

A Theory of Platform-Dependent Low-Level Software (Extended Version)

Marius Nita Dan Grossman Craig Chambers
{marius,djg,chambers}@cs.washington.edu
Department of Computer Science and Engineering
University of Washington

Abstract

The C language definition leaves the sizes and layouts of types partially unspecified. When a C program makes assumptions about type layout, its semantics is defined only on platforms (C compilers and the underlying hardware) on which those assumptions hold. Previous work on formalizing C-like languages has ignored this issue, either by assuming that programs do not make such assumptions or by assuming that all valid programs target only one platform. In the latter case, the platform’s choices are hard-wired in the language semantics.

In this paper, we present a practically-motivated model for a C-like language in which the memory layouts of types is left unspecified. The dynamic semantics is parameterized by a platform’s layout policy and makes manifest the consequence of platform-dependent (i.e., unspecified) steps. A type-and-effect system produces a *layout constraint*: a logic formula encoding layout conditions under which the program is memory-safe. We prove that if a program type-checks, it is memory-safe on all platforms satisfying its constraint.

Based on our theory, we have implemented a tool that discovers unportable layout assumptions in C programs. Our approach should generalize to other kinds of platform-dependent assumptions.

1 Introduction

In recent years, research has demonstrated many ways to improve the quality of low-level software (typically written in C) by using programming-language and program-analysis technology. Such work has detected safety violations (array-bounds errors, dangling-pointer dereferences, uninitialized memory, etc.), enforced temporal protocols, and provided new languages and compilers that support reliable systems programming. The results are an important and practical success for programming-language theory. However, there remains a crucial and complementary set of complications that this paper begins to address:

The memory-safety of a C program often depends on assumptions that hold for some but not all compilers and machines.

Examples of assumptions include how `struct` values are laid out in memory (including padding), the size of values, and alignment restrictions on memory accesses. To our knowledge, existing work on safe low-level code (see Section 6) either (1) checks or simply assumes full layout portability (e.g., that the input program is unaffected by structure padding) or (2) checks the input program assuming a particular C platform, making no guarantee for other platforms.

Requiring that code makes no platform-dependent assumptions (e.g., by enforcing poorly understood and informally specified [23] restrictions on C programs) is too strict because low-level code often has inherently non-portable parts. An impractical solution is to rewrite large legacy applications in fully portable languages or to use perfect libraries that abstract all platform dependencies. Such high-level approaches ignore legacy issues, can be a poor match for low-level code, and assume that language or library implementations are available for an ever-increasing number of computing platforms.

Conversely, allowing implicit platform-dependent assumptions can lead to pernicious defects that lie dormant until one uses a platform violating the assumptions. Whereas defects like dangling-pointer dereferences are largely independent of the language implementation, so testing or verification on the “old platform” can find many of them, defects like assuming two `struct` types have similar data layouts are not. The results can be severe. Conceptually simple tasks like porting an application from a 32-bit machine to a 64-bit machine become expensive and error-prone. Software tested on widely available platforms can break when run on novel hardware such as embedded systems. Widely used compilers cannot change data-representation strategies without breaking legacy code that implicitly relies on undocumented behavior. Section 2.2 discusses some specific real-world examples.

In practice, C programmers identify and isolate platform dependencies manually. They may retest on each platform or use ad hoc tool support, such as compiler flags and lint-like technology, relying on informal knowledge of how target platforms lay out data. In contrast, the work presented in this paper informs the development of a principled, fully automatic tool that discovers particular unportable assumptions in C code.

More generally, we develop a semantics for low-level software that has an explicit and separable notion of a platform. Without such a notion, formal models or language-based tools for C face the same dilemma as the C code they are designed to help: either they apply only to fully portable code or they assume C implementation details that do not hold on all platforms. With our approach, one can “plug in” a platform description into a generic framework for the operational semantics. We advocate our general approach for making any work on C semantics (or semantics for any language with some unspecified behavior) robust to platform dependencies.

Moreover, we use logic formulas to describe a program’s platform-dependent assumptions. Such descriptions give a formal definition to the idea of “semi-portability” — a piece of software or an analysis may be correct given some assumptions, i.e., portable to those platforms on which the assumptions hold. Software and analyses could use these formulas as documentation for their assumptions. In our work, we extract them automatically but conservatively from C code.

1.1 Overview of Our Approach

The key to our formal model is isolating the notion of *platform*, which we can think of as an oracle that answers queries about how types are laid out in memory. In our model, a platform has two roles: (1) as a parameter to the operational semantics, and (2) as something we can describe with a *layout constraint*. The key insight is this: Given a program P , we can state that P is memory-safe on all platforms satisfying a layout constraint S , which is statically extracted from P .

To see how platforms work as parameters to the operational semantics, suppose we have a pointer dereference $*e$. The number of bytes accessed depends on the size of the type of e , and this size is determined by the platform. Therefore, our operational semantics has the form $\Pi \vdash P \rightarrow P'$ where Π is a platform and P is a program state. That way, the dereference rule can use Π to guide the memory access (and become stuck if Π deems the access misaligned).

As for layout constraints, they are formulas in a first-order theory in which platforms are models. For example, the constraint “ $\text{access}(4, 8) \wedge \text{size}(\text{long}) = 8$ ” is modeled by any platform in which values of type `long` occupy 8 bytes and 8-byte loads of 4-byte aligned data are allowed.¹ Per convention, we write $\Pi \models S$ when platform Π models constraint S .

Now given a program P we can try to find a constraint S such that if $\Pi \models S$, then the abstract machine does not get stuck when running P given Π . For our operational semantics, that means it will not treat an integer as a pointer, read beyond the end of a `struct` value, perform an improperly aligned memory access, etc. Finding an S that describes exactly the set of platforms on which the program does not get stuck is trivially undecidable, so a sound approximation is warranted. In our theory, we take a very conservative approach: A type system for source programs produces S using no flow-sensitivity or alias information. Our tool uses points-to information, but the general setup remains the same.

¹This example constraint is slightly simplified; see Section 3.4.

The key metatheoretic result is showing that the S our system produces is indeed sound, i.e., its only models are platforms on which the program does not get stuck. For the proof, we define a second type system for program states. This second type system, which exists only to show safety, is parameterized by a Π like the dynamic semantics. Our type-safety argument then has two parts:

1. The second type system and operational semantics enjoy the conventional preservation and progress properties.
2. If the first type system produces S given P , then P type-checks in the second type system for any Π such that $\Pi \models S$.

1.2 Contributions and Caveats

To our knowledge, this work is the first to consider describing a *set of platforms* on which a low-level program can run safely. At a more detailed level, our development clarifies several points:

- We can define a sound type system for a language with partially unspecified type layouts. The soundness theorem is proved once and for all instead of once for each platform.
- Layout-portability questions are reduced to pointer-cast questions, namely, “when can a pointer to a τ_1 be treated as a pointer to a τ_2 ?” This question, clearly akin to subtyping, depends on the platform.
- Platform constraints can be described in a first-order theory and extracted statically from the program.
- There should be a notion of “sensible” platform, meaning platforms on which every cast-free program cannot get stuck.

For tractability, the formal model considers only a small expression language inspired by C. It has many relevant features, including structs, heap allocation, and taking the address of fields, but we omit some relevant features (e.g., bit-fields), and many irrelevant ones (e.g., functions and `goto`). We also make some simplifying assumptions in our definition of platform. In particular, we assume all pointers have the same size and that alignment restrictions depend on the size, but not the type, of data. We see no fundamental problems extending our model in these directions.

The formal model directly informs the implementation of an automatic tool we wrote to detect layout-portability problems in C programs. In addition to producing a constraint describing the platforms on which the program is portable, the tool checks the constraint against a set of platforms of interest and outputs informative warnings when the constraint is violated. No access to the platforms of interest is needed.

1.3 Outline

Section 2 presents examples of platform-dependent code and the constraints describing their layout assumptions. Section 3 presents our core formal model, including the definition of platforms, our first-order theory, the dynamic and static semantics of our language, and our soundness theorem. Section 4 describes important extensions to the base model. Section 5 describes our tool. The last two sections discuss related work and conclude.

2 Examples

Section 2.1 presents several tiny examples of C code to explain issues of layout portability and relevant platform constraints. Section 2.2 complements this “tutorial” with actual platforms, systems, and coping strategies related to these concepts.

2.1 Small Code Fragments

Example 1: Accessing Memory

`(*e).f`

A memory access such as `(*e).f` reads or writes s bytes at some alignment a . If e has type `struct T*` and the f field has type τ , then s is the *size* of τ and a is the greatest common divisor of the *alignment* of `struct T` and the *offset* of f .

Platforms choose sizes, alignments, and offsets such that cast-free programs do not fail. For example, if a machine prohibits 8-byte accesses on 4-byte alignments, a compiler might put pad bytes before f fields or break 8-byte accesses into two 4-byte accesses. In the latter case, the compiler “supports” 8-byte accesses on 4-byte alignments.

In this paper, we assume platforms include an *access* function of type $\text{Int} \rightarrow \text{Int} \rightarrow \text{Bool}$, as well as *size* and *alignment* functions that map types to integers. Our example `(*e).f` therefore induces the constraint $\text{access}(a, s)$ where a and s are defined above. However, this constraint assumes e actually evaluates to a pointer with alignment a and a τ at the right offset. The constraints for cast expressions must ensure this.

Example 2: Prefix

```
struct S1 {    struct D1 {
  int* f1;      int* g1;
  int* f2;      int* g2;
  int* f3;      };
};
```

A cast from `struct S1*` to `struct D1*` requires that `struct D1` has a less stringent alignment than `struct S1`, and for each field in `struct D1` there is a field of compatible type in `struct S1` at the same offset. In this case, the C standard requires every platform to meet these constraints (and for $g1$ and $f1$ to have offset 0 and $g2$ to have the same offset as $f2$). Our formal model captures this notion precisely.

Example 3: Flattening and Alignment

```
struct S2 {                                struct D2 {
  int* f1;                                  int* g1;
  struct {int* f2; double f3;} f4;          int* g2;
};                                           };
```

A cast from `struct S2*` to `struct D2*` has similar constraints as in Example 2, but this time the C standard provides no guarantee. In fact, some platforms put pad bytes before $f4$ because of alignment restrictions. Our system will generate constraints preventing such a representation mismatch if the program has this cast.

Example 4: Suffix

```
struct S3 {    struct D3 {
  int* f1;      int* g1;
  int* f2;      double g2;
  double f3;    };
};
struct S3* x = ...;
struct D3* y = (struct D3 *)(&(x->f2));
```

The cast in the initializer for y above is a situation where the source and destination types both point to an `int*` followed by a `double`. However, a platform with 4-byte pointers, 8-byte doubles, and 8-byte alignment of doubles cannot support this cast because `struct D3` has more padding. Platforms without padding can allow this cast, even though $\&x \rightarrow f2$ has type `int**` in C.

Example 5: Arrays

```
struct S4 { int f1; short f2; };
struct D4 { int g1; };

struct S4 * x = ...;
struct D4  y = ((struct D4 *)x)[7];
```

Previous examples implicitly assumed the destination pointer was not used as an array (i.e., it was used as a pointer to exactly one struct). On many platforms that add trailing padding after field `f2`, a cast from `struct S4*` to `struct D4*` is legal. However, the cast above is broken due to intermittent pad bytes in the layout pointed to by `x`. This issue is orthogonal to array-bounds violations; we must reject the cast even if `x` points to more than 7 elements. Section A.4 extends our model to distinguish pointers to single-objects from pointers to arrays.

Example 6: Deep Subtyping

```
struct S4 {
  struct {int* f1; int* f2;} *f3;
};
struct D4 {
  const struct {int* g1;} *g2;
};
```

A cast from `struct S4*` to `struct D4*` is safe only because `const` qualifies the type of `g2`. As expected, read-only access permits more casts (i.e., fewer implementation constraints) for the same reasons covariant subtyping is sound on read-only fields. Section 4.3 adds this orthogonal feature to our model.

Example 7: Skipping Pad Bytes

```
struct S6 {  struct D6 {
  int* f1;    int* g1;
  int* f2;    double g2;
  double f3; };
};
```

A cast from `struct S6*` to `struct D6*` might appear safe if pointers are 4 bytes, `struct S6` has no padding, and `struct D6` has 4 bytes of padding before `g2`. However, an assignment such as `*p=*q` where `p` and `q` have type `struct D6*` may overwrite the pad bytes (which thanks to casting actually need to hold a pointer). Section 4.4 extends our model to let implementations indicate on a per-pad basis how they implement assignment to `struct` values.

Example 8: Safe-But-Inequivalent Implementations

```
struct S7 {  struct D7 {
  long f1;    short g1;
};           short g2;
};
```

Assuming there is no padding, that `long` is twice the size of `short`, and that there are no misaligned accesses, a cast from `struct S7*` to `struct D7*` is safe. No misaligned memory access or treating an integer as pointer can result. However, endianness can cause different implementations to behave differently. We leave such notions of equivalence to future work, focusing here only on safety, which we believe will still prove incredibly useful in writing semi-portable code and debugging ported applications. In particular, by preventing reading beyond the end of a struct, we detect many no-padding assumptions.

$$\begin{aligned}
\tau &::= \text{short} \mid \text{long} \mid \tau* \mid N \\
t &::= N\{\overline{\tau} \overline{f}\} \\
e &::= s \mid l \mid x \mid e = e \mid e.f \mid (\tau*)e \mid *(\tau*)(e) \mid (\tau*)&e \rightarrow f \\
&\quad \mid \text{new } \tau \mid e; e \mid \text{if } e \ e \ e \mid \text{while } e \ e \mid \tau \ x; \ e
\end{aligned}$$

Figure 1: Source-Language Syntax: A program has the form $\bar{t}; e$

2.2 Practical Scenarios

The scope of the platform dependency problem is not precisely known because defects can lie dormant until one changes hardware or compiler. Therefore, like date bugs (such as the famous Y2K problem of last decade), offending code can be difficult to locate and fix.²

The LinuxARM project, a port of Linux to the ARM embedded processor, provides compelling evidence that defects are subtle and widespread. The ARM compiler gives all structs at least 4-byte alignment whereas the original Linux implementation (gcc and x86) uses less alignment for structs containing only `short` and `char` fields. To quote [5]:

At this point, several years of fixing alignment defects in Linux packages have reduced the problems in the most common packages. Packages known to have had alignment defects are: Linux kernel; binutils; cpio; RPM; Orbit (part of Gnome); X Windows. This list is *very* incomplete.³

They also note that defects sometimes lead to alignment traps, but sometimes lead to silent data corruption. Kernel developers are basically told to, “be careful” [26].

Ports to 64-bit platforms provide another evidence source. Some vendors do little more than suggest using lint-like technology, such as gcc’s `-Wpadded` flag for reporting when a struct type has padding [22]. However, others find that aggressive warning levels produce so much information for legacy code that they recommend using multiple independent compilers and looking only at lines for which they all produce warnings [27].

3 Core Language

This section develops a formal model that can explain examples 1–4 from Section 2. We define an appropriate core language, a definition of *platform*, a dynamic semantics, a first-order theory that constrains platforms, a type-and-effect system to produce constraints, and the type-soundness result, acquired via a lower-level, platform-dependent type system. Figure 8 on page 15 summarizes the model’s judgments.

3.1 Idealized Syntax

For source programs, we consider a small subset of C with some convenient syntactic changes, as defined in Figure 1. Most significantly, we omit functions and make all terms expressions. A program is a sequence of struct definitions (\bar{t}) and an expression (e) to be evaluated. (We consistently write \bar{x} for a sequence of elements from syntactic category x and \cdot for the empty sequence. We also write x^i for a length i sequence.) Type definitions have global scope, allowing mutually recursive types.

Types τ include short and long (for two sizes of data), pointers ($\tau*$), and struct types (N rather than the more verbose `struct N`). As in C, all pointers (levels of indirection) are explicit. A struct definition (t) names the type and gives a sequence of fields. For simplicity, we assume all field names in a program are disjoint. Several expression forms are identical to C, including short and long constants (s and l ; we leave their exact form unspecified), variables (x), assignments ($e = e$), field access ($e.f$), and pointer casts ($(\tau*)e$).

²Furthermore, semantics-based Y2K solutions [15] serve as inspiration for us, though portability bugs fortunately do not share a worldwide deadline.

³Emphasis in original.

i, a, o	$\in \mathbb{N}$
b	$::= 0 \mid 1 \mid \dots \mid 255$
w	$::= b \mid \text{uninit} \mid \ell+i$
e	$::= \dots \mid \bar{w}$
v	$::= \bar{w}$
α	$::= [a, o]$
H	$::= \cdot \mid H, \ell \mapsto v, \alpha$

Figure 2: Syntax extensions for run-time behavior

For pointer dereference $(*(\tau*)(e))$ and pointing to a field $((\tau*)&e \rightarrow f)$, it is a technical convenience to require a type annotation in the syntax. (Our particular choice happens to correspond to C’s syntax.) Dereference in C *is* type-directed (if e has type $\tau*$, then $*e$ reads `sizeof(τ)` bytes); our type decoration makes this explicit. The cast in address-of-field expressions helps us support “suffix casts” as in Example 4 of Section 2.

The remaining expression forms are for memory allocation or control flow. `new τ` heap-allocates uninitialized space to hold a τ ; it is less verbose than `malloc(sizeof(τ))`. $(e; e)$ is sequence. $(if\ e\ e\ e)$ is a conditional, branching on whether the first subexpression is 0. To avoid distinguishing statements from expressions, a while-loop evaluates to a number if it terminates. Finally, $(\tau\ x; e)$ creates a local variable x of type τ bound in e .

As defined in Section 3.3, program evaluation depends on a platform Π and modifies a heap H . Because \bar{t} does not change during evaluation, we write $\Pi; \bar{t} \vdash H; e \rightarrow H'; e'$ for one evaluation step. Rather than define a *translation* (i.e., a compiler) from e to a lower-level platform-dependent language, we *extend* e with new lower-level forms. This equivalent approach of consulting the platform lazily (at run-time) simplifies the metatheory while fully exposing the intricacies of platform dependencies.

Figure 2 defines the syntactic extensions for run-time expressions and heaps. A value v is a sequence of atomic values w , which can be initialized bytes b , uninitialized bytes `uninit`, or pointers $\ell+i$. A pointer is a label ℓ and an offset i because the heap maps labels to value sequences, so a pointer into the middle of a value has a non-zero offset. This heap model is higher level than assembly language but low enough for middle pointers, suffix casts, etc. In other words, it is ideal for modeling layout dependencies.

Heaps also map labels to alignments. We model alignments as pairs $[a, o]$ where o is an offset from alignment a . Typically o is 0, e.g., $[4, 0]$ describes 4-byte aligned pointers. Supporting offsets adds some precision, e.g., if we add 2 bytes to a $[4, 0]$ pointer we get $[4, 2]$ and adding 2 more bytes gives $[4, 4]$, which is isomorphic to $[4, 0]$. Without offsets, adding 2 bytes to a 4-byte aligned pointer would produce a 2-byte aligned pointer. Section 3.5 describes a subalignment relation precisely.

3.2 Platforms

Before we show the semantics of our language, we need to introduce a precise notion of platform, as platforms play central roles in both the dynamic and static semantics. A platform (Π) is a record of functions with the following components, summarized in Figure 3:

- A translation of types into a lower-level representation ($\bar{\sigma}$), described below. We write $\Pi.xtype(\bar{t}, \tau)$ for the $\bar{\sigma}$ corresponding to the translation of a type τ assuming type definitions \bar{t} .
- An *alignment* function ($\Pi.align$) returns the alignment α used to allocate space for a τ .
- An *offset* function ($\Pi.offset$) takes a field f and returns the number of bytes from the beginning of the nearest enclosing struct to the field f .
- An *access* function takes an alignment α and a size i and returns true if accessing i bytes at an alignment α is not an error.

$$\sigma ::= \text{byte} \mid \text{pad}[i] \mid \text{ptr}_\alpha(\bar{\sigma}) \mid \text{ptr}_\alpha(N)$$

$$\begin{aligned} \Pi.xtype(\bar{t}, \tau) &= \bar{\sigma} \\ \Pi.align(\bar{t}, \tau) &= \alpha \\ \Pi.offset(\bar{t}, f) &= i \\ \Pi.access(\alpha, i) &= \{\text{true}, \text{false}\} \\ \Pi.ptrsize &= i \\ \Pi.xliteral(s) &= \bar{b} \\ \Pi.xliteral(l) &= \bar{b} \end{aligned}$$

Figure 3: Platforms and low-level types

- The size of pointers ($\Pi.ptrsize$) is a constant i .⁴
- An *xliteral* function translates integer literals into byte sequences. The layout of values in memory is platform-dependent.

The *access* function is typically associated with hardware and the other components with compilers, but a platform comprises all components. It is clear a “sensible” platform cannot define its components in isolation (e.g., the type translation must mind the access function); our constraint language will let us define these restrictions precisely.

Low-level types (the target of $\Pi.xtype$) are $\bar{\sigma}$, a sequence of σ . For example, if long is four bytes, the translation is byte^4 . The type $\text{pad}[i]$ represents i bytes of padding. The type $\text{ptr}_\alpha(\bar{\sigma})$ describes pointers to data described by $\bar{\sigma}$ at alignment α . As a technical point, we disallow the type N for low-level types except for the form $\text{ptr}_\alpha(N)$. This restriction simplifies type equalities without restricting platforms or disallowing recursive types.

3.3 Dynamic Semantics

The dynamic semantics is a small-step rewrite system for expressions, parameterized by a platform and a sequence of type declarations. Figure 4 holds the full definition for $\Pi; \bar{t} \vdash H; e \rightarrow H'; e'$. It is defined via evaluation contexts for conciseness. As in C, the left side of assignments (called left-expressions) are evaluated differently from other expressions (called right-expressions). Therefore, we have two sorts of contexts (L and R) defined by mutual induction and a different sort of primitive reduction ($\xrightarrow{\text{L}}$ and $\xrightarrow{\text{R}}$) for each sort of context [17]. In particular, $\text{R}[e]$ is a right-context R containing a right-hole filled by e and $\text{R}[e]_l$ is a right-context R containing a left-hole filled by e . Each context contains exactly one right-hole or exactly one left-hole, but not both.

Most primitive reductions depend on Π , but let us first dispense with those that do not. D-CAST shows that casts have no run-time effect. D-SEQ is typical. D-IF and D-IFT are typical except we treat 0 as false (as in C) and other byte-sequences as true. D-WHILE is a typical small-step unrolling; we make the arbitrary choice that a terminating loop produces some s literal nondeterministically.

D-NEW extends the heap with a new label holding uninitialized data. The platform determines the alignment and size of the new space, with the latter computed by applying the auxiliary size function to the translation of the allocated type. The resulting value $\ell+0$ is a pointer to the beginning of the space. The type system does not prevent getting stuck due to uninitialized data; this issue is orthogonal. D-LET has the same hypotheses as D-NEW. Because memory management is not our concern, we use heap allocation even for local variables. We substitute $*(\tau*)(\ell+0)$ for x in the resulting expression.

D-DEREF reads data from the heap and the resulting expression is the data. In particular, it extracts a sequence “from the middle” of $H(\ell)$. This sequence is from offset j (where the expression before the step is

⁴This is a slight simplification since a C implementation could use different sizes for different pointers.

$R ::= [\cdot]_r \mid L = e \mid *(\tau*)(\ell+i) = R \mid R.f \mid *(\tau*)(R) \mid (\tau*)R \mid R; e \mid (\tau*)&R \rightarrow f \mid \text{if } R \ e \ e$
 $L ::= [\cdot]_l \mid L.f \mid *(\tau*)(R)$
 $D ::= \Pi; \bar{t}$

	$\frac{D \vdash H; e \xrightarrow{\tau} H'; e'}{D \vdash H; R[e]_r \rightarrow H'; R[e']_r}$	$\frac{D \vdash H; e \xrightarrow{\tau} H'; e'}{D \vdash H; R[e]_l \rightarrow H'; R[e']_l}$
D-CAST	D-SEQ	D-WHILE
$\frac{}{D \vdash H; (\tau*)\bar{w} \xrightarrow{\tau} H; \bar{w}}$	$\frac{}{D \vdash H; (v; e) \xrightarrow{\tau} H; e}$	$\frac{}{D \vdash H; \text{while } e_1 \ e_2 \xrightarrow{\tau} H; \text{if } e_1 \ (e_2; \text{while } e_1 \ e_2) \ s}$
D-IF	D-IFT	
$\frac{}{D \vdash H; \text{if } 0^i \ e_1 \ e_2 \xrightarrow{\tau} H; e_2}$	$\frac{b_1 \dots b_i \neq 0^i}{D \vdash H; \text{if } (b_1 \dots b_i) \ e_1 \ e_2 \xrightarrow{\tau} H; e_1}$	
D-NEW	D-LET	
$\frac{\ell \notin \text{Dom}(H) \quad \Pi.\text{align}(\bar{t}, \tau) = \alpha \quad \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}}{\Pi; \bar{t} \vdash H; \text{new } \tau \xrightarrow{\tau} (H, \ell \mapsto \text{uninit}^i, \alpha); \ell+0}$	$\frac{\ell \notin \text{Dom}(H) \quad \Pi.\text{align}(\bar{t}, \tau) = \alpha \quad \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \quad \text{size}(\Pi, \bar{\sigma}) = i}{\Pi; \bar{t} \vdash H; \tau \ x; \ e \xrightarrow{\tau} (H, \ell \mapsto \text{uninit}^i, \alpha); e\{*(\tau*)(\ell+0)/x\}}$	
D-DEREF		
$\frac{H(\ell) = \bar{w}_1 \bar{w}_2 \bar{w}_3, [a, o] \quad \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.\text{access}([a, o + j], k) \quad \text{size}(\Pi, \bar{w}_1) = j \quad \text{size}(\Pi, \bar{w}_2) = \text{size}(\Pi, \bar{\sigma}) = k}{\Pi; \bar{t} \vdash H; *(\tau*)(\ell+j) \xrightarrow{\tau} H; \bar{w}_2}$		
D-ASSIGN		
$\frac{H(\ell) = \bar{w}_1 \bar{w}_2 \bar{w}_3, [a, o] \quad \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.\text{access}([a, o + j], k) \quad \text{size}(\Pi, \bar{w}_1) = j \quad \text{size}(\Pi, \bar{w}_2) = \text{size}(\Pi, \bar{\sigma}) = k \quad \text{size}(\Pi, \bar{w}) = k}{\Pi; \bar{t} \vdash H; (*(\tau*)(\ell+j)) = \bar{w} \xrightarrow{\tau} (H, \ell \mapsto \bar{w}_1 \bar{w}_2 \bar{w}_3, \alpha); \bar{w}}$		
D-FADDR	D-FETCHL	
$\frac{\Pi.\text{offset}(f) = j'}{\Pi; \bar{t} \vdash H; (\tau*)&(\ell+j) \rightarrow f \xrightarrow{\tau} H; \ell+(j+j')}$	$\frac{\Pi.\text{offset}(f) = j' \quad N\{\dots \tau_2 \ f \dots\} \in \bar{t}}{\Pi; \bar{t} \vdash H; (*(\tau_1*)(\ell+j)).f \xrightarrow{\tau_1} H; *(\tau_2*)(\ell+(j+j'))}$	
D-FETCH	D-SHORT	D-LONG
$\frac{N\{\dots \tau \ f \dots\} \in \bar{t} \quad \Pi.\text{offset}(\bar{t}, f) = \text{size}(\Pi, \bar{w}_1) \quad \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \quad \text{size}(\Pi, \bar{\sigma}) = \text{size}(\Pi, \bar{w}_2)}{\Pi; \bar{t} \vdash H; \bar{w}_1 \bar{w}_2 \bar{w}_3.f \xrightarrow{\tau} H; \bar{w}_2}$	$\frac{\Pi.\text{xliteral}(s) = \bar{b}}{\Pi; \bar{t} \vdash H; s \xrightarrow{\tau} H; \bar{b}}$	$\frac{\Pi.\text{xliteral}(l) = \bar{b}}{\Pi; \bar{t} \vdash H; l \xrightarrow{\tau} H; \bar{b}}$
$\text{size}(\Pi, \sigma) = \begin{cases} 1 & \text{if } \sigma = \text{byte} \\ i & \text{if } \sigma = \text{pad}[i] \\ \Pi.\text{ptrsize} & \text{if } \sigma \in \{\text{ptr}_\alpha(N), \text{ptr}_\alpha(\bar{\sigma})\} \end{cases}$	$\text{size}(\Pi, w) = \begin{cases} 1 & \text{if } w \in \{b, \text{uninit}\} \\ \Pi.\text{ptrsize} & \text{if } w = \ell+i \end{cases}$	
$\text{size}(\Pi, \sigma_1 \dots \sigma_n) = \sum_{i=1}^n \text{size}(\Pi, \sigma_i)$	$\text{size}(\Pi, w_1 \dots w_n) = \sum_{i=1}^n \text{size}(\Pi, w_i)$	

Figure 4: Dynamic Semantics

$*(\tau*)(\ell + j)$) to $j + k$ (where k is the size of the translation of τ). If it is not possible to “carve up” $H(\ell)$ in this way, then the rule does not apply and the machine is stuck. As expected, we also use $\Pi.access$ to model alignment constraints on the memory access.

