already supports contextual compilation to a certain extent.

We also would like to support contextual compilation of multi-threaded programs. As we have discussed in Sec. 2.4.2, the existing solutions invented ad hoc mechanisms to cope with the global `nextblock` which prevent further compilation to a realistic machine model in which each thread has its own contiguous and private stack. In the rest of this section, we show that by further instantiating supports with multiple stack trees, we can grow the stacks individually without interference with each other, thereby eliminating the problems with `nextblock`. Furthermore, by enriching supports

with multiple abstract stacks following the idea of Stack-Aware CompCert [Wang et al. 2019], we are able to compile multi-threaded programs onto multi-stack machine models. These ideas form a complete solution to thread-safe contextual compilation, which we shall discuss in detail below.

5.2 Stack-Aware Nominal CompCert

We first extend the instance of Nominal CompCert in Sec. 4 to support compilation with a single and finite stack by incorporating the key ideas in Stack-Aware CompCert [Wang et al. 2019]. Stack-Aware CompCert explicitly manages the call stack by adding a data type called *abstract stack* to memory states. The abstract stack records the history of memory consumption incurred by stack allocation and maintains fine-grained *stack permissions*. By exploiting that information, Stack-Aware CompCert achieves contextual compilation of single-threaded C programs into an assembly language that is aware of a single and finite stack.

In our structured nominal memory model, the abstract stack can be readily absorbed into the support. Moreover, we decide to drop stack permissions from the abstract stack and deal with them separately in the future. To understand this design choice, note that there are two different ways to enforce stack permissions in the literature. One is to embed them explicitly in memory models like in Stack-Aware CompCert, and the other is to separately enforce them as part of the reasoning framework (e.g., simulation conventions in CompCertO [Koenig and Shao 2021] and enriched memory injection in CompCertM [Song et al. 2020]). We notice that in the former approach, the complexity for reasoning about stack permissions permeates the entire proof development, while with the latter approach it affects only the specification and composition of simulation proofs at the top level. Therefore, we decide to follow the second approach in this work. Without stack permissions, our extensions provide weaker support to contextual compilation than the original Stack-Aware CompCert does. On the other hand, we are able to prove preservation of stack consumption by using a much simpler technique thanks to the absence of stack permissions (as we shall see below). We leave as future work the investigation of techniques for combining CompCertO or CompCertM with Stack-Aware Nominal CompCert to support stack permissions without explicitly enforcing them.

Now, our abstract stack only contains information about stack consumption by each function call. Like the original Stack-Aware CompCert, it is organized into *stages* of *abstract frames* where a stage corresponds to a continuous sequence of tailcalls. It is defined as follows:

Definition 5.1 (Abstract Stack). Abstract frames are records of the type $\text{frame} := \{\text{fsize} : \mathbb{Z}; \text{fsize} \geq 0\}$ where *fsize* contain the sizes of concrete stack frames. A stage is a list of abstract frames $\text{stage} := \text{list frame}$ allocated by a sequence of tailcalls where its head is the frame for the active tailcall and the remaining frames have been deallocated by previous tailcalls in the sequence. Finally, an abstract stack is a list of stages $\text{stackadt} := \text{list stage}$.

Given an abstract stack *a*, its size is the summation of sizes of all the frames in *a*. We introduce a constant `MAX_STACK` that denotes the maximum stack size and enforces finiteness of the stack. We now extend the support type defined in Sec. 4.2.3 with an abstract stack:

```
Record sup := mkSUP { global: list ident; stack: stree; astack: stackadt; }.
```

The operations for managing the structure of the abstract stack are the following: `push_stage : mem → mem` pushes a new stage onto the abstract stack. It is invoked when a regular call happens. `record : mem → frame → [mem]` pushes a new frame into the topmost stage of the active stack. It succeeds only if the stack size after pushing this frame does not exceeds `MAX_STACK`. It is invoked either when a regular call or a tailcall happens. Note that only the topmost frame in a stage is “alive.” When a tailcall records a new frame in the topmost stage, the previous frame becomes “dead”,

Fig. 14. Effects of the Operations on the Abstract Stack

corresponding to the parent frame deallocated by the tail call. $\text{pop_stage} : \text{mem} \rightarrow [\text{mem}]$ is invoked when a regular call followed by a sequence of tailcall returns. It simply pops the topmost stage from the stack. For example, Fig. 14 illustrates the effects of these operations applied to the abstract stack where a regular call to g in function f is followed by a tailcall to h which finally returns.

With the enriched memory model, we update the semantics of every language of Nominal CompCert by inserting push_stage , record and pop_stage operations accordingly. Like in Stack-Aware CompCert, the semantics of a program P is also parameterized by an oracle stackspace that provides the sizes of concrete stack frames for each function, which we shall denote as $[[P, \text{stackspace}]]$. In this semantics, a function call to f queries stackspace for the concrete frame size of f and allocates a new frame of this size; the execution gets stuck if such an allocation overflows the stack (i.e., gets an undefined behavior if the consumed stack space exceeds MAX_STACK).

