

Supplementary Material for: Reconciling High-level Optimizations and Low-level Code with Twin Memory Allocation

Anonymous Author(s)

1 Full Semantics

1.1 Syntax

Figure 1 shows syntax of IR with some instructions omitted for brevity.

1.2 Memory

In our semantics, memory *Mem* is defined as a tuple of a current “time”, a partial function from block ids to memory blocks, and a partial function from call ids to call times. We use $M(l)$ to refer to memory block l , and $M(cid)$ to refer to the call time of call cid . Memory allocation/deallocation (call `malloc()`, `alloca`, `call free()`) increments the time of the memory. The time never decreases during the execution of a program.

The memory has two parameters: First, a partial function `ptrsz` saying for each address space, if it exists, how large the pointer is. Address space 0 has to exist. Second, the parameter `memtwins` says how many memory ranges are allocated for each block. The details of this “twin allocation” will be discussed later.

Memory Block. A memory block is defined as a tuple (t, r, n, a, c, P) . t is a *tag* indicating which instruction was used to allocate this block: For memory blocks allocated by `alloca`, the tag is `stack`; for `malloc` it is `heap`; for global variables, it is `global` tag; and for functions (i.e., the target of function pointers), it is `function`.

r is a pair of timestamps defining the *lifetime* of the block. If $r = (s, e)$, the block is alive in the time range $[s, e)$. When a block is newly allocated, r is assigned (s, ∞) where s is the current time. If the block is freed, r is changed to (s, e) where e is the current time. We say that l is alive, or $\text{alive}_M(l)$, if its lifetime has not ended.

n is the *size* of the block in bytes. a is the *alignment* of the block in bytes. When a logical block becomes a concrete block, its integer address $(P(s))$, which we will discuss shortly must be divisible by a .

c is the *content* of the block, stored as a sequence of n bytes.

P stores, for each address space, the integer addresses of the beginning of the block. These addresses are assigned on allocation. For all address spaces s and twin indices i ,

$P(s)_i + n$ should not exceed the maximal integer address of address space s . For example, if address space 1 has size 2^{16} , then we must have $P(1)_i + n < 2^{16}$. Furthermore, in address space 0, the first and last address (0 and $2^{\text{ptrsz}(0)} - 1$) must not be used for any block. For every address space, the first address ($P(s)_0$, or just $P(s)$ for short) is the base address of the block on the physical machine. The remaining addresses (at indices $1, \dots, \text{memtwins} - 1$) are the base addresses of the *twin blocks*. For every block allocated via `malloc` or `alloc`, we actually reserve *several* blocks of the same size in the address space. This lets us prove that it is impossible for others to correctly “guess” the address that the block has been allocated at. The twin blocks’ addresses are not used anywhere in the semantics. However, we demand that all the address ranges covered by the alive blocks of an address space are disjoint: For all address spaces s and all l_1, l_2, i_1, i_2 , if $(l_1, i_1) \neq (l_2, i_2)$ and $\text{alive}_M(l_1)$ and $\text{alive}_M(l_2)$, then $[M(l_1).P(s)_{i_1}, M(l_1).P(s)_{i_1} + P(l_1).n)$ is disjoint from $[M(l_2).P(s)_{i_2}, M(l_2).P(s)_{i_2} + P(l_2).n)$. Furthermore, the base address of one alive block must not be in the address range covered by another alive block: For all address spaces s and all l_1, l_2, i_1, i_2 , if $(l_1, i_1) \neq (l_2, i_2)$ and $\text{alive}_M(l_1)$ and $\text{alive}_M(l_2)$, then $M(l_1).P(s)_{i_1} \notin [M(l_2).P(s)_{i_2}, M(l_2).P(s)_{i_2} + P(l_2).n)$. This second condition is required to handle 0-sized blocks.

$\text{convert}(s, i, s')$ is a partial function that maps an integer address i from address space s to s' . If there exists $P(s) = i$ and $P(s') = i'$, we have for all offsets $o \leq n$ that $\text{convert}(s, i + o, s') = i' + o$ and $\text{convert}(s', i', s) = i + o$.

Memory Addresses. There are two kinds of memory addresses (besides `poison`): logical addresses $\text{Log}(l, o, s)$, and physical addresses $\text{Phy}(o, s, l, cid)$. Both track their address space to be able to detect partial loads of a pointer, and to detect address space punning on load.

A logical memory address is of the shape $\text{Log}(l, o, s)$, where l is a block id, o is a byte offset from the beginning of the block and s is its address space. An offset o is an *inbounds offset* of l , written $\text{inbounds}_M(l, o)$, if o is non-negative and not larger than size of the block, i.e., $0 \leq o \leq n$. The offset one-past-the-end is explicitly allowed. Because the last address of address space 0 is never allocated, we know that computing on inbounds addresses can never overflow. The offset is *strictly inbounds*, written $\text{strict_inbounds}_M(l, o)$, if

```

111      stmt ::= reg=inst | br op, label, label | br label
112           | call ty funcname (ty op[, ty op]*) fnattr | store ty op, ty* op
113           | ret void | ret ty op | unreachable
114
115      inst ::= load ty, ty* op, align op | alloca ty, align op
116           | ptrtoint ty op to ty | inttoptr ty op to ty
117           | addrspacecast ty op to ty
118           | getelementptr [inbounds] ty op[, [inrange] ty op]*
119           | select ty op, ty op, ty op | freeze ty op
120           | psub ty op op | icmp cond, op, op
121           | phi ty [op, label], ..., [op, label]
122           | call ty funcname (ty op[, ty op]*) fnattr
123
124           | aop [nuw] [nsw] ty op, op // arithmetic ops
125           | div [exact] ty op, op // divisions
126           | lop ty op, op // logical ops
127           | castop ty op to ty // casting ops
128           | ...
129
130      op ::= reg | constant | poison
131      bty ::= isz | ty * addrspac n // base type
132      vty ::= <sz × bty> // vector type
133      sty ::= { ty [, ty]* } // struct type
134      aty ::= [sz × ty] // array type
135      ty ::= bty | vty | sty | aty
136
137      funcname ::= malloc | free | ...
138      fnattr ::= alwaysinline | readnone | readonly | writeonly ...
139      cond ::= eq | ne | ugt | uge | ult | ule | sgt | sge | slt | sle
140      fcond ::= oeq | ogt | olt | ole | one | ord | ueq | uno | ...
141      aop ::= add | sub | mul | shl
142      div ::= udiv | sdiv | urem | srem | lshr | ashr
143      lop ::= and | or | xor
144      castop ::= trunc | zext | sext | fptrunc | fpext | fptoui | fptosi | ...

```

Figure 1. Syntax of LLVM IR (unrelated instructions omitted for brevity)

it is inbounds and it does not point beyond the end of the block, i.e., $0 \leq o < n$.