D-ASSIGN has exactly the same hypotheses as D-DEREF plus the requirement that the right-hand side be a value equal in size to the value being replaced in the heap. The resulting heap differs only from offset j to offset $j + k$ of $H(\ell)$.

D-FADDR takes a pointer value and increases its offset by the offset of the field f , which is defined by Π . D-FETCHL, the one primitive reduction in left contexts, is similar, but we also have to change a type to reflect that $e.f$ refers to less memory than e . A “left-value” (i.e., a terminal left-expression) looks like $*(\tau*)(\ell+j)$.

D-FETCH uses the offset and size information from Π to project a subsequence of a value. We do not use the access function here because we are not accessing the heap.⁵

Finally, D-SHORT and D-LONG use the platform directly to translate literals to byte-sequences.

Several of the rules require computing the size of a value \bar{w} or a type $\bar{\sigma}$. Figure 4 includes these platform-dependent functions.

There are many ways to get stuck in the dynamic semantics, especially in the presence of arbitrary, unchecked pointer casts. To characterize memory-safe programs and the platforms on which they will not become stuck, we need a way to write down the platform-dependent layout assumptions made by a program and then a way to extract the assumptions from the program. The next two sections address these issues.

3.4 Constraint Language

To define a sound type system for our language, we need to limit what platforms we consider. That is, “ P does not get stuck” makes no sense, but “ P run on platform Π does not get stuck” does. We use first-order logic to give a syntactic representation to a set of platforms; a formula S represents the platforms that model it, i.e., the set $\{\Pi \mid \Pi \models S\}$.

The syntax for formulas S is a first-order theory with (1) arithmetic, (2) sorts for aspects of our language (including fields f , types τ , low-level types $\bar{\sigma}$, etc.), and (3) function symbols relevant to platform-dependencies. Figure 5 defines these function symbols and their interpretations. These interpretations induce the full definition of $\Pi \models S$ as usual (e.g., $\Pi \models S_1 \wedge S_2$ iff $\Pi \models S_1$ and $\Pi \models S_2$).

Consider two example formulas:

- $\forall \tau, \bar{t}. \text{access}(\text{align}(\bar{t}, \tau), \text{size}(\text{xtype}(\bar{t}, \tau)))$
- Let \bar{t}_0 abbreviate:
 - $N_1\{\text{short } f_1 \text{ short } f_2 \text{ short } f_3\}$
 - $N_2\{\text{short } g_1 \text{ short } g_2\}$
 in the formula:
 - $\text{subtype}(\bar{t}_0, \text{xtype}(\bar{t}_0, N_1*), \text{xtype}(\bar{t}_0, N_2*))$.

The first formula says every type must have a size and alignment that allows memory to be accessed. Without this constraint, a program like $(\tau \ x; x = e)$ could get stuck because D-LET uses the alignment $\Pi.align(\bar{t}, \tau)$ for the space allocated for x . The second formula requires a low-level subtyping relationship between two pointer types. This is the constraint our static semantics generates for a cast like in Example 2 from Section 2.

These examples demonstrate the two flavors of formulas that arise in practice. First, there are constraints that every “sensible” platform would satisfy. We are not interested in other platforms, but stating these requirements as a constraint is much simpler than revisiting our definition of platforms. Second, there are constraints that we do not expect every platform to satisfy. Our static semantics produces a formula for these extra assumptions that a particular program makes.

⁵On actual machines, large values do not fit in registers so alignment remains a concern. We could model this by treating field access as an address-of-field computation followed by a dereference. However, the computation that produced the v in $v.f$ must have done a properly aligned memory access, so if v has the right type, then the more complicated treatment of field-access also would not have failed for any sensible platform.

<u>syntax</u>	<u>interpretation under Π</u>	<u>defined in</u>
$xtype(\bar{t}, \tau)$	$\Pi.xtype(\bar{t}, \tau)$	Figure 3
$align(\bar{t}, \tau)$	$\Pi.align(\bar{t}, \tau)$	
$offset(\bar{t}, f)$	$\Pi.offset(\bar{t}, f)$	
$access(\alpha, i)$	$\Pi.access(\alpha, i)$	
$xliteral(s)$	$\Pi.xliteral(s)$	
$xliteral(l)$	$\Pi.xliteral(l)$	
$size(\bar{\sigma})$	$size(\Pi, \bar{\sigma})$	Figure 4
$size(\bar{w})$	$size(\Pi, \bar{w})$	
$subtype(\bar{t}, \bar{\sigma}_1, \bar{\sigma}_2)$	$\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$	Figure 6
$subalign(\alpha_1, \alpha_2)$	$\vdash \alpha_1 \leq \alpha_2$	

Figure 5: Function Symbols for the First-Order Theory

The sensibility clauses we assume for type safety are straightforward to enumerate. We collect them in a constraint, called $S_{sensible}$, defined as the conjunction of the following formulas:

1. Size and alignment allows access of all types:
 $\forall \tau, \bar{t}. access(align(\bar{t}, \tau), size(xtype(\bar{t}, \tau)))$
2. Translation of literals respects the translation of their types:
 $\forall s, l, \bar{t}. size(xliteral(s)) = size(xtype(\bar{t}, short))$
 $\wedge size(xliteral(l)) = size(xtype(\bar{t}, long))$
3. Greater alignment does not restrict access:
 $\forall \alpha_1, \alpha_2, i. (access(\alpha_1, i) \wedge subalign(\alpha_1, \alpha_2)) \Rightarrow access(\alpha_2, i)$
4. Translation of τ^* respects the alignment and translation of τ :
 $\forall \tau, \bar{t}. subtype(\bar{t}, ptr_{align(\bar{t}, \tau)}(xtype(\bar{t}, \tau)), xtype(\bar{t}, \tau^*))$
5. Struct translation respects the offset and alignment of each field:
 $\forall \bar{t}, \tau, f, \bar{\sigma}, N.$
 $(N\{\dots \tau f \dots\} \in \bar{t} \wedge (xtype(\bar{t}, \tau) = \bar{\sigma}) \Rightarrow$
 $(\exists \bar{\sigma}_1, \bar{\sigma}_2, \alpha, o, i.$
 $xtype(\bar{t}, N) = \bar{\sigma}_1 \bar{\sigma} \bar{\sigma}_2 \wedge size(\bar{\sigma}_1) = offset(\bar{t}, f) = o$
 $\wedge align(\bar{t}, N) = [a, o] \wedge subalign([a, o + o'], align(\bar{t}, \tau))))$

This gives rise to a precise notion of sensible platform:

Definition 1 *A platform Π is sensible if $\Pi \models S_{sensible}$.*

The sensibility constraints are necessary for portable code in the sense that without them, cast-free programs could get stuck. The C standard also allows other assumptions that we can write in our logic but that our safety theorem need not assume. Here are just two examples:

- The first field always has offset 0:
 $\forall f, \bar{t}, \tau, N. (N\{\tau f \dots\} \in \bar{t}) \Rightarrow offset(\bar{t}, f) = 0$
- Fields are in order and do not overlap:
 $\forall \tau_1, f_1, \tau_2, f_2, N. N\{\dots \tau_1 f_1 \dots \tau_2 f_2 \dots\} \in \bar{t} \Rightarrow$
 $(offset(\bar{t}, f_1) + size(xtype(\bar{t}, \tau_1)) \leq offset(\bar{t}, f_2))$

$\frac{\text{PTR} \quad \vdash \alpha_1 \leq \alpha_2}{D \vdash \text{ptr}_{\alpha_1}(\bar{\sigma}_1 \bar{\sigma}_2) \leq \text{ptr}_{\alpha_2}(\bar{\sigma}_1)}$	$\frac{\text{UNROLL} \quad \Pi.xtype(\bar{t}, N) = \bar{\sigma}}{\Pi; \bar{t} \vdash \text{ptr}_\alpha(N) \leq \text{ptr}_\alpha(\bar{\sigma})}$	$\frac{\text{ROLL} \quad \Pi.xtype(\bar{t}, N) = \bar{\sigma}}{\Pi; \bar{t} \vdash \text{ptr}_\alpha(\bar{\sigma}) \leq \text{ptr}_\alpha(N)}$	
$\frac{\text{PAD} \quad \text{size}(\Pi, \sigma) = i}{\Pi; \bar{t} \vdash \sigma \leq \text{pad}[i]}$	$\frac{\text{ADD}}{\Pi; \bar{t} \vdash \text{pad}[i]\text{pad}[j] \leq \text{pad}[i+j]}$	$\frac{\text{SEQ} \quad D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2 \quad D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_4}{D \vdash \bar{\sigma}_1 \bar{\sigma}_3 \leq \bar{\sigma}_2 \bar{\sigma}_4}$	$\frac{\text{REFL}}{D \vdash \bar{\sigma} \leq \bar{\sigma}}$
$\frac{\text{TRANS} \quad D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2 \quad D \vdash \bar{\sigma}_2 \leq \bar{\sigma}_3}{D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_3}$	$\frac{\text{ALIGN-BASE} \quad a_1 = a_2 \times i}{\vdash [a_1, o] \leq [a_2, o]}$	$\frac{\text{ALIGN-OFFSET} \quad o_1 \equiv o_2 \pmod a}{\vdash [a, o_1] \leq [a, o_2]}$	$\frac{\text{ALIGN-TRANS} \quad \vdash \alpha_1 \leq \alpha_2 \quad \vdash \alpha_2 \leq \alpha_3}{\vdash \alpha_1 \leq \alpha_3}$

Figure 6: Physical Subtyping and Alignment Subtyping

3.5 Static Semantics and Constraint Generation

With constraints in hand, we can now define a type-and-effect system where the main judgment $\bar{t}; \Gamma \vdash e : \tau; S$ gives a constraint S suitable for e . Because it is the pointer casts in e that give rise to the constraints, this type system needs an expressive notion of pointer subtyping. Therefore, we consider subtyping (Figure 6) before describing the typing judgments for expressions (Figure 7).

Subtyping: We use a subtyping relation on low-level types to formalize when data described by $\bar{\sigma}$ can also be described by $\bar{\sigma}'$, and hence when pointer casts are safe. This notion has been called *physical subtyping* because it relies on actual memory layouts. Because we take a byte-for-byte view of memory, our notion of physical subtyping is richer than those defined in prior work [9, 36, 31], which are at the level of ground C types instead of bytes. For example, our definition allows casting a `struct S{short x; short y}*` to a `struct D{long a; }*` on many platforms, whereas prior definitions forbid it. The rules for our judgment $\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$ appear in Figure 6.

As expected in a language with mutation, pointer types have invariant subtyping (rule `PTR`). However, we do allow forgetting fields under a pointer type as this corresponds to restricting access to a prefix of the data previously accessible. This encodes the core concept behind casts like Example 2 in Section 2. We also allow assuming a less restrictive alignment. For example, a 4-byte aligned value can be safely treated as if it were 2- or 1-byte aligned.

Although we allow sequence-shortening under pointer types, it is *not* correct to allow shortening as a subtyping rule because a supertype should have the same size as a subtype (we can prove our rules have this property by induction on a subtyping derivation). This fact may seem odd to readers not used to subtyping in a language with explicit pointers. It is why C correctly disallows casts between struct types (as opposed to pointers to structs).

Rules `UNROLL` and `ROLL` witness the equivalence between a struct name and its definition. Recall we restrict a type N to occur under pointers.

Rules `PAD` and `ADD` let us forget about the form of data (not under a pointer) without forgetting its size. Note that we disallow $\Pi; \bar{t} \vdash \text{pad}[i+j] \leq \text{pad}[i]\text{pad}[j]$ to prohibit accessing “part of a pointer”, which would cause the abstract machine to get stuck. Rule `SEQ` lifts subtyping to sequences.

As usual, subsumption is sound for right-expressions but unsound for left-expressions. The static semantics enforces this restriction by disallowing casts as left-expressions.

Static Semantics: The static semantics is shown in Figure 7. The judgments $\bar{t}; \Gamma \vdash e : \tau; S$ and $\bar{t}; \Gamma \vdash e : \tau; S$ (for right- and left-expressions respectively) produce types as usual, but also layout constraints S . This constraint is a conjunction of the layout assumptions the program is making. An alternative approach could parameterize the typing judgment by a platform instead of outputting a constraint, to essentially perform

$\frac{\text{S-SHORT}}{\bar{t}; \Gamma \vdash s : \text{short}; \text{true}}$	$\frac{\text{S-LONG}}{\bar{t}; \Gamma \vdash l : \text{long}; \text{true}}$	$\frac{\text{S-NEW}}{\bar{t}; \Gamma \vdash \text{new } \tau : \tau*; \text{true}}$	$\frac{\text{S-VAR}}{\Gamma(x) = \tau}}{\bar{t}; \Gamma \vdash x : \tau; \text{true}}$
$\frac{\text{S-ASSN}}{\bar{t}; \Gamma \vdash e_1 : \tau; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2}}{\bar{t}; \Gamma \vdash e_1 = e_2 : \tau; S_1 \wedge S_2}$	$\frac{\text{S-FETCH}}{\bar{t}; \Gamma \vdash e : N; S \quad N\{\dots \tau f \dots\} \in \bar{t}}}{\bar{t}; \Gamma \vdash e.f : \tau; S}$		
$\frac{\text{S-SEQ}}{\bar{t}; \Gamma \vdash e_1 : \tau'; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2}}{\bar{t}; \Gamma \vdash e_1; e_2 : \tau; S_1 \wedge S_2}$		$\frac{\text{S-DEREF}}{\bar{t}; \Gamma \vdash e : \tau*; S}}{\bar{t}; \Gamma \vdash *(\tau*)(e) : \tau; S}$	
$\frac{\text{S-CAST}}{\bar{t}; \Gamma \vdash e : \tau_1*; S_1}}{\bar{t}; \Gamma \vdash (\tau*)e : \tau*; S_1 \wedge \text{subtype}(\bar{t}, \text{xtype}(\bar{t}, \tau_1*), \text{xtype}(\bar{t}, \tau*))}$			
$\frac{\text{S-FADDR}}{\bar{t}; \Gamma \vdash e : N*; S_1 \quad N\{\dots \tau_1 f \dots\} \in \bar{t}}}{\bar{t}; \Gamma \vdash (\tau*)(\&e \rightarrow f) : \tau*; S_1 \wedge \exists \bar{\sigma}_1, \bar{\sigma}_2, a, o. \text{xtype}(\bar{t}, N*) = \text{ptr}_{[a, o]}(\bar{\sigma}_1 \bar{\sigma}_2) \wedge \text{offset}(f) = \text{size}(\bar{\sigma}_1) \wedge \text{subtype}(\text{ptr}_{[a, o + \text{offset}(f)]}(\bar{\sigma}_2), \text{xtype}(\bar{t}, \tau*))}$			
$\frac{\text{S-IF}}{\bar{t}; \Gamma \vdash e_1 : \text{long}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2 \quad \bar{t}; \Gamma \vdash e_3 : \tau; S_3}}{\bar{t}; \Gamma \vdash \text{if } e_1 \ e_2 \ e_3 : \tau; S_1 \wedge S_2 \wedge S_3}$			$\frac{\text{S-WHILE}}{\bar{t}; \Gamma \vdash e_1 : \text{long}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2}}{\bar{t}; \Gamma \vdash \text{while } e_1 \ e_2 : \text{short}; S_1 \wedge S_2}$
$\frac{\text{S-DECL}}{\bar{t}; \Gamma, x : \tau_1 \vdash e : \tau_2; S}}{\bar{t}; \Gamma \vdash \tau_1 x; e : \tau_2; S}$			
$\frac{\text{S-VARL}}{\Gamma(x) = \tau}}{\bar{t}; \Gamma \vdash x : \tau; \text{true}}$	$\frac{\text{S-DEREFL}}{\bar{t}; \Gamma \vdash e : \tau*; S}}{\bar{t}; \Gamma \vdash *(\tau*)(e) : \tau; S}$	$\frac{\text{S-FETCHL}}{N\{\dots \tau f \dots\} \in \bar{t} \quad \bar{t}; \Gamma \vdash e : N; S}}{\bar{t}; \Gamma \vdash e.f : \tau; S}$	

Figure 7: Static Semantics (letting $\Gamma ::= \cdot \mid \Gamma, x:\tau$)

platform-dependent type-checking. This approach can be recovered from ours: we can check the constraint against a platform.

The interesting rules are S-CAST and S-FADDR because the constraint S_{sensible} in Section 3.4 suffices to ensure other expression forms (such as dereferences and assignments) cannot fail due to a platform dependency. The constraints directly describe the implicit assumptions made in Examples 2, 3, and 4 in Section 2. The key insight here is that we can allow a pointer cast to assign an arbitrary type to an expression, but the cast will only be deemed legal on platforms that model the associated layout constraint. Recall $\text{subtype}(\bar{t}, \bar{\sigma}_1, \bar{\sigma}_2)$ is the logical formula corresponding to $\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$. The S-FADDR constraint is much more complicated because it must state that there is some sequence of fields starting at the offset of field f that can be viewed as a τ .

Absent from this formal type system is support for downcasts, which are obviously important in practice. To support safe downcasts, we would just need to invert the direction of the subtyping constraint generated by the cast and employ existing techniques [31, 24] to ensure that the casted value actually has the type dictated by the cast.

3.6 Metatheory and Low-Level Static Semantics

Safety: Ideally, our type-safety result would claim that running a well-typed program on a “sensible” platform that also models the program’s constraint would never lead to a stuck state. That is, given $\bar{t}; \cdot \vdash e : \tau; S$, Π is sensible, $\Pi \models S$, and $\Pi; \bar{t} \vdash \cdot; e \rightarrow^* H; e'$ (where \rightarrow^* is the reflexive, transitive closure of \rightarrow), either e' is a value or there exists H', e'' such that $\Pi; \bar{t} \vdash H; e' \rightarrow^* H'; e''$.

However, it is possible that the abstract machine can get stuck by accessing uninitialized data. Because preventing uninitialized accesses is not our focus, we relax our safety guarantee to admit that e' might also be *legally stuck*. An expression e is legally stuck if e is of the form $R[ls]$, or $R[ls]_i$, where

$$ls ::= \text{if } (\bar{w}_1 \text{ uninit } \bar{w}_2) e \mid * (\tau *) (\text{uninit}^i) \mid (\tau *) \& \text{uninit}^i \rightarrow f$$

Our memory-safety proof employs a low-level type system that captures the relevant invariants that evaluation preserves. The main judgment of this type system has the form $\Pi; \bar{t}; \Psi; \Gamma \vdash e : \bar{\sigma}$ where Ψ gives a type to the heap.⁶ This system has implicit subsumption, which is necessary for a step via D-CAST to preserve typing:

$$\frac{\Pi; \bar{t}; \Psi; \Gamma \vdash e : \bar{\sigma}_1 \quad \Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2}{\Pi; \bar{t}; \Psi; \Gamma \vdash e : \bar{\sigma}_2}$$

Like in the source-level type system, we also have a judgment for left-expressions $(\Pi; \bar{t}; \Psi; \Gamma \vdash e : \bar{\sigma}, \alpha)$. This judgment does not have a subsumption rule, but does produce an alignment α for the location to which e will evaluate.

Many of the low-level typing rules have hypotheses that refer directly to the platform. For example, the rule for type-checking dereferences is:

$$\frac{\Pi; \bar{t}; \Psi; \Gamma \vdash e : \text{ptr}_\alpha(\bar{\sigma}_1 \bar{\sigma}_2) \quad \Pi. \text{xtype}(\bar{t}, \tau) = \bar{\sigma}_1 \quad \Pi. \text{access}(\alpha, \text{size}(\Pi, \bar{\sigma}_1))}{\Pi; \bar{t}; \Psi; \Gamma \vdash * (\tau *) (e) : \bar{\sigma}_1}$$

See Appendix A for the complete system, which includes rules for run-time forms (such as \bar{w}) and heaps.

The dynamic semantics and low-level type system enjoy the usual type soundness property (modulo legally stuck states), proven with the aid of the usual progress and preservation lemmas.

Theorem 2 (Low-Level Type Soundness) *If $\Pi; \bar{t}; \cdot; \vdash e : \bar{\sigma}$ and $\Pi; \bar{t} \vdash \cdot; e \rightarrow^* H'; e'$, then $H'; e'$ is not stuck on Π .*

The connection between the static semantics and the low-level type system is concisely stated by this theorem:

Theorem 3 *If $\bar{t}; \Gamma \vdash e : \tau; S$, Π is sensible, $\Pi \models S$, and $\Pi. \text{xtype}(\bar{t}, \tau) = \bar{\sigma}$, then $\Pi; \bar{t}; \cdot; \vdash e : \bar{\sigma}$.*

The proof, by induction on the derivation of $\bar{t}; \Gamma \vdash e : \tau; S$, uses the definition of $\Pi \models S$ in many cases. For example, a source derivation ending in S-DEREF can produce a low-level derivation ending in the dereference rule above because sensible platforms model $\text{access}(\text{align}(\bar{t}, \tau), \text{size}(\text{xtype}(\bar{t}, \tau)))$. Indeed, the proof of Theorem 3 ensures our definition of S_{sensible} is sufficient.

Last but not least, we state the key theorem that a program will not get stuck on any platform on which its layout assumptions hold:

Theorem 4 (Layout Portability) *If $\bar{t}; \Gamma \vdash e : \tau; S$, Π is sensible, $\Pi \models S$, and $\Pi; \bar{t} \vdash \cdot; e \rightarrow^* H'; e'$, then $H'; e'$ is not stuck on Π .*

This is a corollary to Theorems 2 and 3.

⁶ $\Psi ::= \cdot \mid \Psi, \ell \mapsto \bar{\sigma}, \alpha$

Dynamic Semantics:

$\Pi; \bar{t} \vdash H; e \rightarrow H'; e'$	small step
$\Pi; \bar{t} \vdash H; e \xrightarrow{r} H'; e'$	primitive right-step
$\Pi; \bar{t} \vdash H; e \xrightarrow{l} H'; e'$	primitive left-step

High-Level Static Semantics:

$\Pi \models S$	platform Π models formula S
$\vdash \alpha_1 \leq \alpha_2$	subtyping on alignments
$\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$	platform-dependent subtyping
$\bar{t}; \Gamma \vdash_r e : \tau; S$	typing for right-expressions
$\bar{t}; \Gamma \vdash_l e : \tau; S$	typing for left-expressions

Low-Level Static Semantics:

$\Pi; \bar{t}; \Psi; \Gamma \vdash_r e : \bar{\sigma}$	typing for right-expressions
$\Pi; \bar{t}; \Psi; \Gamma \vdash_l e : \bar{\sigma}, \alpha$	typing for left-expressions

Figure 8: Summary of Judgments

Cast-Free Portability: The constraint produced by our type system is fairly expressive, ruling out only platforms for which some cast in the program would make no sense. To make this intuition precise, we prove that for the right definition of “cast-free”, a cast-free program does not get stuck on any sensible platform.

Definition 5 (*Cast-Free*) A program $\bar{t}; e$ is cast-free if:

- No expressions of the form $(\tau^*)e'$ occur in e .
- For every expression of the form $(\tau^*)&e' \rightarrow f$ in e , the type τ is the type of f . That is, $N\{\dots\tau f\dots\} \in \bar{t}$.

The second point allows taking the address of a field but requires the resulting type to be the type of the field (rather than allowing a platform-dependent suffix cast). The key theorem is as follows:

Theorem 6 If $\bar{t}; e$ is cast-free and $\bar{t}; \cdot \vdash_r e : \tau; S$, then $S_{sensible} \Rightarrow S$.

The intuition that only casts threaten layout portability is captured by the following theorem, a corollary to Theorems 4 and 6:

Theorem 7 (*Cast-Free Layout Portability*) If $\bar{t}; e$ is cast-free, $\bar{t}; \cdot \vdash_r e : \tau; S$, Π is sensible, and $\Pi; \bar{t} \vdash \cdot; e \rightarrow^* H'; e'$, then $H'; e'$ is not stuck on Π .

4 Extensions

This section sketches how the core model we have developed is flexible enough to be extended with some other relevant features of C and its platforms. We focus first on arrays because they are ubiquitous and require restricting our subtyping definition.

4.1 Arrays

As Example 5 demonstrates, a subtyping rule for pointers that drops a suffix of pointed-to fields (rule PTR in Figure 6) is unsound if the pointer may be used as a pointer to an array. Therefore, extending our model with arrays is important and requires some otherwise unnecessary restrictions. Figure 9 defines this extension formally.

Rather than conservatively assume all pointers may point to arrays, the types distinguish pointers to one object (τ^* as already defined) from pointers to arrays ($\tau^{*\omega}$; the ω just distinguishes it from τ^*). This

Syntax: $\tau ::= \dots \mid \tau^{*\omega}$ $e ::= \dots \mid \text{new } \tau[e] \mid \&((\tau^{*\omega})(e))[e]$ $R ::= \dots \mid \text{new } \tau[R] \mid \&((\tau^{*\omega})(R))[e] \mid \&((\tau^{*\omega})(\ell+i))[R]$ $\sigma ::= \dots \mid \text{ptr}_\alpha^\omega(\bar{\sigma}) \mid \text{ptr}_\alpha^\omega(N)$	Platforms: $\Pi.\text{val}(\bar{b}) = i$
--	--

Dynamic semantics:

$\frac{\text{D-NEWARR} \quad \begin{array}{l} \ell \notin \text{Dom}(H) \\ \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \\ \Pi.\text{val}(\bar{b}) = j \geq 0 \end{array} \quad \begin{array}{l} \Pi.\text{align}(\bar{t}, \tau) = \alpha \\ \text{size}(\Pi, \bar{\sigma}) = i \end{array}}{\Pi; \bar{t} \vdash H; \text{new } \tau[\bar{b}] \xrightarrow{c} H, \ell \mapsto \text{uninit}^{i \times j}, \alpha; \ell+0}$	$\frac{\text{D-ARRELT} \quad \begin{array}{l} \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \\ \text{size}(\Pi, \bar{\sigma}) = j \\ \Pi.\text{val}(\bar{b}) = k \end{array} \quad \begin{array}{l} H(\ell) = \bar{w}, \alpha \\ 0 \leq (i + j \times k) < \text{size}(\Pi, \bar{w}) \end{array}}{\Pi; \bar{t} \vdash H; \&((\tau^{*\omega})(\ell+i))[\bar{b}] \xrightarrow{c} H; \ell+(i + j \times k)}$
--	--

Sensibility constraint: size is a multiple of alignment

$$\forall \tau, \bar{t}. \exists i, a, o. \text{size}(\bar{t}, \text{xtype}(\bar{t}, \tau)) = i \times a \wedge \text{align}(\bar{t}, \tau) = [a, o]$$

Subtyping and static semantics:

$\frac{\text{ARR} \quad \bar{\sigma}_1 = \bar{\sigma}_2^i \quad \vdash \alpha_1 \leq \alpha_2}{\Pi; \bar{t} \vdash \text{ptr}_{\alpha_1}^\omega(\bar{\sigma}_1) \leq \text{ptr}_{\alpha_2}^\omega(\bar{\sigma}_2)}$	$\frac{\text{S-NEWARR} \quad \bar{t}; \Gamma \vdash e : \text{long}; S}{\bar{t}; \Gamma \vdash \text{new } \tau[e] : \tau^{*\omega}; S}$	$\frac{\text{S-ARRELT} \quad \bar{t}; \Gamma \vdash e_1 : \tau^{*\omega}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \text{long}; S_2}{\bar{t}; \Gamma \vdash \&((\tau^{*\omega})(e_1))[e_2] : \tau^*; S_1 \wedge S_2}$
---	--	---

Figure 9: Additions for Arrays

dichotomy is common in safe C-like languages [31, 24], can be approximated via static analysis over C code, and is necessary to identify what platform assumptions are due only to arrays. The low-level types (σ) make the same distinction.

We add two right-expression forms. First, $\text{new } \tau[e]$ creates a pointer to a heap-allocated array of length e . (Because e will evaluate to a byte-sequence \bar{b} , a platform must interpret \bar{b} as an integer; we use $\Pi.\text{val}$ for this conversion.) The dynamic rule D-NEWARR is exactly like D-NEW except it creates enough space at $H(\ell)$ for the array. Our type system does not prevent $\text{new } \tau[e]$ from being stuck if e has uninitialized bytes or e is negative.⁷

Second, $\&((\tau^{*\omega})(e_1))[e_2]$ is more easily read as $\&e_1[e_2]$; the size of τ guides the dynamic semantics like it does with pointer dereferences. This form produces a pointer to one array element, which can be dereferenced or assigned through. The dynamic rule D-ARRELT produces the pointer $\ell+(i + j \times k)$ where the array begins at $\ell+i$, elements have size j , and e_2 evaluates to k . However, the two hypotheses on the right perform a *run-time bounds check*; our type system does not prevent this check from failing and therefore the machine being stuck.⁸ The bounds-check on $\&((\tau^{*\omega})(e_1))[e_2]$ ensures an ensuing dereference can never fail.

With this economical addition of arrays, we can design constraints and subtyping such that the only failures are bounds-checks. A key issue is alignment: Given the alignment of e_1 , how can we know the alignment of $\&((\tau^{*\omega})(e_1))[e_2]$ without statically constraining the value of e_2 ? The solution taken by every sensible C platform is to ensure the size of τ is a multiple of its alignment; see Figure 9 for the formal constraint. That way, $\&((\tau^{*\omega})(e_1))[e_2]$ is at least as aligned as e_1 . Assuming this constraint, the typing rules for the new expression forms add nothing notable.