Then, we replay the proofs of Stack-Aware CompCert by establishing preservation of semantics for every pass. This guarantees that if the source semantics are defined (i.e., the stack is not overflowed), then so are the target semantics. The key to these proofs is to show at any matching point of execution, the stack consumption in the target is no greater than in the source. The only complication appears when verifying advanced optimizations such as tailcall recognition and inlining. Because inlining is performed after tailcall recognition, it may lift certain tailcalls back to regular calls, causing complicated changes in stack consumption. We handle this problem by assuming every tailcall consumes stack space like a regular call up to inlining, and inserting an identity transformation (called RTL_{mach}) after inlining for shifting from this assumption to an accurate account of stack consumption by tailcalls. This greatly simplifies the proofs of preservation of stack consumption. However, it also means the abstract stack may be out-of-sync with the actual stack tree before inlining. This is exactly why we choose to keep the abstract stack and the stack tree separate in the definition of supports.

By composing the proofs for its individual passes, we derive the correctness of Stack-Aware Nominal CompCert stated as follows:

THEOREM 5.2 (CORRECTNESS OF STACK-AWARE NOMINAL COMPCERT). *Let $C_{\text{sa-nominal-compcert}}$ be Stack-Aware Nominal CompCert that compiles C programs into assembly programs with an abstract stack. We have*

$$\forall P_s, P_t, C_{\text{sa-nominal-compcert}}(P_s) = [P_t, \text{stackspace}] \implies [[P_t, \text{stackspace}]] \leq [[P_s, \text{stackspace}]].$$

where stackspace is an instance of the oracle (a mapping from functions to sizes of stack frames) generated by the compiler at the Mach level where the concrete layouts of frames are fixed.

5.3 Multi-Stack CompCert

Multi-Stack CompCert is a straightforward extension of Stack-Aware Nominal CompCert by 1) further compiling assembly programs with abstract stacks into those operating over a contiguous and finite stack, known as RealAsm programs, and 2) enriching the support with multiple copies of stack trees and abstract stacks. These extensions guarantee that the target code operates on a realistic machine model where each thread has its own *finite*, *contiguous* and *private* stack.

$$\begin{array}{c}
\frac{L[A] \vdash_R M : L_1[A] \quad L[A] \vdash_R N : L_2[A]}{L[A] \vdash_R M \oplus N : L_1[A] \oplus L_2[A]} \text{HComp} \\
\frac{L_1[A] \vdash_R M : L_2[A] \quad L_2[A] \vdash_S N : L_3[A]}{L_1[A] \vdash_{R \circ S} M \oplus N : L_3[A]} \text{VComp} \\
\frac{L_1[A] \vdash_R M : L_2[A] \quad L_1[B] \vdash_R M : L_2[B] \quad \text{compat}(L_2[A], L_2[B])}{L_1[A \cup B] \vdash_R M : L_2[A \cup B]} \text{PComp}
\end{array}$$

Fig. 15. Composition Rules in CCAL

The compilation of assembly programs with abstract stacks into RealAsm programs follows the approach in Stack-Aware CompCert: we first merge the stack frames in CompCert’s Asm into a single stack, we then eliminate the pseudo instructions for managing stack frames to get to RealAsm.

The enrichment of supports is achieved by defining the support type as follows:

```
Record sup := mk sup { global: list ident; stack: list stree; astack: list stackadt; sid: nat}.
```

Compared to the sup type of Stack-Aware Nominal CompCert, it now has a list of stack trees and a list of abstract stacks. The field sid denotes index to the stack which the running thread operates on. The functions get_stacktree and get_abststack are then defined for accessing the stack tree and the abstract stack indexed by sid, respectively.

With the above abstraction, the entire development of Stack-Aware Nominal CompCert is lifted to work with multiple stacks, resulting in Multi-Stack CompCert. This is possible because we assume sid is bound to a thread which cannot by itself switch to a different stack. Instead, context switching must be completed through external mechanisms (e.g., Certified Concurrent Abstraction Layers [Gu et al. 2018]). With this assumption, a thread simply treats all the stacks except for its own one as part of the contextual memory. Its semantics is completely unchanged except for the uses of get_stacktree and get_abststack for accessing its own stack. In the end, we get the correctness theorem of Multi-Stack CompCert which is exactly like Theorem. 5.2 except for the updated semantics and the target language (i.e., RealAsm).

Note that, Multi-Stack CompCert can be very easily proved correct because the modification to contextual stacks is completely irrelevant to the operations on the focused stack. This is in stark contrast with the situation in the original CompCert where such modification affects nextblock at a global scope. These developments illustrate the simplicity and power of nominal memory model in formalizing contextual memory, which we shall further exploit below.

5.4 Thread-Safe Contextual Compilation

To demonstrate the effectiveness of Multi-Stack CompCert in verified compilation of multi-threaded programs, we apply it to solve the problem of compiling and linking Concurrent Certified Abstraction Layers (CCAL) as described in Sec. 2.4.2.