Physical addresses are of the shape $\text{Phy}(p, s, I, cid)$ where o is the physical address, i.e., it is an offset starting at address 0×0 . s is the address space, and I and cid are additional constraints which should be met when the pointer is dereferenced. I is a set of integer addresses which should be inbounds addresses of the dereferencing memory block when the physical pointer is dereferenced. cid is a CallId enforcing that the physical pointer cannot access memory blocks created inside the function call. I and cid are both used for supporting more alias analysis rules. They are not used in pointer comparison, pointer subtraction, and pointer to integer casting, but address space casting may update I . For simplicity, we write $\text{Phy}(o, s)$ whenever I is an empty set and cid is **None**.

To describe cid , we first introduce the concept of a *call id*. A call id is a natural number that is uniquely assigned to

each function call. For each function call, the time at which it occurs is maintained in the partial map $\text{CallID} \rightarrow \text{Time}$ that is part of the memory. If a physical pointer is passed to a function call and cid of the physical pointer is either **None** or a call that has already returned, cid is updated to the call id of the new function call. Otherwise, cid does not change. Inside the function call, even if a physical pointer points to some memory block l , dereferencing the physical pointer is UB if the beginning of the lifetime of l is not earlier than call time of cid . Escaping the physical pointer (e.g., storing it into a global variable or returning it at the end of the function) does not change cid . After the function call is returned, all physical pointers having the call id act as if their cid are **None**. In other words, even if a pointer with a cid is returned back to its caller, it is no longer restricted in how it can be used.

221	Num(<i>sz</i>)	::= { <i>i</i> 0 ≤ <i>i</i> < 2 ^{<i>sz</i>} }	276
222	Time	::= ℕ	277
223	BlockID	::= ℕ	278
224	CallID	::= ℕ	279
225	Mem	::= Time × (BlockID → Block) × (CallID → Time ⊔ None)	280
226	AddrSpace	::= ℕ	281
227	Block	::= { (<i>t, r, n, a, c, P</i>) <i>t</i> ∈ { stack, heap, global, function } ∧ <i>r</i> ∈ (Time × (Time ⊔ {∞})) ∧ <i>n</i> ∈ ℕ ∧ <i>a</i> ∈ ℕ ∧ <i>c</i> ∈ Byte ^{<i>n</i>} ∧ <i>P</i> ∈ ((<i>s</i> ∈ AddrSpace) → Num(ptrsz(<i>s</i>)) ^{memtwins}) }	282
228			283
229			284
230	LogAddr(<i>s</i>)	::= { Log(<i>l, o, s</i>) <i>l</i> ∈ BlockID ∧ <i>o</i> ∈ Num(ptrsz(<i>s</i>)) }	285
231	PhyAddr(<i>s</i>)	::= { Phy(<i>o, s, I, cid</i>) <i>o</i> ∈ Num(ptrsz(<i>s</i>)) ∧ <i>I</i> ⊂ Num(ptrsz(<i>s</i>)) ∧ <i>cid</i> ∈ CallID ⊔ {None} }	286
232			287
233	Addr(<i>s</i>)	::= LogAddr(<i>s</i>) ⊔ PhyAddr(<i>s</i>)	288
234	[[<i>isz</i>]]	::= Num(<i>sz</i>) ⊔ { poison }	289
235	[[⟨ <i>sz</i> × <i>ty</i> ⟩]]	::= { 0, . . . , <i>sz</i> − 1 } → [[<i>ty</i>]]	290
236	[[<i>ty</i> * addrspace(<i>s</i>)]]	::= Addr(<i>s</i>) ⊔ { poison }	291
237	Name	::= { % <i>x</i> , % <i>y</i> , . . . }	292
238	Reg	::= Name → { (<i>ty, v</i>) <i>v</i> ∈ [[<i>ty</i>]] }	293
239	Byte	::= Bit ⁸	294
240	Bit	::= [[i1]] ⊔ AddrBit	295
241	AddrBit	::= { (<i>p, i</i>) ∃ <i>s. p</i> ∈ Addr(<i>s</i>) ∧ (0 ≤ <i>i</i> < ptrsz(<i>s</i>)) }	296
242			297
243			298
244			299
245			300

Figure 2. Semantic domain

To actually perform a memory access of size $sz > 0$ through a pointer p , it must be *dereferencable* as some block-offset-pair, written $\text{deref}_M(p, sz, l, o)$. If p is a logical pointer $\text{Log}(l, o, 0)$ and $\text{alive}_M(l)$ and $\text{inbounds}_M(l, o + sz)$, then we have $\text{deref}_M(p, sz, l, o)$. If p is a physical pointer $\text{Phy}(p, 0, I, cid)$, there must be a block l and an offset o such that $M(l).P(0) + o = p$ and $\text{alive}_M(l)$ and $\text{inbounds}_M(l, o + sz)$ and moreover, for all $p' \in I$, we must have $\text{inbounds}_M(l, p' - M(l).P(0))$ (i.e., all these p' are inbounds of the same block) and if $cid \neq \text{None}$ and $M(cid) \neq \text{None}$, then $M(l).b < M(cid)$ (i.e., the block was allocated before the function call identified by cid began, and that function call is still ongoing). If all these requirements are satisfied, we have $\text{deref}_M(p, sz, l, o)$. Notice that l and o are uniquely determined even for physical pointers due to memory blocks being disjoint.

The NULL pointer of address space s is defined as $\text{Phy}(0, s, \emptyset, \text{None})$. Defining NULL pointer as physical pointer allows folding $\text{inttoptr}(0)$ into NULL ($\text{inttoptr}(x)$ is an instruction that casts integer x to pointer), and replacing p by NULL if $p == \text{NULL}$ holds. Also LLVM can optimize NULL + idx into $\text{inttoptr}(\text{idx})$.

Values. $[[ty]]$ denotes the set of values of type ty . An integer value of type isz is either a concrete number i within range $0 \leq i < 2^{sz}$, or **poison**. A pointer value of type ' $ty * \text{addrspace}(s)$ ' is defined as either a logical address $\text{Log}(l, o, s)$, a physical address $\text{Phy}(o, s, I, cid)$, or **poison**. There is no distinction between pointer values of different types (i32* , i64* , ..), but there is a distinction between

pointer values of different address spaces.¹ This is needed to make sure that load punning cannot be used to perform an address space cast, and because the size of a pointer may depend on its address space. The register file Reg maps a name to a type and a value of that type.

