

A Theory of Implementation-Dependent Low-Level Software

Marius Nita

Dan Grossman

Craig Chambers

{maris,djg,chambers}@cs.washington.edu

Department of Computer Science and Engineering
University of Washington

UW-CSE Technical Report 2006-10-01

Abstract

This technical report is an extended version of the recent paper *A Theory of Implementation-Dependent Low-Level Software* by Nita, Grossman, and Chambers. The report contains additional discussion of the work and full proofs of the theorems described in the paper. The differences between the report and the paper are more precisely outlined in Appendix B. If the reader is interested only in the safety proof, please skip to Appendix A.

We present a theory for describing implementation-dependent assumptions that a program in a C-like language might make, such as the size and alignment of data. We define a static analysis to encode such assumptions in a constraint that describes language implementations (i.e., compilers and architectures) on which a program is memory-safe. More specifically, the constraint produced by the analysis is a formula in first-order logic and implementations are models for the logic. By defining an abstract machine in such a way that it takes an implementation as a parameter, we can prove the analysis is sound. We use this foundation to explain some common but non-portable coding practices and the poorly understood assumptions they are implicitly making.

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 layed 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 5) either (1) checks or simply assumes full portability (e.g., that the input program is unaffected by structure padding) or (2) checks the input program assuming a particular language implementation (making no guarantee for a different compiler or architecture).

Requiring full portability for all code (e.g., by enforcing poorly understood and informally specified [22] 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 implementation 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, implicitly relying on some language-implementation assumptions can lead to pernicious defects that lie dormant until one uses an implementation violating the assumptions. Whereas defects like dangling-pointer dereferences are largely independent of the language implementation (so testing or verification on the “old machine” 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 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 relied on “undocumented behavior”. Section 2.2 discusses some specific real-world examples.

Common practice confronts this “somewhat but not completely portable” dilemma by manually isolating and minimizing implementation-dependent assumptions. For example, the Linux source code has an `arch` subdirectory; avoiding assumptions in the rest of the kernel is left to programmer discipline. As another example, a garbage collector for a high-level language might assume pointers are four bytes and aim to be correct for any implementation satisfying this assumption. Similarly, run-time system code for accessing object headers may make assumptions about the layout of `struct` values. In all cases, the code is *semi-portable*, meaning it is—by design—correct for many but not all compilers and machines.

To help support semi-portable programming, we have begun building a semantics-based porting tool for C. When finished, it will input a C program and output a *description of a set of implementations* on which the program “makes sense” (and the source-code locations that influence the description). Defining the implementation-description language has forced us to give precise meaning to poorly understood platform-specific issues. The result is a foundation not only for a porting tool, but for any analysis seeking to account for implementation dependencies. While our focus is C, where incorrect assumptions can violate memory safety, many safe languages also have implementation-dependent behavior for which our approach of isolating the implementation-definition should apply.

1.1 Key Questions

Many compilers and low-level language tools have had to wrestle with defining semi-portability and implementation dependencies, but the issues have not been isolated and considered rigorously. In hindsight, the key questions are:

- How can we define the notion of an *implementation* to expose issues relevant to semi-portable software without exposing a full translation to assembly language?
- Can we analyze a C program statically to determine (conservatively) what implementation-dependent assumptions must hold for it to be memory-safe?
- What soundness guarantee can we provide? Why does memory-safety (unfortunately) not imply that a program behaves equivalently on different implementations?
- How can we prove a soundness theorem for a source-code analysis given that the code runs on an implementation-dependent low-level machine?

To answer these questions, we have built a novel model for a small-but-relevant C-level expression language.

1.2 Approach

The key to our formal model is isolating the notion of an *implementation*. An implementation has two roles: (1) as a parameter to the operational semantics, and (2) as something we can describe with a portability constraint. The actual definition of an implementation includes things like how to determine the offset of a `struct` field and what alignment restrictions a memory access must obey. The details are fully described in this work, but they are less important than the general insight: by parameterizing the operational semantics

by an implementation, we can take a program P and state that a property (such as safety) holds for P on implementations satisfying a constraint S . That is, P is semi-portable as described by S .

To see how implementations 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 implementation. Therefore, our operational semantics has the form $impl \vdash P \rightarrow P'$ where $impl$ is an implementation and P is a program state. That way, the dereference rule can use $impl$ to guide the memory access (and become stuck if $impl$ deems the access misaligned).

As for portability constraints, they are formulas in a theory of first-order logic that we interpret by having implementations serve as models. For example, the constraint “ $\text{access}(4, 8) \wedge \text{size}(\text{long}) = 8$ ” is modeled by any implementation in which values of type `long` occupy 8 bytes and 8-byte loads of 4-byte aligned data are allowed.¹

Now given a program P we can try to find a constraint S such that if $impl \models S$, then the abstract machine does not get stuck when running P given $impl$. 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 “safe” implementations is trivially undecidable, so a sound approximation (all models are safe, but not all safe implementations are models) is warranted. In this work, we take a very conservative approach: A type system for source programs produces S (which one can view as an effect), using no flow-sensitivity or alias information. In practice, a code-analysis tool will use a more sophisticated analysis, but the set-up (given P , produce an S) will remain the same.

The key metatheoretic result is showing that the S our system produces is indeed sound. To do so, we define a second type system for program states. This second type system, which exists only to prove safety, is parameterized by an $impl$ just 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 (not including orthogonal issues like array-bounds violations and uninitialized memory, which we could prevent at the cost of losing our focus on semi-portability).
2. If the first type system produces S given P , then P type-checks in the second type system for any $impl$ such that $impl \models S$.

Together, these lemmas ensure the program cannot get stuck when run on any model of S .

1.3 Contributions and Caveats

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

- Most semi-portability questions can be 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, often depends on the implementation.
- A theory in first-order logic can describe implementation constraints.
- There should be a notion of “sensible” implementation, meaning implementations on which every program without casts cannot get stuck. The constraint language lets us write a formula for which the models are the “sensible” implementations.
- A program’s safety might depend on a pointer not being treated as an array or a pointer not being written through.

For tractability, the model we present considers only a small expression language inspired by C. It has many relevant features, including structs (and assignment of struct values), 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

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

and `goto`). We see no fundamental problems extending our model in such directions. We also present a definition of implementation that is slightly simpler than one may encounter in practice. In particular, we assume all pointers have the same size and that alignment restrictions do not depend on the destination of a memory access (e.g., we do not distinguish floating-point operations). Again, extending our definition of implementation is straightforward.

1.4 Outline

Section 2 presents examples of semi-portable code and the constraints describing their assumptions. Section 3 presents our core formal model, including the definition of implementations, our first-order theory, the dynamic and static semantics of our language, and our soundness theorem. Section 4 describes how to extend the model in several directions, most importantly to support arrays. 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 semi-portability and relevant implementation constraints. Section 2.2 complements this “tutorial” with actual platforms, systems, and coping strategies related to these concepts.

2.1 Small Code Fragments

Example 0: Accessing Memory

```
(*e).f
```

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

Implementations 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 implementation (which comprises compiler and machine) “supports” 8-byte accesses on 4-byte alignments. In this paper, we assume implementations include an *access* function of type `Int → Int → Bool`, as well as *size* and *alignment* functions that map types to integers.²

Our example `(*e).f` therefore induces the constraint `access(n,m)` where m and n are defined above. However, this constraint assumes `e` actually evaluates to a pointer with alignment m and a τ at the right offset. The constraints for cast expressions must ensure this.

Example 1: 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 implementation to meet these constraints (and for `g1` and `f1` to have offset 0 and `g2` to have the same offset as `f2`), but our purpose is to capture these and less portable notions precisely.

²We actually generalize the notion of alignment to include an offset, e.g., `[8,0]` would describe an 8-byte aligned address and `[8,2]` would describe a 2-byte offset from an 8-byte aligned address.

Example 2: Flattening and Alignment

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

struct D2 {
    int* g1;
    int* g2;
};
```

A cast from `struct S2*` to `struct D2*` has similar constraints as in Example 1, but this time the C standard provides no guarantee. In fact, some implementations put pad bytes before `f4` because of alignment constraints and an assumption that all `struct` types are defined at top-level. Our system will generate constraints preventing such a representation mismatch if the program has this cast.

Example 3: Suffix

```
struct S3 {
    int* f1;
    int* f2;
    double f3;
};

struct D3 {
    int* g1;
    double g2;
};

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, an implementation with 4-byte pointers, 8-byte doubles, and 8-byte alignment of doubles cannot support this cast because `struct D3` has more padding. Implementations without padding can allow this cast, even though `&x->f2` has type `int**` in C.

Example 4: Arrays and Prefixes

```
struct S1 * x = ...;
struct D1 y = ((struct D1 *)x)[7];
```

Example 1 (and 2–3) implicitly assumed the destination pointer was not used as an array (i.e., it was used as a pointer to one and not more than one `struct D1`). This issue is orthogonal to array-bounds violations; we must reject the cast in Example 4 even if `x` points to more than 7 elements. Section 4.1 therefore extends our model to make the necessary distinction between pointers to single-objects and pointers to arrays.

Example 5: 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 4.2 adds this orthogonal feature to our model.

Example 6: 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 4.3 extends our model to let implementations indicate on a per-pad basis how they implement assignment to `struct` values.

Example 7: 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.

Interestingly, once one focuses on safety, certain constraints are actually unnecessary even though correct C implementations must support them. For example, type-compatible fields of a `struct` could have the same offsets (like a `union`). So omitting constraints like “fields have disjoint offsets” does not ruin our safety result.

2.2 Practical Scenarios

The scope of the portability problem is not precisely known because defects lie dormant until one changes hardware or compiler (or at least compiler flags). Therefore, like date bugs (such as the famous Y2K problem of last decade), offending code can be difficult to locate and fix.³

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 [2]:

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.⁴

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” [25].

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 [21]. 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 [26].

³Furthermore, semantics-based Y2K solutions [13] serve as inspiration for us, though portability bugs fortunately do not share a worldwide deadline.

⁴Emphasis in original.

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

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

A third example comes from the common practice of writing garbage collectors for high-level languages in C. If objects are allocated in-line by generated assembly code but the garbage collector accesses them via a C struct, there are unchecked and often undocumented assumptions that the code generator and C code make about each other. In principle, our constraint language could enable documenting the relevant assumptions, and a tool could automatically test a C implementation against the documentation.

3 Core Language

This section develops a formal model that can explain examples 0–3 from Section 2. We define an appropriate core language, a definition of *implementation*, a dynamic semantics, a first-order theory that constrains implementations, a type-and-effect system to produce constraints, and a type-soundness result (acquired via a second, lower-level type system). Figure 8 on page 15 summarizes the judgments defined for the model.

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 (τ^*), 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 ($(\tau^*)e$).

For pointer dereference ($*(\tau^*)(e)$) and pointing to a field ($(\tau^*)&\tau^* \rightarrow e.f$), 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 τ^* , then $*e$ reads `sizeof`(τ) bytes); our type decoration makes this explicit. The cast in address-of-field expressions lets us support “suffix casts” as in Example 3 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, `τ x; e` creates a local variable x of type τ bound in e . Because memory management is not our concern, the dynamic semantics uses heap allocation even for local variables.

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

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

Figure 2: Syntax extensions for run-time behavior

$$\begin{aligned}
\sigma &::= \text{byte} \mid \text{pad}[i] \mid \text{ptr}_\alpha(\bar{\sigma}) \mid \text{ptr}_\alpha(N) \\
\text{impl.xtype}(\bar{t}, \tau) &= \bar{\sigma} \\
\text{impl.align}(\bar{t}, \tau) &= \alpha \\
\text{impl.offset}(\bar{t}, f) &= i \\
\text{impl.access}(\alpha, i) &= \{\text{true}, \text{false}\} \\
\text{impl.ptrsize} &= i \\
\text{impl.xliteral}(s) &= \bar{b} \\
\text{impl.xliteral}(l) &= \bar{b}
\end{aligned}$$

Figure 3: Implementations and low-level types

Figure 2 defines the syntactic extensions for run-time expressions and heaps. A value v is a sequence of “small 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 we roughly model the heap as a mapping from labels to values (which recall are 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 “just right” for modeling semi-portable C.