5.4.1 Challenges in Compiling and Linking CCAL Objects. To understand the underlying challenges, we first give a more detailed introduction to CCAL. CCAL is a generalization of Certified Abstraction Layers [Gu et al. 2015] for building concurrent programs in a layered style. A certified concurrent layer L provides shared and private memory states and primitive operations for manipulating them. Concurrent objects are built on top of a lower layer L and provide implementations of a higher layer L' . From the point view of the user of L , there is no need to worry about how execution of objects running on L interleaves in a concurrent setting, because the shared primitives are already abstracted into atomic operations. A predicate of the form $L[A] \vdash_R M : L'[A]$ formally describes

Fig. 16. Realistic Linking of CCAL Objects

the implementation M of the layer L' running on top of the layer L . Given the domain of threads $\mathbb{D}$, it holds iff concurrent execution of M with a focused subset of threads $A \subseteq \mathbb{D}$ on L backward simulates the layer L' where R is the invariant of the simulation. Because A may be only a subset of all threads, M is an open multi-threaded program whose execution trace may be interleaved with that of unknown threads in the context. With the devices above, it is possible to *horizontally*, *vertically* and *parallelly* compose concurrent layers, as described by the rules in Fig. 15.

The common pattern of developing verified concurrent programs using CCAL is to start with source modules focusing on a single thread, i.e., layer implementations of the form $L[i] \vdash_R M : L'[i]$ where i is the id of the focused thread, then compile these modules as sequential programs, and finally compose the generated objects at the assembly level.

For the above approach to work, it is essential that the compilation is *thread safe*, i.e., when compiling modules focusing on a single thread, the compiler C indeed preserves CCAL:

$$\forall i L L' M R, L[i] \vdash_R M : L'[i] \implies L[i] \vdash_R C(M) : L'[i].$$

As we have discussed in Sec. 2.4.2, thread-safeness as currently implemented has serious problems in that it relies on introduction of dummy blocks and does not support compilation to realistic machine models.

5.4.2 Compilation and Linking of CCAL Objects onto Multi-Stack Machine Models. We make immediate use of Multi-Stack CompCert to overcome the above challenges. Given a CCAL object $L[i] \vdash_R M : L'[i]$ written in C, we assign a unique stack i to M in Multi-Stack CompCert. Then, M represents a sequential program that is ignorant of its context in a multi-stack memory model, exactly matching the requirement of source programs for Multi-Stack CompCert. By individually compiling CCAL objects with Multi-Stack CompCert, we can reduce them to RealAsm programs that look like sequential assembly codes and that are ignorant of the existence of other threads. At the target level, the stacks are finite and contiguous and the thread-local data are allocated by adjusting the stack pointers to private stacks.

To link the generated RealAsm programs, we observe that because the heap is managed by primitives in abstraction layers in CCAL (the same as our treatment of heaps as discussed in Sec. 4.1), no new memory blocks will ever be allocated after the initial allocation of global variables and the finite stacks. Therefore, there is no need to allocate dummy blocks or to synchronize any memory structure across threads. Furthermore, we observe that the finite stacks for threads are separated from each other from the beginning and the operations on them never interfere with each other. Therefore, stack merging becomes trivial. For example, the problematic situation described in Fig. 5 becomes the natural merging of multiple private stacks as described in Fig. 16.

Table 1. Proof Effort for Nominal CompCert Relative to CompCert 3.8

Files	CC 3.8	NCC	Additions(+)	% of Additions	Changes(*/+)	% of Changes
Memory.v	3838	4030	192	5.00%	372	9.69%
Memtype.v	866	893	27	3.12%	52	6.00%
Globalenvs.v	1640	1665	25	1.52%	128	7.80%
SimplLocalsproof.v	1971	1992	21	1.07%	143	7.26%
Cminorgenproof.v	1903	1931	28	1.47%	203	10.67%
Inliningproof.v	1161	1178	17	1.46%	178	15.33%
ValueAnalysis.v	1805	1847	42	2.33%	159	8.81%
Unusedglobproof.v	1259	1271	12	0.95%	67	5.32%
Total	136682	137338	656	0.48%	1396	1.02%

With the above changes, we are able to combine Multi-Stack CompCert with CCAL to—for the first time—realize the development of verified concurrent objects operating on a realistic machine model with multiple contiguous and finite stacks.

6 EVALUATION

In this section, we discuss the proof effort for our developments. We first present an evaluation of the changes introduced into Nominal CompCert relative to the original CompCert. It illustrates the conciseness of the extension for obtaining Nominal CompCert, which is a consequence of the natural generalization of CompCert’s memory model to the nominal memory model. We then present an evaluation of the proof effort for further extending Nominal CompCert as described in Sec. 4 and Sec. 5. It demonstrates that these extensions only require moderate effort to develop, thanks to the generality of the interfaces for the nominal memory model and, on top of that, the ability to exploit sophisticated instantiations of the nominal memory model to verify individual compilation passes on an as-needed basis.