Byte and Bit denote the set of values that one byte or bit can hold, respectively. A byte can hold 8 bits. A bit can hold either a value of type **i1** (i.e., 0, 1, or **poison**) or the i th bit of a pointer value p . Storing to memory involves converting a value to an array of bits. $ty\Downarrow(v)$ is a function that converts value v to bits. Similarly, loading a value from memory involves converting bits to a value. $ty\Uparrow(b)$ is a function that converts bits b to a value of type ty .

$$\begin{aligned} ty\Downarrow &\in [[ty]] \rightarrow \text{Bit}^{\text{bitwidth}(ty)} \\ ty\Uparrow &\in \text{Bit}^{\text{bitwidth}(ty)} \rightarrow [[ty]] \end{aligned}$$

To convert an individual bit, we define a partial function $\text{getbit } vi$ that returns i th bit of a value v with base type bty . If v is an integer, $\text{getbit } vi$ returns i th bit of the integer v . If the integer v is **poison**, all its bits are **poison**; otherwise all bits are either 0 or 1. If v is a pointer, $\text{getbit } vi$ returns either **poison** if v is **poison** or a pair (p, i) which is an element of AddrBit denoting the i th bit of a non-poison pointer p . We

¹This matches LLVM's plans to move to a ptr type: <https://lists.llvm.org/pipermail/llvm-dev/2015-February/081822.html>.

define all bits of **poison** to be **poison**, regardless of its type.

```

getbit ∈  $\llbracket bty \rrbracket \rightarrow \mathbb{N} \rightarrow \text{Bit}$ 
getbit  $n\ i$  = platform dependent
              where  $n \in \llbracket isz \rrbracket, 0 \leq i < sz$ 
getbit  $p\ i$  =  $(p, i)$ 
              where  $p \in \llbracket ty * \text{addr space}(s) \rrbracket,$ 
               $0 \leq i < \text{ptrsz}(s)$ 

```

Further details of these definitions will be given together with the load/store instructions.

1.3 Instructions

In this subsection, we introduce the semantics of the IR instructions. Instruction ι updates the register file $R \in \text{Reg}$ and the memory $M \in \text{Mem}$, denoted $R, M \xrightarrow{\iota} R', M'$. Note that program text and program counter are omitted in state because every operation explained in this section increments the program counter by one and does not change the program text. The aforementioned global map from call id to call time is omitted as well. For semantics of branch instructions, we follow earlier work [?].

The value $\llbracket op \rrbracket_R$ of an operand op in register file R is given by:

```

 $\llbracket r \rrbracket_R = R(r)$  // register
 $\llbracket C \rrbracket_R = C$  // constant
 $\llbracket \text{poison} \rrbracket_R = \text{poison}$  // poison

```

We propose to add one new instruction to LLVM: **psub**.² This instruction takes two pointers p_1, p_2 and returns $p_1 - p_2$ as an integer. Currently LLVM uses ‘sub (ptrtoint p1), (ptrtoint p2)’ to subtract two pointers. This is already correct in this semantics, but using **psub** can improve compiler’s optimization power.³

Integer↔Pointer Casting. We formally define the semantics of **ptrtoint** and **inttoptr** instructions. Figure 3 shows semantics of **ptrtoint** and **inttoptr**. There are two auxiliary functions $\text{cast2int}_M(l, o, s)$ and $\text{cast2ptr}(o, s)$:

```

 $\text{cast2int}_M(l, o, s) = (P(s) + o) \% 2^{\text{ptrsz}(s)}$  where
 $M(l) = (t, r, n, a, c, P)$ 
 $\text{cast2ptr}(o, s) = \text{Phy}(o, s, \emptyset, \text{None})$ 

```

Casting from a logical pointer to integer, or $\text{cast2int}_M(l, o, s)$, yields an integer $P(s) + o$ based on block l . If $P(s) + o$ overflows the size of the address space, it wraps around to 2’s complement. (This can only happen if the pointer is not inbounds.) The semantics of **ptrtoint** is easily represented by cast2int . ‘**ptrtoint** $\text{Log}(l, o, s)$ ’ computes $\text{cast2int}_M(l, o, s)$ and returns it. ‘**ptrtoint** $\text{Phy}(o, s, l, cid)$ ’ simply returns o . If the size of the destination type is isz is

larger than the size of the source type, it is zero-extended. If it is smaller than it, most significant bits are truncated.

Casting from an integer to a pointer, or $\text{cast2ptr}(o, s)$, returns a physical pointer with no provenance information. One option here is to add o to l upon casting, making sure that if this pointer is ever dereferenced, it is still in the block that it started out in. However, that would invalidate replacing p by $\text{inttoptr}(\text{ptrtoint}(p))$.⁴

‘**inttoptr** $isz\ o\ \text{to}\ ty * \text{addr space}(s)$ ’ is equivalent to $\text{cast2ptr}(o, s)$ if the size of the source type isz is equivalent to the size of the destination type $ty * \text{addr space}(s)$. If the size of the source type is larger, high bits of o are truncated. If the size of the source type is smaller, o is zero-extended.

In this semantics, **inttoptr** and **ptrtoint** instructions are allowed to freely move around, be removed, or be introduced.

Address-Space Casting. LLVM IR is a general-purpose intermediate language, and it can be used to compile programs for GPUs as well. The address space of a GPU typically disjoint from the one of the CPU, and moreover, many have multiple address spaces themselves. A programmer can choose which memory to use for allocation. To handle that, LLVM IR tracks address space of a pointer in its type. A pointer in one address space can be casted to a pointer in another address space using **addrspacecast**.

Figure 4 shows formal semantics of **addrspacecast**. If the given pointer is a physical pointer $\text{Phy}(o, s, l, cid)$, the instruction translates both the offset o and the inbounds offsets l using $\text{convert}(s, _, s')$. If the result of convert is not defined (the function is partial, after all), the result is **poison**. If converting any offset in l fails, it is **poison** as well. **addrspacecast** is a capturing operation, as is **ptrtoint**.

If the given pointer is a logical pointer $\text{Log}(l, o, s)$, its block id is maintained, and the new offset is calculated as follows. First, the pointer is casted to integer using cast2int_M . Next, the integer is converted into the corresponding integer address in s' using convert . Finally, offset is calculated by getting relative offset from l in s' .

The reason why the calculation of offset is complex is due to a possible overflow. Let’s assume that the size of address space 1 is 4, i.e., there are 16 bytes, and the size of address space 2 is 5, i.e., there are 32 bytes. Also, let’s assume that $\text{convert}(1, x, 2) = x$ (identity function). Finally, let’s assume that the beginning of a block l is 8 in both address spaces (it must be the same because convert is the identity function). Then, pointer $p = \text{Log}(l, 15, 1)$ will have integer address $(8 + 15) \% 16 = 7$, but **addrspacecast** $p\ \text{to}\ 2 = \text{Log}(l, 15, 2)$ will return $(8 + 15) \% 32 = 23$. This breaks our property that any pointer p can be replaced with **inttoptr**(**ptrtoint**(p)).

²Currently **psub** is implemented as an intrinsic function, `@llvm.psub`. For simplicity, it is represented as an instruction in this document.

³SPEC CPU2017 has up to 2% speedup.