Finally but most importantly, we consider subtyping for pointer-to-array types. Analogues of UNROLL and ROLL are sound for types of the form $\text{ptr}_\alpha^\omega(\bar{\sigma})$, but PTR must be replaced with a more restrictive rule. Therefore, ARR requires the element type of the subtype to be the element type of the supertype repeated i

⁷In C, e is unsigned, but large allocations due to conversion from negative numbers are a well-known cause of defects.

⁸This check disallows pointing just past the end of the array, unlike C.

$$e ::= \dots \mid \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n$$

$$\begin{array}{c} \text{D-PCASE} \\ \frac{\Pi \models S_k}{\Pi; \bar{t} \vdash (\text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n) \xrightarrow{\tau} e_k} \end{array} \qquad \begin{array}{c} \text{S-PCASE} \\ \frac{\forall 1 \leq i \leq n . \bar{t}; \Gamma \vdash e_i : \tau; S'_i}{\bar{t}; \Gamma \vdash \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \tau; \bigwedge_{i=1}^n (S_i \Rightarrow S'_i) \wedge \bigvee_{i=1}^n S_i} \end{array}$$

$$\begin{array}{c} \text{L-PCASE} \\ \frac{\forall 1 \leq i \leq n . \text{if } \Pi \models S_i \text{ then } \Pi; \bar{t}; \Psi; \Gamma \vdash e_i : \bar{\sigma} \quad \exists 1 \leq i \leq n . \Pi \models S_i}{\Pi; \bar{t}; \Psi; \Gamma \vdash \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \bar{\sigma}} \end{array}$$

Figure 10: Additions for Platform Selection

times. This is more lenient than strict invariance. For example, it supports the platform-dependent idiom of treating an array of: `struct { short i1; short i2; short i3; short i4; }`; as an array of `short`. We have proven safety given this subtyping rule (Appendix A).

The `ARR` rule does *not* support subtyping such as:

$D \vdash \text{ptr}_\alpha^\omega(\text{byte byte byte}) \leq \text{ptr}_\alpha^\omega(\text{byte byte})$. A cast requiring this subtyping makes sense if the pointed-to-array has an element count divisible by 6, else it is memory-safe but probably a bug since the target type will “forget” the last byte in the array. We have not extended our formal model with arrays of known size, but we see no problems doing so. Such arrays are common in C, particularly with multidimensional arrays (all but one dimension must have known size), which is why CCured [31] allows casts like this.

4.2 Platform Selection

In practice, programs written in low-level languages can selectively run code based on features of the underlying platform. For example, in the following snippet, `lb` has type `long*`, `ib` points to a buffer filled with integers, and the program needs to treat `ib` as if it were an array of `long`:

```
if (sizeof(int) == sizeof(long))
    lb = (long*)ib;
else
    lb = convert(ib);
```

If the size of `int` equals the size of `long`, the buffer can be directly used at type `long*`. Otherwise, the function `convert` allocates a new `long*` buffer, copies the elements from `ib` into it, and assigns it into `lb`. The test offers a common-case short path, avoiding a copy on many platforms while remaining portable.

Figure 10 shows the additions to our model to support this idiom. The `pcase` form has n branches, each guarded by a constraint. Given a platform Π , `pcase` steps to a branch whose guard is true under Π (shown in `D-PCASE`). The output constraint in the high-level typing rule (`S-PCASE`) encodes two important properties. The first conjunct requires the constraint for each branch’s body to hold only if the constraint guarding the branch holds. In particular, an implication holds vacuously on platforms not modeling the guard. The second conjunct demands that at least one of the guards is true. The low-level typing rule (`L-PCASE`) similarly demands that at least one guard is true under Π and its corresponding expression type-checks under Π . Using `pcase`, the previous example can be written as:

```
pcase
    size(xtype(int)) = size(xtype(long)) => lb=(long*)ib
    size(xtype(int)) != size(xtype(long)) => lb=convert(ib)
```

In addition to the idiom exemplified above, `pcase` effectively explains *selective compilation*, where the preprocessor is used to compile code on a platform-dependent basis.

4.3 Read-Only Pointers

We do not allow subtyping under pointer types because a pointer can be dereferenced on the left side of an assignment. In C, `const τ*` describes pointers that cannot be used to write to the pointed-to data (though the lack of qualifier polymorphism [16] causes `const` to be used rarely and often removed via unsafe casts).

Adding this new flavor of pointer type to our model (for both high-level and low-level types) is straightforward:

- There are no changes to the dynamic semantics.
- S-DEREF_L must require a non-`const` pointer; S-DEREF can allow a `const` or non-`const` pointer; and S-FADDR must produce a type with the same qualifier as its subexpression.
- For subtyping rules PTR, UNROLL, and ROLL, we can add versions where the two types are both `const`.
- Finally, we add two new subtyping rules to show that `const` permits deep subtyping and allows less access than non-`const`:

$$\frac{D \vdash \bar{\sigma} \leq \bar{\sigma}'}{D \vdash \text{const ptr}_\alpha(\bar{\sigma}) \leq \text{const ptr}_\alpha(\bar{\sigma}')} \quad \frac{}{D \vdash \text{ptr}_\alpha(\bar{\sigma}) \leq \text{const ptr}_\alpha(\bar{\sigma})}$$

This addition is synergistic with more expressive recursive subtyping, discussed below.

4.4 Byte-Skipping

Our dynamic semantics assumes that an assignment copies all bytes of the right-hand value into the corresponding heap location. Actually, C implementations may choose to *skip* pad bytes. In practice, skipping reduces the amount of memory written but can increase the number of store instructions.

Example 7 showed a contrived situation where skipping could allow more subtyping. Conversely, although we did not add equality on structures to our expression language, skipping can lead to equality failing because some bytes remain uninitialized (in practice, holding unpredictable bits) despite the struct value being initialized. Though these issues are probably rare, our model is flexible enough to handle them and shed light on the meaning of skipping. To summarize the changes:

- Add a form `skip[i]` to σ and let $D \vdash \text{pad}[i] \leq \text{skip}[i]$.
- Change D-ASSN to “merge” the new value with the old one by skipping over any “skip bytes” as indicated by $\Pi.xtype(\bar{t}, \tau)$. That is, we use the type translation to indicate where skipping does and does not occur.
- The new skip type allows additional subtyping under pointers:

$$\frac{D \vdash \bar{\sigma}_2 \leq \text{skip}[i]}{D \vdash \text{ptr}_\alpha(\bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3) \leq \text{ptr}_\alpha(\bar{\sigma}_1 \text{skip}[i] \bar{\sigma}_3)}$$

4.5 Recursive Subtyping

Our definition of named struct types (\bar{t}) allows recursive types, but for simplicity our definition of physical subtyping is overly restrictive. For example, given two isomorphic structs defining linked-lists of integers, $N_1\{\text{long } f_1; N_1 * f_2\}$ and $N_2\{\text{long } g_1; N_2 * g_2\}$, we might have $\Pi.xtype(\bar{t}, N_1) = \text{byte}^4 \text{ptr}_\alpha(N_1)$ and $\Pi.xtype(\bar{t}, N_2) = \text{byte}^4 \text{ptr}_\alpha(N_2)$. We would be unable to show $\Pi; \bar{t} \vdash \text{byte}^4 \text{ptr}_\alpha(N_1) \leq \text{byte}^4 \text{ptr}_\alpha(N_2)$.

Sound solutions to this sort of limitation are well-known [6]. For this example, one could maintain a context of valid subtyping assumptions and to derive $\text{ptr}_\alpha(N_1) \leq \text{ptr}_\alpha(N_2)$ one can show the translation of N_1 is a subtype of the translation of N_2 while assuming $\text{ptr}_\alpha(N_1) \leq \text{ptr}_\alpha(N_2)$.

While our formal exposition opts for simplicity over expressiveness, our tool favors the latter and supports proper recursive subtyping.

5 Implementation

To give our theory a practical outlet, we have implemented a bug-finding tool that is directly inspired by the formal model. In general, a tool based on the model should extract a layout constraint S from a program and present informative warnings based on S . We see several ways to achieve this. For example:

- Directly explain S to the user, *simplifying* it into a legible form and connecting it to relevant locations in the code.
- Check S against a set of platforms. That is, the user chooses a set of platforms of interest Π_1, \dots, Π_n , and warnings are reported whenever $\Pi_i \not\models S$, for $1 \leq i \leq n$.

Our tool takes the latter approach; we leave the former to future work. Implementing the formal model poses two main challenges:

1. The type system relies on syntactic types to produce constraints, using no flow or alias information. While convenient for the exposition and metatheory, such simplicity is too imprecise.
2. Real C programs contain many downcasts that are safe at run-time. In a simple implementation of the model, these downcasts would result in spurious warnings.

To address (1), our tool replaces the simple type system with a points-to analysis. Using points-to information yields much more precise constraints and correctly handles “round-trip” casts through `void*`. For example, `D*t=(D*)(void*)e`, where `e` points to a value of type `S*`, is properly identified as a cast from `S*` to `D*`.

To address (2), we structured the tool to assume a standard porting scenario: The program has been developed and tested on a platform called the *host*, and the programmer is now interested in porting it to one or more *targets*. We assume that if a pointer cast is legal on the host, it should also be legal on the target. If a cast is illegal on both the host and the target, we do not report a warning. Our experience suggests that (a) if a cast is legal on the host and illegal on the target, it is highly likely to be a real bug, and (b) if a cast is illegal on both the host and the target, then it is highly *unlikely* to be a bug. The warning is likely a result of imprecision in the static analysis. In essence, our host-target scenario is an effective false-positive filter.

While it remains future work to use the tool to investigate thoroughly the extent of layout-portability bugs in C software, preliminary experience suggests it is a valuable addition to the developer’s toolset when porting or programming with portability in mind. The rest of this section describes the tool’s architecture and a case study that discovered a previously unknown bug.

5.1 Tool Overview

The tool has two main components. The *cast gatherer* is a static analysis that takes a C program and outputs a list of pointer casts that may occur at run-time. The *cast analyzer* inputs this list of casts, generates and checks their corresponding constraints, and outputs a list of warnings with code locations.

The cast gatherer uses an interprocedural points-to analysis⁹ to determine which memory layouts an expression may point to at run-time. A first pass analyzes all `malloc`¹⁰ sites and records a map of associated program points and their allocation-time types in a *type table*. In addition to allocation sites, program points for local variables that participate in pointer casts are also entered in the type table. A second pass analyzes each pointer cast `(τ^*)e`. For each entry (p, τ^*) in the type table, if `e` may-alias the allocation expression at program point p (hence `e` may point to memory with run-time type τ^*), a pointer cast from τ^* to τ^* is recorded.

The list of pointer casts output by the cast gatherer is passed to the cast analyzer along with a set of *platform descriptions* chosen by the user: one for the host and one or more for each target. A platform

⁹We use the points-to analysis that ships with CIL [1].

¹⁰A command-line flag allows specifying names of user-defined allocators.

```

data_link.c:196: scat_element * ==> iovec *
  Host (Gcc/32-bit X86):
    Src: ptr_4(ptr_1(b) bbbb)
    Dest: ptr_4(ptr_1(b) bbbb)
  Target (Gcc/LP-64):
    Src: ptr_8(ptr_1(b) bbbb----)
    Dest: ptr_8(ptr_1(b) bbbbbbbb)
events.c:150: sp_time * ==> timeval *
  Host (Gcc/32-bit X86):
    Src: ptr_4(bbbb bbbb)
    Dest: ptr_4(bbbb bbbb)
  Target (Gcc/LP-64):
    Src: ptr_8(bbbbbbbb bbbbbbbb)
    Dest: ptr_8(bbbbbbbb bbbb----)

```

Figure 11: Tool Output on Spread

description is a module, written by us in Caml, that implements a platform’s layout policy by defining exactly the platform functions in our model (recall Figure 3). We have implemented many such platform descriptions, most of which represent real platforms, and some of which implement imagined platforms that are still within the C language specification.

For each pointer cast, the cast analyzer generates and checks the corresponding constraint. First, the constraint is checked against the host description. If it is true (hence the cast is legal on the host), it is checked against each target. If it is false against any of the targets, the relevant warnings (described below) are output. The constraint checker was implemented manually and is specialized to our current set of constraints. It queries the physical subtyping relation, for which we have implemented an algorithm, and the platform description functions.

While a whole-program analysis is necessary for soundness, the tool can be run on a subset of the program at the expense of coverage. Also, it is easy to plug in a different cast gatherer. For example, we have experimented with a dynamic analysis that produces exact results per run.

5.2 Case Study

We ran our tool on Spread [3], a high-performance messaging service intended for use by distributed applications. We chose Spread because (a) it contained a reported layout portability bug [2], and (b) it is intended to be portable. Our tool issued two warnings, one of which was the known bug, and the other a new bug that to our knowledge has not been reported on the developer mailing lists. No false positives were reported.

A relevant subset of the tool’s output is given in Figure 11. The rest of the output was about these same casts at different locations in the code. The two warnings (the first of which is the known bug) are issued in the context of a conventional “Gcc on 32-bit X86” host description and a so-called “LP-64” target, on which integers are 32 bits and `long` and pointers are 64 bits. Each warning lists the bad pointer cast and its location, followed by the layouts of the source and destination type on each platform. The layouts are displayed in an ASCII version of the language for σ in Figure 3, where ‘b’ stands for byte, ‘-’ for pad[1], ‘ptr_4(b)’ for ptr_[4,0](byte), etc.

In the first case, we learn that the two types are laid out identically on the host, but on the target, the source type has four bytes of trailing padding where the destination type contains data bytes. The types involved in the cast are as follows:

```

struct scat_element { char *buf; int len; }
struct iovec        { char *buf; size_t len; }

```

The cast makes an assumption that `sizeof(int)` is equal to `sizeof(size_t)`. On LP-64 platforms, however, `int` is 4 bytes and `size_t` is 8 bytes, and the cast leads to a bigger-than-expected value in `iovec`'s `len` field and hence out-of-bounds `buf` accesses.

In the second warning, the target layout contains trailing padding where the host does not. The two types are:

```
struct sp_time { long sec; long usec; }
struct timeval { time_t tv_sec;
                suseconds_t tv_usec; }
```

On both the host and target, `time_t` is a type alias for `long` and `suseconds_t` is a type alias for `int`. The cast leads to truncation of the `usec` field, only half of the bytes being visible through `tv_usec`. If the target machine is big-endian, most or all of the relevant data is lost, being mapped to the sequence of padding.

This bug is particularly interesting, because in addition to being impossible to detect on the host via testing, it is also very difficult to detect on the target. Since the `tv_usec` field is off by some number of nanoseconds, finding this bug by testing literally depends on the time of day. Our tool detects it statically and without access to the target platform.

6 Previous Work

To our knowledge, previous work considering platform-dependent layout assumptions or low-level type-safety has taken one of the following approaches:

- Consider only one platform or one complete “bit-by-bit” data description.
- Check that a C program is portable, giving a compile-time error or fail-stop run-time termination if it is not.
- Assume that a C program is portable, i.e., extend the C compiler’s view that behavior for unportable programs is undefined.

The first approach is not helpful for writing code that works correctly on a set of platforms. The second approach suffices for applications that should be written in a higher-level language. The third approach relegates some issues to work such as ours, much as we relegate some issues like array-bounds errors to other work.

6.1 Assuming a Platform

Most closely related is the “physical type-checking” work of Chandra *et al.* [9, 36], which motivated our work considerably. Their tool classifies C casts as “upcasts”, “downcasts”, or “neither”, reporting a warning for the last possibility. Their notion of physical subtyping is at a higher level than ours (requiring that types, offsets, and field names match), and they neither parameterize their system by a platform nor produce descriptions of sets of platforms. Checking code against a new platform would require reverification and changing their tool. They present no metatheory validating their approach.

CCured [31], a memory-safe C platform, includes physical type-checking to reduce the number of casts that require run-time checks. That is, CCured permits casts that work in practice but are not allowed by the C standard. The allowed casts are safe under a padding strategy used by common C compilers for the x86 architecture, which covers some but certainly not all platforms. An interesting avenue for future work is to reimplement the CCured type analysis engine using our platform-as-parameter approach, thus making its memory-safety guarantee portable.

Work on typed assembly language and proof-carrying code [30, 29, 13, 10, 21] clearly needs a low-level view of memory. Such projects can establish that certifying compilers produce code that cannot get stuck due to uninitialized memory, unaligned memory access, segmentation faults, etc. In particular, work on

allocation semantics [34, 4] has taken a lower-level view than our formalism by treating addresses as integers and exposing that pointer arithmetic can move between adjacent data objects. These approaches provide less help for writing platform-dependent code because verification of type-safety is repeated for each platform. In practice, defining a new platform is an enormous amount of work.

Various program analyses for C (e.g., [37, 28]) have represented structs and unions with explicit bytes and offsets by assuming one particular platform.

6.2 Safe C

Memory-safe dialects or implementations of C, such as CCured [31, 32, 12], Deputy [11], SAFECode [14], and Cyclone [24, 19, 18], do not solve the platform dependency problem. Rather, they may reject (at compile-time) or terminate (at run-time) programs that attempt platform-dependent operations, or they may support only certain platforms (e.g., certain C compilers as back-ends). These approaches are fine for fully portable code or code that is correct assuming particular compilers.

Furthermore, the implementations of these systems include run-time systems (automatic memory managers, type-tag checkers, etc.) that themselves make platform-dependent assumptions! For example, Cyclone assumes 32-bit integers and pointers, and making this code more portable is a top request from actual users.

6.3 Formalizing C

Recent work by Leroy *et al.* [25, 8] uses Coq to prove a C compiler correct. Their operational semantics for C distinguishes left and right expressions much as we do. However, their source language omits structs, avoiding many alignment and padding issues, and their metatheory proves correctness only for correct source programs, presumably saying nothing about platform-dependent code.

Norrish’s HOL formalization of C [33] includes structs and uses a global namespace mapping struct names to sequences of typed fields, like our work. However, he purposely omits padding and alignment from his formalism. He has no separable notion of a platform; instead he models platform choices as nondeterminism.

6.4 Low-Level Code without C

Our work has been C-centric, whereas other projects have started with languages at higher levels of abstraction and added bit-level views for low-level programming. (See [7, 20] for just two recent examples.) We believe this complementary approach would benefit from our constraint-based view rather than choosing just between completely high-level types and completely low-level ones.

C-- [35] makes data representation and alignment explicit, but C-- is not appropriate for writing platform-dependent code. Rather, it is a low-level language designed as a target for compiling high-level languages. It has explicit padding on data (a compiler inserts bits where desired) and explicit alignment on all memory accesses.¹¹ Incorrect alignment is an unchecked run-time error. The purpose of C-- is to handle back-end code-generation issues for a compiler; it is still expected that the front-end compiler will generate different (but similar) code for each platform and provide a run-time system, probably written in C.

7 Conclusions and Future Work

This work has developed a formal description of platform dependencies in low-level software. The key insight is a semantic definition of “platform” that directs a low-level operational semantics *and* models a syntactic constraint that we can produce via static analysis on a source program. We have proven soundness for a small core language and a simple static analysis, and extended the approach to account for arrays and other language features. Giving platforms a clear identity in our framework clarifies a number of poorly understood issues. The formal model directly informs the implementation of a useful tool that finds and reports layout portability problems in C programs.

¹¹Syntactically, an omitted alignment is taken to be n for an n -byte access.

The technique of platform-as-parameter can apply broadly since high-level languages also have platform-defined behavior. As examples, SML programs may depend on the size of `int`, Scheme programs may depend on evaluation order, and Java programs may depend on fair thread-scheduling.

Looking to the future, we have begun investigating the issue of checking that a C program has observationally equivalent behavior on a set of platforms. This notion of portability needs to take into account issues such as such as endianness, integer overflow, and perhaps floating-point roundoff behavior. We believe that our notion of platform selection described in Section 4.2 (`pcase`) is instrumental in such a framework. The idea is to translate a non-preprocessed C program into a version of C with a `pcase` primitive, where preprocessor directives and `if` statements encoding platform selection are translated to corresponding `pcase` statements. Then, we can extract from this program the necessary proof obligations under which all `pcase` branches are equivalent when executed on their respective platforms. Since `pcase` branches can be arbitrary code, checking their equivalence is undecidable. We envision discharging the proof obligations in an interactive proof environment.

A key feature of the work in this paper is that it detects a relevant subset of portability problems *fully automatically*. A portability checker requiring interactive theorem proving is not accessible to most C programmers.

References

- [1] CIL: Infrastructure for C program analysis and transformation. <http://hal.cs.berkeley.edu/cil/>.
- [2] Spread: Portability bug on Solaris 8. <http://commedia.cnds.jhu.edu/pipermail/spread-users/2002-November/001185.html>.
- [3] The Spread Toolkit. <http://www.spread.org>.
- [4] A. Ahmed and D. Walker. The logical approach to stack typing. In *International Workshop on Types in Language Design and Implementation*, 2003.
- [5] *The ARMLinux Book Online*, Chapter 10. 2005. <http://www.aleph1.co.uk/armlinux/book>.
- [6] R. M. Amadio and L. Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4), 1993.
- [7] D. F. Bacon. Kava: a Java dialect with a uniform object model for lightweight classes. *Concurrency and Computation: Practice and Experience*, 15(3-5), 2003.
- [8] S. Blazy, Z. Dargaye, and X. Leroy. Formal verification of a C compiler front-end. In *14th International Symposium on Formal Methods*, 2006.
- [9] S. Chandra and T. Reps. Physical type checking for C. In *Workshop on Program Analysis for Software Tools and Engineering*, 1999.
- [10] J. Chen, D. Wu, A. W. Appel, and H. Fang. A provably sound TAL for back-end optimization. In *ACM Conference on Programming Language Design and Implementation*, 2003.
- [11] J. Condit, M. Harren, Z. Anderson, D. Gay, and G. Necula. Dependent types for low-level programming. In *European Symposium on Programming*, 2007.
- [12] J. Condit, M. Harren, S. McPeak, G. Necula, and W. Weimer. CCured in the real world. In *ACM Conference on Programming Language Design and Implementation*, 2003.
- [13] K. Crary. Toward a foundational typed assembly language. In *30th ACM Symposium on Principles of Programming Languages*, 2003.

- [14] D. Dhurjati, S. Kowshik, and V. Adve. SAFECode: Enforcing alias analysis for weakly typed languages. In *ACM Conference on Programming Language Design and Implementation*, 2006.
- [15] P. H. Eidorff, F. Henglein, C. Mossin, H. Niss, M. H. Sørensen, and M. Tofte. AnnoDomini: from type theory to year 2000 conversion tool. In *26th ACM Symposium on Principles of Programming Languages*, 1999.
- [16] J. Foster, M. Fähndrich, and A. Aiken. A theory of type qualifiers. In *ACM Conference on Programming Language Design and Implementation*, 1999.
- [17] D. Grossman. Type-safe multithreading in Cyclone. In *International Workshop on Types in Language Design and Implementation*, 2003.
- [18] D. Grossman. Quantified types in imperative languages. *ACM Transactions on Programming Languages and Systems*, 28(3), 2006.
- [19] D. Grossman, G. Morrisett, T. Jim, M. Hicks, Y. Wang, and J. Cheney. Region-based memory management in Cyclone. In *ACM Conference on Programming Language Design and Implementation*, 2002.
- [20] T. Hallgren, M. P. Jones, R. Leslie, and A. Tolmach. A principled approach to operating system construction in Haskell. In *10th ACM International Conference on Functional Programming*, 2005.
- [21] N. A. Hamid, Z. Shao, V. Trifonov, S. Monnier, and Z. Ni. A syntactic approach to foundational proof-carrying code. *Journal of Automated Reasoning*, 31(3–4), 2003.
- [22] IBM. Developing embedded software for the IBM PowerPC 970FX processor. Application Note 970, IBM, 2004. <http://www.ibm.com/chips/techlib/>.
- [23] *ISO/IEC 9899:1999, International Standard—Programming Languages—C*. International Standards Organization, 1999.
- [24] T. Jim, G. Morrisett, D. Grossman, M. Hicks, J. Cheney, and Y. Wang. Cyclone: A safe dialect of C. In *USENIX Annual Technical Conference*, 2002.
- [25] X. Leroy. Formal certification of a compiler back-end. In *33rd ACM Symposium on Principles of Programming Languages*, 2006.
- [26] R. Love. *Linux Kernel Development, Second Edition*. Novell Press, 2005. Page 328.
- [27] B. Martin, A. Rettinger, and J. Singh. Multiplatform porting to 64 bits. *Dr. Dobbs's Journal*, 2005.
- [28] A. Miné. Field-sensitive value analysis of embedded c programs with union types and pointer arithmetics. In *Conference on Language, Compilers, and Tool Support for Embedded Systems*, 2006.
- [29] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly language. *ACM Transactions on Programming Languages and Systems*, 21(3), 1999.
- [30] G. Necula. Proof-carrying code. In *24th ACM Symposium on Principles of Programming Languages*, 1997.
- [31] G. Necula, J. Condit, M. Harren, S. McPeak, and W. Weimer. CCured: Type-safe retrofitting of legacy software. *ACM Transactions on Programming Languages and Systems*, 27(3), 2005.
- [32] G. Necula, S. McPeak, and W. Weimer. CCured: Type-safe retrofitting of legacy code. In *29th ACM Symposium on Principles of Programming Languages*, 2002.
- [33] M. Norrish. *C formalised in HOL*. PhD thesis, University of Cambridge, 1998.

- [34] L. Petersen, R. Harper, K. Crary, and F. Pfenning. A type theory for memory allocation and data layout. In *30th ACM Symposium on Principles of Programming Languages*, 2003.
- [35] N. Ramsey, S. P. Jones, and C. Lindig. The C-- language specification version 2.0, 2005. <http://www.cminusminus.org/extern/man2.pdf>.
- [36] M. Siff, S. Chandra, T. Ball, K. Kunchithapadam, and T. Reps. Coping with type casts in C. In *7th European Software Engineering Conference 7th ACM Symposium on the Foundations of Software Engineering*, 1999.
- [37] R. P. Wilson and M. S. Lam. Efficient context-sensitive pointer analysis for C programs. In *ACM Conference on Programming Language Design and Implementation*, 1995.

A Proofs

This appendix is organized as follows. Section A.2 (page 30) develops the proof of memory safety relative to the low-level type system. Section A.3 (page 48) develops the connection between the two type systems and proves the key layout portability theorems. Sections A.4 (page 51) and A.5 (page 57) show how the metatheory extends for the array and platform selection extensions, respectively.

A.1 Preliminaries

Figure 12 gives the syntax for our C-like source language. Figure 14 gives an extended, “low-level” version of this language, where types are replaced by physical layouts, and runtime values are added. The syntax for heaps, heap typings, and evaluation contexts is also shown. Figure 15 gives the dynamic semantics for the low-level language and Figure 16 defines a set of value and type size functions used throughout the low-level semantics. Figures 17 and 18 give the static semantics for the low-level language. Figure 19 gives the static semantics for the high-level language.

We use the notation \bar{x} to denote a sequence of objects drawn from the syntactic class x , the notation x^i to denote a sequence of length i , and \cdot to denote the empty sequence. For space-saving purposes, we occasionally write two inference rules $\frac{D}{x}$ and $\frac{D}{y}$ as $\frac{D}{x}$ when they have identical hypotheses.

Platforms Our dynamic and low-level static semantics are parameterized by *platforms*. A platform is an oracle that guides the semantics in various ways. Each platform is a record of six functions, as follows:

- *xtype* : $\bar{t} \times \tau \rightarrow \bar{\sigma}$: Translates a high-level type into a low-level type.
- *align* : $\bar{t} \times \tau \rightarrow \alpha$: Gives the alignment of type τ .
- *offset* : $\bar{t} \times f \rightarrow \mathbb{N}$: Gives the offset of field f . Without loss of generality, we assume that all fields in \bar{t} are uniquely named.
- *access* : $\alpha \times \mathbb{N} \rightarrow \{\text{true}, \text{false}\}$: $\Pi.\text{access}(\alpha, k)$ determines whether it is okay to access a memory chunk of size k at alignment α on platform Π .
- *xlit* : $\{s, l\} \rightarrow \bar{w}$: Translates a literal expression into a sequence of bytes.
- *ptrsize* : \mathbb{N} : Gives the size of pointers. We limit our focus to platforms on which all pointers have the same size.

Constraint Language The high-level semantics (Figure 19) is a type and effect system that, in addition to a type, outputs a constraint S . The language used to express S is a multi-sort first-order logic in which each of the $\bar{t}, \tau, \bar{\sigma}, f, \alpha, \mathbb{N}$ syntactic classes is assigned a corresponding sort. Moreover, the logic is enriched with function symbols corresponding to each platform component plus a couple of others. Platforms are models in this theory; we write $\Pi \models S$ when Π satisfies formula S . The (\models) judgment is entirely ordinary, defined inductively over the structure of S . For example, $\Pi \models S_1 \wedge S_2$ exactly when $\Pi \models S_1$ and $\Pi \models S_2$; and $\Pi \models \forall x : s. S$ exactly when for all objects t of sort s , $S[t/x]$ is true, where $S[t/x]$ denotes the capture-avoiding substitution of t for x into S . We generally omit explicit sorts on quantified variables; it is clear what is meant from the context.