6.1 Proof Effort for Nominal CompCert

It took us 1 person month to develop Nominal CompCert, with much of the effort devoted to understanding the details in the proofs for CompCert’s passes, especially for those changing the memory structure. We give the statistics of our Coq implementation of Nominal CompCert (NCC) relative to CompCert 3.8 (CC 3.8) in Table. 1. Columns 2 and 3 show the lines of code (LOC) for each file (counted by using `coqwc`) in CompCert 3.8 and Nominal CompCert, respectively. Column 4 shows the LOC that were *added*, while column 5 shows them in percentage relative to CompCert 3.8. To better reflect on the actual proof development, we further show the LOC that were *modified* or *added* and their percentage relative to CompCert 3.8 in columns 6 and 7. Due to space limitations, this table only lists files with substantial (more than 5%) changes in LOC.

As we can see in Table. 1, the total amount of additional LOC is quite small (about 650 lines and 0.5%), with only a few lines for each file except for `Memory.v`. The substantial additions include 1) an implementation of the nominal memory model as described in Sec. 3.2 and 2) explicit proofs about stack pointer and valid blocks as described in Sec. 3.3. The total amount of changes is also relatively small (about 1.4k lines and 1% changes). It is worth noting that, except for the substantial additions mentioned above, the remaining changes are mostly trivial substitutions of `nextblock` with support that do not increase the LOC. From the above statistics, we can see that our implementation of Nominal CompCert is indeed lightweight.

Table 2. Proof Effort for Extensions Relative to Nominal CompCert

Files	NCC	NCC-SMS			SA-NCC			MS-CC		
	LOC	LOC	+	%	LOC	+	%	LOC	+	%
Memory.v	4030	5485	1455	36.1%	6599	2569	63.7%	6861	2831	70.2%
Globalenvs.v	1665	1803	138	8.29%	1834	169	10.2%	1848	183	11.0%
SimplLocalsproof.v	1992	2397	405	20.3%	2509	517	26.0%	2533	541	27.2%
Cminorngenproof.v	1931	2358	427	22.1%	2518	587	30.4%	2536	605	31.3%
Unusedglobproof.v	1271	1478	207	16.3%	1544	273	21.5%	1571	300	23.6%
Stackingproof.v	1823	2007	184	10.1%	2059	236	12.9%	2128	305	16.7%
Total	137338	140787	3449	2.51%	145278	7940	5.78%	151833	14495	10.6%

6.2 Proof Effort for Extending Nominal CompCert

Table 2 shows the statistics for Nominal CompCert with Structured Memory Space (NCC-SMS) introduced in Sec. 4 and Stack-Aware Nominal CompCert (SA-NCC) and Multi-Stack CompCert (MS-CC) introduced in Sec. 5. For each extension, it shows the LOC for the representative files, and the additional LOC (column +) and their percentages relative to Nominal CompCert (column %).

6.2.1 Proof Effort for NCC-SMS. It took us about 2 person months to implement Nominal CompCert with Structured Memory Space. A major addition is the implementation of the memory model with structured space together with the properties of newly added memory operations (in `Memory.v` and `Globalenvs.v`). On top of this, we have improved the proofs for the four non-optimizing passes in `SimplLocalsproof.v`, `Cminorngenproof.v`, `Unusedglobproof.v`, and `Stackingproof.v`, by making them more intuitive as described in Sec. 4. This requires moderate modification to the original files (about 10-22% additional LOC) as shown in Table 2. Besides these changes, we have also re-verified all the remaining compilation passes of CompCert based on the new memory model. The whole development amounts to about 3.5k additional LOC, which is only about 2.5% more LOC on top of Nominal CompCert. We are able to finish it with moderate effort because that, thanks to the abstraction provided by Nominal CompCert, we can reuse most of the old proofs. Note that because our proofs are still based on the old ones they are *not* simpler than before, even though the structural injections do help make the essence of memory transformations explicit. To get simpler proofs, we will need to rewrite the proofs from scratch, which is left for future work.

6.2.2 Proof Effort for SA-NCC and MS-CC. It took us an additional 2 person months to implement Stack-Aware Nominal CompCert. The main effort includes further enriching the nominal memory model and proving preservation of stack consumption for every compilation pass. These changes amount to about 8k more LOC (relative to Nominal CompCert) as shown in Table 2. They are significantly less than the 21k additional LOC (relative to CompCert 3.0.1) for the original Stack-Aware CompCert [Wang et al. 2019]. The main reasons are that 1) we have dropped stack permissions, 2) we have implemented the abstract stack as part of the support type, with which the relation between source and target stack consumptions can be easily absorbed into memory injections, and 3) we have invented a novel technique (introduced in Sec. 5.2) for approximating stack consumptions before inlining and introduced an identity transformation after inlining to convert this over approximation into accurate stack consumptions. Finally, Multi-Stack CompCert is obtained from Stack-Aware Nominal CompCert by adding about 500 LOC for enriching the support type with multiple stacks and about 6k LOC for extending the compilation chain to `RealAsm`, which took us another 1 person month.