Actually, heaps also map labels to alignments $[a, o]$ which means the address ℓ is $o \bmod a$. (Typically o is zero.) As the next section shows, an implementation specifies an alignment for each type; allocating an object of type τ produces a fresh heap location with this alignment. Note that if ℓ has alignment $[a, o]$, then $\ell+i$ has alignment $\text{gcd}(a, a+i+o)$ where gcd computes the greatest-common-divisor.

3.2 Implementations

An implementation (*impl*) has several components, summarized in Figure 3. Together these components suffice to guide an abstract machine parameterized by an implementation.

- A translation of types into a lower-level representation ($\bar{\sigma}$), described below. We write $\text{impl.xtype}(\bar{t}, \tau)$ for the $\bar{\sigma}$ corresponding to the translation of a type τ (assuming type definitions \bar{t}).
- An *alignment* function (impl.align) returns the alignment α used to allocate space for a τ .
- An *offset* function (impl.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 and a size and returns true if accessing size-bytes of memory at the alignment is not an error. We write $\text{impl.access}(\alpha, i)$ if this function returns true.
- The size of pointers (impl.ptrsize) is a constant i . This is a slight simplification since a C implementation could use different sizes for different pointers.
- A *literal* function (impl.xliteral) translates integer literals into byte sequences.

The access function is typically associated with hardware and the other components with compilers, but an implementation comprises all components. It is clear a “sensible” implementation 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 *impl.xtext*) are $\bar{\sigma}$, a sequence of σ . For example, if long is four bytes, the translation is byte byte byte byte, i.e., byte^4 . The type $\text{pad}[i]$ represents i bytes of padding (data of unknown type but known size). The type $\text{ptr}_\alpha(\bar{\sigma})$ describes pointers to data described by $\bar{\sigma}$ at alignment α (i.e., α is the alignment of the pointed-to data). 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 implementations or disallowing recursive types.

3.3 Dynamic Semantics

The dynamic semantics is a small-step rewrite system for expressions, parameterized by an implementation and a sequence of type definitions. Figure 4 holds the full definition for $\text{impl}; \bar{t} \vdash H; e \rightarrow H'; e'$. 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{L} and \xrightarrow{R}) for each sort of context [16]. In particular, $\text{R}[e]_r$ 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 *impl*, 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. We do not restrict the length of the byte-sequences.⁵ 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 implementation 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 is much like D-NEW (its has identical hypotheses) because we model local variables by heap allocating them. The resulting expression is the substitution of $\ast(\tau\ast)(\ell+0)$ for x .

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 $\ast(\tau\ast)(\ell+j)$) 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 *impl.access* to model alignment constraints on the memory access.

D-ASSIGN has the exact 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 *impl*. 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 $\ast(\tau\ast)(\ell+j)$.

D-FETCH uses the offset and size information from *impl* to project a subsequence of a value. We do not use the access function here because we are not accessing the heap.⁶

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

⁵We also disallow pointer values in the sequence, but allowing them would cause no problems.

⁶On 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 implementation.

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

$D ::= \text{impl}; \bar{t}$

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

Figure 4: Dynamic Semantics

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

3.4 First-Order Formulas

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

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

Consider two example constraints, which are just 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 that 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 $impl.\text{align}(\bar{t}, \tau)$ for the space allocated for x . The second formula requires a low-level subtyping relationship between two pointer types (see Section 3.5). This is the constraint our static semantics generates for a cast like in Example 1 from Section 2.

These examples also demonstrate the two flavors of formulas that arise in practice. First, there are constraints that every “sensible” implementation would satisfy. We are not interested in other implementations, but stating these requirements syntactically is much simpler than revisiting our definition of implementations. Second, there are constraints that describe semi-portable assumptions, i.e., we do not expect every implementation to satisfy them. Our static semantics produces a formula describing the assumptions of this form that a particular program makes.

The “sensible” constraints we assume for type safety are straightforward to enumerate and eminently justifiable:

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

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

Figure 5: Function Symbols for the First-Order Theory

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

Figure 6: Physical Subtyping (and Subtyping on Alignments)

These constraints are necessary for portable code in the sense that without them certain 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 three examples:

- long is at least as big as short:
 $\forall \bar{t}. xtype(\bar{t}, \text{long}) \geq xtype(\bar{t}, \text{short})$
- The first field always has offset 0:
 $\forall f, \bar{t}, \tau. (N\{\tau f \dots\} \in \bar{t}) \Rightarrow offset(\bar{t}, f) = 0$
- Fields are in order and do not overlap:
 $\forall \tau_1, f_1, \tau_2, f_2. N\{\dots \tau_1 f_1 \dots \tau_2 f_2 \dots\} \in \bar{t} \Rightarrow$
 $(offset(\bar{t}, f_1) + size(xtype(\bar{t}, \tau_1)) < offset(\bar{t}, f_2))$

3.5 Physical Subtyping

As in prior work [6, 35, 31], we use a subtyping relation on low-level types to formalize that data described by $\bar{\sigma}$ can also be treated as a $\bar{\sigma}'$. This notion has been called *physical subtyping* because it relies on actual memory layouts. The rules for our judgment $impl; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$ appear in Figure 6 (recall we abbreviate $impl; \bar{t}$ as D when their form is unimportant).

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 1 in Section 2. We also

allow assuming a less restrictive alignment (see the rules for $\vdash \alpha_1 \leq \alpha_2$), which also can only restrict how a pointer can be used.

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. We allow such use of N in order to support recursive types.

Rules `PAD` and `ADD` let us forget about the form of data (not under a pointer) without forgetting its size.⁷

Rule `SEQ` lifts subtyping to sequences. As usual, subtyping is reflexive and transitive.

As usual, subsumption (explicit or implicit) is sound for right-expressions but unsound for left-expressions. (For example, in Java, given $e_1=e_2$, one may use subsumption on e_2 but not on e_1 .) The static semantics enforces this restriction by disallowing casts as left-expressions.

3.6 Static Semantics and Constraint Generation

The preceding definitions of constraints and subtyping provide what we need to define a static semantics for source programs (Figure 7). The judgments $\bar{t}; \Gamma \vdash_r e : \tau; S$ and $\bar{t}; \Gamma \vdash_l e : \tau; S$ (for right- and left-expressions respectively) produce types as usual, but also formulas S . This formula is just a conjunction of the semi-portable assumptions the program may be making.

The only interesting rules are `S-CAST` and `S-FADDR` because the “sensible” constraints in Section 3.4 suffice to ensure other expression forms (such as dereferences and assignments) cannot fail due to an implementation dependency. The constraints directly describe the implicit assumptions made in Examples 1, 2, and 3 in Section 2. The `S-FADDR` constraint is much more complicated because it is not required in general that every subsequence of fields have an alignment appropriate for treating it as a type.

By using separate but mutually recursive typing rules for left-expressions and right-expressions we ensure left-expressions conform to a restricted grammar (as in C), namely expressions of the form x , $*(\tau*)(e_1)$, and $e_2.f$ where e_2 is itself a left-expression. Because subsumption is explicit (via casts) and casts are not left-expressions, there is no subsumption for left-expressions. Note the “cast” in dereference expressions is not really a cast (`S-DEREF` requires a type equality); the explicit type indicates a run-time size used by `D-DEREF`.

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 to ensure that the casted value actually has the result of the cast. Techniques used in existing safe-C approaches [31, 23] include implicit run-time type information, explicit discriminated unions, (bounded) parametric polymorphism, etc.

3.7 Metatheory and Low-Level Static Semantics

Safety:

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

However, this claim is not quite true because our type system does not prevent trying to use uninitialized data. Therefore, we must relax our safety guarantee to admit that e' might also be “legally stuck,” which

⁷For technical reasons we cannot allow $impl; \bar{t} \vdash \text{pad}[i+j] \leq \text{pad}[i]\text{pad}[j]$; it could allow a program to access “part of a pointer”. On one level this is sound (it could not do anything with the result because of type $\text{pad}[i]$), but our abstract machine would get stuck when trying to extract part of a pointer.

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

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

$impl; \bar{t} \vdash H; e \rightarrow H'; e'$	small step in dynamic semantics
$impl; \bar{t} \vdash H; e \xrightarrow{r} H'; e'$	primitive right-step reduction in dynamic semantics
$impl; \bar{t} \vdash H; e \xrightarrow{l} H'; e'$	primitive left-step reduction in dynamic semantics
$impl \models S$	implementation $impl$ models formula S
$\vdash \alpha_1 \leq \alpha_2$	subtyping on alignments
$impl; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$	low-level implementation-dependent subtyping
$\bar{t}; \Gamma \vdash_r e : \tau; S$	static semantics for right-expressions
$\bar{t}; \Gamma \vdash_l e : \tau; S$	static semantics for left-expressions
$impl; \bar{t}; \Psi; \Gamma \vdash_r e : \bar{\sigma}$	low-level typing for right-expressions
$impl; \bar{t}; \Psi; \Gamma \vdash_l e : \bar{\sigma}, \alpha\alpha$	low-level typing for left-expressions

Figure 8: Summary of Judgments

we define as expressions of the form $R[stuck]_r$ or $R[stuck]_l$ where:

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

The proof, available in the appendix, employs a “low-level, run-time type system” that captures the relevant invariants that evaluation preserves. The main judgment of this type system has the form $impl; \bar{t}; \Psi; \Gamma \vdash_r e : \bar{\sigma}$ where Ψ gives a type to the heap.⁸ This system has implicit subsumption, which is necessary for a step via D-CAST to preserve typing:

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

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

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

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

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

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

If $\bar{t}; \Gamma \vdash_r e : \tau; S$, $impl \models S$ and $impl$ is sensible, and $impl.xtype(\bar{t}, \tau) = \bar{\sigma}$, then $impl; \bar{t}; \Psi; \Gamma \vdash_r e : \bar{\sigma}$.

The proof, by induction on the derivation of $\bar{t}; \Gamma \vdash_r e : \tau; S$, uses the definition of \models 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 implementations model $\text{access}(\text{align}(\bar{t}, \tau), \text{size}(xtype(\bar{t}, \tau)))$.

Given the lemma above, showing preservation and progress [36] (modulo legally stuck states) for the dynamic semantics relative to the low-level type system suffices to establish type safety.

Cast-Free Portability:

Having safety rely on a portability constraint (the S in $\bar{t}; \Gamma \vdash_r e : \tau; S$) can be viewed as weak, since S could be difficult to establish, even unsatisfiable. However, we can formalize the intuitive notion that only casts

⁸ $\Psi ::= \cdot \mid \Psi, \ell \mapsto \bar{\sigma}, \alpha$

Syntax:

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

Implementations:

$$\text{impl.val}(\bar{b}) = i$$

Dynamic semantics:

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

Sensibility constraint: size is a multiple of alignment

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

Subtyping and static semantics:

$ \begin{array}{c} \text{ARR} \\ \bar{\sigma}_1 = \bar{\sigma}_2^i \quad \vdash \alpha_1 \leq \alpha_2 \\ \hline \text{impl}; \bar{t} \vdash \text{ptr}_{\alpha_1}^\omega(\bar{\sigma}_1) \leq \text{ptr}_{\alpha_2}^\omega(\bar{\sigma}_2) \end{array} $	$ \begin{array}{c} \text{S-NEWARR} \\ \bar{t}; \Gamma \vdash e : \text{long}; S \\ \hline \bar{t}; \Gamma \vdash \mathbf{new} \tau[e] : \tau^{*\omega}; S \end{array} $	$ \begin{array}{c} \text{S-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} $
---	---	---

Figure 9: Extensions for Arrays

can threaten portability. That is, for the right definition of “cast-free,” if e is cast-free and $\bar{t}; \Gamma \vdash e : \tau; S$, then every sensible impl models S . It then follows from safety that e is portable (i.e., it will not get stuck on any sensible implementation).