Figure 13 summarizes the function symbols and their interpretation under Π . For an example involving function symbols, suppose we have a constraint $S = (\text{align}(\bar{t}, N) = [a, o] \wedge \text{align}(\bar{t}, \tau) = [a, o + \text{offset}(\bar{t}, f)])$ and want to know whether a platform Π models S ($\Pi \models S$). Π models S iff $\Pi \models \text{align}(\bar{t}, N) = [a, o]$ and $\Pi \models \text{align}(\bar{t}, \tau) = [a, o + \text{offset}(\bar{t}, f)]$. To satisfy the former conjunct, $\Pi.\text{align}(\bar{t}, N) = [a, o]$ must be true. To satisfy the latter, $\Pi.\text{align}(\bar{t}, \tau) = [a, o + \Pi.\text{offset}(\bar{t}, f)]$ must be true.

(types)	τ	$::=$	short long τ^* N
(declarations)	t	$::=$	$N\{\overline{\tau} f\}$
(expressions)	e	$::=$	$s \mid l \mid x \mid e = e \mid e.f \mid *(\tau^*)(e) \mid \mathbf{new} \tau \mid (\tau^*)e \mid (\tau^*)&e \rightarrow f$
			$\mid e; e \mid \mathbf{if} e e e \mid \mathbf{while} e e \mid \tau x; e$
(contexts)	Γ	$::=$	$\overline{x:\tau}$

Figure 12: Syntax for the Source Language

symbol	interpretation under Π
$xtype(\overline{t}, \tau)$	$\Pi.xtype(\overline{t}, \tau)$
$align(\overline{t}, \tau)$	$\Pi.align(\overline{t}, \tau)$
$offset(\overline{t}, f)$	$\Pi.offset(\overline{t}, f)$
$access(\alpha, i)$	$\Pi.access(\alpha, i)$
$xlit(s)$	$\Pi.xlit(s)$
$xlit(l)$	$\Pi.xlit(l)$
$ptrsize$	$\Pi.ptrsize$
$size(\overline{\sigma})$	$size(\Pi, \overline{\sigma})$
$size(\overline{w})$	$size(\Pi, \overline{w})$
$subtype(\overline{t}, \overline{\sigma}_1, \overline{\sigma}_2)$	$\Pi; \overline{t} \vdash \overline{\sigma}_1 \leq \overline{\sigma}_2$
$subalign(\alpha_1, \alpha_2)$	$\vdash \alpha_1 \leq \alpha_2$

Figure 13: Constraint Language Function Symbols

Definition 8 (Sensible Platforms)

A platform Π is said to be *sensible* if

1. $\forall \tau. \Pi.access(\Pi.align(\overline{t}, \tau), size(\Pi, \Pi.xtype(\overline{t}, \tau)))$.
2. $\exists i. \Pi.xtype(\overline{t}, \text{short}) = \text{byte}^i$ and $\forall s \exists \overline{b}. \Pi.xlit(s) = \overline{b}$ and $size(\Pi, \overline{b}) = i$.
 $\exists i. \Pi.xtype(\overline{t}, \text{long}) = \text{byte}^i$ and $\forall s \exists \overline{b}. \Pi.xlit(l) = \overline{b}$ and $size(\Pi, \overline{b}) = i$.
3. $\forall \overline{t}. \text{if } N\{\dots \tau f \dots\} \in \overline{t}$ and $\Pi.xtype(\overline{t}, \tau) = \overline{\sigma}_2$ then $\exists \overline{\sigma}_1, \overline{\sigma}_3, a, o, \alpha$ such that $\Pi.xtype(\overline{t}, N) = \overline{\sigma}_1 \overline{\sigma}_2 \overline{\sigma}_3$ where $\Pi.offset(\overline{t}, f) = size(\Pi, \overline{\sigma}_1)$, $\Pi.align(\overline{t}, N) = [a, o]$, $\Pi.align(\overline{t}, \tau) = \alpha$, and $\vdash [a, o + \Pi.offset(\overline{t}, f)] \leq \alpha$.

The reason we have the “subalignment” requirement here, rather than equality, is to admit more implementations without sacrificing soundness; otherwise we would be unnecessarily restrictive. For example, suppose $\Pi.align(\overline{t}, N) = [8, 0]$, $\Pi.offset(\overline{t}, f) = 8$, and $\Pi.align(\overline{t}, \tau) = [4, 0]$. Then, clause (3) would require τ to be 8-byte aligned (alignment $[8, 0 + \Pi.offset(\overline{t}, f)] \equiv [8, 8] \equiv [8, 0]$), which is unreasonable. It is okay for τ to be 4-byte, 2-byte, or 1-byte aligned.

4. $\forall \tau \exists \alpha, \overline{\sigma}. \Pi.xtype(\overline{t}, \tau^*) = \text{ptr}_\alpha(\overline{\sigma})$ and $\Pi.xtype(\overline{t}, \tau) = \overline{\sigma}$ and $\Pi.align(\overline{t}, \tau) = \alpha$.
5. If $\Pi.access(\alpha, i)$ and $\vdash \alpha' \leq \alpha$ then $\Pi.access(\alpha', i)$.

		$\sigma ::= \text{byte} \mid \text{pad}[i] \mid \text{ptr}_\alpha(\bar{\sigma}) \mid \text{ptr}_\alpha(N)$	
(types)			
(alignments)		$\alpha ::= [a, o]$	
(atomic values)		$w ::= b \mid \ell+i \mid \text{uninit}$	
(bytes)		$b ::= 0 \mid 1 \mid \dots \mid 255$	
(expressions)		$e ::= \dots \mid \bar{w}$	
(values)		$v ::= \bar{w}$	
(heaps)		$H ::= \frac{\ell \mapsto v, \alpha}{\ell : \bar{\sigma}, \alpha}$	
(heap typings)		$\Psi ::= \ell : \bar{\sigma}, \alpha$	
		$D ::= \Pi; \bar{t}$	
		$C ::= \Psi; \Gamma$	
		$i, j, k, a, o \in \mathbb{N}$	
(left contexts)		$L ::= [\cdot]_L \mid L.f \mid *(\tau*)(R)$	
(right contexts)		$R ::= [\cdot]_R \mid L = e \mid *(\tau*)(\ell+i) = R \mid R.f \mid *(\tau*)(R) \mid (\tau*)R$	
		$\mid R; e \mid (\tau*)&R \rightarrow f \mid \text{if } R \ e \ e$	

Figure 14: Syntax for the Full Language (Extends Source Language) and Eval Contexts

Definition 9 (Substitution)

$$\begin{aligned}
s\{e/x\} &= s \\
l\{e/x\} &= l \\
x\{e/x\} &= e \\
y\{e/x\} &= y \\
\bar{w}\{e/x\} &= \bar{w} \\
(e_1 = e_2)\{e/x\} &= e_1\{e/x\} = e_2\{e/x\} \\
(e_1.f)\{e/x\} &= (e_1\{e/x\}).f \\
(\tau)(e_1)\{e/x\} &= *(\tau*)((e_1\{e/x\})) \\
(\text{new } \tau)\{e/x\} &= \text{new } \tau \\
((\tau*)e_1)\{e/x\} &= (\tau*)(e_1\{e/x\}) \\
((\tau*)&e_1 \rightarrow f)\{e/x\} &= (\tau*)&(e_1\{e/x\}) \rightarrow f \\
(e_1; e_2)\{e/x\} &= e_1\{e/x\}; e_2\{e/x\} \\
(\text{if } e_1 \ e_2 \ e_3)\{e/x\} &= \text{if } (e_1\{e/x\}) \ (e_2\{e/x\}) \ (e_3\{e/x\}) \\
(\text{while } e_1 \ e_2)\{e/x\} &= \text{while } (e_1\{e/x\}) \ (e_2\{e/x\}) \\
(\tau \ y; \ e_1)\{e/x\} &= \tau \ z; \ (e_1\{z/y\})\{e/x\} \quad (z \text{ fresh})
\end{aligned}$$

Definition 10 (Legal Stuck State) *Let*

$$\begin{aligned}
rstuck &::= \text{if } \bar{w}_1 \ \text{uninit} \ \bar{w}_2 \ e \ e \mid *(\tau*)(\text{uninit}^i) \mid (\tau*)&\text{uninit}^i \rightarrow f \\
lstuck &::= *(\tau*)(\text{uninit}^i)
\end{aligned}$$

A state $H; e$ is legally stuck if one of the following is true:

- $e = R[lstuck]_i$,
- $e = R[rstuck]_i$,
- $e = L[lstuck]_i$,
- $e = L[rstuck]_i$,

Definition 11 (Illegal Stuck State)

A state $H; e$ is illegally stuck on a platform Π and declarations \bar{t} if $H; e$ is not legally stuck, e is not a value, and there exist no H' and e' such that $\Pi; \bar{t} \vdash H; e \rightarrow H'; e'$.

(D-CAST)	$D \vdash H; (\tau^*)\bar{w}$	\xrightarrow{r}	$H; \bar{w}$
(D-SHORT)	$D \vdash H; s$	\xrightarrow{r}	$H; \bar{b}$ if $\Pi.xlit(s) = \bar{b}$
(D-LONG)	$D \vdash H; l$	\xrightarrow{r}	$H; \bar{b}$ if $\Pi.xlit(l) = \bar{b}$
(D-SEQ)	$D \vdash H; (v; e)$	\xrightarrow{r}	$H; e$
(D-IFIF)	$D \vdash H; \text{if } 0^i e_1 e_2$	\xrightarrow{r}	$H; e_2$
(D-IFT)	$D \vdash H; \text{if } b_1 \dots b_i e_1 e_2$	\xrightarrow{r}	$H; e_1$ if $b_1 \dots b_i \neq 0^i$
(D-WHILE)	$D \vdash H; \text{while } e_1 e_2$	\xrightarrow{r}	$H; \text{if } e_1 (e_2; \text{while } e_1 e_2) s$
(D-NEW)	$\Pi; \bar{t} \vdash H; \text{new } \tau$	\xrightarrow{r}	$H, \ell \mapsto \text{uninit}^i, \alpha; \ell+0$ if $\ell \notin \text{Dom}(H)$ $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ $\text{size}(\Pi, \bar{\sigma}) = i$ $\Pi.align(\bar{t}, \tau) = \alpha$
(D-DECL)	$\Pi; \bar{t} \vdash H; \tau x; e$	\xrightarrow{r}	$H, \ell \mapsto \text{uninit}^i, \alpha; e\{*(\tau^*)(\ell+0)/x\}$ if ... same as above ...
(D-FETCH)	$\Pi; \bar{t} \vdash H; \bar{w}_1 \bar{w}_2 \bar{w}_3.f$	\xrightarrow{r}	$H; \bar{w}_2$ if $\Pi.offset(\bar{t}, f) = \text{size}(\Pi, \bar{w}_1)$ $N\{\dots \tau f \dots\} \in \bar{t}$ $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ $\text{size}(\Pi, \bar{\sigma}) = \text{size}(\Pi, \bar{w}_2)$
(D-FADDR)	$\Pi; \bar{t} \vdash H; (\tau^*)\&\ell+j \rightarrow f$	\xrightarrow{r}	$H; \ell+(j + \Pi.offset(\bar{t}, f))$
(D-DEREF)	$\Pi; \bar{t} \vdash H; *(\tau^*)(\ell+j)$	\xrightarrow{r}	$H; \bar{w}_2$ if $H(\ell) = \bar{w}_1 \bar{w}_2 \bar{w}_3, [a, o]$ $\text{size}(\Pi, \bar{w}_1) = j$ $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ $\text{size}(\Pi, \bar{w}_2) = \text{size}(\Pi, \bar{\sigma}) = k$ $\Pi.access([a, o + j], k)$
(D-ASSN)	$\Pi; \bar{t} \vdash H; *(\tau^*)(\ell+j) = \bar{w}$	\xrightarrow{r}	$H, \ell \mapsto \bar{w}_1 \bar{w}_2 \bar{w}_3, [a, o]; \bar{w}$ if ... same as above ... $\text{size}(\Pi, \bar{w}) = \text{size}(\Pi, \bar{w}_2)$
(D-FETCHL)	$\Pi; \bar{t} \vdash H; (*(\tau_1^*)(\ell+j)).f$	\xrightarrow{l}	$H; *(\tau_2^*)(\ell+(j+j'))$ if $\Pi.offset(f) = j'$ $N\{\dots \tau_2 f \dots\} \in \bar{t}$
D-STEP L	$\frac{D \vdash H; e \xrightarrow{l} H'; e'}{D \vdash H; R[e]_l \rightarrow H'; R[e']_l}$	D-STEP R	$\frac{D \vdash H; e \xrightarrow{r} H'; e'}{D \vdash H; R[e]_r \rightarrow H'; R[e']_r}$

Figure 15: Dynamic Semantics

size	:	$(\Pi \times \sigma) \rightarrow \mathbb{N}$
size(Π, σ)	=	$\begin{cases} 1 & \text{if } \sigma = \text{byte} \\ i & \text{if } \sigma = \text{pad}[i] \\ \Pi.\text{ptrsize} & \text{if } \sigma \in \{\text{ptr}_\alpha(N), \text{ptr}_\alpha(\bar{\sigma})\} \end{cases}$
size	:	$(\Pi \times \bar{\sigma}) \rightarrow \mathbb{N}$
size($\Pi, \sigma_1 \dots \sigma_n$)	=	$\sum_{i=1}^n \text{size}(\Pi, \sigma_i)$
size	:	$(\Pi \times w) \rightarrow \mathbb{N}$
size(Π, w)	=	$\begin{cases} 1 & \text{if } w \in \{b, \text{uninit}\} \\ \Pi.\text{ptrsize} & \text{if } w = \ell+i \end{cases}$
size	:	$(\Pi \times \bar{w}) \rightarrow \mathbb{N}$
size($\Pi, w_1 \dots w_n$)	=	$\sum_{i=1}^n \text{size}(\Pi, w_i)$

Figure 16: Size Functions

Definition 12 (Heap Typing Extension)

A heap typing Ψ *extends* a heap typing Ψ' iff there exists Ψ'' such that $\Psi = \Psi' \Psi''$.

A.2 Low-Level Memory Safety

Lemma 13 (Weakening)

1. (a) If $D; \Psi; \Gamma \vdash_r e : \bar{\sigma}$ then $D; \Psi\Psi'; \Gamma\Gamma' \vdash_r e : \bar{\sigma}$.
 (b) If $D; \Psi; \Gamma \vdash e : \bar{\sigma}, \alpha$ then $D; \Psi\Psi'; \Gamma\Gamma' \vdash e : \bar{\sigma}, \alpha$.
2. If $D; \Psi \vdash H : \Psi'$ then $D; \Psi\Psi'' \vdash H : \Psi'$.

Proof:

1. By simultaneous induction on the assumed typing derivations, using the facts that if $\ell \in \text{Dom}(\Psi)$ then $(\Psi\Psi')(\ell) = \Psi(\ell)$, and if $x \in \text{Dom}(\Gamma)$ then $(\Gamma\Gamma')(x) = \Gamma(x)$.
2. By straightforward induction on the assumed heap typing derivation, using the result of part (1).

Lemma 14 (Heap Canonical Forms)

If $D; \Psi \vdash H : \Psi$ and $\Psi(\ell) = \bar{\sigma}, \alpha$, then $\exists \bar{w}$ such that $H(\ell) = \bar{w}, \alpha$ and $D; \Psi; \cdot \vdash_r \bar{w} : \bar{\sigma}$.

Proof: By straightforward induction on the assumed heap typing derivation.

Lemma 15 (Uninit Type)

If $\text{size}(\Pi, \bar{\sigma}) = i$ then $\Pi; \bar{t}; \Psi; \cdot \vdash_r \text{uninit}^i : \bar{\sigma}$.

Proof: By induction on the length of $\bar{\sigma}$. In the base case, $\bar{\sigma} = \cdot$ and the claim follows immediately. In the inductive case, $\bar{\sigma} = \sigma\bar{\sigma}'$ where by induction we have $\text{size}(\Pi, \bar{\sigma}') = j$ and

1 $\Pi; \bar{t}; \Psi; \cdot \vdash_r \text{uninit}^j : \bar{\sigma}'$

If $\text{size}(\Pi, \sigma) = k$ then we can derive

2 $\Pi; \bar{t}; \Psi; \cdot \vdash_r \text{uninit}^k : \sigma$

$\frac{\text{L-SHORT}}{\frac{\Pi.xtype(\bar{t}, \text{short}) = \text{byte}^i}{\Pi; \bar{t}; C \vdash s : \text{byte}^i}}$	$\frac{\text{L-LONG}}{\frac{\Pi.xtype(\bar{t}, \text{long}) = \text{byte}^i}{\Pi; \bar{t}; C \vdash l : \text{byte}^i}}$	$\frac{\text{L-VARR}}{\frac{\Gamma(x) = \tau \quad \Pi.xtype(\bar{t}, \tau) = \bar{\sigma}}{\Pi; \bar{t}; \Psi; \Gamma \vdash x : \bar{\sigma}}}$
$\frac{\text{L-VARL}}{\frac{\Gamma(x) = \tau \quad \Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.align(\bar{t}, \tau) = \alpha}{\Pi; \bar{t}; \Psi; \Gamma \vdash x : \bar{\sigma}, \alpha}}$		$\frac{\text{L-ASSN}}{\frac{D; C \vdash e_1 : \bar{\sigma}, \alpha \quad D; C \vdash e_2 : \bar{\sigma}}{D; C \vdash e_1 = e_2 : \bar{\sigma}}}$
$\frac{\text{L-FETCHR}}{\frac{\Pi; \bar{t}; C \vdash e : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3 \quad \Pi.offset(f) = \text{size}(\Pi, \bar{\sigma}_1) \quad N\{\dots \tau f \dots\} \in \bar{t} \quad \Pi.xtype(\bar{t}, \tau) = \bar{\sigma}_2}{\Pi; \bar{t}; C \vdash e.f : \bar{\sigma}_2}}$		
$\frac{\text{L-FETCHL}}{\frac{\Pi; \bar{t}; C \vdash e : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a, o] \quad \Pi.offset(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_1) = o' \quad N\{\dots \tau f \dots\} \in \bar{t} \quad \Pi.align(\bar{t}, N) = \alpha \quad \vdash [a, o] \leq \alpha \quad \Pi.xtype(\bar{t}, \tau) = \bar{\sigma}_2}{\Pi; \bar{t}; C \vdash e.f : \bar{\sigma}_2, [a, o + o']}}$		
$\frac{\text{L-DEREF}\{R,L\}}{\frac{\Pi; \bar{t}; C \vdash e : \text{ptr}_\alpha(\bar{\sigma}_1 \bar{\sigma}_2) \quad \Pi.xtype(\bar{t}, \tau) = \bar{\sigma}_1 \quad \Pi.access(\alpha, \text{size}(\Pi, \bar{\sigma}_1))}{\frac{\Pi; \bar{t}; C \vdash *(\tau*)(e) : \bar{\sigma}_1}{\Pi; \bar{t}; C \vdash *(\tau*)(e) : \bar{\sigma}_1, \alpha}}}$		
$\frac{\text{L-NEW}}{\frac{\Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.align(\bar{t}, \tau) = \alpha}{\Pi; \bar{t}; C \vdash \text{new } \tau : \text{ptr}_\alpha(\bar{\sigma})}}$	$\frac{\text{L-CAST}}{\frac{\Pi; \bar{t}; C \vdash e : \text{ptr}_\alpha(\bar{\sigma})}{\Pi; \bar{t}; C \vdash (\tau*)e : \text{ptr}_\alpha(\bar{\sigma})}}$	
$\frac{\text{L-FADDR}}{\frac{\Pi; \bar{t}; C \vdash e : \text{ptr}_{[a, o_1]}(\bar{\sigma}_1 \bar{\sigma}_2) \quad \Pi.offset(f) = \text{size}(\Pi, \bar{\sigma}_1)}{\Pi; \bar{t}; C \vdash (\tau*)&e \rightarrow f : \text{ptr}_{[a, o_2]}(\bar{\sigma}_2)}}$	$\frac{\text{L-SEQ}}{\frac{D; C \vdash e_1 : \bar{\sigma}' \quad D; C \vdash e_2 : \bar{\sigma}}{D; C \vdash e_1; e_2 : \bar{\sigma}}}$	
$\frac{\text{L-DECL}}{\frac{D; \Psi; \Gamma, x : \tau \vdash e : \bar{\sigma}}{D; \Psi; \Gamma \vdash \tau x; e : \bar{\sigma}}}$	$\frac{\text{L-IF}}{\frac{\Pi; \bar{t}; C \vdash e_1 : \text{byte}^i \quad \Pi; \bar{t}; C \vdash e_2 : \bar{\sigma} \quad \Pi; \bar{t}; C \vdash e_3 : \bar{\sigma}}{\Pi; \bar{t}; C \vdash \text{if } e_1 \ e_2 \ e_3 : \bar{\sigma}}}$	
$\frac{\text{L-WHILE}}{\frac{\Pi; \bar{t}; C \vdash e_1 : \text{byte}^j \quad \Pi; \bar{t}; C \vdash e_2 : \bar{\sigma}_2 \quad \Pi.xtype(\bar{t}, \text{short}) = \text{byte}^i}{\Pi; \bar{t}; C \vdash \text{while } e_1 \ e_2 : \text{byte}^i}}$		
$\frac{\text{L-VALUE}}{\frac{D; \Psi \Vdash \bar{w}_1 : \sigma_1 \quad D; \Psi; \Gamma \vdash \bar{w}_2 : \bar{\sigma}_2}{D; \Psi; \Gamma \vdash \bar{w}_1 \bar{w}_2 : \sigma_1 \bar{\sigma}_2}}$	$\frac{\text{L-VALUEEMPTY}}{D; \Psi \vdash \cdot : \cdot}$	$\frac{\text{L-SUB}}{\frac{D; C \vdash e : \bar{\sigma}_1 \quad D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2}{D; C \vdash e : \bar{\sigma}_2}}$

Figure 17: Static Semantics for the Low-Level Language

Heap Typing

$$\frac{\text{HT-AX}}{D; \Psi \vdash \dots}$$

$$\frac{\text{HT-INF} \quad D; \Psi \vdash H : \Psi' \quad D; \Psi; \cdot \vdash v : \bar{\sigma}}{D; \Psi \vdash H, \ell \mapsto v, \alpha : \Psi', \ell : \bar{\sigma}, \alpha}$$

Alignment Subtyping

$$\frac{\text{ALST-OFF} \quad o_1 \equiv o_2 \pmod{a}}{\vdash [a, o_1] \leq [a, o_2]}$$

$$\frac{\text{ALST-BASE} \quad a_1 = a_2 \times k}{\vdash [a_1, o] \leq [a_2, o]}$$

$$\frac{\text{ALST-TRANS} \quad \vdash \alpha_1 \leq \alpha_2 \quad \vdash \alpha_2 \leq \alpha_3}{\vdash \alpha_1 \leq \alpha_3}$$

Physical Subtyping

$$\frac{\text{ST-UNROLL} \quad \Pi.xtype(\bar{t}, N) = \bar{\sigma}}{\Pi; \bar{t} \vdash \text{ptr}_\alpha(N) \leq \text{ptr}_\alpha(\bar{\sigma})}$$

$$\frac{\text{ST-ROLL} \quad \Pi.xtype(\bar{t}, N) = \bar{\sigma}}{\Pi; \bar{t} \vdash \text{ptr}_\alpha(\bar{\sigma}) \leq \text{ptr}_\alpha(N)}$$

$$\frac{\text{ST-PTR} \quad \vdash \alpha_1 \leq \alpha_2}{D \vdash \text{ptr}_{\alpha_1}(\bar{\sigma}_1 \bar{\sigma}_2) \leq \text{ptr}_{\alpha_2}(\bar{\sigma}_1)}$$

$$\frac{\text{ST-PAD} \quad \text{size}(\Pi, \sigma) = i}{\Pi; \bar{t} \vdash \sigma \leq \text{pad}[i]}$$

$$\frac{\text{ST-PADADD}}{\Pi; \bar{t} \vdash \text{pad}[i]\text{pad}[j] \leq \text{pad}[i+j]}$$

$$\frac{\text{ST-REFL}}{D \vdash \bar{\sigma} \leq \bar{\sigma}}$$

$$\frac{\text{ST-SEQ} \quad D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2 \quad D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_4}{D \vdash \bar{\sigma}_1 \bar{\sigma}_3 \leq \bar{\sigma}_2 \bar{\sigma}_4}$$

$$\frac{\text{ST-TRANS} \quad D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2 \quad D \vdash \bar{\sigma}_2 \leq \bar{\sigma}_3}{D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_3}$$

Value Typing

$$\frac{\text{LW-BYTE}}{D; \Psi \Vdash b : \text{byte}}$$

$$\frac{\text{LW-LBL} \quad \Psi(\ell) = \bar{\sigma}_1 \bar{\sigma}_2, [a, o] \quad i = \text{size}(\Pi, \bar{\sigma}_1)}{\Pi; \bar{t}; \Psi \Vdash \ell + i : \text{ptr}_{[a, o+i]}(\bar{\sigma}_2)}$$

$$\frac{\text{LW-UNINIT1}}{D; \Psi \Vdash \text{uninit} : \text{byte}}$$

$$\frac{\text{LW-UNINIT2} \quad \Pi.\text{ptrsize} = i}{\Pi; \bar{t}; \Psi \Vdash \text{uninit}^i : \text{ptr}_\alpha(\bar{\sigma})}$$

Figure 18: Auxiliary Relations

If σ is one of $(\text{byte}, \text{ptr}_\alpha(\bar{\sigma}))$, this is immediate from LW-UNINIT1 and LW-UNINIT2. For $\text{pad}[i]$ we can use LW-UNINIT1, ST-PAD, and ST-PADADD k times. Plugging (1,2) into L-VALUE yields the desired conclusion. Finally, if σ is $\text{ptr}_\alpha(N)$, we can use LW-UNINIT2 to derive $\Pi; \bar{t}; \Psi; \cdot \vdash \text{uninit}^i : \text{ptr}_\alpha(\bar{\sigma})$ where $\Pi.xtype(\bar{t}, N) = \bar{\sigma}$, then use ST-ROLL and L-SUB to assign type $\text{ptr}_\alpha(N)$ to uninit^i .

Lemma 16 (Constant-Size Subtyping)

If $D \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$ then $\text{size}(\Pi, \bar{\sigma}_1) = \text{size}(\Pi, \bar{\sigma}_2)$.

Proof: By induction on the assumed subtyping derivation, by cases on the last rule used.

- ST-PADADD and ST-REFL are immediate.
- ST-ROLL, ST-UNROLL, ST-PTR: Two pointers have the same size.
- ST-PAD: $\text{size}(\Pi, \sigma) = \text{size}(\Pi, \text{pad}[i]) = i$.

H-SHORT $\frac{}{\bar{t}; \Gamma \vdash i : \text{short}; \top}$	H-LONG $\frac{}{\bar{t}; \Gamma \vdash d : \text{long}; \top}$	H-NEW $\frac{}{\bar{t}; \Gamma \vdash \text{new } \tau : \tau*; \top}$	H-VAR{R,L} $\frac{\Gamma(x) = \tau}{\bar{t}; \Gamma \vdash x : \tau; \top}$
H-ASSN $\frac{\bar{t}; \Gamma \vdash e_1 : \tau; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2}{\bar{t}; \Gamma \vdash e_1 = e_2 : \tau; S_1 \wedge S_2}$	H-FETCHR $\frac{N\{\dots \tau f \dots\} \in \bar{t} \quad \bar{t}; \Gamma \vdash e : N; S}{\bar{t}; \Gamma \vdash e.f : \tau; S}$		
H-FETCHL $\frac{N\{\dots \tau f \dots\} \in \bar{t} \quad \bar{t}; \Gamma \vdash e : N; S}{\bar{t}; \Gamma \vdash e.f : \tau; S}$	H-SEQ $\frac{\bar{t}; \Gamma \vdash e_1 : \tau'; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2}{\bar{t}; \Gamma \vdash e_1; e_2 : \tau; S_1 \wedge S_2}$	H-DEREF{R,L} $\frac{\bar{t}; \Gamma \vdash e : \tau*; S}{\bar{t}; \Gamma \vdash *(\tau*)(e) : \tau; S}$	
H-CAST $\frac{\bar{t}; \Gamma \vdash e : \tau*; S}{\bar{t}; \Gamma \vdash (\tau')e : \tau'; S \wedge \text{subtype}(\text{xtype}(\bar{t}, \tau*), \text{xtype}(\bar{t}, \tau'))}$			
H-FADDR $\frac{N\{\dots \tau' f \dots\} \in \bar{t} \quad \bar{t}; \Gamma \vdash e : N*; S}{\bar{t}; \Gamma \vdash (\tau*)(\&e \rightarrow f) : \tau*; S \wedge \exists \bar{\sigma}_1, \bar{\sigma}_2, a, o \quad \begin{array}{l} \text{xtype}(\bar{t}, N*) = \text{ptr}_{[a,o]}(\bar{\sigma}_1 \bar{\sigma}_2) \\ \wedge \text{subtype}(\bar{t}, \text{ptr}_{[a,o+\text{offset}(\bar{t}, f)]}(\bar{\sigma}_2), \text{xtype}(\bar{t}, \tau*)) \\ \wedge \text{offset}(\bar{t}, f) = \text{size}(\bar{\sigma}_1) \end{array}}$			
H-IF $\frac{\bar{t}; \Gamma \vdash e_1 : \text{long}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2 \quad \bar{t}; \Gamma \vdash e_3 : \tau; S_3}{\bar{t}; \Gamma \vdash \text{if } e_1 \text{ } e_2 \text{ } e_3 : \tau; S_1 \wedge S_2 \wedge S_3}$	H-WHILE $\frac{\bar{t}; \Gamma \vdash e_1 : \text{long}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2}{\bar{t}; \Gamma \vdash \text{while } e_1 \text{ } e_2 : \text{short}; S_1 \wedge S_2}$		
H-DECL $\frac{\bar{t}; \Gamma, x : \tau_1 \vdash e : \tau_2; S}{\bar{t}; \Gamma \vdash \tau_1 x; e : \tau_2; S}$			

Figure 19: Static semantics for high level language

- **ST-SEQ:** By induction, $\text{size}(\Pi, \bar{\sigma}_1) = \text{size}(\Pi, \bar{\sigma}_2)$ and $\text{size}(\Pi, \bar{\sigma}_3) = \text{size}(\Pi, \bar{\sigma}_4)$. It follows that $\text{size}(\Pi, \sigma_1) + \text{size}(\Pi, \bar{\sigma}_3) = \text{size}(\Pi, \sigma_2) + \text{size}(\Pi, \bar{\sigma}_4)$.
- **ST-TRANS:** By induction twice, $\text{size}(\Pi, \bar{\sigma}_1) = \text{size}(\Pi, \bar{\sigma}_2) = \text{size}(\Pi, \bar{\sigma}_3)$.