7 RELATED WORK

Nominal techniques [Gabbay and Pitts 2002; Pitts 2016] have been widely used to define the semantics of formal calculi with binders (e.g., λ -calculus, π -calculus) using inductive definitions of nominal sets [Pitts 2013]. They have also been used in the game-semantics community [Abramsky et al. 2004; Laird 2004; Murawski and Tzevelekos 2016] to define the nominal games which led to full abstraction results for languages with dynamic generative behaviors, such as ν -calculus [Abramsky et al. 2004], higher-order concurrency [Laird 2006], ML references [Murawski and Tzevelekos 2011], and Interface Middleweight Java [Murawski and Tzevelekos 2014, 2021]. Urban and Berghofer [Urban and Berghofer 2006; Urban and Tasson 2005] have implemented a “nominal datatype” package (in the Isabelle/HOL proof assistant) which has been used to mechanize formal proofs and operational semantics using nominal techniques. However, none of these have attempted to address the memory semantics and verified compilation of low-level C-like languages.

CompCert uses a unified memory model [Leroy et al. 2012; Leroy and Blazy 2008] for all of its compiler intermediate and target languages. The memory model treats global variables, and heap and stack objects as separate memory blocks to enforce isolation. It supports bound-checking (to give semantics to “undefined behaviors”) and uses the `Vundef` value to denote results from ill-defined loads. CompCert memory block identifiers are kept relatively abstract and can be renamed, deleted, or injected as sub-blocks of bigger blocks while preserving the observable behaviors of programs. While the abstract block identifiers in CompCert seem to be suitable for a nominal treatment, for a long time, it was not clear how it should be done and what benefits it would bring. CompCert’s memory-injection-based simulation proofs are also quite challenging so it is unclear whether nominal techniques can actually help improve the existing proofs.

What we have shown in this paper is that the nominal techniques can indeed be used to both generalize and simplify the CompCert memory model in a really clean way. Furthermore, the nominal extension is backward compatible in that regardless what representations we use for the block identifiers (“names”), the rest of the compiler (including all the memory-injection-heavy compilation phases and their proofs) will remain valid. The separation of “global-variables-” and “stack-” supports allows us to improve the memory-injection-related proofs, but the overall CompCert memory remains as a mapping from block identifiers to values as before (so the existing proofs would still work). This is in contrast to previous attempts [Ramananandro et al. 2015; Wang et al. 2019] in which they either had to introduce a “stack tag” for each block [Ramananandro et al. 2015] or add a separate stack component [Wang et al. 2019]; both required a major overhaul over the memory-injection-related proofs in CompCert.

There exists abundant work on extending CompCert to support various memory structures for verified compilation of open and concurrent programs. We shall make comparisons with them from the following two perspectives:

Stack-Awareness. Stack-Awareness means the support of a memory model with an explicit notion of stack memory, in particular, how close the memory model of the final target language is to the actual memory model used by concurrent machine or assembly code.

There has been previous work on translating the unbounded memory of CompCert into some kind of finite memory. CompCertS [Besson et al. 2017] supports low-level manipulation of pointer values in a finite memory space; CompCert-TSO [Sevcík et al. 2013] builds in a notion of finite memory into all levels of CompCert for certifying concurrency in compilation. A lower-level semantics for CompCert assembly that models pointers as 32-bit values for verifying the peephole optimization is defined by Mullen et al. [2016].

None of the above work tries to model an explicit stack in memory. Quantitative CompCert (QCC) [Carbonneaux et al. 2014] does present a more stack-aware view of CompCert; it extends

event traces with call/return events, and then reasons about stack consumption in terms of the *same* structure of such call and return events on traces of the source and target programs. The Cerco C compiler [Amadio et al. 2014] uses a similar approach to QCC for reasoning about stack consumption and uses a different backend from that of CompCert. Because the exact same call/return events are needed throughout the compilation, they cannot support inlining and tailcall optimizations which are essential to the performance of the generated code.

Stack-Aware CompCert [Wang et al. 2019] is the first extension to CompCert that explicitly models a finite stack and that supports all the optimization passes in CompCert. Its further extension (known as CompCertMC) also supports merging of stack frames into a finite and contiguous stack and elimination of pseudo instructions for stack manipulation. The target code CompCertMC generates is very close to machine code. CASCompCert [Jiang et al. 2019] reasons about what they call the footprints of programs, and in particular reasons about memory ownership. In its memory model, an infinite set of memory locations is used for allocating stack frames. It seems challenging to use their framework to merge stack blocks into a finite and contiguous stack.

To support Stack-Awareness, all the above work requires intrusive and sometimes ad hoc changes to CompCert’s memory model. On the other hand, Nominal CompCert provides a well-defined and flexible interface for supporting Stack-Awareness through instantiations of block ids and supports. As we have shown before, a complete call stack can be seamlessly integrated into Nominal CompCert with Structured Memory Space, and compilation to multiple and contiguous machine stacks can be supported without intrusive changes through the development of Stack-Aware Nominal CompCert and Multi-Stack CompCert.

