Polymorphic Type Inference for Machine Code^{*}

Matthew Noonan Alexey Loginov David Cok

GrammaTech, Inc.

Ithaca NY, USA

{mnoonan,alexey,dcok}@grammatech.com

Abstract

For many compiled languages, source-level types are erased very early in the compilation process. As a result, further compiler passes may convert type-safe source into type-unsafe machine code. Type-unsafe idioms in the original source and type-unsafe optimizations mean that type information in a stripped binary is essentially nonexistent. The problem of recovering high-level types by performing type inference over stripped machine code is called *type reconstruction*, and offers a useful capability in support of reverse engineering and decompilation.

In this paper, we motivate and develop a novel type system and algorithm for machine-code type inference. The features of this type system were developed by surveying a wide collection of common source- and machine-code idioms, building a catalog of challenging cases for type reconstruction. We found that these idioms place a sophisticated set of requirements on the type system, inducing features such as recursively-constrained polymorphic types. Many of the features we identify are often seen only in expressive and powerful type systems used by high-level functional languages.

Using these type-system features as a guideline, we have developed Retypd: a novel static type-inference algorithm for machine code that supports recursive types, polymorphism, and subtyping. Retypd yields more accurate inferred types than existing algorithms, while also enabling new capabilities

such as reconstruction of pointer const annotations with 98% recall. Retypd can operate on weaker program representations than the current state of the art, removing the need for high-quality points-to information that may be impractical to compute.

Categories and Subject Descriptors F.3.2 [Logics and Meanings of Programs]: Semantics of Programming Languages; D.2.7 [Software Engineering]: Distribution, Maintenance, and Enhancement; D.3.3 [Programming Languages]: Language Constructs and Features; F.4.3 [Mathematical Logic and Formal Languages]: Formal Languages

Keywords Reverse Engineering, Type Systems, Polymorphism, Static Analysis, Binary Analysis, Pushdown Automata

1. Introduction

In this paper we introduce Retypd, a machine-code type-inference tool that finds **regular types** using **pushdown** systems. Retypd includes several novel features targeted at improved types for reverse engineering, decompilation, and high-level program analyses. These features include:

- Inference of most-general type schemes (§5)
- Inference of recursive structure types (Figure 2)
- Sound analysis of pointer subtyping (§3.3)
- Tracking of customizable, high-level information such as purposes and typedef names (§3.5)
- Inference of type qualifiers such as const (§6.4)
- No dependence on high-quality points-to data (§6)
- More accurate recovery of source-level types (§6)

Retypd continues in the tradition of SecondWrite [10] and TIE [17] by introducing a principled static type-inference algorithm applicable to stripped binaries. Diverging from previous work on machine-code type reconstruction, we use a rich type system that supports polymorphism, mutable references, and recursive types. The principled type-inference phase is followed by a second phase that uses heuristics to “downgrade” the inferred types to human-readable C types before display. By factoring type inference into two phases, we can sequester unsound heuristics and quirks of the C type systems from the sound core of the type-inference engine. This adds a degree of freedom to the design space so that we may leverage a relatively complex type system during type

^{*}This research was developed with funding from the Defense Advanced Research Projects Agency (DARPA). The views, opinions, and/or findings contained in this material are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

DISTRIBUTION A. Approved for public release; distribution unlimited.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

PLDI’16, June 13–17, 2016, Santa Barbara, CA, USA
© 2016 ACM. 978-1-4503-4261-2/16/06...\$15.00
<http://dx.doi.org/10.1145/2908080.2908119>

analysis, yet still emit familiar C types for the benefit of the reverse engineer.

Retypd operates on an intermediate representation (IR) recovered by automatically disassembling a binary using GrammaTech’s static analysis tool CodeSurfer® for Binaries [5]. By generating type constraints from a TSL-based abstract interpreter [18], Retypd can operate uniformly on binaries for any platform supported by CodeSurfer, including x86, x86-64, and ARM.