Lemma 17 (Value-Type Size)

1. If $\Pi; \bar{t}; \Psi \Vdash \bar{w} : \bar{\sigma}$ then $\text{size}(\Pi, \bar{w}) = \text{size}(\Pi, \bar{\sigma})$.
2. If $\Pi; \bar{t}; \Psi; \Gamma \vdash \bar{w} : \bar{\sigma}$ then $\text{size}(\Pi, \bar{w}) = \text{size}(\Pi, \bar{\sigma})$.

Proof:

1. By inspection of the assumed typing derivation:

- **LW-BYTE:** $\text{size}(\Pi, b) = \text{size}(\Pi, \text{byte}) = 1$.
- **LW-LBL:** $\text{size}(\Pi, \ell+i) = \text{size}(\Pi, \text{ptr}_\alpha(\bar{\sigma}_2)) = \Pi.\text{ptrsize}$.

- Lw-UNINIT1: $\text{size}(\Pi, \text{uninit}) = \text{size}(\Pi, \text{byte}) = 1$.
 - Lw-UNINIT2: $\text{size}(\Pi, \text{uninit}^i) = \text{size}(\Pi, \text{ptr}_\alpha(\bar{\sigma})) = \Pi.\text{ptrsize} = i$.
2. By induction on the assumed typing derivation, by cases on the last rule used (all but 3 are impossible):
- L-VALUEEMPTY: trivial; both sizes are 0.
 - L-VALUE: we have $\Pi; \bar{t}; \Psi; \Gamma \vdash \bar{w}_1 \bar{w}_2 : \sigma_1 \bar{\sigma}_2$ where inversion gives $\Pi; \bar{t}; \Psi \Vdash \bar{w}_1 : \sigma_1$ and $\Pi; \bar{t}; \Psi; \Gamma \vdash \bar{w}_2 : \bar{\sigma}_2$. Part (1) of the theorem gives $\text{size}(\Pi, \bar{w}_1) = \text{size}(\Pi, \sigma_1)$. The induction hypothesis gives $\text{size}(\Pi, \bar{w}_2) = \text{size}(\Pi, \bar{\sigma}_2)$. It follows that $\text{size}(\Pi, \bar{w}_1 \bar{w}_2) = \text{size}(\Pi, \bar{\sigma}_1 \bar{\sigma}_2)$.
 - L-SUB: Follows from induction and the Constant-Size Subtyping Lemma.

Lemma 18 (Sequence Typing)

1. If $D; C \vdash \bar{w}_1 : \bar{\sigma}_1$ and $D; C \vdash \bar{w}_2 : \bar{\sigma}_2$, then $D; C \vdash \bar{w}_1 \bar{w}_2 : \bar{\sigma}_1 \bar{\sigma}_2$.
2. If $D; C \vdash \bar{w}_1 : \bar{\sigma}_1, \dots, D; C \vdash \bar{w}_n : \bar{\sigma}_n$, then $D; C \vdash \bar{w}_1 \dots \bar{w}_n : \bar{\sigma}_1 \dots \bar{\sigma}_n$.

Proof:

1. By induction on the derivation of $D; C \vdash \bar{w}_1 : \bar{\sigma}_1$, by cases on the last rule used (all but 3 are impossible):
 - L-VALUEEMPTY: trivial; the second assumption is what we need.
 - L-VALUE: Then $\bar{w}_1 = \bar{w}_a \bar{w}_b$ and induction (applied to \bar{w}_b and \bar{w}_2) and L-VALUE suffice.
 - L-SUB: By inversion $D; C \vdash \bar{w}_1 : \bar{\sigma}_3$ and $D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_1$. So by induction $D; C \vdash \bar{w}_1 \bar{w}_2 : \bar{\sigma}_3 \bar{\sigma}_2$. So with $D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_1$, ST-SEQ, ST-REFL, and L-SUB we can derive $D; C \vdash \bar{w}_1 \bar{w}_2 : \bar{\sigma}_1 \bar{\sigma}_2$.
2. By induction on n . Cases $n = 0$ and $n = 1$ are trivial. For $n > 1$, use induction and part(1).

Lemma 19 (Subtyping Partition)

If $D \vdash \bar{\sigma}_1 \leq \sigma_{21} \dots \sigma_{2n}$, then $\exists \bar{\sigma}_{11}, \dots, \bar{\sigma}_{1n}$ such that $\bar{\sigma}_1 = \bar{\sigma}_{11} \dots \bar{\sigma}_{1n}$ and for all $1 \leq i \leq n$, $D \vdash \bar{\sigma}_{1i} \leq \sigma_{2i}$.

Proof: By induction on the assumed subtyping derivation, by cases on the last rule used:

- ST-ROLL, ST-UNROLL, ST-PTR, ST-PAD, ST-PADADD: Immediate because $n = 1$ so $\bar{\sigma}_1 = \bar{\sigma}_{11}$.
- ST-REFL: Trivial, let $\bar{\sigma}_{1i} = \sigma_{2i}$.
- ST-SEQ: Follows from two invocations of the induction hypothesis and the *union* of the partitions they prove exist.
- ST-TRANS: Follows from two invocations of the induction hypothesis and *composing* the partitions they prove exist.

Lemma 20 (Subtyping Split)

If $D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_1 \bar{\sigma}_2$, then there exists $\bar{\sigma}'_1$ and $\bar{\sigma}'_2$ such that $\bar{\sigma}_3 = \bar{\sigma}'_1 \bar{\sigma}'_2$, $D \vdash \bar{\sigma}'_1 \leq \bar{\sigma}_1$, and $D \vdash \bar{\sigma}'_2 \leq \bar{\sigma}_2$.

Proof: Let $\bar{\sigma}_1 = \sigma_{11} \dots \sigma_{1n}$ and $\bar{\sigma}_2 = \sigma_{21} \dots \sigma_{2m}$. Then by the Subtyping Partition Lemma, there exist $\bar{\sigma}_{11}, \dots, \bar{\sigma}_{1n}, \bar{\sigma}_{21}, \dots, \bar{\sigma}_{2m}$ such that $\bar{\sigma}_3 = \bar{\sigma}_{11} \dots \bar{\sigma}_{1n} \bar{\sigma}_{21} \dots \bar{\sigma}_{2m}$, $D \vdash \bar{\sigma}_{11} \leq \sigma_{11}, \dots, D \vdash \bar{\sigma}_{1n} \leq \sigma_{1n}, D \vdash \bar{\sigma}_{21} \leq \sigma_{21}, \dots, D \vdash \bar{\sigma}_{2m} \leq \sigma_{2m}$. So letting $\bar{\sigma}'_1 = \bar{\sigma}_{11} \dots \bar{\sigma}_{1n}$ (respectively $\bar{\sigma}'_2 = \bar{\sigma}_{21} \dots \bar{\sigma}_{2m}$), we can use n (respectively m) uses of ST-SEQ and ST-TRANS to derive what we need.

Lemma 21 (Typing Partition)

If $D; C \vdash \bar{w} : \sigma_1 \dots \sigma_n$, then $\exists \bar{w}_1, \dots, \bar{w}_n$ such that $\bar{w} = \bar{w}_1 \dots \bar{w}_n$ and for all $1 \leq i \leq n$, $D; C \vdash \bar{w}_i : \sigma_i$.

Proof: By induction on the assumed typing derivation, by cases on the last rule used (all but 3 are impossible):

1. L-VALUEEMPTY: trivial because $n = 0$.
2. L-VALUE: Then $\bar{w} = \bar{w}_1\bar{w}_2$, $D; \Psi \vdash_{\bar{w}} \bar{w}_1 : \sigma_1$ follows from inversion, and the other results follow from inversion and induction.
3. L-SUB: Then by inversion there exists $\sigma'_1 \dots \sigma'_m$ such that $D; C \vdash_{\bar{r}} \bar{w} : \sigma'_1 \dots \sigma'_m$ and $D \vdash \sigma'_1 \dots \sigma'_m \leq \sigma_1 \dots \sigma_n$. So by induction there exists $\bar{w}'_1, \dots, \bar{w}'_m$ such that $\bar{w} = \bar{w}'_1 \dots \bar{w}'_m$ and $1 \leq i \leq m$, $D; C \vdash_{\bar{r}} \bar{w}'_i : \sigma'_i$. And by the Subtyping Partition Lemma, there exist $\bar{\sigma}'_1, \dots, \bar{\sigma}'_n$ such that $\sigma'_1 \dots \sigma'_m = \bar{\sigma}'_1 \dots \bar{\sigma}'_n$ and for all $1 \leq i \leq n$, $D \vdash \bar{\sigma}'_i \leq \sigma_i$. Therefore, we can use the $\sigma'_i \dots \sigma'_j$ that is $\bar{\sigma}'_k$ and the corresponding $\bar{w}'_i \dots \bar{w}'_j$ and use the Sequence Typing Lemma to conclude $D; C \vdash_{\bar{r}} \bar{w}'_i \dots \bar{w}'_j : \bar{\sigma}'_k$ and then L-SUB to conclude $D; C \vdash_{\bar{r}} \bar{w}'_i \dots \bar{w}'_j : \sigma_k$. So letting $\bar{w}'_i \dots \bar{w}'_j = \bar{w}_k$ suffices.

Lemma 22 (Value Split)

1. If $D; C \vdash_{\bar{r}} \bar{w} : \bar{\sigma}_1\bar{\sigma}_2$, then $\exists \bar{w}_1, \bar{w}_2$ such that $\bar{w} = \bar{w}_1\bar{w}_2$, $D; C \vdash_{\bar{r}} \bar{w}_1 : \bar{\sigma}_1$, and $D; C \vdash_{\bar{r}} \bar{w}_2 : \bar{\sigma}_2$.
2. If $D; C \vdash_{\bar{r}} \bar{w} : \bar{\sigma}_1\bar{\sigma}_2\bar{\sigma}_3$, then $\exists \bar{w}_1, \bar{w}_2, \bar{w}_3$ such that $\bar{w} = \bar{w}_1\bar{w}_2\bar{w}_3$, $D; C \vdash_{\bar{r}} \bar{w}_1 : \bar{\sigma}_1$, $D; C \vdash_{\bar{r}} \bar{w}_2 : \bar{\sigma}_2$, and $D; C \vdash_{\bar{r}} \bar{w}_3 : \bar{\sigma}_3$.

Proof:

1. Let $\bar{\sigma}_1 = \sigma_{11} \dots \sigma_{1n}$ and $\bar{\sigma}_2 = \sigma_{21} \dots \sigma_{2m}$. Then by the Typing Partition Lemma there exists $\bar{w}_{11}, \dots, \bar{w}_{1n}, \bar{w}_{21}, \dots, \bar{w}_{2m}$ such that $\bar{w} = \bar{w}_{11} \dots \bar{w}_{1n} \bar{w}_{21} \dots \bar{w}_{2m}$ and $D; C \vdash_{\bar{r}} \bar{w}_{11} : \sigma_{11}, \dots, D; C \vdash_{\bar{r}} \bar{w}_{1n} : \sigma_{1n}, D; C \vdash_{\bar{r}} \bar{w}_{21} : \sigma_{21}, \dots, D; C \vdash_{\bar{r}} \bar{w}_{2m} : \sigma_{2m}$. So letting $\bar{w}_1 = \bar{w}_{11} \dots \bar{w}_{1n}$ and $\bar{w}_2 = \bar{w}_{21} \dots \bar{w}_{2m}$, two uses of the Sequence Typing Lemma provide what we need.
2. Letting $\bar{\sigma}'_2 = \bar{\sigma}_2\bar{\sigma}_3$, part 1 ensures there exist \bar{w}_1 and \bar{w}'_2 such that $\bar{w} = \bar{w}_1\bar{w}'_2$, $D; C \vdash_{\bar{r}} \bar{w}_1 : \bar{\sigma}_1$, and $D; C \vdash_{\bar{r}} \bar{w}'_2 : \bar{\sigma}'_2$. Applying part 1 again to the last conclusion provides $D; C \vdash_{\bar{r}} \bar{w}_2 : \bar{\sigma}_2$ and $D; C \vdash_{\bar{r}} \bar{w}_3 : \bar{\sigma}_3$. (Note we can extend this to n by a trivial induction should we need it).

Lemma 23 (Value-Type Split) *If*

$$\begin{aligned} & \Pi; \bar{t}; C \vdash_{\bar{r}} \bar{w}_1\bar{w}_2 : \bar{\sigma}_1\bar{\sigma}_2 \\ & \text{size}(\Pi, \bar{w}_1) = \text{size}(\Pi, \bar{\sigma}_1) \end{aligned}$$

then

$$\begin{aligned} & \Pi; \bar{t}; C \vdash_{\bar{r}} \bar{w}_1 : \bar{\sigma}_1 \\ & \Pi; \bar{t}; C \vdash_{\bar{r}} \bar{w}_2 : \bar{\sigma}_2 \end{aligned}$$

Proof: If $\bar{w}_1 = \cdot$, then $\bar{\sigma}_1 = \cdot$ and the results follow from L-VALUEEMPTY and the assumption. So assume $\bar{w}_1 \neq \cdot$ and proceed by induction on the assumed typing derivation, by cases on the last rule applied (all but 2 are impossible):

- L-VALUE: By inversion and the assumptions, $\bar{w}_1 = \bar{w}_{11}\bar{w}_{12}$, $\bar{\sigma}_1 = \sigma_{11}\bar{\sigma}_{12}$,

$$\mathbf{1} \quad \Pi; \bar{t}; \Psi \vdash_{\bar{w}} \bar{w}_{11} : \sigma_{11}$$

$$\mathbf{2} \quad \Pi; \bar{t}; \Psi; \cdot \vdash_{\bar{r}} \bar{w}_{12}\bar{w}_2 : \bar{\sigma}_{12}\bar{\sigma}_2$$

From the assumption, we know $\text{size}(\Pi, \bar{w}_{11}\bar{w}_{12}) = \text{size}(\Pi, \sigma_{11}\bar{\sigma}_{12})$. By the Value-Type Size Lemma, we have $\text{size}(\Pi, \bar{w}_{11}) = \text{size}(\Pi, \sigma_{11})$. It follows that

$$\mathbf{3} \quad \text{size}(\Pi, \bar{w}_{12}) = \text{size}(\Pi, \bar{\sigma}_{12})$$

From (2,3), the induction hypothesis yields

4 $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_{12} : \bar{\sigma}_{12}$

5 $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_2 : \bar{\sigma}_2$

Plugging (1,4) into L-VALUE gives $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_{11}\bar{w}_{12} : \sigma_{11}\bar{\sigma}_{12}$, which, together with (5), gives the desired conclusion.

- L-SUB: By inversion there exists a $\bar{\sigma}_3$ such that $D; C \vdash \bar{w}_1\bar{w}_2 : \bar{\sigma}_3$ and $D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_1\bar{\sigma}_2$. So by the Subtyping Split Lemma, there exists $\bar{\sigma}'_1$ and $\bar{\sigma}'_2$ such that $\bar{\sigma}_3 = \bar{\sigma}'_1\bar{\sigma}'_2$, $D \vdash \bar{\sigma}'_1 \leq \bar{\sigma}_1$, and $D \vdash \bar{\sigma}'_2 \leq \bar{\sigma}_2$. By the Constant-Size Subtyping Lemma $\text{size}(\Pi, \bar{\sigma}'_1) = \text{size}(\Pi, \bar{\sigma}_1)$, so the assumption $\text{size}(\Pi, \bar{w}_1) = \text{size}(\Pi, \bar{\sigma}_1)$ ensures $\text{size}(\Pi, \bar{w}'_1) = \text{size}(\Pi, \bar{\sigma}'_1)$. So by induction $\Pi; \bar{t}; C \vdash \bar{w}_1 : \bar{\sigma}'_1$ and $\Pi; \bar{t}; C \vdash \bar{w}_2 : \bar{\sigma}'_2$. So with one use of L-SUB each, we can derive $\Pi; \bar{t}; C \vdash \bar{w}_1 : \bar{\sigma}_1$ and $\Pi; \bar{t}; C \vdash \bar{w}_2 : \bar{\sigma}_2$.

Lemma 24 (Subsequence Replacement)

1. If $D; C \vdash \bar{w}_1\bar{w}_2 : \bar{\sigma}_1\bar{\sigma}_2$ and $D; C \vdash \bar{w}_2 : \bar{\sigma}_2$ and $D; C \vdash \bar{w}_3 : \bar{\sigma}_2$ then $D; C \vdash \bar{w}_1\bar{w}_3 : \bar{\sigma}_1\bar{\sigma}_2$.
2. If $D; C \vdash \bar{w}_1\bar{w}_2 : \bar{\sigma}_1\bar{\sigma}_2$ and $D; C \vdash \bar{w}_1 : \bar{\sigma}_1$ and $D; C \vdash \bar{w}_3 : \bar{\sigma}_1$ then $D; C \vdash \bar{w}_3\bar{w}_2 : \bar{\sigma}_1\bar{\sigma}_2$.

Proof: Follows from the Value Split and Sequence Typing Lemmas.

Lemma 25 (Alignment Subtyping Reflexivity)

For all $[a, o]$, we can derive $\vdash [a, o] \leq [a, o]$.

Proof: Can be derived with either ALST-OFF ($o \equiv o \pmod{a}$) or ALST-BASE ($a = a \times 1$).

Lemma 26 (Alignment Subtyping Form)

If $\vdash [a, o + i] \leq \alpha$ then $\exists a', o'$ such that $\alpha = [a', o' + i]$.

Proof: By induction on the assumed derivation.

Lemma 27 (Alignment Addition)

$\vdash [a, o] \leq [a', o']$ if and only if $\vdash [a, o + k] \leq [a', o' + k]$.

Proof: Both directions proceed by induction on the assumed derivation, using the fact that $o \equiv o' \pmod{a}$ iff $o + k \equiv o' + k \pmod{a}$.

Lemma 28 (Subtyping Type Form)

- If $\Pi; \bar{t} \vdash \bar{\sigma}' \leq \text{ptr}_\alpha(\bar{\sigma})$ then $\exists \bar{\sigma}'', \alpha'$ such that $\bar{\sigma}' = \text{ptr}_{\alpha'}(\bar{\sigma}\bar{\sigma}'')$ or $\bar{\sigma}' = \text{ptr}_{\alpha'}(N)$ where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}\bar{\sigma}'$, and $\vdash \alpha' \leq \alpha$.
- If $\Pi; \bar{t} \vdash \bar{\sigma}' \leq \text{ptr}_\alpha(N)$ and $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}$ then $\exists \bar{\sigma}'', \alpha'$ such that $\bar{\sigma}' = \text{ptr}_{\alpha'}(\bar{\sigma}\bar{\sigma}'')$ or $\bar{\sigma}' = \text{ptr}_{\alpha'}(N')$ where $\Pi.\text{xtype}(\bar{t}, N') = \bar{\sigma}\bar{\sigma}'$, and $\vdash \alpha' \leq \alpha$.

Proof: By simultaneous induction on the assumed subtyping derivations, by cases on the last rule used. The cases ST-PAD and ST-PADADD cannot occur.

- ST-ROLL: $\bar{\sigma}' = \text{ptr}_\alpha(N)$ and inversion gives $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}$. We satisfy the property with $\bar{\sigma}'' = \cdot$ and $\alpha' = \alpha$.
- ST-UNROLL: $\bar{\sigma}' = \text{ptr}_\alpha(\bar{\sigma}\bar{\sigma}'')$ where inversion gives $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}$. We satisfy the property with $\bar{\sigma}'' = \cdot$ and $\alpha' = \alpha$.
- ST-PTR: Follows from inspection and inversion.
- ST-REFL: Immediate, with $\alpha' = \alpha$, $N' = N$, and $\bar{\sigma}'' = \cdot$.

- ST-SEQ: We have $\Pi; \bar{t} \vdash \bar{\sigma}_1 \bar{\sigma}_3 \leq \bar{\sigma}_2 \bar{\sigma}_4$. It must be the case that either $\bar{\sigma}_1 = \bar{\sigma}_2 = \cdot$ or $\bar{\sigma}_3 = \bar{\sigma}_4 = \cdot$, because pointer types cannot be broken into subsequences. The property follows by induction.
- ST-TRANS: We have $\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_3$ where inverting gives $\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$ and $\Pi; \bar{t} \vdash \bar{\sigma}_2 \leq \bar{\sigma}_3$. There are two possibilities for $\bar{\sigma}_3$:
 1. $\bar{\sigma}_3 = \text{ptr}_\alpha(N)$ where $\Pi.\text{xtyp}(\bar{t}, N) = \bar{\sigma}'$. By induction on the second subderivation, we get either
 - $\bar{\sigma}_2 = \text{ptr}_{\alpha'}(\bar{\sigma}'\bar{\sigma}'')$, where $\vdash \alpha' \leq \alpha$. By induction on the first subderivation, we have either $\bar{\sigma}_1 = \text{ptr}_{\alpha''}(\bar{\sigma}'\bar{\sigma}''\bar{\sigma}''')$ or $\bar{\sigma}_1 = \text{ptr}_{\alpha''}(N')$ where $\Pi.\text{xtyp}(\bar{t}, N') = \bar{\sigma}'\bar{\sigma}''\bar{\sigma}'''$ and $\vdash \alpha'' \leq \alpha'$, and by ALST-TRANS, $\vdash \alpha'' \leq \alpha$, which is what we want.
 - $\bar{\sigma}_2 = \text{ptr}_{\alpha'}(N')$ where $\Pi.\text{xtyp}(\bar{t}, N') = \bar{\sigma}'\bar{\sigma}''$ and $\vdash \alpha' \leq \alpha$. By induction on the first subderivation, we have either $\bar{\sigma}_1 = \text{ptr}_{\alpha''}(\bar{\sigma}'\bar{\sigma}''\bar{\sigma}''')$ or $\bar{\sigma}_1 = \text{ptr}_{\alpha''}(N)$ where $\Pi.\text{xtyp}(\bar{t}, N) = \bar{\sigma}'\bar{\sigma}''\bar{\sigma}'''$ and $\vdash \alpha'' \leq \alpha'$, and by ALST-TRANS, $\vdash \alpha'' \leq \alpha$, which is what we want.
 2. $\bar{\sigma}_3 = \text{ptr}_\alpha(\bar{\sigma}')$. Proceeds exactly as above.

Lemma 29 (Canonical Forms)

1. If $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \cdot$ then $\bar{w} = \cdot$.
2. If $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{byte}$ then $\bar{w} = b$ or $\bar{w} = \text{uninit}$.
3. If $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{byte}^k$ then $\bar{w} = w_1 \dots w_k$ where $w_i = b$ or $w_i = \text{uninit}$ for $i \in \{1, \dots, k\}$.
4.
 - If $\Pi; \bar{t}; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_\alpha(\bar{\sigma}_1)$ then $\exists \bar{\sigma}_0, \bar{\sigma}_2$ s.t. either
 - $\bar{w} = \text{uninit}^i$ where $\Pi.\text{ptrsize} = i$.
 - $\bar{w} = \ell+i$ and $\alpha = [a, o+i]$ and $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2, \alpha'$ where $\vdash \alpha' \leq [a, o]$ and $\text{size}(\Pi, \bar{\sigma}_0) = i$.
 - If $\Pi; \bar{t}; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_\alpha(N)$ where $\Pi.\text{xtyp}(\bar{t}, N) = \bar{\sigma}_1$, then $\exists \bar{\sigma}_0, \bar{\sigma}_2$ s.t. either
 - $\bar{w} = \text{uninit}^i$ where $\Pi.\text{ptrsize} = i$.
 - $\bar{w} = \ell+i$ and $\alpha = [a, o+i]$ and $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2, \alpha'$ where $\vdash \alpha' \leq [a, o]$ and $\text{size}(\Pi, \bar{\sigma}_0) = i$.

Proof:

1. By induction on the typing derivation. Last rule applied is either L-VALUEEMPTY or L-SUB. The former follows immediately. If the latter, the only subtype of \cdot is \cdot , so the property follows from induction.
2. By induction on the derivation. Last rule applied is either L-VALUE or L-SUB. In the former case, by inversion we have $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{byte}$. Only the LW-BYTE or LW-UNINIT1 can apply, meaning either $\bar{w} = b$ or $\bar{w} = \text{uninit}$. In the latter case, byte has no subtypes other than itself, so the property follows from induction.
3. By induction on the type derivation, by cases on the last rule used. If the last rule is L-VALUE, we have $\bar{w} = w\bar{w}'$ and inversion gives $D; \Psi \vdash_{\bar{\tau}} w : \text{byte}$ and $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w}' : \text{byte}^{k-1}$. By induction, the property holds for the latter derivation. By part (2) of the theorem, the property holds for the former derivation. If the last rule applied is L-SUB, the property holds by induction, since byte^i has no subtypes other than itself.
4. By simultaneous induction on the assumed typing derivations, by cases on the last rule applied. Reflexivity of alignment subtyping (Lemma 25) is used implicitly. There are only two possibilities:
 - L-VALUE: There are two sub-cases:
 - Assume $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_\alpha(\bar{\sigma}_1)$ Inversion gives $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_\alpha(\bar{\sigma}_1)$. Quick inspection reveals that only the rules LW-LBL or LW-UNINIT2 can lead to this conclusion. In the latter case, clearly $\bar{w} = \text{uninit}^i$. In the former case, inversion yields the desired result, with $\bar{\sigma}_2 = \cdot$ and $\alpha = [a, o+i]$ and $\alpha' = [a, o]$.

- Assume $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_{[a, o+i]}(N)$, where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}_1$. Vacuous; cannot be derived via L-VALUE.
- L-SUB: There are two sub-cases:
 - Assume $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_{\alpha}(\bar{\sigma}_1)$. Inversion gives $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \bar{\sigma}$ and $D \vdash \bar{\sigma} \leq \text{ptr}_{\alpha}(\bar{\sigma}_1)$. By the Subtyping Type Form Lemma, $\bar{\sigma}$ is one of
 - * $\text{ptr}_{\alpha'}(N)$ where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}_1 \bar{\sigma}_k$ for some $\bar{\sigma}_k$, and $\vdash \alpha' \leq \alpha$. Applying the IH to $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_{\alpha'}(N)$ we get that either $\bar{w} = \text{uninit}^i$ or $\bar{w} = \ell+i$ and $\alpha' = [a, o+i]$ and $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_k \bar{\sigma}_2, \alpha''$ where $\vdash \alpha'' \leq [a, o]$ and $\text{size}(\Pi, \bar{\sigma}_0) = i$. It remains to be shown that α is of the form $[a', o'+i]$ and $\vdash \alpha'' \leq [a', o']$. We know $\vdash [a, o+i] \leq \alpha$. By the Alignment Subtyping Form Lemma, $\alpha = [a', o'+i]$ for some a' and o' . By the Alignment Addition Lemma, $\vdash [a, o] \leq [a', o']$. We also know $\vdash \alpha'' \leq [a, o]$, so by ALST-TRANS, $\vdash \alpha'' \leq [a', o']$, as desired.
 - * $\text{ptr}_{\alpha}(\bar{\sigma}_1 \bar{\sigma}_k)$ for some $\bar{\sigma}_k$. Proceeds like the above case.
 - Assume $D; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \text{ptr}_{\alpha}(N)$, where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}_1$. Similar to the above case.

Lemma 30 (Context Reordering)

If $D; \Gamma, x_1 : \tau_1, x_2 : \tau_2 \vdash_{\bar{\tau}} e : \bar{\sigma}$ then $D; \Gamma, x_2 : \tau_2, x_1 : \tau_1 \vdash_{\bar{\tau}} e : \bar{\sigma}$.

Proof: Straightforward induction.