To be precise, a program $\bar{t}; e$ is cast-free if:

- The only occurrences of $(\tau^*)e'$ in e are in expressions of the form $*((\tau^*)e')$ and $(\tau^*)\&\tau^* \rightarrow e'f$.
- For every expression of the form $(\tau^*)\&\tau^* \rightarrow e'f$ 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 an implementation-dependent suffix cast).

4 Extensions

This section sketches how the core model we have developed is flexible enough to be extended with other relevant features of C and its implementations. We focus primarily on arrays because they are ubiquitous and require restricting our subtyping definition. Other extensions we consider merely permit additional subtyping.

4.1 Arrays

As Example 4 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 ($\tau*$ as already defined) from pointers to arrays ($\tau^{*\omega}$; the ω just distinguishes it from $\tau*$). This dichotomy is common in safe C-like languages [31, 23], can be approximated via static analysis over C code, and is necessary to identify what implementation assumptions are due only to arrays. The low-level types (σ) make the same distinction.

We add two right-expression forms. First, `new $\tau[e]$` creates a pointer to a heap-allocated array of “length” e . (Because e will evaluate to a byte-sequence \bar{b} , an implementation must interpret \bar{b} as an integer; we use *impl.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 `new $\tau[e]$` from being stuck if e has uninitialized bytes or e is negative.⁹

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 then be dereferenced or assigned through. (Adding a form that produces a $\tau^{*\omega}$ type instead of a $\tau*$ would work fine, but we omit it here.) 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.¹⁰ By performing the bounds-check on $\&((\tau^{*\omega})(e_1))[e_2]$, we ensure an ensuing dereference can never fail.

With this economical addition of arrays to the abstract machine, we can design constraints and subtyping such that the only failures are bounds-checks (not unaligned memory accesses or treating bytes as pointers). 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 ? (This issue does not arise with $(\tau*)\&\tau* \rightarrow ef$ precisely because the offset of f is known statically.) The solution taken by every sensible C implementation 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. Interestingly, this sensibility constraint is not needed without arrays.

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 times (for some i). This is more lenient than strict invariance. For example, it supports the semi-portable practice 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. (See 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.

An alternate approach could allow such subtyping for pointers to arrays of unknown size, but change the run-time behavior of casts to check that the array length is an appropriate multiple. In any case, our model provides an excellent starting point for considering subtle variations of subtyping and how it relies on implementation dependencies, which is exactly the goal of our work.

4.2 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 $\tau*$` describes pointers that cannot be used to write to the pointed-to data (though

⁹In C, e is unsigned, but large allocations due to conversion from negative numbers are a well-known defect source.

¹⁰This check disallows pointing just past the end of the array, unlike C.

the lack of qualifier polymorphism [15] 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` 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`:

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

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

4.3 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 6 showed a contrived example 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 $\text{impl.xtype}(\bar{t}, \tau)$. 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:

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

4.4 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 $\text{impl.xtype}(\bar{t}, N_1) = \text{byte}^4 \text{ptr}_\alpha(N_1)$ and $\text{impl.xtype}(\bar{t}, N_2) = \text{byte}^4 \text{ptr}_\alpha(N_2)$. We would be unable to show $\text{impl}; \bar{t} \vdash \text{byte}^4 \text{ptr}_\alpha(N_1) \leq \text{byte}^4 \text{ptr}_\alpha(N_2)$.

Sound solutions to this sort of limitation are well-known [3]. 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 $\text{ptr}_\alpha(N_1) \leq \text{ptr}_\alpha(N_2)$.

The issues regarding subtyping recursive types are the usual ones. We have opted for simplicity of presentation over expressiveness. In an actual tool, one should prefer expressiveness.

5 Previous Work

To our knowledge, previous work considering implementation-dependent data-layout assumptions or low-level type-safety has taken one of a few approaches:

- Consider only one or two implementations (e.g., verify that one compiler to one architecture produces type-safe code) 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.
- Restrict C, e.g., prohibit dynamic memory allocation.

The first approach is less helpful for writing semi-portable code because we must reverify the code for each implementation. The second approach is good only for (parts of) applications that should be written in a higher-level language. The third and fourth approaches essentially relegate some issues to work such as ours, much as we relegate some issues like array-bounds errors to other work [9, 12].

5.1 Assuming an Implementation

Most closely related is the “physical type-checking” work of Chandra et al. [6, 35], which motivated our work considerably. Their tool classifies C casts as “upcasts”, “downcasts”, or “neither”, reporting a warning for the last possibility. They take a byte-for-byte view of memory for a low-level type system, but they neither parameterize their system by an implementation nor produce descriptions of sets of implementations. Hence checking code against a new implementation would require reverification and changing their tool. They present no metatheory validating their approach.

CCured [31], a memory-safe C implementation, 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 implementations. The formal model establishing CCured’s soundness shares similarities with our work (particularly the subtyping between arrays and single-objects), but it lacks a distinct notion of *implementation*, alignment constraints, local variables with `struct` types, memory allocation, `const`, recursive types, etc.

Work on typed assembly language and proof-carrying code [29, 28, 10, 7, 20] 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 [33, 1] 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 semi-portable code because verification of type-safety is repeated for each implementation. In practice, defining a new implementation is an enormous amount of work.

For external data such as network packets [27] or *ad hoc* data streams [14], one can produce parsers directly from explicit bit-by-bit type declarations. While this is a robust solution for external data, it is less appropriate for describing implementation decisions for program data. For semi-portable code, we want to specify some data-representation constraints, but not fix every bit.

5.2 Safe C

Memory-safe dialects or implementations of C, such as Cyclone [23, 18, 17], CCured [31, 30, 8], and SAFE-Code [11], do not solve the semi-portability problem. Rather, they may reject (at compile-time) or terminate

(at run-time) programs that attempt implementation-dependent operations, or they may support only certain implementations (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 all these systems include “run-time systems” (automatic memory managers, type-tag checkers, etc.) that are themselves semi-portable! For example, the Cyclone run-time system assumes 32-bit integers and pointers, and making this code more portable is a top request from actual users.

5.3 Formalizing C

Recent work by Leroy et al. [24, 5] uses Coq to prove a C compiler correct. Their (large-step) 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 implementation-dependent code).

Norrish’s formalization of C [32] uses HOL and includes structs. Like in our work, he uses a global namespace mapping struct names to sequences of typed fields. However, he purposely omits padding and alignment from his formalism. He has no separable notion of an implementation; instead he models implementation choices as non-determinism.

5.4 Low-Level Code without C

Our work has been somewhat C-centric, whereas other projects have started with languages at higher levels of abstraction and added bit-level views for low-level programming. (See [4, 19] 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. That is, our model provides a foundation for semi-portable types, which are useful for building low-level but retargetable systems in any language.

C-- [34] makes data representation and alignment explicit, but C-- is not a platform for writing semi-portable 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 just 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.

6 Conclusions

This work has developed a formal description of implementation-dependencies in low-level software. The key insight is a semantic definition of “implementation” 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 implementations a clear identity in our framework clarifies a number of poorly understood issues.

The technique of implementation-as-parameter can apply broadly since high-level languages also have implementation-defined behavior. As examples, SML programs may depend on the size of `int`, Scheme programs may depend on evaluation order, or Java programs may depend on fair thread-scheduling. A program property relying on such assumptions would be stronger than memory-safety (which is always assured); the point is describing implementations via constraints would let us identify a program’s assumptions.

Our next step is to complete a practical tool to determine implementation assumptions of real C code. We expect the simple static analysis formalized here will prove too imprecise. However, the problem decomposition in this work still applies; we just replace the “high-level type system” with a more accurate

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

producer of an implementation constraint. In fact, for a bug-finding porting tool, we may wish to sacrifice soundness to reduce false positives.

Another area for future work is a stronger guarantee, namely that a program has similar (ideally equivalent) behavior on a set of implementations. Tackling this problem requires arithmetic reasoning in the presence of overflow and different endiannesses.

7 Acknowledgments

Matthew Fluet, Greg Morrisett, and David Walker provided excellent feedback on the presentation.

References

- [1] A. Ahmed and D. Walker. The logical approach to stack typing. In *ACM International Workshop on Types in Language Design and Implementation*, pages 74–85, New Orleans, LA, Jan. 2003.
- [2] *The ARMLinux Book Online*, Chapter 10. May 2005.
<http://www.aleph1.co.uk/armlinux/book>.
- [3] R. M. Amadio and L. Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4):575–631, Sept. 1993.
- [4] D. F. Bacon. Kava: a Java dialect with a uniform object model for lightweight classes. *Concurrency and Computation: Practice and Experience*, 15(3–5):185–206, 2003.
- [5] S. Blazy, Z. Dargaye, and X. Leroy. Formal verification of a C compiler front-end. In *14th International Symposium on Formal Methods*, Aug. 2006.
- [6] S. Chandra and T. Reps. Physical type checking for C. In *ACM Workshop on Program Analysis for Software Tools and Engineering*, pages 66–75, Toulouse, France, Sept. 1999.
- [7] 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*, pages 208–219, San Diego, CA, June 2003.
- [8] 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*, pages 232–244, June 2003.
- [9] P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. The ASTRÉE Analyser. In *European Symposium on Programming*, volume 3444 of *Lecture Notes in Computer Science*, pages 21–30. Springer-Verlag, Apr. 2005.
- [10] K. Crary. Toward a foundational typed assembly language. In *30th ACM Symposium on Principles of Programming Languages*, pages 198–212, New Orleans, LA, Jan. 2003.
- [11] D. Dhurjati, S. Kowshik, and V. Adve. SAFECode: Enforcing alias analysis for weakly typed languages. In *ACM Conference on Programming Language Design and Implementation*, Ottawa, Canada, June 2006.
- [12] N. Dor, M. Rodeh, and M. Sagiv. CSSV: towards a realistic tool for statically detecting all buffer overflows in C. In *ACM Conference on Programming Language Design and Implementation*, pages 155–167, San Diego, CA, June 2003.
- [13] 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*, pages 1–14, San Antonio, TX, Jan. 1999.