⁴This enables replacing p with `NULL` if $p == \text{NULL}$ is given. Also we can make optimizers like GVN insert this if needed, although we didn’t utilize this replacement in our prototypes.

$$\begin{array}{c}
441 \quad (\iota = \text{"r = ptrtoint ty * addrspace(s) op to isz"}) \quad (\iota = \text{"r = ptrtoint ty * addrspace(s) op to isz"}) \quad (\iota = \text{"r = ptrtoint ty * addrspace(s) op to isz"}) \\
442 \quad \frac{\text{Log}(l, o, s) = \llbracket op \rrbracket_R \quad \text{cast2int}_M(l, o, s) = j}{R, M \xrightarrow{\iota} R[r \mapsto j\%2^{sz}], M} \quad \frac{\text{Phy}(o, s, I, cid) = \llbracket op \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto o], M} \quad \frac{\text{poison} = \llbracket op \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto \text{poison}], M} \\
443 \\
444 \\
445 \quad (\iota = \text{"r = inttoptr isz op to ty * addrspace(s)}) \quad (\iota = \text{"r = inttoptr isz op to ty * addrspace(s)}) \\
446 \quad \frac{\text{poison} \neq \llbracket op \rrbracket_R \quad \text{cast2ptr}(\llbracket op \rrbracket_R, s) = p}{R, M \xrightarrow{\iota} R[r \mapsto p], M} \quad \frac{\text{poison} = \llbracket op \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto \text{poison}], M} \\
447 \\
448 \\
449 \\
450
\end{array}$$

Figure 3. Semantics of ptrtoint, inttoptr

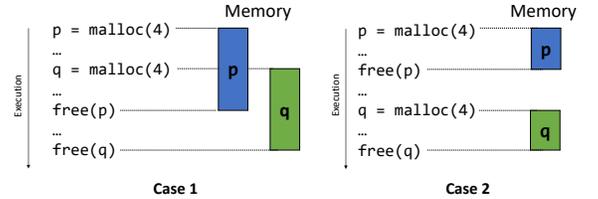
$$\begin{array}{c}
451 \quad (\iota = \text{"r = addrspacecast ty}_1 * \text{addrspace}(s_1) \text{op to ty}_2 * \text{addrspace}(s_2)") \\
452 \quad \frac{\text{Phy}(o, s_1, I, cid) = \llbracket op \rrbracket_R \quad o' = \text{convert}(s_1, o, s_2) \quad I' = \{ \text{convert}(s_1, i, s_2) \mid i \in I \}}{R, M \xrightarrow{\iota} R[r \mapsto \text{Phy}(o', s_2, I', cid)], M} \\
453 \\
454 \\
455 \\
456 \quad (\iota = \text{"r = addrspacecast ty}_1 * \text{addrspace}(s_1) \text{op to ty}_2 * \text{addrspace}(s_2)") \\
457 \quad \frac{\text{Log}(l, o, s_1) = \llbracket op \rrbracket_R \quad o' = (\text{convert}(s_1, \text{cast2int}_M(l, o, s_1), s_2) - \text{cast2int}_M(l, 0, s_2))\%2^{\text{ptrsz}(s_2)}}{R, M \xrightarrow{\iota} R[r \mapsto \text{Log}(l, o', s_2)], M} \\
458 \\
459 \\
460 \\
461 \\
462
\end{array}$$

Figure 4. Semantics of addrspacecast

Pointer Comparison. Now we define the semantics of pointer comparison. `icmp` compares two pointers in the same address space, and returns a value of type `i1`. In this semantics, `icmp` can freely move across any other operations like `call malloc()`, `free()`. Also, `icmp` on two logical pointers does not capture their integer addresses. In address space 0, `icmp NULL, p` also does not escape `p` if `p` is inbounds address, because integer address of `p` is always positive if it is inbounds.⁵

The definition of `icmp eq p1 p2` is as follows.

1. If `p1, p2` are both logical addresses $\text{Log}(l_1, o_1, s)$ and $\text{Log}(l_2, o_2, s)$, we first check whether their block ids are same, e.g., $l_1 = l_2$. If they are same, the comparison is equivalent to $o_1 = o_2$. If $l_1 \neq l_2$, the comparison can evaluate to `false`. However, there are also some cases where comparison is non-deterministic, i.e., it can evaluate to either `false` or `true`. This is the case if either one of the offsets is not strictly inbounds, i.e., $\neg(0 \leq o_1 < n_1) \vee \neg(0 \leq o_2 < n_2)$, or if the lifetimes of the two blocks do not overlap. In other words, the result is only *guaranteed* to be `false` if both offsets are strictly inbounds and the lifetimes overlap. These are sufficient conditions to ensure that the bit representations of the two pointers on the hardware differ.



This figure visualizes two cases where (1) lifetimes overlap, and (2) lifetimes do not overlap. If lifetimes overlap, the two blocks never have overlapping memory addresses. Therefore comparison on pointers from each of these blocks yields `false` if the offsets are strictly inbounds. Notably, the result of the comparison does not depend on whether `p` or `q` has already been freed. However, if their lifetimes are disjoint, `p` and `q` may overlap their addresses, and comparison on two pointers is nondeterministic value.⁶ Note that whether the two blocks overlap or not is determined when the second `malloc` is called. The result of the comparison does not depend on whether a block is still allocated, so `icmp eq` is allowed to freely move across `free`.

2. If `p1, p2` are both physical addresses $\text{Phy}(o_1, s, I_1, cid_1)$, $\text{Phy}(o_2, s, I_2, cid_2)$: the result is equivalent to $o_1 = o_2$.
3. If `p1 = Phy`(o_1, s, I_1, cid_1) and `p2 = Log`(l_2, o_2, s) or vice versa, the result is equivalent to $o_1 = \text{cast2int}_M(l_2, o_2, s)$.

The rules for comparing logical pointers allow `'p+n == q'` to be folded into `'false'`, which is an optimization currently

⁵Note that LLVM is rather conservative, so it assumes that `p == NULL` does not capture only if `p` is some system memory allocating function.

⁶This is the case of <http://lists.llvm.org/pipermail/llvm-dev/2017-April/112009.html>.

performed by gcc/llvm. Figure 5 shows the formal rules for ‘icmp eq’.

icmp ne $p1, p2$ is simply defined as a negation of **icmp eq** $p1, p2$. This enables free conversion between $p = q$ and $!(p != q)$.

icmp ule $p1, p2$ (or $p <= q$) is defined as follows.

1. If $p1 = \text{Log}(l, o_1, s)$ and $p2 = \text{Log}(l, o_2, s)$, we check whether the offsets o_1 and o_2 are inbounds. This condition is needed because $l + o_1$ or $l + o_2$ can overflow at runtime. If this is the case, the result is $o_1 \leq_u o_2$. Otherwise, we allow non-deterministic choice.
2. If $p1 = \text{Log}(l_1, o_1, s), p2 = \text{Log}(l_2, o_2, s)$ and $l_1 \neq l_2$, the result is nondeterministic choice.
3. If $p1 = \text{Phy}(o_1, s, I_1, cid_1)$ and $p2 = \text{Phy}(o_2, s, I_2, cid_2)$, the result is $o_1 \leq_u o_2$.
4. If $p1 = \text{Phy}(o_1, s, I_1, cid_1)$ and $p2 = \text{Log}(l_2, o_2, s)$, it is $o_1 \leq_u \text{cast2int}_M(l_2, o_2, s)$.

Figure 6 shows the rules for ‘icmp ule’. The semantics of ‘icmp ult’ is defined in a similar way to ‘icmp ule’.