Lemma 31 (Substitution Preserves Types)

Suppose $\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}_2$, $\Pi.\text{align}(\bar{t}, \tau) = \alpha$, $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} e_2 : \bar{\sigma}_2$, and $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} e_2 : \bar{\sigma}_2, \alpha$.

1. If $\Pi; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\bar{\tau}} e_1 : \bar{\sigma}_1$, then $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} e_1\{e_2/x\} : \bar{\sigma}_1$.
2. If $\Pi; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\bar{\tau}} e_1 : \bar{\sigma}_1, \alpha$, then $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} e_1\{e_2/x\} : \bar{\sigma}_1, \alpha$.

Proof: By simultaneous induction the assumed typing derivations, by cases on the last rule used.

- L-SHORT, L-LONG, L-NEW, L-VALUEEMPTY, L-VALUE: trivial because substitution and x are irrelevant.
- L-VARR: We assume $\Pi; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\bar{\tau}} y : \bar{\sigma}_1$ and it follows that $\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}_1$. If $x = y$ then $\bar{\sigma}_1 = \bar{\sigma}_2$ and $y\{e_2/x\} = e_2$, and so $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} e_2 : \bar{\sigma}_1$ by assumption, as desired. If $x \neq y$, then $y\{e_2/x\} = y$ and we can drop x from the context and still derive $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} y : \bar{\sigma}_1$.
- L-VARL: We assume $\Pi; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\bar{\tau}} y : \bar{\sigma}_1, \alpha$ and it follows that $\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}_1$ and $\Pi.\text{align}(\bar{t}, \tau) = \alpha$. If $x = y$ then $\bar{\sigma}_1 = \bar{\sigma}_2$ and $y\{e_2/x\} = e_2$, and so $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} e_2 : \bar{\sigma}_1, \alpha$ by assumption, as desired. If $x \neq y$, then $y\{e_2/x\} = y$ and we can drop x from the context and still derive $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\bar{\tau}} y : \bar{\sigma}_1, \alpha$.
- L-ASSN: We assume $\Pi; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\bar{\tau}} e = e' : \bar{\sigma}_1$. By inversion and induction (once for left-typing and once for right-typing),
 - 1 $\Pi; \bar{t}; \Gamma \vdash_{\bar{\tau}} e\{e_2/x\} : \bar{\sigma}_1, \alpha$
 - 2 $\Pi; \bar{t}; \Gamma \vdash_{\bar{\tau}} e'\{e_2/x\} : \bar{\sigma}_1$

The desired property follows from an application of L-ASSN to these facts and from the definition of substitution.

- L-FETCH{L,R}, L-DEREF{L,R}, L-CAST, L-FADDR, L-SEQ, L-IF, L-WHILE: Inversion, induction applied to each subexpression, and plugging the results back into the rule.
- L-DECL: We assume $\Pi; \bar{t}; \Gamma, x : \tau \vdash_{\bar{\tau}} \tau' y; e : \bar{\sigma}_1$. By inversion, $\Pi; \bar{t}; \Gamma, x : \tau, y : \tau' \vdash_{\bar{\tau}} e : \bar{\sigma}_1$. By the Context Reordering Lemma, $\Pi; \bar{t}; \Gamma, y : \tau', x : \tau \vdash_{\bar{\tau}} e : \bar{\sigma}_1$. By Weakening, $\Pi; \bar{t}; \Gamma, y : \tau' \vdash_{\bar{\tau}} e_2 : \bar{\sigma}_2$. By induction, $\Pi; \bar{t}; \Gamma, y : \tau' \vdash_{\bar{\tau}} e\{e_2/x\} : \bar{\sigma}_1$. By D-DECL, $\Pi; \bar{t}; \Gamma \vdash_{\bar{\tau}} \tau' y; e\{e_2/x\} : \bar{\sigma}_1$.

- L-SUB: Inversion, induction, and an application of L-SUB.

Lemma 32 (Subject Reduction)

Suppose Π is a sensible platform. If

$$\begin{aligned} & \Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma} \text{ (or } \Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma}, \alpha) \\ & \Pi; \bar{t} \vdash H; e \xrightarrow{\cdot} H'; e' \text{ (or } \Pi; \bar{t} \vdash H; e \xrightarrow{\cdot} H'; e') \\ & \Pi; \bar{t}; \Psi \vdash H : \Psi \end{aligned}$$

then there exists Ψ' , extending Ψ , such that

$$\begin{aligned} & \Pi; \bar{t}; \Psi'; \cdot \vdash e' : \bar{\sigma} \text{ (or } \Pi; \bar{t}; \Psi'; \cdot \vdash e' : \bar{\sigma}, \alpha) \\ & \Pi; \bar{t}; \Psi' \vdash H' : \Psi' \end{aligned}$$

Proof: Proof proceeds by induction on the assumed typing derivation. The only inductive case is L-SUB. We assume $\Psi = \Psi'$ and $H = H'$ unless otherwise stated.

- L-VAR{L,R}, L-VALUEEMPTY, L-VALUE: Vacuous; variables and values cannot step.
- L-CAST: The step must have been derived by D-CAST and the result follows from inversion.
- L-SHORT: The step must have been derived by D-SHORT and we have $e' = \bar{b}$ where $\Pi.xlit(s) = \bar{b}$. From L-SHORT, $\bar{\sigma} = \text{byte}^i$. Because Π is sensible (clause 2), we know $\text{size}(\Pi, \bar{b}) = i$. We can apply LW-BYTE and L-VALUE i times to derive $\Pi; \bar{t}; \Psi; \cdot \vdash e' : \bar{\sigma}$.
- L-LONG: Similar to L-SHORT, where the step was derived by D-LONG.
- L-SEQ: The step was derived using D-SEQ and the result follows from inversion.
- L-IF: The step was derived using D-IFT or D-IFF and the result follows from inversion.
- L-WHILE: The step was derived using D-WHILE and we have $\Pi; \bar{t}; \Psi; \cdot \vdash \text{while } e_1 e_2 : \text{byte}^i$ and $e' = \text{if } e_1 (e_2; \text{while } e_1 e_2) s$. Inversion gives

- 1 $\Pi; \bar{t}; \Psi; \cdot \vdash e_1 : \text{byte}^j$
- 2 $\Pi; \bar{t}; \Psi; \cdot \vdash e_2 : \bar{\sigma}_2$
- 3 $\Pi.xtext(\bar{t}, \text{short}) = \text{byte}^i$

From the assumed typing and (2), L-SEQ gives

- 4 $\Pi; \bar{t}; \Psi; \cdot \vdash (e_2; \text{while } e_1 e_2) : \text{byte}^i$

From L-SHORT, we have

- 5 $\Pi; \bar{t}; \Psi; \cdot \vdash s : \text{byte}^i$

Plugging (1,4,5) into L-IF yields $\Pi; \bar{t}; \Psi; \cdot \vdash e' : \text{byte}^i$.

- L-NEW: The step was derived using D-NEW. We assume $\Pi; \bar{t}; \Psi; \cdot \vdash \text{new } \tau : \text{ptr}_\alpha(\bar{\sigma})$ and have $e' = \ell+0$ and $H' = H, \ell \mapsto \text{uninit}^i, \alpha$. Let $\Psi' = \Psi, \ell : \bar{\sigma}, \alpha$, so

- 1 $\Psi'(\ell) = \bar{\sigma}, \alpha$

From the side conditions of the assumed reduction, we have

- 2 $\Pi.xtext(\bar{t}, \tau) = \bar{\sigma}$
- 3 $\text{size}(\Pi, \bar{\sigma}) = i$
- 4 $\Pi.align(\bar{t}, \tau) = \alpha$

Note that the α and $\bar{\sigma}$ here are the same as those mentioned in the assumption, from inversion on L-NEW. From (1,3), LW-LBL followed by L-VALUE give $\Pi; \bar{t}; \Psi'; \cdot \vdash e' : \text{ptr}_\alpha(\bar{\sigma})$, which satisfies the first part of the claim.

To satisfy the second part of the claim, we first apply the Heap Weakening Lemma to the third assumption to get

5 $\Pi; \bar{t}; \Psi' \vdash H : \Psi$

From (3) and the Uninit Type Lemma, we have

6 $\Pi; \bar{t}; \Psi'; \cdot \vdash \text{uninit}^i : \bar{\sigma}$ (6)

Plugging (5,6) into the HT-INF rule yields $\Pi; \bar{t}; \Psi' \vdash H' : \Psi'$.

- L-DECL: The step was derived using D-DECL. We assume $\Pi; \bar{t}; \Psi; \cdot \vdash \tau x; e : \bar{\sigma}_1$, and we have $e' = e\{*(\tau*)(\ell+0)/x\}$ and $H' = H, \ell \mapsto \text{uninit}^i, \alpha$, where the side conditions of the assumed reduction give

1 $\Pi.\text{ctype}(\bar{t}, \tau) = \bar{\sigma}$

2 $\Pi.\text{align}(\bar{t}, \tau) = \alpha$

By inversion on the assumed typing, we have

3 $\Pi; \bar{t}; \Psi; \cdot, x : \tau \vdash e : \bar{\sigma}_1$

Let $\Psi' = \Psi, \ell : \bar{\sigma}, \alpha$, so

4 $\Psi'(\ell) = \bar{\sigma}, \alpha$

Because the platform is sensible (clause 5), we know $\Pi.\text{access}(\Pi.\text{align}(\bar{t}, \tau), \text{size}(\Pi, \Pi.\text{ctype}(\bar{t}, \tau)))$. In other words, using (1,2),

5 $\Pi.\text{access}(\alpha, \text{size}(\Pi, \bar{\sigma}))$

With (1,4,5), we can perform the following forward derivation:

$$\frac{\frac{\frac{\Psi'(\ell) = \bar{\sigma}, \alpha}{\Pi; \bar{t}; \Psi'; \cdot \vdash \ell+0 : \text{ptr}_\alpha(\bar{\sigma})} \quad \Pi; \bar{t}; \Psi'; \cdot \vdash \cdot : \cdot}{\Pi; \bar{t}; \Psi'; \cdot \vdash \ell+0 : \text{ptr}_\alpha(\bar{\sigma})} \quad \Pi.\text{ctype}(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.\text{access}(\alpha, \text{size}(\Pi, \bar{\sigma}))}{\Pi; \bar{t}; \Psi'; \cdot \vdash *(\tau*)(\ell+0) : \bar{\sigma}}$$

Notice that under the same assumptions, we can also derive $\Pi; \bar{t}; \Psi'; \cdot \vdash *(\tau*)(\ell+0) : \bar{\sigma}, \alpha$. (L-DEREFR and L-DEREF have identical hypotheses.) With this, we can apply Substitution Preserves Types Lemma to the above conclusion and (1,2,3) to get $\Pi; \bar{t}; \Psi; \cdot \vdash e' : \bar{\sigma}_1$. The demonstration of the second part of the claim is identical to the D-NEW case.

- L-FETCHL: The step was derived using D-FETCHL; we assume $\Pi; \bar{t}; \Psi; \cdot \vdash *(\tau_1*)(\ell+j).f : \bar{\sigma}_2, [a, o + j']$ and have $e' = *(\tau*)(\ell+(j+j'))$ where

1 $\Pi.\text{offset}(\bar{t}, f) = j'$

By inversion we obtain

2 $\Pi; \bar{t}; \Psi; \cdot \vdash *(\tau_1*)(\ell+j) : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a, o]$

3 $\Pi.\text{offset}(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_1)$

4 $N\{\dots \tau f \dots\} \in \bar{t}$

$$5 \quad \Pi.xtype(\bar{t}, \tau) = \bar{\sigma}_2$$

$$6 \quad \Pi.align(\bar{t}, N) = [a_N, o_N]$$

$$7 \quad \vdash [a, o] \leq [a_N, o_N]$$

Inverting (2) (note this is a left-typing so there is no subsumption) gives

$$8 \quad \Pi; \bar{t}; \Psi; \cdot \vdash \ell + j : \text{ptr}_{[a, o]}(\bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4)$$

Canonical Forms gives

$$9 \quad \Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5, [a', o']$$

$$10 \quad \text{size}(\Pi, \bar{\sigma}_0) = j$$

$$11 \quad o = o'' + j$$

$$12 \quad \vdash [a', o'] \leq [a, o'']$$

From (1,3), we have $j' = \text{size}(\Pi, \bar{\sigma}_1)$. From this and (10), it follows that $j + j' = \text{size}(\Pi, \bar{\sigma}_0 \bar{\sigma}_1)$. Plugging this along with (9) into Lw-LBL, and the result into L-VALUE, gives

$$13 \quad \Pi; \bar{t}; \Psi; \cdot \vdash \ell + (j + j') : \text{ptr}_{[a', o'' + j + j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5)$$

From the Alignment Addition Lemma on (12), we get $\vdash [a', o'' + j + j'] \leq [a, o'' + j + j']$. Plugging this into ST-PTR, we get $\Pi; \bar{t} \vdash \text{ptr}_{[a', o'' + j + j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5) \leq \text{ptr}_{[a, o'' + j + j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5)$. Plugging this and (13) into L-SUB, we get

$$14 \quad \Pi; \bar{t}; \Psi; \cdot \vdash \ell + (j + j') : \text{ptr}_{[a, o'' + j + j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5).$$

It remains to be shown that $\Pi.access([a, o'' + j + j'], \text{size}(\Pi, \bar{\sigma}_2))$. This, plugged into L-DEREF together with (5,14) would yield $\Pi; \bar{t}; \Psi; \cdot \vdash *(\tau*)(\ell + (j + j')) : \bar{\sigma}_2, [a, o'' + j + j']$, as desired.

To prove this, suppose $\Pi.align(\bar{t}, \tau) = \alpha$. By clause (1) of platform sensibility we have $\Pi.access(\alpha, \bar{\sigma}_2)$. By clause (3), we have $\vdash [a_N, o_N + j'] \leq \alpha$ (where $[a_N, o_N]$ is the alignment of N as established in (7)). By clause (5), we then have $\Pi.access([a_N, o_N + j'], \text{size}(\Pi, \bar{\sigma}_2))$. From (7) and the Alignment Addition Lemma, we get $\vdash [a, o + j'] \leq [a_N, o_N + j']$, which, with (11), can be rewritten as $\vdash [a, o'' + j + j'] \leq [a_N, o_N + j']$. Again, by clause (5) of platform sensibility, we have $\Pi.access([a, o'' + j + j'], \text{size}(\Pi, \bar{\sigma}_2))$, which concludes the proof case.

- L-FETCHR: The step was derived using D-FETCH, so we assume $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_1 \bar{w}_2 \bar{w}_3.f : \bar{\sigma}_2$ and $e' = \bar{w}_2$. The step side condition gives

$$1 \quad \Pi.offset(\bar{t}, f) = \text{size}(\Pi, \bar{w}_1)$$

$$2 \quad \text{size}(\Pi, \bar{\sigma}_2) = \text{size}(\Pi, \bar{w}_2)$$

By inversion on the assumed typing we have

$$3 \quad \Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_1 \bar{w}_2 \bar{w}_3 : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$$

$$4 \quad \Pi.offset(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_1)$$

Note that the metavariable $\bar{\sigma}_2$ is bound to the same type sequence in both the step side condition and inversion on the assumed typing, because $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}_2$ is given by both.

Applying the Value-Type Split Lemma to (1,3) gives

$$5 \quad \Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_2 \bar{w}_3 : \bar{\sigma}_2 \bar{\sigma}_3$$

Applying the Value-Type Split Lemma to (2,5) yields $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_2 : \bar{\sigma}_2$.

- L-FADDR: The step was derived using D-FADDR; we assume $\Pi; \bar{t}; \Psi; \cdot \vdash (\tau^*) \& \ell + j \rightarrow f : \text{ptr}_{[a, o_2]}(\bar{\sigma}_2)$ and $e' = \ell + (j + \Pi.\text{offset}(\bar{t}, f))$. By inversion on the assumed typing, we get

- 1 $\Pi; \bar{t}; \Psi; \cdot \vdash \ell + j : \text{ptr}_{[a, o_1]}(\bar{\sigma}_1 \bar{\sigma}_2)$
- 2 $\Pi.\text{offset}(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_1)$
- 3 $o_2 = o_1 + \Pi.\text{offset}(\bar{t}, f) = o_1 + \text{size}(\Pi, \bar{\sigma}_1)$

Canonical Forms gives

- 4 $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a', o'_0]$
- 5 $o_1 = o_0 + j$
- 6 $j = \text{size}(\Pi, \bar{\sigma}_0)$
- 7 $\vdash [a', o'_0] \leq [a, o_0]$

Adding j to both sides of (2) and making use of (6), we obtain $j + \Pi.\text{offset}(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_0 \bar{\sigma}_1)$. We can plug this fact along with (4) into Lw-LBL to get

$$\Pi; \bar{t}; \Psi; \cdot \vdash \ell + (j + \Pi.\text{offset}(\bar{t}, f)) : \text{ptr}_{[a', o'_0 + j + \Pi.\text{offset}(\bar{t}, f)]}(\bar{\sigma}_2 \bar{\sigma}_3)$$

By the Alignment Addition Lemma on (7), $\vdash [a', o'_0 + j + \Pi.\text{offset}(\bar{t}, f)] \leq [a, o_0 + j + \Pi.\text{offset}(\bar{t}, f)]$. Plugging this into ST-ALN1 and the result into L-SUB, we get

$$\Pi; \bar{t}; \Psi; \cdot \vdash \ell + (j + \Pi.\text{offset}(\bar{t}, f)) : \text{ptr}_{[a, o_0 + j + \Pi.\text{offset}(\bar{t}, f)]}(\bar{\sigma}_2 \bar{\sigma}_3)$$

Notice however, by way of (3,5), that $o_2 = o_1 + \Pi.\text{offset}(\bar{t}, f) = o_0 + j + \Pi.\text{offset}(\bar{t}, f)$, so we get

$$\Pi; \bar{t}; \Psi; \cdot \vdash \ell + (j + \Pi.\text{offset}(\bar{t}, f)) : \text{ptr}_{[a, o_2]}(\bar{\sigma}_2 \bar{\sigma}_3)$$

A final application of L-SUB to this and $\Pi; \bar{t} \vdash \text{ptr}_{[a, o_2]}(\bar{\sigma}_2 \bar{\sigma}_3) \leq \text{ptr}_{[a, o_2]}(\bar{\sigma}_2)$ (derivable via ST-Ptr) yields the desired result.

- L-DEREF: Vacuous; no primitive step rule can apply.
- L-DEREF: The step was derived using D-DEREF. We assume $\Pi; \bar{t}; \Psi; \cdot \vdash *(\tau^*)(\ell + j) : \bar{\sigma}_1$ and $e' = \bar{w}_2$. By inversion on L-DEREF, we obtain

- 1 $\Pi; \bar{t}; \Psi; \cdot \vdash \ell + j : \text{ptr}_\alpha(\bar{\sigma}_1 \bar{\sigma}_2)$
- 2 $\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}_1$

The D-DEREF side condition gives

- 3 $H(\ell) = \bar{w}_1 \bar{w}_2 \bar{w}_3, [a', o']$
- 4 $\text{size}(\Pi, \bar{w}_1) = j$
- 5 $\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}_1$
- 6 $\text{size}(\Pi, \bar{w}_2) = \text{size}(\Pi, \bar{\sigma}_1) = k$
- 7 $\Pi.\text{access}([a', o' + j], k)$

From (1), Canonical Forms gives

- 8 $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a', o']$
- 9 $j = \text{size}(\Pi, \bar{\sigma}_0)$
- 10 $\alpha = [a, o + i]$

11 $\vdash [a', o'] \leq [a, o]$

From the assumption $\Pi; \bar{t}; \Psi \vdash H : \Psi$ and (8), the Heap Canonical Forms Lemma gives $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_1 \bar{w}_2 \bar{w}_3 : \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$. Since $\text{size}(\Pi, \bar{w}_1) = \text{size}(\Pi, \bar{\sigma}_0) = j$ (by (4,9)), the Value-Type Split Lemma gives $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_2 \bar{w}_3 : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$. Since $\text{size}(\Pi, \bar{w}_2) = \text{size}(\Pi, \bar{\sigma}_1)$ (by (6)), the same lemma gives $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_2 : \bar{\sigma}_1$, as desired.

- L-ASSN: The step must have been derived using D-ASSN. We assume $\Pi; \bar{t}; \Psi; \cdot \vdash *(\tau*)(\ell+j) = \bar{w} : \bar{\sigma}_1$ and $e' = \bar{w}$, where $H(\ell) = \bar{w}_1 \bar{w}_2 \bar{w}_3, [a, o]$. By inversion on L-ASSN, we get $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w} : \bar{\sigma}_1$, which satisfies the first part of the claim.

Identically to the L-FETCHR case, we use the Canonical Forms, Heap Canonical Forms, and Value-Type Size Lemmas to establish that $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_2 : \bar{\sigma}_1$ and $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a, o]$. Using the Subsequence Replacement and Value-Type Split Lemmas, we get $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_1 \bar{w}_2 : \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$. Plug this and the assumption $\Pi; \bar{t}; \Psi \vdash H : \Psi$ into HT-INF to obtain $\Pi; \bar{t}; \Psi \vdash H' : \Psi'$, where $\Psi' = \Psi, \ell \mapsto \Psi(\ell)$. Using Weakening, we obtain $\Pi; \bar{t}; \Psi' \vdash H' : \Psi'$, as desired.

- L-SUB: Follows from inversion, induction, and an application of L-SUB.

Lemma 33 (Replacement)

- If $D; \Psi; \cdot \vdash_{\text{r}} R[e]_r : \bar{\sigma}$, then $\exists \bar{\sigma}'$ s.t. $D; \Psi; \cdot \vdash e : \bar{\sigma}'$ and if $D; \Psi; \cdot \vdash e' : \bar{\sigma}'$ then $D; \Psi; \cdot \vdash_{\text{r}} R[e']_r : \bar{\sigma}$.
- If $D; \Psi; \cdot \vdash_{\text{l}} L[e]_r : \bar{\sigma}, \alpha$, then $\exists \bar{\sigma}'$ s.t. $D; \Psi; \cdot \vdash e : \bar{\sigma}'$ and if $D; \Psi; \cdot \vdash e' : \bar{\sigma}'$ then $D; \Psi; \cdot \vdash_{\text{l}} L[e']_r : \bar{\sigma}, \alpha$.
- If $D; \Psi; \cdot \vdash_{\text{r}} R[e]_l : \bar{\sigma}$, then $\exists \bar{\sigma}'$ s.t. $D; \Psi; \cdot \vdash e : \bar{\sigma}', \alpha$ and if $D; \Psi; \cdot \vdash e' : \bar{\sigma}', \alpha$ then $D; \Psi; \cdot \vdash_{\text{r}} R[e']_l : \bar{\sigma}$.
- If $D; \Psi; \cdot \vdash_{\text{l}} L[e]_l : \bar{\sigma}, \alpha$, then $\exists \bar{\sigma}'$ s.t. $D; \Psi; \cdot \vdash e : \bar{\sigma}', \alpha$ and if $D; \Psi; \cdot \vdash e' : \bar{\sigma}', \alpha$ then $D; \Psi; \cdot \vdash_{\text{l}} L[e']_l : \bar{\sigma}, \alpha$.

Proof: Proof proceeds by simultaneous structural induction on R and L, by cases on their top-level syntactic form. In the base cases we have $R = [\cdot]_{\text{R}}$ or $L = [\cdot]_{\text{L}}$ and the property follows immediately. All inductive cases follow directly from the inductive hypothesis and typing rules. For example, consider the case when $R = (L = e_2)$. We assume $D; \Psi; \cdot \vdash_{\text{r}} L[e]_r = e_2 : \bar{\sigma}$ and proceed by cases on the last rule applied in order to reach this conclusion. All but two are impossible:

- L-ASSN: By inversion we get

- 1 $D; \Psi; \cdot \vdash_{\text{l}} L[e]_r : \bar{\sigma}, \alpha$
- 2 $D; \Psi; \cdot \vdash e_2 : \bar{\sigma}$

By induction hypothesis on (1) we have that there exists a $\bar{\sigma}'$ s.t. $D; \Psi; \cdot \vdash e : \bar{\sigma}'$, and also that $D; \Psi; \cdot \vdash_{\text{l}} L[e']_r : \bar{\sigma}, \alpha$. Plugging this along with (2) into L-ASSN, we get $D; \Psi; \cdot \vdash_{\text{r}} L[e']_r = e_2 : \bar{\sigma}$, which is the same as $D; \Psi; \cdot \vdash_{\text{r}} R[e']_r : \bar{\sigma}$.

- L-SUB: Follows from inversion, induction, and an application of L-SUB.

All other cases follow an identical pattern.

Lemma 34 (Preservation) If

- $$\begin{aligned} & \Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma} \text{ (or } \Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma}, \alpha) \\ & \Pi; \bar{t}; \Psi \vdash H : \Psi \\ & \Pi; \bar{t} \vdash H; e \rightarrow H'; e' \end{aligned}$$

then there exist Ψ' , extending Ψ , such that

- $$\begin{aligned} & \Pi; \bar{t}; \Psi'; \cdot \vdash e' : \bar{\sigma} \text{ (or } \Pi; \bar{t}; \Psi'; \cdot \vdash e' : \bar{\sigma}, \alpha) \\ & \Pi; \bar{t}; \Psi' \vdash H' : \Psi' \end{aligned}$$

Proof: Since it is a right-expression, e is of the form $R[e_0]_r$ or $R[e_0]_l$. Consider the former case. It must be the case that the step is of the form $\Pi; \bar{t} \vdash R[e_0]_r \rightarrow R[e'_0]_r$ (the D-STEP_R rule) which can only happen under condition $\Pi; \bar{t} \vdash e_0 \xrightarrow{\cdot} e'_0$. Then:

- Replacement on the assumed typing derivation gives $\exists \bar{\sigma}'$ s.t. $\Pi; \bar{t}; \Psi; \cdot \vdash e_0 : \bar{\sigma}'$.
- Subject Reduction gives $\exists \Psi'$ extending Ψ such that $\Pi; \bar{t}; \Psi'; \cdot \vdash e'_0 : \bar{\sigma}'$.
- Weakening gives $\Pi; \bar{t}; \Psi'; \cdot \vdash R[e_0]_r : \bar{\sigma}$.
- Replacement gives $\Pi; \bar{t}; \Psi'; \cdot \vdash R[e'_0]_r : \bar{\sigma}$.

The case when e is of the form $R[e_0]_l$, proceeds similarly, using the D-STEP_L rule instead.

Lemma 35 (Progress)

Suppose $\Pi; \bar{t}; \Psi \vdash H : \Psi$ and Π is a sensible platform.

- If $\Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma}$ then one of the following is true:
 - e is of the form \bar{w} – an r -value.
 - e is legally stuck on Π and \bar{t} .
 - e can step. I.e., there exist R, e_1, e_2 , and H' , such that
 - * $e = R[e_1]_r$, and $\Pi; \bar{t} \vdash H; e_1 \xrightarrow{\cdot} H'; e_2$, or
 - * $e = R[e_1]_l$, and $\Pi; \bar{t} \vdash H; e_1 \xrightarrow{\cdot} H'; e_2$
- If $\Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma}, \alpha$ then one of the following is true:
 - e is of the form $*(\tau*)(\ell+i)$ – an l -value.
 - e is legally stuck on Π and \bar{t} .
 - e can step. I.e., there exist L, e_1, e_2 , and H' , such that
 - * $e = L[e_1]_r$, and $\Pi; \bar{t} \vdash H; e_1 \xrightarrow{\cdot} H'; e_2$, or
 - * $e = L[e_1]_l$, and $\Pi; \bar{t} \vdash H; e_1 \xrightarrow{\cdot} H'; e_2$

Proof: Proof proceeds by simultaneous induction on the assumed typing derivations, by cases on the last rule used. In the cases where the induction hypothesis applies to a subexpression, there are three cases to consider: (a) the subexpression is a value, (b) the subexpression is legally stuck, or (c) the subexpression can take a step. In case (b), the subexpression is a context (L or R) whose left- or right-hole is plugged by a legally stuck expression. In each case, we can show that the outer expression is also legally stuck. Take for example the outer expression $e = *(\tau *)(e')$. If the subexpression $e' = R[rstuck]_r$, let $R' = *(\tau *)(R)$. Then, $e = R'[rstuck]_r$, which is legally stuck. The cases for left-contexts and left-holes are analogous. In case (c), the induction hypothesis says that the subexpression e' is a context whose hole is plugged by an expression e_1 that can take a primitive step to an expression e_2 . In each case, we can build a context that, when plugged with e_1 , equals the outer expression e ; it follows that e can take a step. Take for example $e = *(\tau *)(e')$. If $e' = R[e_1]_r$, and $\Pi; \bar{t} \vdash H; e_1 \xrightarrow{\cdot} H'; e_2$, we take $R' = *(\tau *)(R)$. Then $e = R'[e_1]_r$, which can take a step. The cases for left-contexts and left-holes are analogous. In the rest of the proof, whenever we apply the induction hypothesis, we consider only the case when the subexpression is a value. The other two cases follow the consistent pattern explained above. We assume $H = H'$ unless otherwise stated.

- L-SHORT: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash s : \text{byte}^i$. Take $R = [\cdot]_R$, $e_1 = s$, $e_2 = b^i$, and D-SHORT applies.
- L-LONG: Analogous to L-SHORT, where D-LONG applies.
- L-VAR $\{L,R\}$: Holds vacuously as variables cannot type-check in an empty Γ .