- [14] K. Fisher, Y. Mandelbaum, and D. Walker. The next 700 data description languages. In *33rd ACM Symposium on Principles of Programming Languages*, pages 2–15, Charleston, SC, Jan. 2006.
- [15] J. Foster, M. Fähndrich, and A. Aiken. A theory of type qualifiers. In *ACM Conference on Programming Language Design and Implementation*, pages 192–203, Atlanta, GA, May 1999.
- [16] D. Grossman. Type-safe multithreading in Cyclone. In *ACM International Workshop on Types in Language Design and Implementation*, pages 13–25, New Orleans, LA, Jan. 2003.
- [17] D. Grossman. Quantified types in imperative languages. *ACM Transactions on Programming Languages and Systems*, 28(3):429–475, May 2006.
- [18] 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*, pages 282–293, Berlin, Germany, June 2002.
- [19] 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*, pages 116–128, Tallinn, Estonia, Sept. 2005.
- [20] 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):191–229, Sept. 2003.
- [21] IBM. Developing embedded software for the IBM PowerPC 970FX processor. Application Note 970, IBM, July 2004. <http://www.ibm.com/chips/techlib/>.
- [22] *ISO/IEC 9899:1999, International Standard—Programming Languages—C*. International Standards Organization, 1999.
- [23] T. Jim, G. Morrisett, D. Grossman, M. Hicks, J. Cheney, and Y. Wang. Cyclone: A safe dialect of C. In *USENIX Annual Technical Conference*, pages 275–288, Monterey, CA, June 2002.
- [24] X. Leroy. Formal certification of a compiler back-end or: Programming a compiler with a proof assistant. In *33rd ACM Symposium on Principles of Programming Languages*, pages 42–54, Jan. 2006.
- [25] R. Love. *Linux Kernel Development, Second Edition*. Novell Press, 2005. Page 328.
- [26] B. Martin, A. Rettinger, and J. Singh. Multiplatform porting to 64 bits. *Dr. Dobbs’s Journal*, Dec. 2005. <http://www.ddj.com/184406427>.
- [27] P. J. McCann and S. Chandra. Packet types: abstract specification of network protocol messages. In *SIGCOMM ’00: Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 321–333, 2000.
- [28] 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):528–569, May 1999.
- [29] G. Necula. Proof-carrying code. In *24th ACM Symposium on Principles of Programming Languages*, pages 106–119, Paris, France, Jan. 1997.
- [30] G. Necula, S. McPeak, and W. Weimer. CCured: Type-safe retrofitting of legacy code. In *29th ACM Symposium on Principles of Programming Languages*, pages 128–139, Portland, OR, Jan. 2002.
- [31] G. C. 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):477–526, May 2005.
- [32] M. Norrish. *C formalised in HOL*. PhD thesis, University of Cambridge, 1998.

- [33] L. Petersen, R. Harper, K. Cray, and F. Pfenning. A type theory for memory allocation and data layout. In *30th ACM Symposium on Principles of Programming Languages*, pages 172–184, New Orleans, LA, Jan. 2003.
- [34] N. Ramsey, S. P. Jones, and C. Lindig. The C-- language specification version 2.0, Feb. 2005. <http://www.cminusminus.org/extern/man2.pdf>.
- [35] M. Siff, S. Chandra, T. Ball, K. Kunchithapadam, and T. Reps. Coping with type casts in C. In *7th European Software Engineering Conference and 7th ACM Symposium on the Foundations of Software Engineering*, pages 180–198, Toulouse, France, Sept. 1999.
- [36] A. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.

A Safety Proof

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 let \cdot denote the empty sequence. Figure 10 gives the syntax for a C-like source language. Figure 12 gives an extended, “low-level” version of this language, where the types are replaced by types at the byte-sequence level, runtime values are added, and syntax for heaps, heap typings, and evaluation contexts is also shown. Figure 13 gives the dynamic semantics for the low-level language and Figure 14 defines a set of value and type size functions used throughout the low-level semantics. Figures 15 and 16 give the static semantics for the low-level language. Figure 17 gives the static semantics for the high-level language.

A.1 Implementations

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

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

A.2 Constraint Language

The high-level semantics (Figure 17) 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 implementation component plus a couple of others. Implementations are models in this theory; we write $\text{impl} \models S$ when impl satisfies formula S . The (\models) judgment is entirely ordinary, defined inductively over the structure of S . For example, $\text{impl} \models S_1 \wedge S_2$ exactly when $\text{impl} \models S_1$ and

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

Figure 10: Syntax for the Source Language

symbol	interpretation under <i>impl</i>
$\text{xtype}(\overline{t}, \tau)$	$\text{impl.xtype}(\overline{t}, \tau)$
$\text{align}(\overline{t}, \tau)$	$\text{impl.align}(\overline{t}, \tau)$
$\text{offset}(\overline{t}, f)$	$\text{impl.offset}(\overline{t}, f)$
$\text{access}(\alpha, i)$	$\text{impl.access}(\alpha, i)$
$\text{xlit}(s)$	$\text{impl.xlit}(s)$
$\text{xlit}(l)$	$\text{impl.xlit}(l)$
ptr_size	impl.ptr_size
$\text{size}(\overline{\sigma})$	$\text{size}(\text{impl}, \overline{\sigma})$
$\text{size}(\overline{w})$	$\text{size}(\text{impl}, \overline{w})$
$\text{subtype}(\overline{t}, \overline{\sigma}_1, \overline{\sigma}_2)$	$\text{impl}; \overline{t} \vdash \overline{\sigma}_1 \leq \overline{\sigma}_2$ is derivable
$\text{subalign}(\alpha_1, \alpha_2)$	$\vdash \alpha_1 \leq \alpha_2$ is derivable

Figure 11: Constraint Language Function Symbols

$\text{impl} \models S_2$; and $\text{impl} \models \forall x : s.S$ 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 11 summarizes the function symbols and their interpretation under *impl*. As a final example, involving function symbols, suppose we have a constraint $S = (\text{align}(\overline{t}, N) = [a, o] \wedge \text{align}(\overline{t}, \tau) = [a, o + \text{offset}(\overline{t}, f)])$ and want to know whether an implementation *impl* satisfies S (alternatively, *impl* models S). *impl* satisfies S iff $\text{impl} \models \text{align}(\overline{t}, N) = [a, o]$ and $\text{impl} \models \text{align}(\overline{t}, \tau) = [a, o + \text{offset}(\overline{t}, f)]$. To satisfy the former conjunct, $\text{impl.align}(\overline{t}, N) = [a, o]$ must be true. To satisfy the latter, $\text{impl.align}(\overline{t}, \tau) = [a, o + \text{impl.offset}(\overline{t}, f)]$ must be true.

A.3 Metatheory

Definition 1 (Sensible Implementations)

An implementation *impl* is said to be *sensible* if

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

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 $\text{impl.align}(\overline{t}, N) = [8, 0]$, $\text{impl.offset}(\overline{t}, f) = 8$, and $\text{impl.align}(\overline{t}, \tau) = [4, 0]$. Then, clause (3) would require τ to be 8-byte aligned (alignment $[8, 0 + \text{impl.offset}(\overline{t}, f)] \equiv [8, 8] \equiv [8, 0]$), which is unreasonable. It is okay for τ to be 4-byte, 2-byte, or 1-byte aligned.

		σ	$::=$	$\text{byte} \mid \text{pad}[i] \mid \text{ptr}_\alpha(\bar{\sigma}) \mid \text{ptr}_\alpha(N)$
(types)		α	$::=$	$[a, o]$
(alignments)		w	$::=$	$b \mid \ell+i \mid \text{uninit}$
(atomic values)		b	$::=$	$0 \mid 1 \mid \dots \mid 255$
(bytes)		e	$::=$	$\dots \mid \bar{w}$
(expressions)		v	$::=$	\bar{w}
(values)		H	$::=$	$\frac{\ell \mapsto v, \alpha}{\ell}$
(heaps)		Ψ	$::=$	$\frac{\ell: \bar{\sigma}, \alpha}{\ell}$
(heap typings)		D	$::=$	$\text{impl}; \bar{t}$
		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 12: Syntax for the Full Language (Extends Source Language) and Eval Contexts

4. $\forall \tau \exists \alpha, \bar{\sigma}. \text{impl.xtype}(\bar{t}, \tau*) = \text{ptr}_\alpha(\bar{\sigma})$ and $\text{impl.xtype}(\bar{t}, \tau) = \bar{\sigma}$ and $\text{impl.align}(\bar{t}, \tau) = \alpha$.

5. If $\text{impl.access}(\alpha, i)$ and $\vdash \alpha' \leq \alpha$ then $\text{impl.access}(\alpha', i)$.

Definition 2 (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 3 (Legal Stuck State) Let

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

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

- $e = R[\text{lstuck}]_l$
- $e = R[\text{rstuck}]_r$
- $e = L[\text{lstuck}]_l$
- $e = L[\text{rstuck}]_r$

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

Figure 13: Dynamic Semantics

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

Figure 14: Size Functions

Definition 4 (Illegal Stuck State)

A state $H; e$ is *illegally stuck* on an implementation $impl$ 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 $impl; \bar{t} \vdash H; e \rightarrow H'; e'$.

Definition 5 (Heap Typing Extension)

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

A.3.1 Type Soundness

Lemma 6 (Weakening)

1. (a) If $D; \Psi; \Gamma \vdash e : \bar{\sigma}$ then $D; \Psi\Psi'; \Gamma\Gamma' \vdash e : \bar{\sigma}$.
(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 7 (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 $D; \Psi; \cdot \vdash \bar{w} : \bar{\sigma}$.

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

Lemma 8 (Uninit Type)

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

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 $\text{size}(impl, \bar{\sigma}') = j$ and

- 1 $impl; \bar{t}; \Psi; \cdot \vdash \text{uninit}^j : \bar{\sigma}'$

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

Figure 15: Static Semantics for the Low-Level Language

 Heap Typing

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

 Alignment Subtyping

$$\frac{\text{ALST-OFF} \quad o_1 \equiv o_2 \pmod{a}}{\vdash [a, o_1] \leq [a, o_2]} \qquad \frac{\text{ALST-BASE} \quad a_1 = a_2 \times k}{\vdash [a_1, o] \leq [a_2, o]} \qquad \frac{\text{ALST-TRANS} \quad \vdash \alpha_1 \leq \alpha_2 \quad \vdash \alpha_2 \leq \alpha_3}{\vdash \alpha_1 \leq \alpha_3}$$

 Physical Subtyping

$$\frac{\text{ST-UNROLL} \quad \text{impl.xtext}(\bar{t}, N) = \bar{\sigma}}{\text{impl}; \bar{t} \vdash \text{ptr}_\alpha(N) \leq \text{ptr}_\alpha(\bar{\sigma})} \qquad \frac{\text{ST-ROLL} \quad \text{impl.xtext}(\bar{t}, N) = \bar{\sigma}}{\text{impl}; \bar{t} \vdash \text{ptr}_\alpha(\bar{\sigma}) \leq \text{ptr}_\alpha(N)} \qquad \frac{\text{ST-PTR} \quad \vdash \alpha_1 \leq \alpha_2}{D \vdash \text{ptr}_{\alpha_1}(\bar{\sigma}_1 \bar{\sigma}_2) \leq \text{ptr}_{\alpha_2}(\bar{\sigma}_1)}$$

$$\frac{\text{ST-PAD} \quad \text{size}(\text{impl}, \sigma) = i}{\text{impl}; \bar{t} \vdash \sigma \leq \text{pad}[i]} \qquad \frac{\text{ST-PADADD}}{\text{impl}; \bar{t} \vdash \text{pad}[i] \text{pad}[j] \leq \text{pad}[i+j]} \qquad \frac{\text{ST-REFL}}{D \vdash \bar{\sigma} \leq \bar{\sigma}}$$

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

 Value Typing

$$\frac{\text{LW-BYTE}}{D; \Psi \Vdash b : \text{byte}} \qquad \frac{\text{LW-LBL} \quad \Psi(\ell) = \bar{\sigma}_1 \bar{\sigma}_2, [a, o] \quad i = \text{size}(\text{impl}, \bar{\sigma}_1)}{\text{impl}; \bar{t}; \Psi \Vdash \ell + i : \text{ptr}_{[a, o+i]}(\bar{\sigma}_2)} \qquad \frac{\text{LW-UNINIT1}}{D; \Psi \Vdash \text{uninit} : \text{byte}}$$

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

Figure 16: Auxiliary Relations

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

$$\mathbf{2} \quad \text{impl}; \bar{t}; \Psi; \cdot \Vdash \text{uninit}^k : \sigma$$

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 $\text{impl}; \bar{t}; \Psi; \cdot \vdash \text{uninit}^i : \text{ptr}_\alpha(\bar{\sigma})$ where $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}$, then use ST-ROLL and L-SUB to assign type $\text{ptr}_\alpha(N)$ to uninit^i .