For all equality/inequality comparisons, the result is **poison** if one or more operands are **poison**.

$$(t = \text{“}r = \text{icmp op ty} * \text{addrspace}(s) \text{ op}_1 \text{ op}_2\text{”})$$

$$\text{poison} = \llbracket \text{op}_1 \rrbracket_R \vee \text{poison} = \llbracket \text{op}_2 \rrbracket_R$$

$$R, M \xrightarrow{t} R[r \mapsto \text{poison}], M$$

Memory Allocation / Deallocation. Memory allocating operations create a new memory block and pick integer base addresses $P(s)$ for each address space they operate in, including some number of twin blocks. Our semantics is parameterized in how many twin blocks are allocated. In address space 0, the base addresses $P(0)_i$ may not be 0 and the last strictly inbounds address $P(0)_i + n - 1$ may not be $2^{\text{ptrsz}(0)} - 1$, i.e., the first and last byte of address space 0 are not allocatable. In particular, this means that $P(0) + n$ cannot overflow. The standard operations **alloca**, **call malloc** only work in address space 0. In order to maintain the memory invariants, all the base addresses must be divisible by the alignment. Furthermore, $P(s)_i + n$ must not exceed the size of s . Finally, for all s , all the $[P(s)_i, P(s)_i + n)$ must be mutually disjoint and disjoint from all existing alive blocks’ ranges. Figure 7 shows semantics of **alloca** and **call malloc**.

alloca creates a new logical block of the size of ty . Every bit of value of a new block is initialized with **poison**. The tag of a new block is **stack** meaning that the block cannot be freed by **free**. It is freed when a function returns.

malloc creates a new logical block of len bytes, or returns NULL nondeterministically. If len is **poison**, it is UB. Similar to **alloca**, every bit is initialized with **poison**. The alignment of blocks created by **malloc** is determined by the ABI (hence platform dependent), and it corresponds to the maximum alignment required for any type. Note that for aggregate

types like struct type / array type only a single block is allocated and the block contains all members of the aggregated type. **malloc(0)** returns NULL.

The pointers returned by **alloca** and **malloc** all have address space 0.

Figure 8 shows semantics of **free()**. **free** invalidates the memory block that ptr refers to by updating its time range. Calling **free** on NULL pointer is a NOP. Otherwise, calling **free** on a pointer p requires $\text{deref}_m(p, 1, l, 0)$ for some block l ; otherwise, it is UB. Notice that the offset must be 0. Moreover, deref only ever holds for pointers of address space 0.

These three operations **alloca**, **malloc**, **free** all increment the time of the memory τ_{cur} .

Address Calculation. The **getelementptr** instruction is used to get the address of a subelement of an aggregate data structure. **getelementptr** does not check whether the block is alive or not. This enables **getelementptr** to freely move across **free** calls.

getelementptr on a logical pointer yields a logical pointer with shifted offset. If the operand is $p = \text{Log}(l, o, s)$, **getelementptr** p, i returns $\text{Log}(l, (o + i')\%2^{\text{ptrsz}(s)}, s)$ where i' is i multiplied by the size of its element type. The **getelementptr** instruction may have the **inbounds** tag, which imposes further requirements on the operands and helps LLVM do further alias analysis. Concretely, it demands that both the base pointer and the returned pointer are inbounds of the block. **getelementptr inbounds** p, i returns **poison** if that is not the case.

getelementptr on a physical pointer yields a physical pointer with shifted offset. If the base pointer is $p = \text{Phy}(o, s, I, cid)$, then **getelementptr** p, i simply returns $\text{Phy}((o + i')\%2^{\text{ptrsz}(s)}, s, I, cid)$ where i' is i multiplied by size of its element type. This operation does not affect I and cid , which enables optimizing **getelementptr** $p, 0$ to p . In the **inbounds** variant, the returned pointer has an updated inbounds set $I' = I \cup \{(o + i')\%2^{\text{ptrsz}(s)}\}$. This allows for further alias analysis even on physical pointers. Also, tracking inbounds addresses and checking them later instead of returning **poison** instantly allows re-ordering of **getelementptr inbounds** and memory allocating/deallocating operations. **getelementptr inbounds** on a physical pointer is **poison** if the added offset overflows.

The formal semantics of **getelementptr** is given in Figure 9.

If the base pointer points to a nested aggregate value, the **getelementptr** instruction may have multiple indexes as its operands. In this case, it is allowed for **getelementptr inbounds** to point past the range of a subtype. For example, ‘`int a[5][5]; int* t=&a[0][7];`’ is translated into

```
%a = alloca [5 x [5 x i32]], align 16
%t = getelementptr inbounds [5 x [5 x i32]]* %a,
    i64 0, i64 0, i64 7
```

$$\begin{array}{c}
661 \\
662 \\
663 \\
664 \\
665 \\
666 \\
667 \\
668 \\
669 \\
670 \\
671 \\
672 \\
673 \\
674 \\
675 \\
676 \\
677 \\
678 \\
679 \\
680 \\
681 \\
682 \\
683 \\
684 \\
685 \\
686 \\
687 \\
688 \\
689 \\
690 \\
691 \\
692 \\
693 \\
694 \\
695 \\
696 \\
697 \\
698 \\
699 \\
700 \\
701 \\
702 \\
703 \\
704 \\
705 \\
706 \\
707 \\
708 \\
709 \\
710 \\
711 \\
712 \\
713 \\
714 \\
715
\end{array}$$

$ \begin{array}{c} (\iota = \text{"r = icmp eq ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \text{ICMP-PTR-LOGICAL-SAME-BLOCK} \\ \frac{\text{Log}(l, o_1, s) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Log}(l, o_2, s) = \llbracket \text{op}_2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto (o_1 = o_2)], M} \end{array} $	$ \begin{array}{c} (\iota = \text{"r = icmp eq ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \text{ICMP-PTR-LOGICAL-DIFFERENT-BLOCK} \\ \frac{\text{Log}(l_1, o_1, s) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Log}(l_2, o_2, s) = \llbracket \text{op}_2 \rrbracket_R \quad l_1 \neq l_2}{R, M \xrightarrow{\iota} R[r \mapsto \text{false}], M} \end{array} $	716 717 718 719 720 721 722 723 724 725 726 727 728
$ \begin{array}{c} (\iota = \text{"r = icmp eq ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \text{ICMP-PTR-LOGICAL-NONDET-TRUE} \\ \frac{\text{Log}(l_1, o_1, s) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Log}(l_2, o_2, s) = \llbracket \text{op}_2 \rrbracket_R \quad l_1 \neq l_2 \quad \neg(0 \leq o_1 < n_1) \vee \neg(0 \leq o_2 < n_2) \vee [b_1, e_1] \cap [b_2, e_2] = \emptyset}{R, M \xrightarrow{\iota} R[r \mapsto \text{true}], M} \end{array} $	$ \begin{array}{c} (\iota = \text{"r = icmp eq ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \text{ICMP-PTR-PHYSICAL} \\ \frac{\text{Phy}(o_1, s, I_1, cid_1) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Phy}(o_2, s, I_2, cid_2) = \llbracket \text{op}_2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto (o_1 = o_2)], M} \end{array} $	729 730 731 732 733 734 735 736 737 738 739 740 741 742 743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767