- L-FETCHR: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash e.f : \bar{\sigma}_2$ where inversion gives

- 1 $\Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$
- 2 $\Pi.offset(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_1)$
- 3 $N\{\dots \tau f \dots\} \in \bar{t}$
- 4 $\Pi.xtextype(\bar{t}, \tau) = \bar{\sigma}_2$

We apply the IH to (1) and consider the case when e is a value, that is $e = \bar{w}$. The Value Split Lemma gives

- 5 $e = \bar{w}_1 \bar{w}_2 \bar{w}_3$

such that $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_j : \bar{\sigma}_j$ for $j \in \{1, 2, 3\}$. The Value-Type Size Lemma gives

- 6 $\text{size}(\Pi, \bar{\sigma}_1) = \text{size}(\Pi, \bar{w}_1)$
- 7 $\text{size}(\Pi, \bar{\sigma}_2) = \text{size}(\Pi, \bar{w}_2)$

From (2,6) we have

- 8 $\Pi.offset(\bar{t}, f) = \text{size}(\Pi, \bar{w}_1)$

(3,4,5,7) allow a step $\Pi; \bar{t} \vdash H; e \xrightarrow{t} H'; \bar{w}_2$ via the D-FETCH rule. Thus, we satisfy the third requirement of the theorem with $R = [\cdot]_R.f$, $e_1 = e$ and $e_2 = \bar{w}_2$.

- L-FETCHL: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash e.f : \bar{\sigma}_2, [a, o + o']$ where inversion gives

- 1 $\Pi; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a, o]$
- 2 $N\{\dots \tau f \dots\} \in \bar{t}$

We apply the IH to (1) and consider the case when e is an l-value, i.e. $e = *(\tau_1*)(\ell+i)$. (2) is a sufficient condition for $e.f$ to take a primitive step via the D-FETCHL rule. We satisfy the third requirement of the theorem with $R = [\cdot]_L.f$, $e_1 = e$ and $e_2 = *(\tau*)(\ell+(i + \Pi.offset(\bar{t}, f)))$.

- L-DEREF: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash *(\tau*)(e) : \bar{\sigma}$. Inverting, we get

- 1 $\Pi; \bar{t}; \Psi; \cdot \vdash e : \text{ptr}_\alpha(\bar{\sigma} \bar{\sigma}_2)$
- 2 $\Pi.xtextype(\bar{t}, \tau) = \bar{\sigma}$
- 3 $\Pi.access(\alpha, \text{size}(\Pi, \bar{\sigma}))$

We apply the IH to (1) and consider the case when e is a value.

Canonical Forms gives $e = \bar{w} = \text{uninit}^i$ or $e = \bar{w} = \ell+i$. In the former case, we have an expression of form $*(\tau*)(\text{uninit}^i)$, which is *legally stuck* so we are done. In the latter case, Canonical Forms also gives

- 4 $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma} \bar{\sigma}_2, [a', o']$
- 5 $\alpha = [a, o + i]$
- 6 $\text{size}(\Pi, \bar{\sigma}_0) = i$
- 7 $\vdash [a', o'] \leq [a, o]$

Heap Canonical Forms on (4) gives

- 8 $H(\ell) = \bar{w}', [a', o']$
- 9 $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}' : \bar{\sigma}_0 \bar{\sigma} \bar{\sigma}_2$

Value Split on (8) gives $\bar{w}' = \bar{w}_0\bar{w}_1\bar{w}_2$ where

10 $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_0 : \bar{\sigma}_0$

11 $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_1 : \bar{\sigma}$

Value-Type Size on (10,11) gives

12 $\text{size}(\Pi, \bar{w}_0) = \text{size}(\Pi, \bar{\sigma}_0) = i$ (from (6))

13 $\text{size}(\Pi, \bar{w}_1) = \text{size}(\Pi, \bar{\sigma})$

From (3,7), the Alignment Addition Lemma, and clause 5 of platform sensibility, we get

14 $\Pi.\text{access}([a', o' + i], \text{size}(\Pi, \bar{\sigma}))$

Notice that (2,8,12,13,14) form sufficient conditions for $H; \bar{w}$ to take a right-step to $H; \bar{w}_1$ under rule D-DEREF, so we satisfy the third requirement of the theorem with $R = [\cdot]_{\mathbb{R}}$, $e_1 = *(\tau*)(\ell+i)$, and $e_2 = \bar{w}_1$.

- L-DEREF: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash *(\tau*)(e) : \bar{\sigma}_1, \alpha$ and by inversion we get $\Pi; \bar{t}; \Psi; \cdot \vdash e : \text{ptr}_{\alpha}(\bar{\sigma}_1\bar{\sigma}_2)$. We apply the IH to this and consider the case when e is a value. Canonical Forms says that $e = \text{uninit}^k$ or $e = \ell+i$. In the former case, $*(\tau*)(\text{uninit}^k)$ is legally stuck and we are done. In the latter case, $*(\tau*)(\ell+i)$ is an l-value.
- L-NEW: The totality of $\Pi.\text{xtype}$, $\Pi.\text{ptrsize}$, and $\Pi.\text{align}$ give sufficient conditions to step.
- L-ASSN: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash e_1 = e_2 : \bar{\sigma}$ where by inversion we have

1 $\Pi; \bar{t}; \Psi; \cdot \vdash e_1 : \bar{\sigma}, \alpha$

2 $\Pi; \bar{t}; \Psi; \cdot \vdash e_2 : \bar{\sigma}$

By induction on (1,2), we have several cases to consider:

- e_1 is legally stuck. In this case we can show that $e_1 = e_2$ is legally stuck, as explained in the proof prelude.
- e_1 is an l-value and e_2 is legally stuck. Then we can show that $e_1 = e_2$ is legally stuck as explained in the prelude.
- e_1 is an l-value and e_2 can take a step. Then we can show that $e_1 = e_2$ can take a step, as explained in the proof prelude.
- e_1 is an l-value ($e_1 = *(\tau*)(\ell+i)$) and e_2 is a value ($e_2 = \bar{w}''$). Inversion on (1) (L-DEREF) gives

3 $\Pi; \bar{t}; \Psi; \cdot \vdash \ell+i : \text{ptr}_{[a, o+i]}(\bar{\sigma}\bar{\sigma}_2)$

4 $\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}$

5 $\Pi.\text{access}([a, o + i], \text{size}(\Pi, \bar{\sigma}))$

Canonical Forms on (3) yields

6 $\Psi(\ell) = \bar{\sigma}_0\bar{\sigma}\bar{\sigma}_2\bar{\sigma}_3, [a', o']$

7 $\vdash [a', o'] \leq [a, o]$

8 $\text{size}(\Pi, \bar{\sigma}_0) = i$

Heap Canonical Forms on (6) gives

9 $H(\ell) = \bar{w}', [a', o']$

10 $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}' : \bar{\sigma}_0\bar{\sigma}\bar{\sigma}_2\bar{\sigma}_3$

The Value Split Lemma on (10) gives

11 $\bar{w}' = \bar{w}_0\bar{w}_1\bar{w}_2\bar{w}_3$

12 $\Pi; \bar{t}; \Psi; \cdot \vdash \bar{w}_1 : \bar{\sigma}$

13 $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \bar{w}_0 : \bar{\sigma}_0$

The Value-Type Size Lemma on (12,13) gives

14 $\text{size}(\Pi, \bar{w}_1) = \text{size}(\Pi, \bar{\sigma})$

15 $\text{size}(\Pi, \bar{w}_0) = \text{size}(\Pi, \bar{\sigma}_0) = i$ (by (8))

From (7) and the Alignment Addition Lemma, we get $\vdash [a', o' + i] \leq [a, o + i]$. From this and (5), clause 5 of platform sensibility gives

16 $\Pi.\text{access}([a', o' + i], \text{size}(\Pi, \bar{\sigma}))$

(9,11,14,15,16) form sufficient conditions for $e_1 = e_2$ to take a step via rule D-ASSN. We satisfy the third part of the theorem with $R = [\cdot]_R$, $H' = H$, $\ell \mapsto \bar{w}_0 \bar{w}'' \bar{w}_2 \bar{w}_3$, where $\ast(\tau\ast)(\ell+i) = \bar{w}''$ steps to \bar{w}'' .

– e_1 can step. Then we can show that $e_1 = e_2$ can step as explained in the prelude.

- L-CAST: Supposing the subexpression is a value, D-CAST applies.
- L-FADDR: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} (\tau\ast)\&e \rightarrow f : \text{ptr}_{[a, o_1 + \Pi.\text{offset}(\bar{t}, f)]}(\bar{\sigma}_2)$ where by inversion we have $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} e : \text{ptr}_{a, o_1}(\bar{\sigma}_1 \bar{\sigma}_2)$. We apply the IH to the latter and consider the case when e is a value. Canonical Forms gives that either $e = \text{uninit}^i$, in which case $(\tau\ast)\&\text{uninit}^i \rightarrow f$ is legally stuck and we are done. Otherwise, $e = \ell+i$ and the expression can step via D-FADDR.
- L-SEQ: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} e_1; e_2 : \bar{\sigma}$ where inversion gives $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} e_1 : \bar{\sigma}'$. Applying the IH to this, if e_1 is a value then $e_1; e_2$ can step by the D-SEQ rule.
- L-DECL: The totality of $\Pi.\text{xtype}$, $\Pi.\text{ptrsize}$, and $\Pi.\text{align}$ give sufficient conditions to step.
- L-IF: We assume $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma}$ if $e_1 e_2 e_3 : \bar{\sigma}$, where inversion gives $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} e_1 : \text{byte}^i$. By induction, if e_1 is a value, Canonical Forms gives $e_1 = w_1 \dots w_i$ where each w_i can be either a b or uninit . If $\exists i$ s.t. $w_i = \text{uninit}$, then the expression is legally stuck and we are done. Otherwise $e_1 = b_1 \dots b_i$, which is a sufficient condition for (if $e_1 e_2 e_3$) to step by D-IFT or D-IFF.
- L-WHILE: D-WHILE applies.
- L-VALUE, L-VALUEEMPTY: The expression in question is already a value.
- L-SUB: Falls out of the induction hypothesis.

Theorem 36 (Type Soundness) *Let (\rightarrow^*) be the reflexive transitive closure of the transition relation (\rightarrow) . Then, if*

1. $\Pi; \bar{t}; \cdot \vdash_{\Gamma} e : \bar{\sigma}$
2. $\Pi; \bar{t} \vdash \cdot; e \rightarrow^* H'; e'$

then $H'; e'$ is not illegally stuck on Π and \bar{t} .

Proof: Proof proceeds by induction on the length of the step sequence. In the case of 0 steps, Progress tells us that $\cdot; e$ is not illegally stuck. In the case of n steps, by induction we have $\Pi; \bar{t} \vdash \cdot; e \rightarrow^{n-1} H''; e''$ and $H''; e''$ not illegally stuck. Preservation tells us that if $\Pi; \bar{t} \vdash H''; e'' \rightarrow H'; e'$, the type of e'' is preserved in the step to e' . It follows from Progress that $H'; e'$ is not illegally stuck.

A.3 Platform Dependencies

Theorem 37 (Constraint Satisfaction)

- If $\bar{t}; \Gamma \vdash e : \tau; S$ and Π is sensible and $\Pi \models S$ and $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$, then $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}$.
- If $\bar{t}; \Gamma \vdash e : \tau; S$ and Π is sensible and $\Pi \models S$ and $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ and $\Pi.align(\bar{t}, \tau) = \alpha$, then $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}, \alpha'$ and $\vdash \alpha' \leq \alpha$.

Proof: Proof proceeds by simultaneous induction on the assumed typing derivations, by cases on the last rule used. Reflexivity of alignment subtyping (Lemma 25) is used implicitly.

- H-SHORT: We assume $\bar{t}; \Gamma \vdash s : \text{short}; \top$. By clause 2 of the definition of a sensible platform, $\Pi.xtype(\bar{t}, \text{short}) = \text{byte}^i$. From this, $\Pi; \bar{t}; \cdot; \Gamma \vdash s : \text{byte}^i$ via L-SHORT.
- H-LONG: Similar to H-SHORT.
- H-NEW: We assume $\bar{t}; \Gamma \vdash \text{new } \tau : \tau*; \top$. Because Π is sensible (clause 4), we know $\Pi.xtype(\bar{t}, \tau*) = \text{ptr}_\alpha(\bar{\sigma})$ where $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ and $\Pi.align(\bar{t}, \tau) = \alpha$. We can plug the latter two facts into L-NEW to obtain $\Pi; \bar{t}; \cdot; \Gamma \vdash \text{new } \tau : \text{ptr}_\alpha(\bar{\sigma})$.
- H-VARR: We assume $\bar{t}; \Gamma \vdash x : \tau; \top$ and $\Gamma(x) = \tau$. Plug this along with $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ into L-VARR to reach the desired conclusion.
- H-VARL: Assume $\bar{t}; \Gamma \vdash x : \tau; \top$ and $\Gamma(x) = \tau$. Plug this along with $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ and $\Pi.align(\bar{t}, \tau) = \alpha$ into L-VARL to reach the desired conclusion.
- H-ASSN: We assume $\bar{t}; \Gamma \vdash e_1 = e_2 : \tau; S_1 \wedge S_2$, where inversion gives $\bar{t}; \Gamma \vdash e_1 : \tau; S_1$ and $\bar{t}; \Gamma \vdash e_2 : \tau; S_2$. Because $\Pi \models S_1 \wedge S_2$, we know by model definition that $\Pi \models S_1$ and $\Pi \models S_2$. Induction on the left subderivation gives $\Pi; \bar{t}; \cdot; \Gamma \vdash e_1 : \bar{\sigma}, \alpha$ where $\bar{\sigma} = \Pi.xtype(\bar{t}, \tau)$. Induction on the right subderivation gives $\Pi; \bar{t}; \cdot; \Gamma \vdash e_2 : \bar{\sigma}$. Plugging these derivations into L-ASSN yields the desired conclusion.
- H-FETCHR: We assume $\bar{t}; \Gamma \vdash e.f : \tau; S$ where inversion gives

- 1 $N\{\dots \tau f \dots\} \in \bar{t}$
- 2 $\bar{t}; \Gamma \vdash e : N; S$

By induction on (2) (with the assumption that $\Pi \models S$), we get

- 3 $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}$
- 4 $\Pi.xtype(\bar{t}, N) = \bar{\sigma}$

Because Π is sensible (clause 3), under assumption (1), we have

- 5 $\Pi.xtype(\bar{t}, N) = \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$
- 6 $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}_2$
- 7 $\text{size}(\Pi, \bar{\sigma}_1) = \Pi.offset(\bar{t}, f)$

It follows from (3,5) that

- 8 $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$

Plugging (1,6,7,8) into L-FETCHR, we obtain $\Pi; \bar{t}; \cdot; \Gamma \vdash e.f : \bar{\sigma}_2$, as desired.

- H-FETCHL: We assume $\bar{t}; \Gamma \vdash e.f : \tau; S$ with $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ and $\Pi.align(\bar{t}, \tau) = [a, o]$. Inversion gives
- 1 $N\{\dots \tau f \dots\} \in \bar{t}$

2 $\bar{t}; \Gamma \vdash e : N; S$

By induction on (2) (under the assumption that $\Pi \models S$), we get

3 $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}', [a', o']$

4 $\Pi.xtype(\bar{t}, N) = \bar{\sigma}'$

5 $\Pi.align(\bar{t}, N) = [a'', o'']$

6 $\vdash [a', o'] \leq [a'', o'']$

By clause 3 of platform sensibility, we get

7 $\bar{\sigma}' = \bar{\sigma}_1 \bar{\sigma} \bar{\sigma}_3$

8 $size(\Pi, \bar{\sigma}_1) = \Pi.offset(\bar{t}, f)$

9 $\vdash [a'', o'' + \Pi.offset(\bar{t}, f)] \leq [a, o]$

Plugging (1,3,5,6,8) along with the assumption $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ into L-FETCHL yields

10 $\Pi; \bar{t}; \cdot; \Gamma \vdash e.f : \bar{\sigma}_2, [a', o' + \Pi.offset(\bar{t}, f)]$

From (6) and the Alignment Addition Lemma we get $\vdash [a', o' + \Pi.offset(\bar{t}, f)] \leq [a'', o'' + \Pi.offset(\bar{t}, f)]$. Plugging this and (9) into ALST-TRANS we get $\vdash [a', o' + \Pi.offset(\bar{t}, f)] \leq [a, o]$. This and (10) form the desired conclusion.

- H-SEQ: Similar to H-ASSN.
- H-DEREF: We assume $\bar{t}; \Gamma \vdash *(\tau*)(e) : \tau; S$ and inversion gives $\bar{t}; \Gamma \vdash e : \tau*; S$. By induction,

1 $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}$

2 $\bar{\sigma} = \Pi.xtype(\bar{t}, \tau*)$

Because Π is sensible (clause 4), we know

3 $\bar{\sigma} = ptr_\alpha(\bar{\sigma}')$

4 $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}'$

5 $\Pi.access(\alpha, size(\Pi, \bar{\sigma}'))$

From (1,3), we get

6 $\Pi; \bar{t}; \cdot; \Gamma \vdash e : ptr_\alpha(\bar{\sigma}')$

Plugging (4,5,6) into L-DEREF, we get $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}'$, which, together with (4), forms our desired conclusion.

- H-DEREF: Similar to H-DEREF.
- H-CAST: We assume $\bar{t}; \Gamma \vdash (\tau*)e : \tau*; S$, where $S = S' \wedge subtype(\bar{t}, xtype(\bar{t}, \tau'*), xtype(\bar{t}, \tau*))$, where by inversion we obtain

1 $\bar{t}; \Gamma \vdash e : \tau'*; S'$

Because $\Pi \models S$, from the definition of \models we get

2 $\Pi \models S'$

Applying the induction hypothesis to (1,2), we get

3 $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}'$

$$4 \quad \Pi.xtype(\bar{t}, \tau^*) = \bar{\sigma}'$$

Interpreting S under Π , we get

$$5 \quad \Pi; \bar{t} \vdash \bar{\sigma}' \leq \bar{\sigma}$$

where $\bar{\sigma} = \Pi.xtype(\bar{t}, \tau^*)$. Plugging (3,5) into L-SUB, we get

$$6 \quad \Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}$$

From (6) and clause 4 of the definition of sensible platforms, we know $\bar{\sigma} = \text{ptr}_\alpha(\bar{\sigma}'')$ where $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}''$. Plugging (5) into L-CAST, we get $\Pi; \bar{t}; \cdot; \Gamma \vdash (\tau^*)e : \text{ptr}_\alpha(\bar{\sigma}'')$, as desired.

- H-FADDR: We assume $\bar{t}; \Gamma \vdash (\tau^*)\&e \rightarrow f : \tau^*; S$, where

$$\begin{aligned} S = S' \wedge \exists \bar{\sigma}_1, \bar{\sigma}_2, a, o \quad & \cdot \quad xtype(\bar{t}, N^*) = \text{ptr}_{[a,o]}(\bar{\sigma}_1 \bar{\sigma}_2) \\ & \wedge \quad \text{subtype}(\bar{t}, \text{ptr}_{[a,o+\text{offset}(\bar{t},f)]}(\bar{\sigma}_2), xtype(\bar{t}, \tau^*)) \\ & \wedge \quad \text{offset}(\bar{t}, f) = \text{size}(\bar{\sigma}_1) \end{aligned}$$

Interpreting this constraint under Π , and by clause 4 of the definition of sensible platforms, we obtain that there exist $\bar{\sigma}_1, \bar{\sigma}_2, a, o$, such that

- 1 $\Pi.\text{align}(\bar{t}, N) = [a, o]$
- 2 $\Pi.xtype(\bar{t}, N^*) = \text{ptr}_{[a,o]}(\bar{\sigma}_1 \bar{\sigma}_2)$
- 3 $\Pi.xtype(\bar{t}, \tau^*) = \text{ptr}_\alpha(\bar{\sigma})$
- 4 $\Pi; \bar{t} \vdash \text{ptr}_{[a,o+\Pi.\text{offset}(\bar{t},f)]}(\bar{\sigma}_2) \leq \text{ptr}_\alpha(\bar{\sigma})$
- 5 $\Pi.\text{offset}(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_1)$

By inversion,

$$6 \quad \bar{t}; \Gamma \vdash e : N^*; S'$$

By definition of the models relation,

$$7 \quad \Pi \models S'$$

By induction on (6,7), we obtain $\Pi; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}'$ and $\Pi.xtype(\bar{t}, N^*) = \bar{\sigma}'$. From these two facts and (1,2), we obtain

$$8 \quad \Pi; \bar{t}; \cdot; \Gamma \vdash e : \text{ptr}_{[a,o]}(\bar{\sigma}_1 \bar{\sigma}_2)$$

We can plug (5,8) into L-FADDR to derive $\Pi; \bar{t}; \cdot; \Gamma \vdash (\tau^*)\&e \rightarrow f : \text{ptr}_{[a,o+\Pi.\text{offset}(\bar{t},f)]}(\bar{\sigma}_2)$. Plugging this and (4) into L-SUB, we obtain $\Pi; \bar{t}; \cdot; \Gamma \vdash (\tau^*)\&e \rightarrow f : \text{ptr}_\alpha(\bar{\sigma})$ (recall that $\text{ptr}_\alpha(\bar{\sigma}) = xtype(\bar{t}, \tau^*)$, from (3)), as desired.

- H-IF: Similar to H-ASSN.
- H-WHILE: Similar to H-ASSN.
- H-DECL: We assume $\bar{t}; \Gamma \vdash \tau_1 x; e : \tau_2; S$, where inversion gives $\bar{t}; \Gamma, x : \tau_1 \vdash e : \tau_2; S$. By induction, $\Pi; \bar{t}; \cdot; \Gamma, x : \tau_1 \vdash e : \bar{\sigma}$ where $\Pi.xtype(\bar{t}, \tau_2) = \bar{\sigma}$. Plugging these facts into L-DECL, we derive $\Pi; \bar{t}; \cdot; \Gamma \vdash \tau_1 x; e : \bar{\sigma}$, as desired.

Theorem 38 (Layout Portability) *If $\bar{t}; \cdot \vdash e : \tau; S$ and Π is sensible, $\Pi \models S$, and $\Pi; \bar{t} \vdash \cdot; e \rightarrow^* H'; e'$, then $H'; e'$ is not illegally stuck on Π and declarations \bar{t} .*

Proof: Corollary to Theorems 36 (Low-Level Type Soundness) and 37 (Constraint Satisfaction).

Definition 39 (Cast-Free Expressions)

An expression e is cast-free if

1. $(\tau^*)e'$ does not occur in e .
2. If $(\tau^*)&e' \rightarrow f$ occurs in e and $N\{\dots\tau' f\dots\} \in \bar{t}$ then $\tau = \tau'$.

Theorem 40 (Cast-Free Typing)

- If $\bar{t}; \Gamma \vdash_{\tau} e : \tau; S$, $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$, and e is cast-free, then for any sensible platform Π , $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} e : \bar{\sigma}$.
- If $\bar{t}; \Gamma \vdash_{\tau} e : \tau; S$, $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$, $\Pi.align(\bar{t}, \tau) = \alpha$, and e is cast-free, then for any sensible platform Π , $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} e : \bar{\sigma}, \alpha'$ and $\vdash \alpha' \leq \alpha$.

Proof: The proof is similar to that of the Constraint Satisfaction theorem, by simultaneous induction on the assumed typing derivations. Notice that the only cases that yield constraints are those for H-CAST and H-FADDR. The rest of the cases are either trivial or follow from the definition of a sensible platform; the corresponding proofs proceed identically as in Constraint Satisfaction. Moreover, the H-CAST case cannot occur, ruled out by the first cast-free requirement. The only remaining interesting case is that for H-FADDR. We assume $\bar{t}; \Gamma \vdash_{\tau} (\tau^*)&e \rightarrow f : \tau^*; S$ where inversion gives

$$1 \quad \bar{t}; \Gamma \vdash_{\tau} e : N^*; S'$$

$$2 \quad N\{\dots\tau f\dots\} \in \bar{t}$$

(Note that the types match as required by assumption (2).) By induction on (1), we obtain

$$3 \quad \Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} e : \text{ptr}_{[a,o]}(\bar{\sigma})$$

where $\Pi.xtype(\bar{t}, N^*) = \text{ptr}_{[a,o]}(\bar{\sigma})$. By platform sensibility clause 4, $\Pi.xtype(\bar{t}, N) = \bar{\sigma}$. By sensibility clause 3, we have

$$4 \quad \bar{\sigma} = \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$$

$$5 \quad \Pi.offset(\bar{t}, f) = \text{size}(\Pi, \bar{\sigma}_1)$$

$$6 \quad \Pi.xtype(\bar{t}, \tau) = \bar{\sigma}_2$$

Plugging (3) into L-SUB with ST-PTR, we get

$$7 \quad \Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} e : \text{ptr}_{[a,o]}(\bar{\sigma}_1 \bar{\sigma}_2)$$

Plugging (5,7) into L-FADDR yields $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} (\tau^*)&e \rightarrow f : \text{ptr}_{[a,o+\Pi.offset(\bar{t},f)]}(\bar{\sigma}_2)$ as desired.

It is important to note that this proof does not rely at all on constraints (we do not assume $\Pi \models S$) and that the sensibility of Π suffices.

Theorem 41 (Cast-Free Sufficiency) *If $\bar{t}; \cdot \vdash_{\tau} e : \tau; S$, e is cast-free, Π is sensible, and $\Pi; \bar{t} \vdash \cdot; e \rightarrow^* H'; e'$, then $H'; e'$ is not illegally stuck on Π and declarations \bar{t} .*

Proof: Corollary to Theorems 36 (Low-Level Type Soundness) and 40 (Cast-Free Typing).

A.4 Array Extension

Figure 20 shows how the syntax and semantics can be augmented with C-like arrays. First, we augment platforms (see Section A.1) with a new component ($\text{val} : \bar{b} \rightarrow \mathbb{N}$) which interprets a sequence of bytes into a number in an platform-prescribed manner. We now show how the definitions, lemmas, and theorems in the previous sections can be extended for arrays.

Extended Definition 42 (Sensible Platforms)

We add two array-related sensibility clauses:

6. $\forall \tau, \bar{t}. \exists a, k, o. \Pi. \text{align}(\bar{t}, \tau) = [a, o]$ and $\Pi. \text{xtype}(\bar{t}, \tau) = \bar{\sigma}$ and $\text{size}(\Pi, \bar{\sigma}) = k \times a$.
7. $\forall \tau \exists \alpha, \bar{\sigma}. \Pi. \text{xtype}(\bar{t}, \tau^{*\omega}) = \text{ptr}_\alpha^\omega(\bar{\sigma})$ and $\Pi. \text{xtype}(\bar{t}, \tau) = \bar{\sigma}$ and $\Pi. \text{align}(\bar{t}, \tau) = \alpha$.

Extended Definition 43 (Legal Stuck State)

We extend the syntax for stuck expressions which can appear in right-holes.

$$\begin{array}{l} \text{rstuck} ::= \dots \\ \quad | \text{new } \tau[\bar{w}_1 \text{ uninit } \bar{w}_2] \\ \quad | \text{new } \tau[\bar{b}] \quad \text{if } \Pi. \text{val}(\bar{b}) < 0 \\ \quad | \&\mathcal{L}((\tau^{*\omega})(\text{uninit}^i))[e] \\ \quad | \&\mathcal{L}((\tau^{*\omega})(\ell+i))[\bar{w}_1 \text{ uninit } \bar{w}_2] \\ \quad | \&\mathcal{L}((\tau^{*\omega})(\ell+i))[\bar{b}] \quad \text{if the first 4 hypotheses of D-ARRELT hold but not the 5th} \end{array}$$
Extended Lemma 44 (Uninit Type)

(No additions.)

Proof: Consider the inductive case, when $\bar{\sigma} = \sigma^{\bar{\sigma}'}$ and $\text{size}(\Pi, \bar{\sigma}') = j$ and by induction, $\Pi; \bar{t}; \Psi; \cdot \vdash \text{uninit}^j : \bar{\sigma}'$. If $\text{size}(\Pi, \sigma) = k$ and $\sigma \in \{\text{ptr}_\alpha^\omega(\bar{\sigma}'), \text{ptr}_\alpha^\omega(N)\}$ then we can derive $\Pi; \bar{t}; \Psi \Vdash \text{uninit}^k : \sigma$ by LW-UNINITARR. We can plug this and $\Pi; \bar{t}; \Psi; \cdot \vdash \text{uninit}^j : \bar{\sigma}'$ into L-VALUE to achieve the desired result.

Extended Lemma 45 (Constant-Size Subtyping)

(No additions.)

Proof: In the array subtyping cases (ST-ROLLARR, ST-UNROLLARR, ST-ARR), two array pointers have the same size (given in Figure 20).

Extended Lemma 46 (Value-Type Size)

(No additions.)

Proof:

- LW-UNINITARR: $\text{size}(\Pi, \text{uninit}^i) = \text{size}(\Pi, \text{ptr}_\alpha^\omega(\bar{\sigma})) = \Pi. \text{ptrsize} = i$.
- LW-LBLARR: $\text{size}(\Pi, \ell+0) = \text{size}(\Pi, \text{ptr}_\alpha^\omega(\bar{\sigma})) = \Pi. \text{ptrsize}$.