Compositional Verification of Compilation. We have shown that, with the nominal memory model and its instances, contextual memory can be separated from internal memory and be reasoned about via a well-defined interface, leading to simpler and more elegant correctness proofs for compiling certain kinds of open programs. We believe that these techniques together with others we have built on top of the nominal memory model could be particularly beneficial in the general context of compositional compiler correctness; we elaborate on this point below.

In CompCert, the complexity of memory injections is largely confined to the correctness proofs of individual passes: injections are existentially quantified as part of the simulation relation and do not appear in the correctness statement itself. This is possible because memory states are not part of the externally observable behavior of programs.

This assumption must be relaxed in compositional extensions of CompCert which model interactions between compilation units. In work such as Compositional CompCert [Stewart et al. 2015], CompCertX [Gu et al. 2015; Wang et al. 2019], CompCertM [Song et al. 2020], CASCompCert [Jiang et al. 2019] and CompCertO [Koenig and Shao 2021], memory states appear as part of the interactions between components. As a consequence, the memory relations used by compilation passes become part of their correctness statements. This makes correctness theorems and the vertical composition of passes much more complex: in CompCertX, memory injections appear in the overall compiler correctness theorem and the passes are composed in an ad-hoc manner; in Compositional CompCert, memory injections must be extended with ownership information and become even more intricate; CompCertO and CompCertM go one step further and introduce Kripke logical relations to deal with the variations in the memory relations used by different compilation passes and facilitate the composition of correctness results.

Our techniques based on the nominal memory model have the potential to significantly simplify the above proofs from the following perspectives. First, the evolving Kripke worlds used in compositional compiler correctness are used to store both memory injections (which can be determined ahead of time using our structural injection functions) and additional permission information

(which could be stored in our model as part of the information associated with block identifiers). As a result, by incorporating our approach, it may be possible to eliminate Kripke worlds entirely, which would significantly simplify the semantic frameworks used in these projects. Moreover, our techniques could help isolate private memory as internal states, for example through the use of a partial commutative monoid structure on memory states which would facilitate splitting and merging private memory regions used by individual components. Finally, they could provide a seamless way to incorporate stack awareness to compositional compilers of this kind.

8 CONCLUSION

In this paper, we have developed a nominal memory model, an elegant generalization of CompCert's block-based memory model that employs the nominal notions of atomic names, supports, and freshness. The nominal memory model comes with an abstraction of block ids and supports. Thus, block ids can be generalized from positive numbers to any desired type for keeping track of names; the notion of valid blocks in a memory m corresponds to the support of m , and allocation must return a name that is simply fresh with respect to the support.

Nominal logics build on the notion of swaps (exchanging two names) and make sure that all definable objects are invariant by swaps. This aspect of normality is partly captured using memory injection in CompCert, but is not fully explored in our paper because clashing of names does not happen in the setting of whole program compilation or compilation with fixed contexts where the assignment of block ids for a given program is fixed. We believe that equivalence of semantics under name swapping is essential for reasoning about general compositional compilation of open modules and concurrent programs, which we plan to investigate in the future.

We have extended CompCert to build Nominal CompCert, with which the original CompCert becomes a specialization obtained by instantiating block ids with positive numbers. We then developed an extension of Nominal CompCert by instantiating block ids and supports with structured types that allow us to distinguish different memory regions. This led to more elegant and modular reasoning about various compiler transformations. We built Stack-Aware Nominal CompCert which uses the nominal idea to improve Stack-Aware CompCert and Multi-Stack CompCert which is the first effort to support compilation of multi-threaded programs all the way to multi-stack machines.

We have left the proofs for inlining and tailcall recognition passes as they are. Because inlining changes the structure of call trees, fixed paths are probably not enough for defining its structural injection function which needs to capture the non-local transformation on stack block ids. We need more information about how inlining of earlier function calls will affect the locations of stack blocks allocated later in the call tree. This can be done by representing stack block ids as "snapshots" of stack trees at allocation time which contain all the historical information for calculating the effects of inlining on stack block ids. Similar observations can be made for tailcall recognition. A complete solution is left for future work.

ACKNOWLEDGMENTS

We would like to thank the anonymous referees for their helpful feedback which improved this paper significantly. This research is based on work supported in part by the National Natural Science Foundation of China (NSFC) under Grant No. 62002217, and by the Natural Science Foundation of the United States (NSF) under Grant No. 1521523, 1763399, 2019285, and 2118851. The third author is a co-founder of and has an equity interest in CertiK Global Ltd. CertiK has licensed Yale University's intellectual property, which is related to the NSF grants 1521523 and 1763399. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.