Lemma 9 (Constant-Size Subtyping)

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

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

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

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} $\frac{\Gamma(x) = \tau}{\bar{t}; \Gamma \vdash x : \tau; \top}$ $\frac{}{\bar{t}; \Gamma \vdash x : \tau; \top}$
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} $\frac{\bar{t}; \Gamma \vdash e : \tau*; S}{\bar{t}; \Gamma \vdash *(\tau*)(e) : \tau; S}$ $\frac{}{\bar{t}; \Gamma \vdash *(\tau*)(e) : \tau; S}$	
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 $\frac{\bar{t}; \Gamma \vdash e_1 : \text{long}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2 \quad \bar{t}; \Gamma \vdash e_3 : \tau; S_3}{\bar{t}; \Gamma \vdash \text{if } e_1 \ e_2 \ e_3 : \tau; S_1 \wedge S_2 \wedge S_3}$	H-WHILE $\frac{\bar{t}; \Gamma \vdash e_1 : \text{long}; S_1 \quad \bar{t}; \Gamma \vdash e_2 : \tau; S_2}{\bar{t}; \Gamma \vdash \text{while } e_1 \ e_2 : \text{short}; S_1 \wedge S_2}$		
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 17: Static semantics for high level language

- ST-PADADD: Immediate.
- ST-REFL: Immediate.
- ST-SEQ: By induction, $\text{size}(\text{impl}, \bar{\sigma}_1) = \text{size}(\text{impl}, \bar{\sigma}_2)$ and $\text{size}(\text{impl}, \bar{\sigma}_3) = \text{size}(\text{impl}, \bar{\sigma}_4)$. It follows that $\text{size}(\text{impl}, \bar{\sigma}_1) + \text{size}(\text{impl}, \bar{\sigma}_3) = \text{size}(\text{impl}, \bar{\sigma}_2) + \text{size}(\text{impl}, \bar{\sigma}_4)$.
- ST-TRANS: By induction twice, $\text{size}(\text{impl}, \bar{\sigma}_1) = \text{size}(\text{impl}, \bar{\sigma}_2) = \text{size}(\text{impl}, \bar{\sigma}_3)$.

Lemma 10 (Value-Type Size)

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

Proof:

1. By inspection of the assumed typing derivation:

- LW-BYTE: $\text{size}(\text{impl}, b) = \text{size}(\text{impl}, \text{byte}) = 1$.
- LW-LBL: $\text{size}(\text{impl}, \ell+i) = \text{size}(\text{impl}, \text{ptr}_\alpha(\bar{\sigma}_2)) = \text{impl.ptr_size}$.
- LW-UNINIT1: $\text{size}(\text{impl}, \text{uninit}) = \text{size}(\text{impl}, \text{byte}) = 1$.
- LW-UNINIT2: $\text{size}(\text{impl}, \text{uninit}^i) = \text{size}(\text{impl}, \text{ptr}_\alpha(\bar{\sigma})) = \text{impl.ptr_size} = i$.

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 $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\bar{r}} \bar{w}_1 \bar{w}_2 : \sigma_1 \bar{\sigma}_2$ where inversion gives $\text{impl}; \bar{t}; \Psi \vdash_w \bar{w}_1 : \sigma_1$ and $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\bar{r}} \bar{w}_2 : \bar{\sigma}_2$. Part (1) of the theorem gives $\text{size}(\text{impl}, \bar{w}_1) = \text{size}(\text{impl}, \sigma_1)$. The induction hypothesis gives $\text{size}(\text{impl}, \bar{w}_2) = \text{size}(\text{impl}, \bar{\sigma}_2)$. It follows that $\text{size}(\text{impl}, \bar{w}_1 \bar{w}_2) = \text{size}(\text{impl}, \bar{\sigma}_1 \bar{\sigma}_2)$.
- L-SUB: Follows from induction and the Constant-Size Subtyping Lemma.

Lemma 11 (Sequence Typing)

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

Proof:

1. By induction on the derivation of $D; C \vdash_{\bar{r}} \bar{w}_1 : \bar{\sigma}_1$, 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 $D; C \vdash_{\bar{r}} \bar{w}_1 : \bar{\sigma}_3$ and $D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_1$. So by induction $D; C \vdash_{\bar{r}} \bar{w}_1 \bar{w}_2 : \bar{\sigma}_3 \bar{\sigma}_2$. So with $D \vdash \bar{\sigma}_3 \leq \bar{\sigma}_1$, ST-SEQ, ST-REFL, and L-SUB we can derive $D; C \vdash_{\bar{r}} \bar{w}_1 \bar{w}_2 : \bar{\sigma}_1 \bar{\sigma}_2$.

2. By induction on n . Cases $n = 0$ and $n = 1$ are trivial. For $n > 1$, use induction and part(1).

Lemma 12 (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 13 (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 14 (Typing Partition)

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

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 $D; C \vdash_{\bar{\tau}} \bar{w} : \sigma'_1 \dots \sigma'_m$ 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 $1 \leq i \leq m$, $D; C \vdash_{\bar{\tau}} \bar{w}'_i : \sigma'_i$. 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 $\sigma'_1 \dots \sigma'_j$ that is $\bar{\sigma}'_k$ and the corresponding $\bar{w}'_1 \dots \bar{w}'_j$ and use the Sequence Typing Lemma to conclude $D; C \vdash_{\bar{\tau}} \bar{w}'_1 \dots \bar{w}'_j : \bar{\sigma}'_k$ and then L-SUB to conclude $D; C \vdash_{\bar{\tau}} \bar{w}'_1 \dots \bar{w}'_j : \sigma_k$. So letting $\bar{w}'_1 \dots \bar{w}'_j = \bar{w}_k$ suffices.

Lemma 15 (Value Split)

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

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 $D; C \vdash_{\bar{\tau}} \bar{w}_{11} : \sigma_{11}, \dots, D; C \vdash_{\bar{\tau}} \bar{w}_{1n} : \sigma_{1n}, D; C \vdash_{\bar{\tau}} \bar{w}_{21} : \sigma_{21}, \dots, D; C \vdash_{\bar{\tau}} \bar{w}_{2m} : \sigma_{2m}$. 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 $\bar{w} = \bar{w}_1 \bar{w}'_2$, $D; C \vdash_{\bar{\tau}} \bar{w}_1 : \bar{\sigma}_1$, and $D; C \vdash_{\bar{\tau}} \bar{w}'_2 : \bar{\sigma}'_2$. Applying part 1 again to the last conclusion provides $D; C \vdash_{\bar{\tau}} \bar{w}_2 : \bar{\sigma}_2$ and $D; C \vdash_{\bar{\tau}} \bar{w}_3 : \bar{\sigma}_3$. (Note we can extend this to n by a trivial induction should we need it).

Non-Lemma: (Type-Split) If $D; C \vdash_{\bar{\tau}} \bar{w}_1 \bar{w}_2 : \bar{\sigma}$, then $\exists \bar{\sigma}_1, \bar{\sigma}_2$ such that $\bar{\sigma} = \bar{\sigma}_1 \bar{\sigma}_2$, $D; C \vdash_{\bar{\tau}} \bar{w}_1 : \bar{\sigma}_1$, and $D; C \vdash_{\bar{\tau}} \bar{w}_2 : \bar{\sigma}_2$.

Non-Proof: This is FALSE. Let $\bar{w}_1 = \text{uninit}^2$, $\bar{w}_2 = \text{uninit}^2$ and $\bar{\sigma} = \text{ptr}_{\alpha}(\cdot)$ when the size of pointers is 4.

Lemma 16 (Value-Type Split) *If*

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

then

$$\begin{aligned} & \text{impl}; \bar{t}; C \vdash_{\bar{\tau}} \bar{w}_1 : \bar{\sigma}_1 \\ & \text{impl}; \bar{t}; C \vdash_{\bar{\tau}} \bar{w}_2 : \bar{\sigma}_2 \end{aligned}$$

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, $\bar{w}_1 = \bar{w}_{11} \bar{w}_{12}$, $\bar{\sigma}_1 = \sigma_{11} \bar{\sigma}_{12}$,
 - 1 $\text{impl}; \bar{t}; \Psi \vdash_{\bar{w}} \bar{w}_{11} : \sigma_{11}$

2 $impl; \bar{t}; \Psi; \cdot \vdash \bar{w}_{12}\bar{w}_2 : \bar{\sigma}_{12}\bar{\sigma}_2$

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

3 $size(impl, \bar{w}_{12}) = size(impl, \bar{\sigma}_{12})$

From (2,3), the induction hypothesis yields

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

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

Plugging (1,4) into L-VALUE gives $impl; \bar{t}; \Psi; \cdot \vdash \bar{w}_{11}\bar{w}_{12} : \sigma_{11}\bar{\sigma}_{12}$, 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 $size(impl, \bar{\sigma}'_1) = size(impl, \bar{\sigma}_1)$, so the assumption $size(impl, \bar{w}_1) = size(impl, \bar{\sigma}_1)$ ensures $size(impl, \bar{w}'_1) = size(impl, \bar{\sigma}'_1)$. So by induction $impl; \bar{t}; C \vdash \bar{w}_1 : \bar{\sigma}'_1$ and $impl; \bar{t}; C \vdash \bar{w}_2 : \bar{\sigma}'_2$. So with one use of L-SUB each, we can derive $impl; \bar{t}; C \vdash \bar{w}_1 : \bar{\sigma}_1$ and $impl; \bar{t}; C \vdash \bar{w}_2 : \bar{\sigma}_2$.

Lemma 17 (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 18 (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 19 (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 20 (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 21 (Subtyping Type Form)

- If $impl; \bar{t} \vdash \bar{\sigma}' \leq \text{ptr}_\alpha(\bar{\sigma})$ then $\exists \bar{\sigma}'', \alpha'$ such that $\bar{\sigma}' = \text{ptr}_{\alpha'}(\bar{\sigma}\bar{\sigma}'')$ or $\bar{\sigma}' = \text{ptr}_{\alpha'}(N)$ where $impl.xtype(\bar{t}, N) = \bar{\sigma}\bar{\sigma}'$, and $\vdash \alpha' \leq \alpha$.
- If $impl; \bar{t} \vdash \bar{\sigma}' \leq \text{ptr}_\alpha(N)$ and $impl.xtype(\bar{t}, N) = \bar{\sigma}$ then $\exists \bar{\sigma}'', \alpha'$ such that $\bar{\sigma}' = \text{ptr}_{\alpha'}(\bar{\sigma}\bar{\sigma}'')$ or $\bar{\sigma}' = \text{ptr}_{\alpha'}(N')$ where $impl.xtype(\bar{t}, N') = \bar{\sigma}\bar{\sigma}'$, 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 $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}$. 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 $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}$. 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 $\text{impl}; \bar{t} \vdash \bar{\sigma}_1\bar{\sigma}_3 \leq \bar{\sigma}_2\bar{\sigma}_4$. 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 $\text{impl}; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_3$ where inverting gives $\text{impl}; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$ and $\text{impl}; \bar{t} \vdash \bar{\sigma}_2 \leq \bar{\sigma}_3$. There are two possibilities for $\bar{\sigma}_3$:
 1. $\bar{\sigma}_3 = \text{ptr}_\alpha(N)$ where $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}'$. 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 $\text{impl.xtext}(\bar{t}, N') = \bar{\sigma}'\bar{\sigma}''\bar{\sigma}'''$ 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 $\text{impl.xtext}(\bar{t}, N') = \bar{\sigma}'\bar{\sigma}''$ 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 $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}'\bar{\sigma}''\bar{\sigma}'''$ 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 22 (Canonical Forms)

1. If $D; \Psi; \cdot \vdash_{\bar{\Gamma}} \bar{w} : \cdot$ then $\bar{w} = \cdot$.
2. If $D; \Psi; \cdot \vdash_{\bar{\Gamma}} \bar{w} : \text{byte}$ then $\bar{w} = b$ or $\bar{w} = \text{uninit}$.
3. If $D; \Psi; \cdot \vdash_{\bar{\Gamma}} \bar{w} : \text{byte}^k$ 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 $\text{impl}; \bar{t}; \Psi; \cdot \vdash_{\bar{\Gamma}} \bar{w} : \text{ptr}_\alpha(\bar{\sigma}_1)$ then $\exists \bar{\sigma}_0, \bar{\sigma}_2$ s.t. either
 - $\bar{w} = \text{uninit}^i$ where $\text{impl.ptr_size} = i$.
 - $\bar{w} = \ell+i$ 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 $\text{size}(\text{impl}, \bar{\sigma}_0) = i$.
 - If $\text{impl}; \bar{t}; \Psi; \cdot \vdash_{\bar{\Gamma}} \bar{w} : \text{ptr}_\alpha(N)$ where $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}_1$, then $\exists \bar{\sigma}_0, \bar{\sigma}_2$ s.t. either
 - $\bar{w} = \text{uninit}^i$ where $\text{impl.ptr_size} = i$.
 - $\bar{w} = \ell+i$ 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 $\text{size}(\text{impl}, \bar{\sigma}_0) = i$.