Extended Lemma 47 (Subtyping Partition)

(No additions.)

Proof: The array subtyping cases (ST-ROLLARR, ST-UNROLLARR, ST-ARR) follow immediately because $n = 1$ so $\bar{\sigma}_1 = \bar{\sigma}_{11}$.

Extended Lemma 48 (Subtyping Type Form)

- If $\Pi; \bar{t} \vdash \bar{\sigma}' \leq \text{ptr}_\alpha^\omega(\bar{\sigma})$ then $\exists \alpha', i$ such that $\bar{\sigma}' = \text{ptr}_{\alpha'}^\omega(\bar{\sigma}^i)$ or $\bar{\sigma}' = \text{ptr}_{\alpha'}^\omega(N)$ where $\Pi. \text{xtype}(\bar{t}, N) = \bar{\sigma}^i$, and $\vdash \alpha' \leq \alpha$.
- If $\Pi; \bar{t} \vdash \bar{\sigma}' \leq \text{ptr}_\alpha^\omega(N)$ and $\Pi. \text{xtype}(\bar{t}, N) = \bar{\sigma}$ then $\exists \alpha', i$ such that $\bar{\sigma}' = \text{ptr}_{\alpha'}^\omega(\bar{\sigma}^i)$ or $\bar{\sigma}' = \text{ptr}_{\alpha'}^\omega(N')$ where $\Pi. \text{xtype}(\bar{t}, N') = \bar{\sigma}^i$, and $\vdash \alpha' \leq \alpha$.

Proof: By simultaneous induction on the assumed typing derivations, by cases on the last rule used. The cases ST-PAD, ST-PADADD, ST-PTR, ST-ROLL, and ST-UNROLL cannot occur.

- ST-ARR: Follows from inspection and inversion.

New Syntax

$$\begin{aligned}
\tau &::= \dots \mid \tau^{*\omega} \\
e &::= \dots \mid \text{new } \tau[e] \mid \&((\tau^{*\omega})(e))[e] \\
R &::= \dots \mid \text{new } \tau[R] \mid \&((\tau^{*\omega})(R))[e] \mid \&((\tau^{*\omega})(\ell+i))[R] \\
\sigma &::= \dots \mid \text{ptr}_\alpha^\omega(\bar{\sigma}) \mid \text{ptr}_\alpha^\omega(N)
\end{aligned}$$

New Dynamic Rules

$$\begin{aligned}
(\text{D-NEWARR}) \quad \Pi; \bar{t} \vdash H; \text{new } \tau[\bar{b}] &\xrightarrow{\tau} H, \ell \mapsto \text{uninit}^{i \times j}, \alpha; \ell+0 \\
&\text{if } \ell \notin \text{Dom}(H) \\
&\quad \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \\
&\quad \text{size}(\Pi, \bar{\sigma}) = i \\
&\quad \Pi.\text{align}(\bar{t}, \tau) = \alpha \\
&\quad \Pi.\text{val}(\bar{b}) = j \geq 0 \\
(\text{D-ARRELT}) \quad \Pi; ts \vdash H; \&((\tau^{*\omega})(\ell+i))[\bar{b}] &\xrightarrow{\tau} H; \ell+(i+j \times k) \\
&\text{if } \Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \\
&\quad \text{size}(\Pi, \bar{\sigma}) = j \\
&\quad \Pi.\text{val}(\bar{b}) = k \\
&\quad H(\ell) = \bar{w}, \alpha \\
&\quad 0 \leq (i+j \times k) < \text{size}(\Pi, \bar{w})
\end{aligned}$$

New High-Level Typing

$$\begin{array}{c}
\text{H-NEWARR} \\
\bar{t}; \Gamma \vdash e : \text{long}; S \\
\hline
\bar{t}; \Gamma \vdash \text{new } \tau[e] : \tau^{*\omega}; S
\end{array}
\qquad
\begin{array}{c}
\text{H-ARRELT} \\
\bar{t}; \Gamma \vdash e_1 : \tau^{*\omega}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \text{long}; S_2 \\
\hline
\bar{t}; \Gamma \vdash \&((\tau^{*\omega})(e_1))[e_2] : \tau^*; S_1 \wedge S_2
\end{array}$$

New Subtyping

$$\begin{array}{c}
\text{ST-ARR} \\
\vdash \alpha_1 \leq \alpha_2 \\
\hline
\Pi; \bar{t} \vdash \text{ptr}_{\alpha_1}^\omega(\bar{\sigma}^i) \leq \text{ptr}_{\alpha_2}^\omega(\bar{\sigma})
\end{array}
\qquad
\begin{array}{c}
\text{ST-UNROLLARR} \\
\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma} \\
\hline
\Pi; \bar{t} \vdash \text{ptr}_\alpha^\omega(N) \leq \text{ptr}_\alpha^\omega(\bar{\sigma})
\end{array}
\qquad
\begin{array}{c}
\text{ST-ROLLARR} \\
\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma} \\
\hline
\Pi; \bar{t} \vdash \text{ptr}_\alpha^\omega(\bar{\sigma}) \leq \text{ptr}_\alpha^\omega(N)
\end{array}$$

New Low-Level Typing

$$\begin{array}{c}
\text{L-NEWARR} \\
\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.\text{align}(\bar{t}, \tau) = \alpha \quad \Pi; \bar{t}; C \vdash e : \bar{\sigma}' \quad \Pi.\text{xtype}(\bar{t}, \text{long}) = \bar{\sigma}' \\
\hline
\Pi; \bar{t}; C \vdash \text{new } \tau[e] : \text{ptr}_\alpha^\omega(\bar{\sigma})
\end{array}$$

$$\begin{array}{c}
\text{L-ARRELT} \\
\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma} \quad \text{size}(\Pi, \bar{\sigma}) = k \times a \quad \Pi; \bar{t}; C \vdash e_1 : \text{ptr}_{[a,o]}^\omega(\bar{\sigma}) \quad \Pi; \bar{t}; C \vdash e_2 : \bar{\sigma}' \quad \Pi.\text{xtype}(\bar{t}, \text{long}) = \bar{\sigma}' \\
\hline
\Pi; \bar{t}; C \vdash \&((\tau^{*\omega})(e_1))[e_2] : \text{ptr}_{[a,o]}^\omega(\bar{\sigma})
\end{array}$$

$$\begin{array}{c}
\text{LW-UNINITARR} \\
\Pi.\text{ptrsize} = i \\
\hline
\Pi; \bar{t}; \Psi \Vdash \text{uninit}^i : \text{ptr}_\alpha^\omega(\bar{\sigma})
\end{array}
\qquad
\begin{array}{c}
\text{LW-LBLARR} \\
\Psi(\ell) = \bar{\sigma}^i, \alpha \\
\hline
\Pi; \bar{t}; \Psi \Vdash \ell+0 : \text{ptr}_\alpha^\omega(\bar{\sigma})
\end{array}$$

New Size Function Case

$$\text{size}(\Pi, \sigma) = \Pi.\text{ptrsize} \text{ if } \sigma \in \{\text{ptr}_\alpha^\omega(\bar{\sigma}), \text{ptr}_\alpha^\omega(N)\}$$

Figure 20: Array Syntax and Semantics

- ST-UNROLLARR: Follows from inspection and inversion, with $\alpha' = \alpha$.
- ST-ROLLARR: Follows from inspection and inversion, with $\alpha' = \alpha$ and $i = 1$.
- ST-REFL: Immediate, with $\alpha' = \alpha$, $N = N'$, and $i = 1$.
- ST-SEQ: We have $\Pi; \bar{t} \vdash \bar{\sigma}_1 \bar{\sigma}_3 \leq \bar{\sigma}_2 \bar{\sigma}_4$. It must be the case that either $\bar{\sigma}_1 = \bar{\sigma}_2 = \cdot$ or $\bar{\sigma}_3 = \bar{\sigma}_4 = \cdot$, since array types cannot be broken into subsequences. The property then follows from induction.
- ST-TRANS: We have $\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_3$ where inversion gives $\Pi; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$ and $\Pi; \bar{t} \vdash \bar{\sigma}_2 \leq \bar{\sigma}_3$. There are two possibilities for $\bar{\sigma}_3$:
 1. $\bar{\sigma}_3 = \text{ptr}_\alpha^\omega(N)$ where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}$. By induction on the second subderivation, we have either
 - $\bar{\sigma}_2 = \text{ptr}_{\alpha'}^\omega(\bar{\sigma}^j)$ where $\vdash \alpha' \leq \alpha$. By induction on the first subderivation, we have either $\bar{\sigma}_1 = \text{ptr}_{\alpha''}^\omega(\bar{\sigma}^{j \times k})$ or $\bar{\sigma}_1 = \text{ptr}_{\alpha''}^\omega(N')$ where $\Pi.\text{xtype}(\bar{t}, N') = \bar{\sigma}^j$ and $\vdash \alpha'' \leq \alpha'$. By ALST-TRANS, $\vdash \alpha'' \leq \alpha$, as required.
 - $\bar{\sigma}_2 = \text{ptr}_{\alpha'}^\omega(N')$ where $\Pi.\text{xtype}(\bar{t}, N') = \bar{\sigma}$ and $\vdash \alpha' \leq \alpha$. By induction on the first subderivation, we have either $\bar{\sigma}_1 = \text{ptr}_{\alpha''}^\omega(\bar{\sigma}^j)$ or $\bar{\sigma}_1 = \text{ptr}_{\alpha''}^\omega(N')$ where $\Pi.\text{xtype}(\bar{t}, N') = \bar{\sigma}^j$. By ALST-TRANS, $\vdash \alpha'' \leq \alpha$, as required.
 2. $\bar{\sigma}_3 = \text{ptr}_\alpha^\omega(\bar{\sigma}')$. Similar to the above case.

Extended Lemma 49 (Canonical Forms)

- If $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \text{ptr}_\alpha^\omega(\bar{\sigma})$ then either
 - $\bar{w} = \text{uninit}^i$ where $\Pi.\text{ptrsize} = i$.
 - $\bar{w} = \ell+0$ and $\exists j, \bar{\sigma}', \alpha'$ such that $\Psi(\ell) = \bar{\sigma}^j, \alpha'$ where $\vdash \alpha' \leq \alpha$.
- If $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \text{ptr}_\alpha^\omega(N)$ and $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}$, then either
 - $\bar{w} = \text{uninit}^i$ where $\Pi.\text{ptrsize} = i$.
 - $\bar{w} = \ell+0$ and $\exists j, \bar{\sigma}', \alpha'$ such that $\Psi(\ell) = \bar{\sigma}^j, \alpha'$ where $\vdash \alpha' \leq \alpha$.

Proof: By simultaneous induction on the assumed derivations, by cases on the last rule applied. All but two are impossible. Reflexivity of the alignment subtype relation is used implicitly (Lemma 25).

- L-VALUE: There are two cases:
 - Assume $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \text{ptr}_\alpha^\omega(\bar{\sigma}_1)$. By inversion we get $\Pi; \bar{t}; \Psi \vdash_{\bar{w}} \bar{w} : \text{ptr}_\alpha^\omega(\bar{\sigma}_1)$. Quick inspection reveals that this conclusion could have only been reached by LW-UNINITARR or LW-LBLARR. In the former case, $\bar{w} = \text{uninit}^i$ where $i = \Pi.\text{ptrsize}$, and in the latter case, $\bar{w} = \ell+0$ where by inversion $\Psi(\ell) = \bar{\sigma}^j, \alpha'$ with $\alpha' = \alpha$.
 - Assume $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \text{ptr}_\alpha^\omega(N)$, where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}_1$. Vacuous; cannot be derived via L-VAL.
- L-SUB: There are two cases:
 - Assume $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \text{ptr}_\alpha^\omega(\bar{\sigma})$. By inversion we get $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \bar{\sigma}''$ and $\Pi; \bar{t} \vdash \bar{\sigma}'' \leq \text{ptr}_\alpha^\omega(\bar{\sigma})$. By the Subtyping Type Form Lemma, $\bar{\sigma}''$ is one of
 - * $\text{ptr}_{\alpha'}^\omega(\bar{\sigma}^i)$ where $\vdash \alpha' \leq \alpha$. Applying the IH to $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \text{ptr}_{\alpha'}^\omega(\bar{\sigma}^i)$, we get that either $\bar{w} = \text{uninit}^k$ or $\bar{w} = \ell+0$ where $\Psi(\ell) = \bar{\sigma}^{i \times j}, \alpha''$ and $\vdash \alpha'' \leq \alpha'$. Plugging $\vdash \alpha'' \leq \alpha'$ and $\vdash \alpha' \leq \alpha$ into ALST-TRANS we get $\vdash \alpha'' \leq \alpha$, as required.
 - * $\text{ptr}_{\alpha'}^\omega(N)$ where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}$ and $\vdash \alpha' \leq \alpha$. By induction and an application of ALST-TRANS, like the above case.
 - Assume $\Pi; \bar{t}; \Psi; \cdot \vdash_r \bar{w} : \text{ptr}_\alpha^\omega(N)$, where $\Pi.\text{xtype}(\bar{t}, N) = \bar{\sigma}$. Similar to the above case.

Extended Lemma 50 (Subject Reduction)

(No additions.)

Proof:

- L-NEWARR:
$$\frac{\Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.align(\bar{t}, \tau) = \alpha \quad \Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \bar{b} : \bar{\sigma}' \quad \Pi.xtype(\bar{t}, \text{long}) = \bar{\sigma}'}{\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \text{new } \tau[\bar{b}] : \text{ptr}_{\alpha}^{\omega}(\bar{\sigma})}$$

The assumed step rule is

$$\Pi; \bar{t} \vdash H; \text{new } \tau[\bar{b}] \xrightarrow{\iota} H, \ell \mapsto \text{uninit}^{i \times j}, \alpha; \ell + 0 \quad \text{if} \quad \begin{array}{l} \ell \notin \text{Dom}(H) \\ \Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \\ \text{size}(\Pi, \bar{\sigma}) = i \\ \Pi.align(\bar{t}, \tau) = \alpha \\ \Pi.\text{val}(\bar{b}) = j \geq 0 \end{array}$$

We have $H' = H, \ell \mapsto \text{uninit}^{i \times j}, \alpha$. Let $\Psi' = \Psi, \ell \mapsto \bar{\sigma}^j, \alpha$, so $\Psi'(\ell) = \bar{\sigma}^j, \alpha$. Plugging this into LW-LBLARR, and the result into L-VALUE, yields $\Pi; \bar{t}; \Psi'; \cdot \vdash_{\Gamma} \ell + 0 : \text{ptr}_{\alpha}^{\omega}(\bar{\sigma})$, which satisfies the first part of the claim.

To satisfy the second part of the claim, we first apply the Heap Weakening Lemma to the third assumption to get

$$\mathbf{1} \quad \Pi; \bar{t}; \Psi' \vdash H : \Psi$$

From the third step side condition and the Uninit Type Lemma, we have $\Pi; \bar{t}; \Psi'; \cdot \vdash_{\Gamma} \text{uninit}^i : \bar{\sigma}$. From the Sequence Typing Lemma applied j times, we have $\Pi; \bar{t}; \Psi'; \cdot \vdash_{\Gamma} \text{uninit}^{i \times j} : \bar{\sigma}^j$. Plugging this and (1) into HT-INF we obtain $\Pi; \bar{t}; \Psi' \vdash H' : \Psi'$, as desired.

- L-ARRFLT:
$$\frac{\Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \quad \text{size}(\Pi, \bar{\sigma}) = k \times a \quad \Pi; \bar{t}; C \vdash_{\Gamma} \bar{b} : \bar{\sigma}' \quad \Pi.xtype(\bar{t}, \text{long}) = \bar{\sigma}'}{\Pi; \bar{t}; C \vdash_{\Gamma} \&((\tau *^{\omega})(\ell + i))[\bar{b}] : \text{ptr}_{[a, o]}^{\omega}(\bar{\sigma})}$$

The assumed step rule is

$$\Pi; \bar{t} \vdash H; \&((\tau *^{\omega})(\ell + i))[\bar{b}] \xrightarrow{\iota} H; \ell + (i + j \times k) \quad \text{if} \quad \begin{array}{l} \Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \\ \text{size}(\Pi, \bar{\sigma}) = j \\ \Pi.\text{val}(\bar{b}) = k \\ H(\ell) = \bar{w}, \alpha' \\ 0 \leq (i + j \times k) < \text{size}(\Pi, \bar{w}) \end{array}$$

Applying the Canonical Forms Lemma to the first typing hypothesis, we learn that $i = 0$ and

- $\mathbf{1} \quad \Psi(\ell) = \bar{\sigma}^h, [a', o']$ for some h
- $\mathbf{2} \quad \vdash [a', o'] \leq [a, o]$

Applying the Heap Canonical Forms Lemma to (1), we get

- $\mathbf{3} \quad H(\ell) = \bar{w}, \alpha'$
- $\mathbf{4} \quad \Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \bar{\sigma}^h$

Applying the Value-Type Size Lemma to (4), we have

- $\mathbf{5} \quad \text{size}(\Pi, \bar{w}) = \text{size}(\Pi, \bar{\sigma}^h)$

From this and the last side condition in the assumed reduction we get

6 $j \times k < \text{size}(\Pi, \bar{\sigma}^h)$

It must be the case then that we can write $\bar{\sigma}^h$ as $\bar{\sigma}^k \bar{\sigma}^g$ (since $\text{size}(\Pi, \bar{\sigma}) = j$, according to the second side condition in the reduction) where $g > 0$. Using (1), we can write $\Psi(\ell) = \bar{\sigma}^k \bar{\sigma}^g, [a', o']$ and we naturally have $\text{size}(\Pi, \bar{\sigma}^k) = k \times \text{size}(\Pi, \bar{\sigma}) = k \times j$. Plugging these two facts into L-LBL and the result into L-VALUE, we get

7 $\Pi; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \ell+(j \times k) : \text{ptr}_{[a', o'+j \times k]}(\bar{\sigma}^g)g$

Using the Alignment Addition Lemma, we can conclude

8 $\vdash [a', o' + j \times k] \leq [a, o + j \times k]$

From the third hypothesis of the typing assumption, we know $j = \text{size}(\Pi, \bar{\sigma}) = a \times k'$ for some k' . This means that $j = 0 \pmod a$. With this, using ALST-OFF, we can conclude $\vdash [a, o + j \times k] \leq [a, o]$. Plugging this and (8) into ALST-TRANS, we get $\vdash [a', o' + j \times k] \leq [a, o]$. Plugging this into ST-PTR we get $\Pi; \bar{t} \vdash \text{ptr}_{[a', o'+j \times k]}(\bar{\sigma}^g) \leq \text{ptr}_{[a, o]}(\bar{\sigma})$. Plugging this along with (7) into L-SUB yields $\Pi; \bar{t}; \Psi; \cdot \vdash \ell+(j \times k) : \text{ptr}_{[a, o]}(\bar{\sigma})$, as desired.

Extended Lemma 51 (Progress)

(No additions.)

Proof:

- L-NEWARR:
$$\frac{\Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \quad \Pi.align(\bar{t}, \tau) = \alpha \quad \Pi; \bar{t}; C \vdash_{\Gamma} e : \bar{\sigma}' \quad \Pi.xtype(\bar{t}, \text{long}) = \bar{\sigma}'}{\Pi; \bar{t}; C \vdash_{\Gamma} \text{new } \tau[e] : \text{ptr}_{\alpha}^{\omega}(\bar{\sigma})}$$

We apply the IH to $\Pi; \bar{t}; C \vdash_{\Gamma} e : \bar{\sigma}'$ and consider the case when e is a value. Clause 2 of platform sensibility gives that $\bar{\sigma}' = \text{byte}^i$. Canonical Forms gives that either $e = \bar{w}_1 \text{uninit } \bar{w}_2$ or $e = \bar{b}$. In the former case, $\text{new } \tau[e]$ is legally stuck. In the latter case, if $\Pi.\text{val}(\bar{b}) < 0$, then $\text{new } \tau[e]$ is again legally stuck. Otherwise it can take a step via rule D-NEWARR.

- L-ARRELT:
$$\frac{\Pi.xtype(\bar{t}, \tau) = \bar{\sigma} \quad \text{size}(\Pi, \bar{\sigma}) = k \times a \quad \Pi; \bar{t}; C \vdash_{\Gamma} e_1 : \text{ptr}_{[a, o]}^{\omega}(\bar{\sigma}) \quad \Pi; \bar{t}; C \vdash_{\Gamma} e_2 : \bar{\sigma}' \quad \Pi.xtype(\bar{t}, \text{long}) = \bar{\sigma}'}{\Pi; \bar{t}; C \vdash_{\Gamma} \&((\tau^{*\omega})(e_1))[e_2] : \text{ptr}_{[a, o]}(\bar{\sigma})}$$

We apply the IH to $\Pi; \bar{t}; C \vdash_{\Gamma} e_1 : \text{ptr}_{[a, o]}^{\omega}(\bar{\sigma})$ and $\Pi; \bar{t}; C \vdash_{\Gamma} e_2 : \bar{\sigma}'$ and consider the case when both e_1 and e_2 are values. Canonical Forms on the former gives that either $e_1 = \text{uninit}^{k_0}$ or $e_1 = \ell+0$ with $\Psi(\ell) = \bar{\sigma}^{k_1}, \alpha'$ where $\vdash \alpha' \leq [a, o]$. In the former case, $\&((\tau^{*\omega})(e_1))[e_2]$ is legally stuck. Suppose the latter case occurs.

Clause 2 of platform sensibility yields $\bar{\sigma}' = \text{byte}^{k_2}$. Canonical Forms gives either $e_2 = \bar{w}_1 \text{uninit } \bar{w}_2$ or $e_2 = \bar{b}$. In the former case, $\&((\tau^{*\omega})(e_1))[e_2]$ is legally stuck. In the latter case, if $\Pi.\text{val}(\bar{b}) < 0$, $\&((\tau^{*\omega})(e_1))[e_2]$ is again legally stuck. Suppose then that $\Pi.\text{val}(\bar{b}) = k$ for $k \geq 0$.

Heap Canonical Forms on $\Psi(\ell) = \bar{\sigma}^{k_1}, \alpha'$ gives $H(\ell) = \bar{w}, \alpha'$. Suppose that $i + \text{size}(\Pi, \bar{\sigma}) \times k \geq \text{size}(\Pi, \bar{w})$. Then $\&((\tau^{*\omega})(e_1))[e_2]$ is legally stuck. Otherwise, $H(\ell) = \bar{w}, \alpha'$ and $i + \text{size}(\Pi, \bar{\sigma}) \times k < \text{size}(\Pi, \bar{w})$ give sufficient conditions to take a step via D-ARRELT.

Extended Lemma 52 (Constraint Satisfaction)

(No additions.)

Proof:

- H-NEWARR:
$$\frac{\bar{t}; \Gamma \vdash_{\Gamma} e : \text{long}; S}{\bar{t}; \Gamma \vdash_{\Gamma} \text{new } \tau[e] : \tau^{*\omega}; S}$$

By induction on $\bar{t}; \Gamma \vdash_{\Gamma} e : \text{long}; S$ we get $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\Gamma} e : \bar{\sigma}'$ where $\Pi.xtype(\bar{t}, \text{long}) = \bar{\sigma}'$. Plugging these two facts along with $\Pi.xtype(\bar{t}, \tau) = \bar{\sigma}$ and $\Pi.align(\bar{t}, \tau) = \alpha$ into L-NEWARR yields $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\Gamma} \text{new } \tau[e] : \text{ptr}_{\alpha}^{\omega}(\bar{\sigma})$. From clause 7 of platform sensibility, we get $\Pi.xtype(\bar{t}, \tau^{*\omega}) = \text{ptr}_{\alpha}^{\omega}(\bar{\sigma})$, as desired.

Syntax:

$$e ::= \dots \mid \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n$$

Dynamic semantics:

$$\text{D-PCASE} \quad \frac{\Pi \models S_k}{\Pi; \bar{t} \vdash (\text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n) \xrightarrow{r} e_k}$$

High-level typing:

$$\text{H-PCASE} \quad \frac{\forall 1 \leq i \leq n . \bar{t}; \Gamma \vdash_r e_i : \tau; S'_i}{\bar{t}; \Gamma \vdash_r \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \tau; \bigwedge_{i=1}^n (S_i \Rightarrow S'_i) \wedge \bigvee_{i=1}^n S_i}$$

Low-level typing:

$$\text{L-PCASE} \quad \frac{\forall 1 \leq i \leq n . \text{if } \Pi \models S_i \text{ then } \Pi; \bar{t}; \Psi; \Gamma \vdash_r e_i : \bar{\sigma} \quad \exists 1 \leq i \leq n . \Pi \models S_i}{\Pi; \bar{t}; \Psi; \Gamma \vdash_r \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \bar{\sigma}}$$

Figure 21: Additions for Platform Selection

- **H-ARRELT:** $\frac{\bar{t}; \Gamma \vdash_r e_1 : \tau^{*\omega}; S_1 \quad \bar{t}; \Gamma \vdash_r e_2 : \text{long}; S_2}{\bar{t}; \Gamma \vdash_r \&((\tau^{*\omega})(e_1))[e_2] : \tau^*; S_1 \wedge S_2}$

We know $\Pi \models S_1 \wedge S_2$ so by definition of (\models) we have $\Pi \models S_1$ and $\Pi \models S_2$. Using this, we can apply the induction hypothesis to the two subderivations, and use clause 7 of platform sensibility to obtain

- 1 $\Pi; \bar{t}; \cdot; \Gamma \vdash_r e_1 : \text{ptr}_{[a,o]}^\omega(\bar{\sigma})$
- 2 $\Pi.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}$
- 3 $\Pi.\text{align}(\bar{t}, \tau) = [a, o]$
- 4 $\Pi; \bar{t}; \cdot; \Gamma \vdash_r e_2 : \bar{\sigma}'$
- 5 $\Pi.\text{xtype}(\bar{t}, \text{long}) = \bar{\sigma}'$

From clause 6 of platform sensibility we learn

- 6 $\text{size}(\Pi, \bar{\sigma}) = k \times a$ for some k

From clause 4 of platform sensibility we have

- 7 $\Pi.\text{xtype}(\bar{t}, \tau^*) = \text{ptr}_{[a,o]}(\bar{\sigma})$

Plugging (1,2,4,5,6) into L-ARRELT yields $\Pi; \bar{t}; \cdot; \Gamma \vdash_r \&((\tau^{*\omega})(e_1))[e_2] : \text{ptr}_{[a,o]}^\omega(\bar{\sigma})$, which together with (7), gives the desired conclusion.

A.5 Platform Selection Extension

The additions for platform selection are given in Figure 21. The metatheory extends as follows:

Extended Lemma 53 (Progress)

(No additions to the statement.)

Proof:

- L-PCASE: We assume $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\tau} \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \bar{\sigma}$. By inversion, we obtain $\Pi \models S_i$ for some i . Using this, we can step to e_i via D-PCASE.

Extended Lemma 54 (Subject Reduction)

(No additions to the statement.)

Proof:

- L-PCASE: We assume $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\tau} \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \bar{\sigma}$ and the step was derived via D-PCASE. By inversion, we get $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\tau} e_i : \bar{\sigma}$ for all i such that $\Pi \models S_i$. By inversion on D-PCASE, we get $\Pi \models S_k$ for a k such that the `pcase` expression stepped to e_k . Therefore, $\Pi; \bar{t}; \Psi; \Gamma \vdash_{\tau} e_k : \bar{\sigma}$, as desired.

Extended Lemma 55 (Constraint Satisfaction)

(No additions to the statement.)

Proof:

- H-PCASE: We assume $\bar{t}; \Gamma \vdash_{\tau} \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \tau; S$ and $\Pi \models S$ and we want to show $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} \text{pcase } S_1 \Rightarrow e_1 \dots S_n \Rightarrow e_n : \bar{\sigma}$ where $\Pi.xtype(\tau) = \bar{\sigma}$.

We know $S = (\bigwedge_{i=1}^n S_i \Rightarrow S'_i) \wedge \bigvee_{i=1}^n S_i$ and by inversion on the assumed typing we get $\bar{t}; \Gamma \vdash_{\tau} e_i : \tau; S'_i$ for all $i \in 1..n$. By definition of (\models) , we have $\Pi \models \bigwedge_{i=1}^n S_i \Rightarrow S'_i$ and $\Pi \models \bigvee_{i=1}^n S_i$. The former means that $\forall 1 \leq i \leq n$, if $\Pi \models S_i$ then $\Pi \models S'_i$. The latter means that $\exists 1 \leq i \leq n$ s.t. $\Pi \models S_i$. Combining the results, we get $\exists 1 \leq i \leq n$ s.t. $\Pi \models S'_i$.

We apply the induction hypothesis to all subderivations $\bar{t}; \Gamma \vdash_{\tau} e_i : \tau; S'_i$ where $\Pi \models S'_i$ to obtain $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} e_i : \bar{\sigma}$. That is, $\forall 1 \leq i \leq n$, if $\Pi \models S'_i$ then $\Pi; \bar{t}; \cdot; \Gamma \vdash_{\tau} e_i : \bar{\sigma}$.

We plug the conclusions of the previous two paragraphs into L-PCASE to obtain the desired result.