REFERENCES

- Samson Abramsky, Dan R. Ghica, Andrzej S. Murawski, C.-H. Luke Ong, and Ian David Bede Stark. 2004. Nominal Games and Full Abstraction for the Nu-Calculus. In *Proc. 19th IEEE Symposium on Logic in Computer Science (LICS'04)*. IEEE Computer Society, 150–159. <https://doi.org/10.1109/LICS.2004.1319609>
- Roberto M. Amadio, Nicolas Ayache, Francois Bobot, Jaap P. Boender, Brian Campbell, Ilias Garnier, Antoine Madet, James McKinna, Dominic P. Mulligan, Mauro Piccolo, Randy Pollack, Yann Régis-Gianas, Claudio Sacerdoti Coen, Ian Stark, and Paolo Tranquilli. 2014. Certified Complexity (CerCo). In *Foundational and Practical Aspects of Resource Analysis*, Ugo Dal Lago and Ricardo Peña (Eds.). LNCS, Vol. 8552. Springer, Cham, 1–18. https://doi.org/10.1007/978-3-319-12466-7_1
- Andrew Appel. 2011. Verified Software Toolchain. In *Proc. 20th European Symposium on Programming (ESOP'11)*, Gilles Barthe (Ed.). LNCS, Vol. 6602. Springer, Saarbrücken, Germany, 1–17. https://doi.org/10.1007/978-3-642-19718-5_1
- Frédéric Besson, Sandrine Blazy, and Pierre Wilke. 2017. CompCertS: A Memory-Aware Verified C Compiler Using Pointer as Integer Semantics. In *Interactive Theorem Proving (ITP'17)*, Mauricio Ayala-Rincón and César A. Muñoz (Eds.). LNCS, Vol. 10499. Springer, Cham, 81–97. https://doi.org/10.1007/978-3-319-66107-0_6
- Quentin Carbonneaux, Jan Hoffmann, Tahina Ramananandro, and Zhong Shao. 2014. End-to-End Verification of Stack-Space Bounds for C Programs. In *Proc. 2014 ACM Conference on Programming Language Design and Implementation (PLDI'14)*. ACM, New York, 270–281. <https://doi.org/10.1145/2594291.2594301>
- Murdoch Gabbay and Andrew M. Pitts. 2002. A New Approach to Abstract Syntax with Variable Binding. *Formal Aspects Comput.* 13, 3-5 (2002), 341–363. <https://doi.org/10.1007/s001650200016>
- Ronghui Gu, Jérémie Koenig, Tahina Ramananandro, Zhong Shao, Xiongnan(Newman) Wu, Shu-Chun Weng, Haozhong Zhang, and Yu Guo. 2015. Deep Specifications and Certified Abstraction Layers. In *Proc. 42nd ACM Symposium on Principles of Programming Languages (POPL'15)*. ACM, New York, 595–608. <https://doi.org/10.1145/2775051.2676975>
- Ronghui Gu, Zhong Shao, Jieung Kim, Xiongnan (Newman) Wu, Jérémie Koenig, Vilhelm Sjöber, Hao Chen, David Costanzo, and Tahina Ramananandro. 2018. Certified Concurrent Abstraction Layers. In *Proc. 2018 ACM Conference on Programming Language Design and Implementation (PLDI'18)*. ACM, New York, 646–661. <https://doi.org/10.1145/3192366.3192381>
- Hanru Jiang, Hongjin Liang, Siyang Xiao, Junpeng Zha, and Xinyu Feng. 2019. Towards Certified Separate Compilation for Concurrent Programs. In *Proc. 40th ACM Conference on Programming Language Design and Implementation (PLDI'19)*. ACM, New York, 111–125. <https://doi.org/10.1145/3314221.3314595>
- Jecheon Kang, Yoonseung Kim, Chung-Kil Hur, Derek Dreyer, and Viktor Vafeiadis. 2016. Lightweight Verification of Separate Compilation. In *Proc. 43rd ACM Symposium on Principles of Programming Languages (POPL'16)*. ACM, New York, 178–190. <https://doi.org/10.1145/2837614.2837642>
- Jérémie Koenig and Zhong Shao. 2021. CompCertO: Compiling Certified Open C Components. In *Proc. 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI'21)*. ACM, New York, 1095–1109. <https://doi.org/10.1145/3453483.3454097>
- James Laird. 2004. A Game Semantics of Local Names and Good Variables. In *Foundations of Software Science and Computation Structures (FOSSACS'04), Held as Part of the Joint European Conferences on Theory and Practice of Software (ETAPS'04)*, Igor Walukiewicz (Ed.). LNCS, Vol. 2987. Springer, 289–303. https://doi.org/10.1007/978-3-540-24727-2_21
- James Laird. 2006. Game Semantics for Higher-Order Concurrency. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS'06)*, S. Arun-Kumar and Naveen Garg (Eds.). LNCS, Vol. 4337. Springer, 417–428. https://doi.org/10.1007/11944836_38
- Xavier Leroy. 2005–2021. The CompCert Verified Compiler. <https://compcert.org/>.
- Xavier Leroy, Andrew W. Appel, Sandrine Blazy, and Gordon Stewart. 2012. *The CompCert Memory Model, Version 2*. Research Report RR-7987. INRIA. 26 pages. <https://hal.inria.fr/hal-00703441>
- Xavier Leroy and Sandrine Blazy. 2008. Formal Verification of a C-like Memory Model and Its Uses for Verifying Program Transformation. *Journal of Automated Reasoning* 41, 1 (2008), 1–31. <https://doi.org/10.1007/s10817-008-9099-0>
- Eric Mullen, Daryl Zuniga, Zachary Tatlock, and Dan Grossman. 2016. Verified Peephole Optimizations for CompCert. In *Proc. 37th ACM Conference on Programming Language Design and Implementation (PLDI'16)*. ACM, New York, NY, USA, 448–461. <https://doi.org/10.1145/2980983.2908109>
- Andrzej S. Murawski and Nikos Tzevelekos. 2011. Game Semantics for Good General References. In *Proc. 26th IEEE Symposium on Logic in Computer Science (LICS'11)*. IEEE Computer Society, 75–84. <https://doi.org/10.1109/LICS.2011.31>
- Andrzej S. Murawski and Nikos Tzevelekos. 2014. Game Semantics for Interface Middleweight Java. In *Proc. 41st ACM Symposium on Principles of Programming Languages (POPL'14)*. ACM, New York, 517–528. <https://doi.org/10.1145/2535838.2535880>
- Andrzej S. Murawski and Nikos Tzevelekos. 2016. Nominal Game Semantics. *Found. Trends Program. Lang.* 2, 4 (2016), 191–269. <https://doi.org/10.1561/25000000017>
- Andrzej S. Murawski and Nikos Tzevelekos. 2021. Game Semantics for Interface Middleweight Java. *J. ACM* 68, 1 (2021), 4:1–4:51. <https://doi.org/10.1145/3428676>