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 $D; \Psi; \cdot \vdash_{\bar{\Gamma}} \bar{w} : \text{byte}$. 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 $D; \Psi \vdash_{\bar{\Gamma}} w : \text{byte}$ and $D; \Psi; \cdot \vdash_{\bar{\Gamma}} \bar{w}' : \text{byte}^{k-1}$. 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 18) is used implicitly. There are only two possibilities:

- L-VALUE: There are two sub-cases:
 - Assume $D; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \text{ptr}_{\alpha}(\bar{\sigma}_1)$. Inversion gives $D; \Psi; \cdot \vdash_{\bar{w}} \bar{w} : \text{ptr}_{\alpha}(\bar{\sigma}_1)$. 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 $D; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \text{ptr}_{[a, o+i]}(N)$, where $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}_1$. Vacuous; cannot be derived via L-VALUE.
- L-SUB: There are two sub-cases:
 - Assume $D; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \text{ptr}_{\alpha}(\bar{\sigma}_1)$. Inversion gives $D; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \bar{\sigma}$ 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 $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}_1 \bar{\sigma}_k$ for some $\bar{\sigma}_k$, and $\vdash \alpha' \leq \alpha$. Applying the IH to $D; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \text{ptr}_{\alpha'}(N)$ 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 $\text{size}(\text{impl}, \bar{\sigma}_0) = i$. 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 $D; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \text{ptr}_{\alpha}(N)$, where $\text{impl.xtext}(\bar{t}, N) = \bar{\sigma}_1$. Similar to the above case.

Lemma 23 (Context Reordering)

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

Proof: Straightforward induction.

Lemma 24 (Substitution Preserves Types)

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

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

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 $\text{impl}; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\Gamma} y : \bar{\sigma}_1$ and it follows that $\text{impl.xtext}(\bar{t}, \tau) = \bar{\sigma}_1$. If $x = y$ then $\bar{\sigma}_1 = \bar{\sigma}_2$ and $y\{e_2/x\} = e_2$, and so $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\Gamma} e_2 : \bar{\sigma}_1$ 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 $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\Gamma} y : \bar{\sigma}_1$.
- L-VARL: We assume $\text{impl}; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\Gamma} y : \bar{\sigma}_1, \alpha$ and it follows that $\text{impl.xtext}(\bar{t}, \tau) = \bar{\sigma}_1$ and $\text{impl.align}(\bar{t}, \tau) = \alpha$. If $x = y$ then $\bar{\sigma}_1 = \bar{\sigma}_2$ and $y\{e_2/x\} = e_2$, and so $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\Gamma} e_2 : \bar{\sigma}_1, \alpha$ 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 $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\Gamma} y : \bar{\sigma}_1, \alpha$.
- L-ASSN: We assume $\text{impl}; \bar{t}; \Psi; \Gamma, x : \tau \vdash_{\Gamma} e = e' : \bar{\sigma}_1$. By inversion and induction (once for left-typing and once for right-typing),
 - 1 $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\Gamma} e\{e_2/x\} : \bar{\sigma}_1, \alpha$
 - 2 $\text{impl}; \bar{t}; \Psi; \Gamma \vdash_{\Gamma} e'\{e_2/x\} : \bar{\sigma}_1$

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

- L-FETCHL: Follows from inversion, induction, and an application of L-FETCHL.
- L-FETCHR: Follows from inversion, induction, and an application of L-FETCHR.
- L-DEREF{L,R}: Inversion, induction, and an application of L-DEREF{L,R}.
- L-CAST: Inversion, induction, and an application of L-CAST.
- L-FADDR: Inversion, induction, and an application of L-FADDR.
- L-SEQ: Inversion, induction applied to each subexpression, and plugging the results into L-SEQ.
- L-IF: Similar to L-SEQ.
- L-WHILE: Similar to L-SEQ.
- L-DECL: We assume $impl; \bar{t}; \Gamma, x : \tau \vdash \tau' y; e : \bar{\sigma}_1$. By inversion, $impl; \bar{t}; \Gamma, x : \tau, y : \tau' \vdash e : \bar{\sigma}_1$. By the Context Reordering Lemma, $impl; \bar{t}; \Gamma, y : \tau', x : \tau \vdash e : \bar{\sigma}_1$. By Weakening, $impl; \bar{t}; \Gamma, y : \tau' \vdash e_2 : \bar{\sigma}_2$. By induction, $impl; \bar{t}; \Gamma, y : \tau' \vdash e\{e_2/x\} : \bar{\sigma}_1$. By D-DECL, $impl; \bar{t}; \Gamma \vdash \tau' y; e\{e_2/x\} : \bar{\sigma}_1$.
- L-SUB: Inversion, induction, and an application of L-SUB.

Lemma 25 (Subject Reduction)

Suppose $impl$ is a sensible implementation. If

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

then there exists Ψ' , extending Ψ , such that

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

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 $impl.xlit(s) = \bar{b}$. From L-SHORT, $\bar{\sigma} = \text{byte}^i$. Because $impl$ is sensible (clause 2), we know $\text{size}(impl, \bar{b}) = i$. We can apply LW-BYTE and L-VALUE i times to derive $impl; \bar{t}; \Psi; \cdot \vdash e' : \bar{\sigma}$.
- 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 $impl; \bar{t}; \Psi; \cdot \vdash \text{while } e_1 e_2 : \text{byte}^i$ and $e' = \text{if } e_1 (e_2; \text{while } e_1 e_2) s$. Inversion gives
 - 1 $impl; \bar{t}; \Psi; \cdot \vdash e_1 : \text{byte}^j$
 - 2 $impl; \bar{t}; \Psi; \cdot \vdash e_2 : \bar{\sigma}_2$

3 $\text{impl.xtype}(\bar{t}, \text{short}) = \text{byte}^i$

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

4 $\text{impl}; \bar{t}; \Psi; \cdot \vdash_{\Gamma} (e_2; \text{while } e_1 \ e_2) : \text{byte}^i$

From L-SHORT, we have

5 $\text{impl}; \bar{t}; \Psi; \cdot \vdash_{\Gamma} s : \text{byte}^i$

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

- L-NEW: The step was derived using D-NEW. We assume $\text{impl}; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \text{new } \tau : \text{ptr}_{\alpha}(\bar{\sigma})$ 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 $\text{impl.xtype}(\bar{t}, \tau) = \bar{\sigma}$

3 $\text{size}(\text{impl}, \bar{\sigma}) = i$

4 $\text{impl.align}(\bar{t}, \tau) = \alpha$

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 $\text{impl}; \bar{t}; \Psi'; \cdot \vdash_{\Gamma} e' : \text{ptr}_{\alpha}(\bar{\sigma})$, 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 $\text{impl}; \bar{t}; \Psi' \vdash H : \Psi$

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

6 $\text{impl}; \bar{t}; \Psi'; \cdot \vdash_{\Gamma} \text{uninit}^i : \bar{\sigma}$ (6)

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

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

1 $\text{impl.xtype}(\bar{t}, \tau) = \bar{\sigma}$

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

By inversion on the assumed typing, we have

3 $\text{impl}; \bar{t}; \Psi; \cdot, x : \tau \vdash_{\Gamma} e : \bar{\sigma}_1$

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

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

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

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

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

$$\frac{\frac{\Psi'(\ell) = \bar{\sigma}, \alpha}{impl; \bar{t}; \Psi'; \cdot \Vdash \ell+0 : ptr_{\alpha}(\bar{\sigma})} \quad impl; \bar{t}; \Psi'; \cdot \Vdash \cdot : \cdot}{impl; \bar{t}; \Psi'; \cdot \Vdash \ell+0 : ptr_{\alpha}(\bar{\sigma})} \quad \frac{}{impl.xtype(\bar{t}, \tau) = \bar{\sigma}} \quad impl.access(\alpha, size(impl, \bar{\sigma}))}{impl; \bar{t}; \Psi'; \cdot \Vdash *(\tau*)(\ell+0) : \bar{\sigma}}$$

Notice that under the same assumptions, we can also derive $impl; \bar{t}; \Psi'; \cdot \Vdash *(\tau*)(\ell+0) : \bar{\sigma}, \alpha$. (L-DEREF and L-DEREF_L have identical hypotheses.) With this, we can apply Substitution Preserves Types Lemma to the above conclusion and (1,2,3) to get $impl; \bar{t}; \Psi; \cdot \Vdash e' : \bar{\sigma}_1$. 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 $impl; \bar{t}; \Psi; \cdot \Vdash *(\tau_1*)(\ell+j).f : \bar{\sigma}_2, [a, o + j']$ and have $e' = *(\tau_1*)(\ell+(j + j'))$ where

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

By inversion we obtain

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

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

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

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

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

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

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

$$8 \quad impl; \bar{t}; \Psi; \cdot \Vdash \ell+j : ptr_{[a, o]}(\bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4)$$

Canonical Forms gives

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

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

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

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

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

$$13 \quad impl; \bar{t}; \Psi; \cdot \Vdash \ell+(j + j') : ptr_{[a', o'+j+j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5)$$

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 $impl; \bar{t} \vdash ptr_{[a', o'+j+j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5) \leq ptr_{[a, o''+j+j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5)$. Plugging this and (13) into L-SUB, we get

$$14 \quad impl; \bar{t}; \Psi; \cdot \Vdash \ell+(j + j') : ptr_{[a, o''+j+j']}(\bar{\sigma}_2 \bar{\sigma}_3 \bar{\sigma}_4 \bar{\sigma}_5).$$

It remains to be shown that $impl.access([a, o'' + j + j'], size(impl, \bar{\sigma}_2))$. This, plugged into L-DEREF_L together with (5,14) would yield $impl; \bar{t}; \Psi; \cdot \Vdash *(\tau_2*)(\ell+(j + j')) : \bar{\sigma}_2, [a, o'' + j + j']$, as desired.

To prove this, suppose $impl.align(\bar{t}, \tau) = \alpha$. By clause (1) of sane implementations we have $impl.access(\alpha, \bar{\sigma}_2)$. By clause (3) of sane implementations, 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) of sane implementations, we then have $impl.access([a_N, o_N + j'], size(impl, \bar{\sigma}_2))$. 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 sane implementations, we have $impl.access([a, o'' + j + j'], size(impl, \bar{\sigma}_2))$, which concludes the proof case.