Figure 5. Semantics of icmp eq

$$\begin{array}{c}
682 \\
683 \\
684 \\
685 \\
686 \\
687 \\
688 \\
689 \\
690 \\
691 \\
692 \\
693 \\
694 \\
695 \\
696 \\
697 \\
698 \\
699 \\
700 \\
701 \\
702 \\
703 \\
704 \\
705 \\
706 \\
707 \\
708 \\
709 \\
710 \\
711 \\
712 \\
713 \\
714 \\
715
\end{array}$$

$ \begin{array}{c} (\iota = \text{"r = icmp ule ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \frac{\text{Log}(l, o_1, s) = \llbracket \text{op}_1 \rrbracket_R \quad \text{inbounds}_M(l, o_1) \quad \text{Log}(l, o_2, s) = \llbracket \text{op}_2 \rrbracket_R \quad \text{inbounds}_M(l, o_2)}{R, M \xrightarrow{\iota} R[r \mapsto (o_1 \leq_u o_2)], M} \end{array} $	$ \begin{array}{c} (\iota = \text{"r = icmp ule ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \frac{\text{Log}(l_1, o_1, s) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Log}(l_1, o_2, s) = \llbracket \text{op}_2 \rrbracket_R \quad l_1 \neq l_2 \quad b \in \{\text{true}, \text{false}\}}{R, M \xrightarrow{\iota} R[r \mapsto b], M} \end{array} $	736 737 738 739 740 741 742 743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770
$ \begin{array}{c} (\iota = \text{"r = icmp ule ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \frac{\text{Phy}(o_1, s, I_1, cid_1) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Phy}(o_2, s, I_2, cid_2) = \llbracket \text{op}_2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto (o_1 \leq_u o_2)], M} \end{array} $	$ \begin{array}{c} (\iota = \text{"r = icmp ule ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \frac{\text{Phy}(p, s, I, cid_1) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Log}(l, o, s) = \llbracket \text{op}_2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto (p \leq_u \text{cast2int}_M(l, o, s))], M} \end{array} $	743 744 745 746 747 748 749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770
$ \begin{array}{c} (\iota = \text{"r = icmp ule ty * addrspac(s) op}_1 \text{ op}_2\text{"}) \\ \frac{\text{Log}(l, o, s) = \llbracket \text{op}_1 \rrbracket_R \quad \text{Phy}(p, s, I, cid_2) = \llbracket \text{op}_2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto (\text{cast2int}_M(l, o, s) \leq_u p)], M} \end{array} $		749 750 751 752 753 754 755 756 757 758 759 760 761 762 763 764 765 766 767 768 769 770

Figure 6. Semantics of icmp ult

$$\begin{array}{c}
701 \\
702 \\
703 \\
704 \\
705 \\
706 \\
707 \\
708 \\
709 \\
710 \\
711 \\
712 \\
713 \\
714 \\
715
\end{array}$$

$ \begin{array}{c} (\iota = \text{"r = alloca ty, align a"}) \\ \frac{n = \text{bytewidth}(ty) \quad u = i(8 \times n) \ll (\text{poison}) \quad l \text{ fresh} \quad P \text{ unallocated physical addresses}}{R, (\tau_{cur}, f, C) \xrightarrow{\iota} R[r \mapsto \text{Log}(l, 0, 0)], (\tau_{cur} + 1, f[l \mapsto (\text{stack}, (\tau_{cur}, \infty), n, a, u, P)], C)} \end{array} $	$ \begin{array}{c} (\iota = \text{"r = call i8* malloc(iptrsz(0) len)}) \\ \frac{n = \llbracket \text{len} \rrbracket_R \quad u = i(8 \times n) \ll (\text{poison}) \quad l \text{ fresh} \quad P \text{ unallocated physical addresses} \quad n > 0}{R, (\tau_{cur}, f, C) \xrightarrow{\iota} R[r \mapsto \text{Log}(l, 0, 0)], (\tau_{cur} + 1, f[l \mapsto (\text{heap}, (\tau_{cur}, \infty), n, a, u, P)], C)} \end{array} $	756 757 758 759 760 761 762 763 764 765 766 767 768 769 770
$ \begin{array}{c} (\iota = \text{"r = call i8* malloc(iptrsz(0) len)}) \\ \frac{-}{R, M \xrightarrow{\iota} R[r \mapsto \text{NULL}], M} \end{array} $		761 762 763 764 765 766 767 768 769 770

Figure 7. Semantics of alloca, call malloc()

In this case, %t is not poison.

Pointer Subtraction. We define a new instruction `psub(ptr1, ptr2)` that calculates $ptr1 - ptr2$. This operation

$$\begin{array}{c}
\text{(} \iota = \text{"call void free(i8* op)} \text{")} \\
\frac{\text{Log}(0, 0, 0) = \llbracket op \rrbracket_R}{R, M \xrightarrow{\iota} R, M} \\
\text{(} \iota = \text{"call void free(i8* op)} \text{")} \\
\frac{\text{Log}(l, 0, 0) = \llbracket op \rrbracket_R \quad M(l) = (\text{heap}, (b, \infty), n, a, c, P)}{R, (\tau_{cur}, f, C) \xrightarrow{\iota} R, (\tau_{cur} + 1, f[l \mapsto (\text{heap}, (b, \tau_{cur}), n, a, c, P)], C)} \\
\text{(} \iota = \text{"call void free(i8* op)} \text{")} \\
\frac{\text{Phy}(o, 0, I, cid) = \llbracket op \rrbracket_R \quad M(l) = (\text{heap}, (b, \infty), n, a, c, P) \quad \text{cast2int}(l, 0) = o \quad b < \text{calltime}(cid)}{R, (\tau_{cur}, f, C) \xrightarrow{\iota} R, (\tau_{cur} + 1, f[l \mapsto (\text{heap}, (b, \tau_{cur}), n, a, c, P)], C)}
\end{array}$$

Figure 8. Semantics of `free()` (cases not mentioned here all raise UB)