- Andrew M. Pitts. 2013. *Nominal Sets: Names and Symmetry in Computer Science*. Cambridge University Press, Cambridge, England.
- Andrew M. Pitts. 2016. Nominal Techniques. *ACM SIGLOG News* 3, 1 (2016), 57–72. <https://doi.org/10.1145/2893582.2893594>
- Tahina Ramananandro, Zhong Shao, Shu-Chun Weng, Jérémie Koenig, and Yuchen Fu. 2015. A Compositional Semantics for Verified Separate Compilation and Linking. In *Proc. 2015 Conference on Certified Programs and Proofs (CPP’15)*. ACM, New York, 3–14. <https://doi.org/10.1145/2676724.2693167>
- Jaroslav Sevcik, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2011. Relaxed-Memory Concurrency and Verified Compilation. In *Proc. 38th ACM Symposium on Principles of Programming Languages (POPL’11)*. ACM, New York, 43–54. <https://doi.org/10.1145/1926385.1926393>
- Jaroslav Sevcik, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2013. CompCertTSO: A Verified Compiler for Relaxed-Memory Concurrency. *J. ACM* 60, 3 (2013), 22:1–22:50. <https://doi.org/10.1145/2487241.2487248>
- Youngju Song, Minki Cho, Dongjoo Kim, Yonghyun Kim, Jeehoon Kang, and Chung-Kil Hur. 2020. CompCertM: CompCert with C-Assembly Linking and Lightweight Modular Verification. *Proc. ACM Program. Lang.* 4, POPL, Article 23 (Jan. 2020), 31 pages. <https://doi.org/10.1145/3371091>
- Gordon Stewart, Lennart Beringer, Santiago Cuellar, and Andrew W. Appel. 2015. Compositional CompCert. In *Proc. 42nd ACM Symposium on Principles of Programming Languages (POPL’15)*. ACM, New York, 275–287. <https://doi.org/10.1145/2676726.2676985>
- Harvey Tuch, Gerwin Klein, and Michael Norrish. 2007. Types, Bytes, and Separation Logic. In *Proc. 34th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL’07)*, Martin Hofmann and Matthias Felleisen (Eds.). ACM, New York, 97–108. <https://doi.org/10.1145/1190216.1190234>
- Christian Urban and Stefan Berghofer. 2006. A Recursion Combinator for Nominal Datatypes Implemented in Isabelle/HOL. In *International Joint Conference on Automated Reasoning (IJCAR’06)*, Ulrich Furbach and Natarajan Shankar (Eds.). LNCS, Vol. 4130. Springer, 498–512. https://doi.org/10.1007/11814771_41
- Christian Urban and Christine Tasson. 2005. Nominal Techniques in Isabelle/HOL. In *International Conference on Automated Deduction (CADE-20)*, Robert Nieuwenhuis (Ed.). LNCS, Vol. 3632. Springer, 38–53. https://doi.org/10.1007/11532231_4
- Yuting Wang, Pierre Wilke, and Zhong Shao. 2019. An Abstract Stack Based Approach to Verified Compositional Compilation to Machine Code. *Proc. ACM Program. Lang.* 3, POPL, Article 62 (Jan. 2019), 30 pages. <https://doi.org/10.1145/3290375>
- Yuting Wang, Xiangzhe Xu, Pierre Wilke, and Zhong Shao. 2020. CompCertELF: Verified Separate Compilation of C Programs into ELF Object Files. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 197 (2020), 28 pages. <https://doi.org/10.1145/3428265>