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

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

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

By inversion on the assumed typing we have

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

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

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 $impl.xtext(\bar{t}, \tau) = \bar{\sigma}_2$ is given by both.

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

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

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

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

$$1 \quad impl; \bar{t}; \Psi; \cdot \vdash \ell + j : ptr_{[a, o_1]}(\bar{\sigma}_1 \bar{\sigma}_2)$$

$$2 \quad impl.offset(\bar{t}, f) = size(impl, \bar{\sigma}_1)$$

$$3 \quad o_2 = o_1 + impl.offset(\bar{t}, f) = o_1 + size(impl, \bar{\sigma}_1)$$

Canonical Forms gives

$$4 \quad \Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a', o'_0]$$

$$5 \quad o_1 = o_0 + j$$

$$6 \quad j = size(impl, \bar{\sigma}_0)$$

$$7 \quad \vdash [a', o'_0] \leq [a, o_0]$$

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

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

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

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

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

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

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

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

1 $impl; \bar{t}; \Psi; \cdot \vdash_{\bar{\tau}} \ell+j : \text{ptr}_{\alpha}(\bar{\sigma}_1 \bar{\sigma}_2)$

2 $impl.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}_1$

The D-DEREF side condition gives

3 $H(\ell) = \bar{w}_1 \bar{w}_2 \bar{w}_3, [a', o']$

4 $\text{size}(impl, \bar{w}_1) = j$

5 $impl.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}_1$

6 $\text{size}(impl, \bar{w}_2) = \text{size}(impl, \bar{\sigma}_1) = k$

7 $impl.\text{access}([a', o' + j], k)$

From (1), Canonical Forms gives

8 $\Psi(\ell) = \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a', o']$

9 $j = \text{size}(impl, \bar{\sigma}_0)$

10 $\alpha = [a, o + i]$

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

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

- L-ASSN: The step must have been derived using D-ASSN. We assume $impl; \bar{t}; \Psi; \cdot \vdash_{\bar{\tau}} *(\tau*)(\ell+j) = \bar{w} : \bar{\sigma}_1$ 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 $impl; \bar{t}; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w} : \bar{\sigma}_1$, 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 $impl; \bar{t}; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w}_2 : \bar{\sigma}_1$ 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 $impl; \bar{t}; \Psi; \cdot \vdash_{\bar{\tau}} \bar{w}_1 \bar{w}_2 : \bar{\sigma}_0 \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$. Plug this and the assumption $impl; \bar{t}; \Psi \vdash H : \Psi$ into H-INF to obtain $impl; \bar{t}; \Psi \vdash H' : \Psi'$, where $\Psi' = \Psi, \ell \mapsto \Psi(\ell)$. Using Weakening, we obtain $impl; \bar{t}; \Psi' \vdash H' : \Psi'$, as desired.

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

Lemma 26 (Replacement)

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

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 $R = [\cdot]_R$ or $L = [\cdot]_L$ 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 $D; \Psi; \cdot \vdash_{\bar{\tau}} L[e]_r = e_2 : \bar{\sigma}$ 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 $D; \Psi; \cdot \vdash_{\bar{\tau}} L[e]_r : \bar{\sigma}, \alpha$

2 $D; \Psi; \cdot \vdash_{\bar{\tau}} e_2 : \bar{\sigma}$

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

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

All other cases follow an identical pattern.

Lemma 27 (Preservation) *If*

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

then there exist Ψ' , extending Ψ , such that

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

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 $\text{impl}; \bar{t} \vdash R[e_0]_r \rightarrow R[e'_0]_r$ (the D-STEPR rule) which can only happen under condition $\text{impl}; \bar{t} \vdash e_0 \xrightarrow{\tau} e'_0$. Then:

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

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

Lemma 28 (Progress)

Suppose $\text{impl}; \bar{t}; \Psi \vdash H : \Psi$ and impl is a sensible implementation.

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

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 $impl; \bar{t} \vdash H; e_1 \xrightarrow{\tau} H'; e_2$, 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 $impl; \bar{t}; \Psi; \cdot \vdash s : \text{byte}^i$. 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 $impl; \bar{t}; \Psi; \cdot \vdash e.f : \bar{\sigma}_2$ where inversion gives

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

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

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

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

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

From (2,6) we have

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

(3,4,5,7) allow a step $impl; \bar{t} \vdash H; e \xrightarrow{\tau} H'; \bar{w}_2$ 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 $impl; \bar{t}; \Psi; \cdot \vdash e.f : \bar{\sigma}_2, [a, o + o']$ where inversion gives
- 1 $impl; \bar{t}; \Psi; \cdot \vdash e : \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3, [a, o]$
 - 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 $e_2 = *(\tau_1*)(\ell+(i + impl.offset(\bar{t}, f)))$.

- L-DEREF: We assume $impl; \bar{t}; \Psi; \cdot \vdash *(\tau_1*)(e) : \bar{\sigma}$. Inverting, we get

- 1 $impl; \bar{t}; \Psi; \cdot \vdash e : ptr_\alpha(\bar{\sigma}\bar{\sigma}_2)$
- 2 $impl.xtext(\bar{t}, \tau) = \bar{\sigma}$
- 3 $impl.access(\alpha, size(impl, \bar{\sigma}))$

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 $size(impl, \bar{\sigma}_0) = i$
- 7 $\vdash [a', o'] \leq [a, o]$

Heap Canonical Forms on (4) gives

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

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

- 10 $impl; \bar{t}; \Psi; \cdot \vdash \bar{w}_0 : \bar{\sigma}_0$
- 11 $impl; \bar{t}; \Psi; \cdot \vdash \bar{w}_1 : \bar{\sigma}$

Value-Type Size on (10,11) gives

- 12 $size(impl, \bar{w}_0) = size(impl, \bar{\sigma}_0) = i$ (from (6))
- 13 $size(impl, \bar{w}_1) = size(impl, \bar{\sigma})$

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

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

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 $R = [\cdot]_R$, $e_1 = *(\tau*)(\ell+i)$, and $e_2 = \bar{w}_1$.

- L-DEREF: We assume $impl; \bar{t}; \Psi; \cdot \vdash *(\tau*)(e) : \bar{\sigma}_1, \alpha$ and by inversion we get $impl; \bar{t}; \Psi; \cdot \vdash e : ptr_\alpha(\bar{\sigma}_1\bar{\sigma}_2)$. 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 $impl.xtext$, $impl.ptr_size$, and $impl.align$ give sufficient conditions to step.
- L-ASSN: We assume $impl; \bar{t}; \Psi; \cdot \vdash e_1 = e_2 : \bar{\sigma}$ where by inversion we have

- 1 $impl; \bar{t}; \Psi; \cdot \vdash e_1 : \bar{\sigma}, \alpha$
- 2 $impl; \bar{t}; \Psi; \cdot \vdash e_2 : \bar{\sigma}$

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 $impl; \bar{t}; \Psi; \cdot \vdash \ell+i : ptr_{[a, o+i]}(\bar{\sigma}\bar{\sigma}_2)$

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

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

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 $size(impl, \bar{\sigma}_0) = i$

Heap Canonical Forms on (6) gives

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

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

The Value Split Lemma on (10) gives

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

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

13 $impl; \bar{t}; \Psi; \cdot \vdash \bar{w}_0 : \bar{\sigma}_0$

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

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

15 $size(impl, \bar{w}_0) = size(impl, \bar{\sigma}_0) = i$ (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 implementation sensibility gives

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

(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 $*(\tau*)(\ell+i) = \bar{w}''$ 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 $impl; \bar{t}; \Psi; \cdot \vdash (\tau*)&e \rightarrow f : ptr_{[a, o_1+impl.offset(\bar{t}, f)]}(\bar{\sigma}_2)$ where by inversion we have $impl; \bar{t}; \Psi; \cdot \vdash e : ptr_{a, o_1}(\bar{\sigma}_1\bar{\sigma}_2)$. 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 $(\tau*)&\text{uninit}^i \rightarrow f$ is legally stuck and we are done. Otherwise, $e = \ell+i$ and the expression can step via D-FADDR.
- L-SEQ: We assume $impl; \bar{t}; \Psi; \cdot \vdash e_1; e_2 : \bar{\sigma}$ where inversion gives $impl; \bar{t}; \Psi; \cdot \vdash e_1 : \bar{\sigma}'$. 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 $impl.xtype$, $impl.ptr_size$, and $impl.align$ give sufficient conditions to step.
- L-IF: We assume $impl; \bar{t}; \Psi; \cdot \vdash \text{if } e_1 e_2 e_3 : \bar{\sigma}$, where inversion gives $impl; \bar{t}; \Psi; \cdot \vdash e_1 : \text{byte}^i$. 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 29 (Type Soundness) *Let (\rightarrow^*) be the reflexive transitive closure of the transition relation (\rightarrow) . Then, if*

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

then $H'; e'$ is not illegally stuck on $impl$ 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 $impl; \bar{t} \vdash \cdot; e \rightarrow^{n-1} H''; e''$ and $H''; e''$ not illegally stuck. Preservation tells us that if $impl; \bar{t} \vdash H''; e'' \rightarrow H'; e'$, the type of e'' is preserved in the step to e' . It follows from Progress that $H'; e'$ is not illegally stuck.

A.3.2 Implementation Dependencies

Theorem 30 (Constraint Satisfaction)

- If $\bar{t}; \Gamma \vdash e : \tau; S$ and $impl$ is sensible and $impl \models S$ and $impl.xtype(\bar{t}, \tau) = \bar{\sigma}$, then $impl; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}$.
- If $\bar{t}; \Gamma \vdash e : \tau; S$ and $impl$ is sensible and $impl \models S$ and $impl.xtype(\bar{t}, \tau) = \bar{\sigma}$ and $impl.align(\bar{t}, \tau) = \alpha$, then $impl; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}, \alpha'$ 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 18) is used implicitly.

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

 - 3 $impl; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}$

$$4 \text{ impl.xtype}(\bar{t}, N) = \bar{\sigma}$$

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

$$5 \text{ impl.xtype}(\bar{t}, N) = \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$$

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

$$7 \text{ size}(\text{impl}, \bar{\sigma}_1) = \text{impl.offset}(\bar{t}, f)$$

It follows from (3,5) that

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

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

- H-FETCHL: We assume $\bar{t}; \Gamma \vdash e.f : \tau; S$ with $\text{impl.xtype}(\bar{t}, \tau) = \bar{\sigma}$ and $\text{impl.align}(\bar{t}, \tau) = [a, o]$. 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 $\text{impl} \models S$), we get

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

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

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

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

By clause 3 of implementation sensibility, we get

$$7 \bar{\sigma}' = \bar{\sigma}_1 \bar{\sigma}_2 \bar{\sigma}_3$$

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

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

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

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

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

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

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

Because *impl* is sensible (clause 4), we know

$$3 \bar{\sigma} = \text{ptr}_\alpha(\bar{\sigma}')$$

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

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

From (1,3), we get

6 $impl; \bar{t}; \cdot; \Gamma \vdash e : \text{ptr}_\alpha(\bar{\sigma}')$

Plugging (4,5,6) into L-DEREF, we get $impl; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}'$, 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 $S = S' \wedge \text{subtype}(\bar{t}, \text{xtype}(\bar{t}, \tau'), \text{xtype}(\bar{t}, \tau*))$, where by inversion we obtain

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

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

2 $impl \models S'$

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

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

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

Interpreting S under $impl$, we get

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

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

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

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

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

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

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

1 $impl.\text{align}(\bar{t}, N) = [a, o]$

2 $impl.\text{xtype}(\bar{t}, N*) = \text{ptr}_{[a,o]}(\bar{\sigma}_1 \bar{\sigma}_2)$

3 $impl.\text{xtype}(\bar{t}, \tau*) = \text{ptr}_\alpha(\bar{\sigma})$

4 $impl; \bar{t} \vdash \text{ptr}_{[a,o+impl.\text{offset}(\bar{t},f)]}(\bar{\sigma}_2) \leq \text{ptr}_\alpha(\bar{\sigma})$

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

By inversion,

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

By definition of the models relation,

7 $impl \models S'$

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

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

We can plug (5,8) into L-FADDR to derive $impl; \bar{t}; \cdot; \Gamma \vdash (\tau*) \& e \rightarrow f : \text{ptr}_{[a,o+impl.offset(\bar{t},f)]}(\bar{\sigma}_2)$. Plugging this and (4) into L-SUB, we obtain $impl; \bar{t}; \cdot; \Gamma \vdash (\tau*) \& e \rightarrow f : \text{ptr}_\alpha(\bar{\sigma})$ (recall that $\text{ptr}_\alpha(\bar{\sigma}) = \text{xtype}(\bar{t}, \tau*)$, 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, $impl; \bar{t}; \cdot; \Gamma, x : \tau_1 \vdash e : \bar{\sigma}$ where $impl.\text{xtype}(\bar{t}, \tau_2) = \bar{\sigma}$. Plugging these facts into L-DECL, we derive $impl; \bar{t}; \cdot; \Gamma \vdash \tau_1 x; e : \bar{\sigma}$, as desired.