$$\begin{array}{c}
\text{(} \iota = \text{"r = getelementptr ty * addrspace(s) op}_1 \text{ isz op}_2 \text{")} \\
\text{GETELEMENTPTR-LOGICAL} \\
\frac{\text{Log}(l, o, s) = \llbracket op_1 \rrbracket_R \quad i = \llbracket op_2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto \text{Log}(l, (o + \text{bytewidth}(ty) * i)\%2^{\text{ptrsz}(s)}, s)], M} \\
\text{(} \iota = \text{"r = getelementptr ty * addrspace(s) op}_1 \text{ isz op}_2 \text{")} \\
\text{GETELEMENTPTR-PHYSICAL} \\
\frac{\text{Phy}(o, s, I, cid) = \llbracket op_1 \rrbracket_R \quad i = \llbracket op_2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto \text{Phy}((o + \text{bytewidth}(ty) * i)\%2^{\text{ptrsz}(s)}, s, I, cid)], M} \\
\text{(} \iota = \text{"r = getelementptr inbounds ty * addrspace(s) op}_1 \text{ isz op}_2 \text{")} \\
\text{GETELEMENTPTR-INBOUNDS-LOGICAL} \\
\frac{\text{Log}(l, o, s) = \llbracket op_1 \rrbracket_R \quad M(l) = (t, r, n, a, c, P) \quad \text{inbounds}_M(l, o) \quad \text{inbounds}_M(l, o + \text{bytewidth}(ty) * i)}{R, M \xrightarrow{\iota} R[r \mapsto (l, o + \text{bytewidth}(ty) * i, s)], M} \\
\text{(} \iota = \text{"r = getelementptr inbounds ty * addrspace(s) op}_1 \text{ isz op}_2 \text{")} \\
\text{GETELEMENTPTR-INBOUNDS-PHYSICAL} \\
\frac{\text{Phy}(o, s, I, cid) = \llbracket op_1 \rrbracket_R \quad i = \llbracket op_2 \rrbracket_R \quad o' = (o + \text{bytewidth}(ty) * i) \quad o + \text{bytewidth}(ty) * i < 2^{\text{ptrsz}(s)}}{R, M \xrightarrow{\iota} R[r \mapsto \text{Phy}(o', s, I \cup \{o, o'\}, cid)], M}
\end{array}$$

Figure 9. Semantics of `getelementptr` (cases not mentioned here are all `poison`)

does not read or write memory. `ptr1` and `ptr2` should have same address space by type checking. Alias analysis can treat `psub` specially so it does not consider the addresses of its operands escaped in some cases. Originally clang used `ptrtoint` and integer arithmetic to emit pointer subtraction; with `psub` we enable a more precise alias analysis.

1. If `ptr1, ptr2` are both logical addresses, they must point to the same logical block; otherwise the result is `poison`.
2. If `ptr1, ptr2` are both physical addresses, say `ptr1 = Phy(o1, s, I1, cid1), ptr2 = Phy(o2, s, I2, cid2)`, the result is `(o1 - o2)%2ptrsz(s)`.
3. If `ptr1` is a logical address and `ptr2` is a physical address or vice versa, say `ptr1 = Log(l1, o1, s), ptr2 =`

`Phy(o2, s, I2, cid2)`, the result is equivalent to `(cast2intM(l1, o1, s) - o2)%2ptrsz(s)`.

Figure 10 shows formal semantics of `psub`. The transformation `(p1 - p2) == 0 → p1 == p2` is valid, but its inverse is not. Similarly, `(p1 - p2) > 0 → p1 > p2` is valid, but its inverse is not. This instruction is implemented as `@llvm.psub` intrinsic function in our prototype.

Load and Store. As mentioned in the beginning of this section, we define two meta operations to support conversion between values of types and low-level bit representation.

$$\begin{array}{l}
ty \Downarrow \in \llbracket ty \rrbracket \rightarrow \text{Bit}^{\text{bitwidth}(ty)} \\
ty \Uparrow \in \text{Bit}^{\text{bitwidth}(ty)} \rightarrow \llbracket ty \rrbracket
\end{array}$$

For base types, `ty` transforms `poison` into the bitvector of all `poison` bits, and defined values into their standard

$$\begin{array}{c}
(\iota = \text{"}r = \mathbf{psub} \text{ } ty * \text{ addrspace}(s) \text{ } op1, \text{ } op2\text{"}) \quad (\iota = \text{"}r = \mathbf{psub} \text{ } ty * \text{ addrspace}(s) \text{ } op1, \text{ } op2\text{"}) \\
\frac{\text{Log}(l, o_1, s) = \llbracket op1 \rrbracket_R \quad \text{Log}(l, o_2, s) = \llbracket op2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto (o_1 - o_2)\%2^{\text{ptrsz}(s)}], M} \quad \frac{\text{Log}(l_1, o_1, s) = \llbracket op1 \rrbracket_R \quad \text{Log}(l_2, o_2, s) = \llbracket op2 \rrbracket_R \quad l_1 \neq l_2}{R, M \xrightarrow{\iota} R[r \mapsto \mathbf{poison}], M} \\
(\iota = \text{"}r = \mathbf{psub} \text{ } ty * \text{ addrspace}(s) \text{ } op1, \text{ } op2\text{"}) \\
\frac{\text{Phy}(o_1, s, I_1, cid_1) = \llbracket op1 \rrbracket_R \quad \text{Phy}(o_2, s, I_2, cid_2) = \llbracket op2 \rrbracket_R}{R, M \xrightarrow{\iota} R[r \mapsto (o_1 - o_2)\%2^{\text{ptrsz}(s)}], M} \\
(\iota = \text{"}r = \mathbf{psub} \text{ } ty * \text{ addrspace}(s) \text{ } op1, \text{ } op2\text{"}) \\
\frac{\text{Log}(l, o_1, s) = \llbracket op1 \rrbracket_R \quad \text{Phy}(o_2, s, I_2, cid_2) = \llbracket op2 \rrbracket_R \quad o' = \text{cast2int}_M(l, o_1, s)}{R, M \xrightarrow{\iota} R[r \mapsto (o' - o_2)\%2^{\text{ptrsz}(s)}], M}
\end{array}$$

Figure 10. Semantics of **psub**

$$\frac{p = \text{Log}(l, o, s) \quad \text{deref}_M(p, sz, l, o) \quad o\%a = 0 \quad M(l) = (t, r, n, a, c, P) \quad b = \text{subst}_{sz}^{8 \times n}(c, o)}{\text{Load}(M, p, sz, a) = b}$$

$$\frac{p = \text{Log}(l, o, s) \quad M(l) = (t, r, n, a, c, P) \quad \text{deref}_M(p, sz, l, o) \quad o\%a = 0 \quad c' = \text{overwrite}_{sz}^{8 \times n}(c, b, o) \quad M' = (\tau_{cur}, f[l \mapsto (t, r, n, a, c', P)], C)}{\text{Store}(M, p, b, a) = M'}$$

Figure 11. Semantics of two auxiliary functions, Load and Store (when p is logical pointer only). If p is a logical pointer and it does not match these two cases, it fails.

$$\text{isz}\uparrow(b) = \begin{cases} n & \text{such that } \forall_{0 \leq i < sz} b[i] = n.i \vee \\ & b[i] = (\text{Phy}(n, 0, \emptyset, \text{None}), i) \\ \mathbf{poison} & \text{if there's no such } n \end{cases}$$

Figure 13. Converting a bit vector to an integer

low-level representation. $\text{getbit } v \ i$ is a function that returns i th bit of a value v . For vector types, $ty\downarrow$ transforms values element-wise, where \uparrow denotes the bitvector concatenation.

$$\begin{array}{l}
\text{isz}\downarrow(v) \text{ or } ty * \text{ addrspace}(s)\downarrow(v) = \lambda i. \text{getbit } v \ i \\
\langle sz \times ty \rangle\downarrow(v) = ty\downarrow(v[0]) \uparrow \dots \\
\quad \uparrow ty\downarrow(v[sz - 1]) \\
\text{where } b = b_0 \uparrow \dots \uparrow b_{sz-1}
\end{array}$$

Figure 12. Converting a value to a bit vector

$\text{isz}\uparrow(b)$ transforms bitwise value b to an integer of type isz . It creates integer n from bits. Notation $n.i$ is used to represent i th bit of non-**poison** integer n . Type punning from pointer

to integer yields **poison** and this explains redundant load-store pair elimination.⁷ One exception is when the pointer is a physical pointer of address space 0. In this case, type punning yields the integer address of the physical pointer. If any bit of b is **poison**, the result of $\text{isz}\uparrow(b)$ is **poison**. Figure 13 shows the definition of $\text{isz}\uparrow(b)$.