Corollary 31 (Semi-Portability)

If $\bar{t}; \cdot \vdash e : \tau; S$ and $impl$ is sensible and $impl \models S$, then $\cdot; e$ is not illegally stuck on implementation $impl$ and declarations \bar{t} .

Definition 32 (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 33 (Cast-Free Sufficiency)

- If $\bar{t}; \Gamma \vdash e : \tau; S$, $impl.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}$, and e is cast-free, then for any sensible implementation $impl$, $impl; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}$.
- If $\bar{t}; \Gamma \vdash e : \tau; S$, $impl.\text{xtype}(\bar{t}, \tau) = \bar{\sigma}$, $impl.\text{align}(\bar{t}, \tau) = \alpha$, and e is cast-free, then for any sensible implementation $impl$, $impl; \bar{t}; \cdot; \Gamma \vdash e : \bar{\sigma}, \alpha'$ 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 implementation; 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 $\bar{t}; \Gamma \vdash (\tau*) \& e \rightarrow f : \tau*; S$ where inversion gives

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

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

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

3 $impl; \bar{t}; \cdot; \Gamma \vdash e : \text{ptr}_{[a,o]}(\bar{\sigma})$

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

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

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

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

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

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

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

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

A.3.3 Array Extension

Figure 18 shows how the syntax and semantics can be augmented with C-like arrays. First, we augment implementations (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 implementation-prescribed manner. We now show how the definitions, lemmas, and theorems in Section A.3 can be extended for arrays.

Extended Definition 34 (Sensible Implementations)

We add two array-related sensibility clauses:

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

Extended Definition 35 (Legal Stuck State)

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

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

Extended Lemma 36 (Uninit Type)

(No additions.)

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

Extended Lemma 37 (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 18).

Extended Lemma 38 (Value-Type Size)

(No additions.)

Proof:

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

Extended Lemma 39 (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 40 (Subtyping Type Form)

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

- If $\text{impl}; \bar{t} \vdash \bar{\sigma}' \leq \text{ptr}_\alpha^\omega(N)$ and $\text{impl.xtype}(\bar{t}, N) = \bar{\sigma}$ then $\exists \alpha', i$ such that $\bar{\sigma}' = \text{ptr}_{\alpha'}^\omega(\bar{\sigma}^i)$ or $\bar{\sigma}' = \text{ptr}_{\alpha'}^\omega(N')$ where $\text{impl.xtype}(\bar{t}, N') = \bar{\sigma}^i$, 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.
- 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 $\text{impl}; \bar{t} \vdash \bar{\sigma}_1 \bar{\sigma}_3 \leq \bar{\sigma}_2 \bar{\sigma}_4$. 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 $\text{impl}; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_3$ where inversion gives $\text{impl}; \bar{t} \vdash \bar{\sigma}_1 \leq \bar{\sigma}_2$ and $\text{impl}; \bar{t} \vdash \bar{\sigma}_2 \leq \bar{\sigma}_3$. There are two possibilities for $\bar{\sigma}_3$:
 1. $\bar{\sigma}_3 = \text{ptr}_\alpha^\omega(N)$ where $\text{impl.xtype}(\bar{t}, N) = \bar{\sigma}$. 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 $\text{impl.xtype}(\bar{t}, N') = \bar{\sigma}^j$ and $\vdash \alpha'' \leq \alpha'$. By ALST-TRANS, $\vdash \alpha'' \leq \alpha$, as required.
 - $\bar{\sigma}_2 = \text{ptr}_{\alpha'}^\omega(N')$ where $\text{impl.xtype}(\bar{t}, N') = \bar{\sigma}$ 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 $\text{impl.xtype}(\bar{t}, N') = \bar{\sigma}^j$. 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 41 (Canonical Forms)

- If $\text{impl}; \bar{t}; \Psi; \cdot \vdash \bar{w} : \text{ptr}_\alpha^\omega(\bar{\sigma})$ then either
 - $\bar{w} = \text{uninit}^i$ where $\text{impl.ptr_size} = i$.
 - $\bar{w} = \ell+0$ and $\exists j, \bar{\sigma}', \alpha'$ such that $\Psi(\ell) = \bar{\sigma}^j, \alpha'$ where $\vdash \alpha' \leq \alpha$.
- If $\text{impl}; \bar{t}; \Psi; \cdot \vdash \bar{w} : \text{ptr}_\alpha^\omega(N)$ and $\text{impl.xtype}(\bar{t}, N) = \bar{\sigma}$, then either
 - $\bar{w} = \text{uninit}^i$ where $\text{impl.ptr_size} = i$.
 - $\bar{w} = \ell+0$ 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 18).

- L-VALUE: There are two cases:
 - Assume $\text{impl}; \bar{t}; \Psi; \cdot \vdash \bar{w} : \text{ptr}_\alpha^\omega(\bar{\sigma}_1)$. By inversion we get $\text{impl}; \bar{t}; \Psi \Vdash \bar{w} : \text{ptr}_\alpha^\omega(\bar{\sigma}_1)$. 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 $i = \text{impl.ptr_size}$, and in the latter case, $\bar{w} = \ell+0$ where by inversion $\Psi(\ell) = \bar{\sigma}^j, \alpha'$ with $\alpha' = \alpha$.
 - Assume $\text{impl}; \bar{t}; \Psi; \cdot \vdash \bar{w} : \text{ptr}_\alpha^\omega(N)$, where $\text{impl.xtype}(\bar{t}, N) = \bar{\sigma}_1$. Vacuous; cannot be derived via L-VAL.
- L-SUB: There are two cases:

- Assume $impl; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \bar{w} : ptr_{\alpha}^{\omega}(\bar{\sigma})$. By inversion we get $impl; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \bar{w} : \bar{\sigma}''$ and $impl; \bar{t} \vdash \bar{\sigma}'' \leq ptr_{\alpha}^{\omega}(\bar{\sigma})$. By the Subtyping Type Form Lemma, $\bar{\sigma}''$ is one of
 - * $ptr_{\alpha'}^{\omega}(\bar{\sigma}^i)$ where $\vdash \alpha' \leq \alpha$. Applying the IH to $impl; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \bar{w} : ptr_{\alpha'}^{\omega}(\bar{\sigma}^i)$, 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.
 - * $ptr_{\alpha'}^{\omega}(N)$ where $impl.\text{xtype}(\bar{t}, N) = \bar{\sigma}$ and $\vdash \alpha' \leq \alpha$. By induction and an application of ALST-TRANS, like the above case.
- Assume $impl; \bar{t}; \Psi; \cdot \vdash_{\Gamma} \bar{w} : ptr_{\alpha}^{\omega}(N)$, where $impl.\text{xtype}(\bar{t}, N) = \bar{\sigma}$. Similar to the above case.

Extended Lemma 42 (Subject Reduction)

(No additions.)

Proof:

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

The assumed step rule is

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

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 $impl; \bar{t}; \Psi'; \cdot \vdash_{\Gamma} \ell+0 : ptr_{\alpha}^{\omega}(\bar{\sigma})$, 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

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

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

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

The assumed step rule is

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

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

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

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

$$3 \quad H(\ell) = \bar{w}, \alpha'$$

$$4 \quad \text{impl}; \bar{t}; \Psi; \cdot \vdash \bar{w} : \bar{\sigma}^h$$

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

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

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

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

It must be the case then that we can write $\bar{\sigma}^h$ as $\bar{\sigma}^k \bar{\sigma}^g$ (since $\text{size}(\text{impl}, \bar{\sigma}) = j$, 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 $\text{size}(\text{impl}, \bar{\sigma}^k) = k \times \text{size}(\text{impl}, \bar{\sigma}) = k \times j$. Plugging these two facts into LW-LBL and the result into L-VALUE, we get

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

Using the Alignment Addition Lemma, we can conclude

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

From the third hypothesis of the typing assumption, we know $j = \text{size}(\text{impl}, \bar{\sigma}) = a \times k'$ 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 $\text{impl}; \bar{t} \vdash \text{ptr}_{[a', o' + j \times k]}(\bar{\sigma}^g) \leq \text{ptr}_{[a, o]}(\bar{\sigma})$. Plugging this along with (7) into L-SUB yields $\text{impl}; \bar{t}; \Psi; \cdot \vdash \ell + (j \times k) : \text{ptr}_{[a, o]}(\bar{\sigma})$, as desired.

Extended Lemma 43 (Progress)

(No additions.)

Proof:

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

We apply the IH to $\text{impl}; \bar{t}; C \vdash e : \bar{\sigma}'$ and consider the case when e is a value. Clause 2 of implementation 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 $\text{impl.val}(\bar{b}) < 0$, then $\text{new } \tau[e]$ is again legally stuck. Otherwise it can take a step via rule D-NEWARR.

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

We apply the IH to $\text{impl}; \bar{t}; C \vdash e_1 : \text{ptr}_{[a, o]}^{\omega}(\bar{\sigma})$ and $\text{impl}; \bar{t}; C \vdash e_2 : \bar{\sigma}'$ 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 implementation 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 $\text{impl.val}(\bar{b}) < 0$, $\&((\tau^{*\omega})(e_1))[e_2]$ is again legally stuck. Suppose then that $\text{impl.val}(\bar{b}) = k$ for $k \geq 0$.

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

Extended Lemma 44 (Constraint Satisfaction)

(No additions.)

Proof:

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

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

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

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

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

From clause 6 of implementation sensibility we learn

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

From clause 4 of implementation sensibility we have

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

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

B Differences Between Report and Paper

This technical report broadly extends the paper with the same title in the following ways:

- The full safety proof is given in Appendix A.
- Section 1.3, omitted from the paper, contains additional discussion regarding the contributions of the work.
- Section 2 contains two additional motivating examples.
- The language given in the report contains one additional syntactic form for “while” loops.
- Section 3.6 contains additional discussion regarding the difference between left- and right-typings in the high-level semantics and the cast-like syntax on the dereference operator.
- Figures 6 (physical subtyping) and 7 (high-level semantics) are complete. The paper omits a subtyping rule for padding from the former and most rules from the latter.

- Section 3.7 contains additional discussion regarding the low-level type system.
- Section 4 contains a detailed discussion of the array extension, and other extensions which are omitted from the paper: read-only pointers, byte skipping, and recursive subtyping.

New Syntax

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

New Dynamic Rules

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

New High-Level Typing

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

New Subtyping

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

New Low-Level Typing

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

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

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

New Size Function Case

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

Figure 18: Array Syntax and Semantics