$ty * \text{ addrspace}(s)\uparrow(b)$ transforms a bitvector b into a pointer of type $ty * \text{ addrspace}(s)$. If b is exactly all the bits of pointer p in the right order, it returns p . If all bits contain a non-**poison** integer, it reconstructs a physical pointer of address space 0 with the corresponding integer address. Otherwise, it returns **poison**. Combined with the definition of $\text{isz}\uparrow(b)$, this allows vectorization of loading heterogeneous aggregates containing both aligned pointers and integer.⁸ Figure 13 shows the definition of $\text{isz}\uparrow(b)$.

For vector types, $ty\uparrow$ transforms bitwise representations element-wise.

$$\langle sz \times ty \rangle\uparrow(b) = \langle ty\uparrow(b_0), \dots, ty\uparrow(b_{sz-1}) \rangle$$

⁷Redundant load-store pair eliminations means removing ' $v = \text{load } i64 \text{ ptr}; \text{ store } v, \text{ ptr}$ '. If reading a logical pointer as integer implicitly casts the pointer, removing this load-store pair is not allowed.

⁸ For example, vectorizing load and store of `struct T{float* a; uintptr_t b}` type as `<2 x i8*>` type is allowed.

$$ty * \text{addrspace}(s) \uparrow(b) = \begin{cases} p & \text{such that } \text{addrspace}(p) = s \wedge \forall_{0 \leq i < 2^{\text{ptrsz}(s)}} b[i] = (p, i) \\ \text{Phy}(n, 0, \emptyset, \text{None}) & \text{such that } s = 0 \wedge \text{isz} \uparrow(b) = n \\ \text{poison} & \text{if there's no such } p \text{ or } n \end{cases}$$

Figure 14. Converting a bit vector to a pointer

Figure 11 is the semantics of two auxiliary functions `Load` and `Store`. $\text{subtr}_{sz}^n(c, o)$ is a partial function $(\text{Bit}^n \times \mathbb{N}) \rightarrow \text{Bit}^{sz}$ which returns $(c[o \times 8], c[o \times 8 + 1], \dots, c[o \times 8 + sz - 1])$. $\text{overwrite}_{sz}^n(c, b, o)$ is a partial function $(\text{Bit}^n \times \text{Bit}^{sz} \times \mathbb{N}) \rightarrow \text{Bit}^n$ that overwrites bits b over c at byte offset o . If a dereferenceable physical pointer $p = \text{Phy}(o, s, I, cid)$ is given, $\text{Load}(M, p, sz, a)$ behaves exactly same as $\text{Load}(M, p', sz, a)$ where p' is a logical pointer $\text{Log}(l, o', s)$, which is uniquely determined for p as described in the early part of this section. If there's no such p' , $\text{Load}(M, p, sz, a)$ fails. $\text{Store}(M, p, b, a)$ on a dereferenceable physical pointer p behaves exactly same as $\text{Store}(M, p', b, a)$, where p' is a logical pointer which is uniquely determined for the physical pointer p . If there's no such p' , $\text{Store}(M, p, b, a)$ fails. $\text{Load}(M, p, sz, a)$ and $\text{Store}(M, p, b, a)$ fail if p is not dereferenceable, regardless of p 's type.

Now we define semantics of `load/store` operations. $\text{Load}(M, p, sz, a)$ returns the bits p points to if it successfully dereferences the pointer with given size sz and alignment a . `load` yields v if $\text{Load}(M, p, sz, a)$ successfully returns a value, or UB if $\text{Load}(M, p, sz, a)$ fails. The store operation $\text{Store}(M, p, b, a)$ successfully stores the bit representation b into the memory M and returns the updated memory if p is dereferenceable with the given alignment a . `store` is UB if $\text{Store}(M, p, b, a)$ fails, or updates memory to M' otherwise.

$$\begin{array}{l} (\iota = \text{"}r = \text{load } ty, ty * op, \text{align } a\text{"}) \\ \text{Load}(M, \llbracket op \rrbracket_R, \text{bitwidth}(ty), a) \text{ fails} \\ \hline R, M \xrightarrow{\iota} \text{UB} \\ \text{Load}(M, \llbracket op \rrbracket_R, \text{bitwidth}(ty)) = v \\ \hline R, M \xrightarrow{\iota} R[r \mapsto ty \uparrow(v)], M \\ (\iota = \text{"}store \text{ } ty \text{ } op_1, ty * op, \text{align } a\text{"}) \\ \text{Store}(M, \llbracket op \rrbracket_R, ty \downarrow(\llbracket op_1 \rrbracket_R), a) \text{ fails} \\ \hline R, M \xrightarrow{\iota} \text{UB} \\ \text{Store}(M, \llbracket op \rrbracket_R, ty \downarrow(\llbracket op_1 \rrbracket_R)) = M' \\ \hline R, M \xrightarrow{\iota} R, M' \end{array}$$

Figure 15. Semantics of load, store

Function Call. `call ty funcname` calls a function with arguments which are given as operands of the instruction. `call` creates a fresh call id cid , and adds (cid, τ_{cur}) where τ_{cur} is current time of memory M to the global call id map. If an argument x is given to a call, register x inside the call has

value $\text{updatecid}(x)$. $\text{updatecid}(x)$ is a function that updates every bit of x to have current cid if possible. ty is type of the argument.

$$\begin{aligned} \text{updatecid}(x) &= ty \uparrow(\text{map}(ty \downarrow(x), \text{updatebit})) \\ \text{updatebit}(b) &= (\text{Phy}(o, s, I, cid), i) \text{ if } b = (\text{Phy}(o, s, I, \text{None}), i) \\ \text{updatebit}(b) &= b \text{ otherwise} \end{aligned}$$

By recording call id in physical pointers, alias analysis can assume that physical pointer which is given as argument never aliases with memory blocks allocated inside the function. Any actual access violating this rule would be UB because of the cid checks performed by `load` and `store`.

Function Return. When a function call with call id cid returns, its entry in the global call ID map gets changed to `None`, indicating that the call has ended.

Non-memory Operations. For the remaining operations we follow earlier work [?]. All operations on `poison` unconditionally return `poison` except `phi` and `select`. The instruction `freeze(isz op)` non-deterministically chooses an arbitrary non-`poison` value of the given type if op is `poison`. Otherwise, it is a NOP. Branching on `poison` is immediate UB. `Select` yields `poison` iff the condition is `poison` or the selected value is `poison`.