During the development of Retypd, we carried out an extensive investigation of common machine-code idioms in compiled C and C++ code that create challenges for existing type-inference methods. For each challenging case, we identified requirements for any type system that could correctly type the idiomatic code. The results of this investigation appear in §2. The type system used by Retypd was specifically designed to satisfy these requirements. These common idioms pushed us into a far richer type system than we had first expected, including features like recursively constrained type schemes that have not previously been applied to machine-code type inference.

Due to space limitations, details of the proofs and algorithms appear in the appendices, which are available in the online version of this paper at arXiv:1603.05495 [22]. Scripts and data sets used for evaluation also appear there.

2. Challenges

There are many challenges to carrying out type inference on machine code, and many common idioms that lead to sophisticated demands on the feature set of a type system. In this section, we describe several of the challenges seen during the development of Retypd that led to our particular combination of type-system features.

2.1 Optimizations after Type Erasure

Since type erasure typically happens early in the compilation process, many compiler optimizations may take well-typed machine code and produce functionally equivalent but ill-typed results. We found that there were three common optimization techniques that required special care: the use of a variable as a syntactic constant, early returns along error paths, and the re-use of stack slots.

Semi-syntactic constants: Suppose a function with signature `void f(int x, char* y)` is invoked as `f(0, NULL)`. This will usually be compiled to x86 machine code similar to

```
xor  eax, eax
push  eax      ; y := NULL
push  eax      ; x := 0
call  f
```

This represents a code-size optimization, since `push eax` can be encoded in one byte instead of the five bytes needed to push an immediate value (0). We must be careful that the type variables for x and y are not unified; here, `eax` is being

<pre>T * get_T(void) { S * s = get_S(); if (s == NULL) { return NULL; } T * t = S2T(s); return t; }</pre>	<pre>get_T: call get_S test eax, eax jz local_exit push eax call S2T add esp, 4 local_exit: ret</pre>
---	--

Figure 1. A common fortuitous re-use of a known value.

used more like a syntactic constant than a dynamic value that should be typed.

Fortuitous re-use of values: A related situation appears in the common control-flow pattern represented by the snippet of C and the corresponding machine code in Figure 1. Note that on procedure exit, the return value in `eax` may have come from either the return value of `S2T` or from the return value of `get_S` (if `NULL`). If this situation is not detected, we will see a false relationship between the incompatible return types of `get_T` and `get_S`.

Re-use of stack slots: If a function uses two variables of the same size in disjoint scopes, there is no need to allocate two separate stack slots for those variables. Often the optimizer will reuse a stack slot from a variable that has dropped out of scope. This is true even if the new variable has a different type. This optimization even applies to the stack slots used to store formal-in parameters, as in Figure 2; once the function’s argument is no longer needed, the optimizer can overwrite it with a local variable of an incompatible type.

More generally, we cannot assume that the map from program variables to physical locations is one-to-one. We cannot even make the weaker assumption that the program variables inhabiting a single physical location at different times will all belong to a single type.

We handle these issues through a combination of type-system features (subtyping instead of unification) and program analyses (reaching definitions for stack variables and trace partitioning [21]).

2.2 Polymorphic Functions

We discovered that, although not directly supported by the C type system, most programs define or make use of functions that are effectively polymorphic. The most well-known example is `malloc`: the return value is expected to be immediately cast to some other type T^* . Each call to `malloc` may be thought of as returning some pointer of a different type. The type of `malloc` is effectively *not* $\text{size_t} \rightarrow \text{void}^*$, but rather $\forall \tau. \text{size_t} \rightarrow \tau^*$.

The problem of a polymorphic `malloc` could be mitigated by treating each call site p as a call to a distinct function `mallocp`, each of which may have a distinct return type T_p^* . Unfortunately, it is not sufficient to treat a handful of special functions like `malloc` this way: it is common to see binaries

<pre>#include <stdlib.h> struct LL { struct LL * next; int handle; }; int close_last(struct LL * list) { while (list->next != NULL) { list = list->next; } return close(list->handle); }</pre>	<pre>close_last: push ebp mov ebp, esp sub esp, 8 mov edx, dword [ebp+arg_0] jmp loc_8048402 loc_8048400: mov edx, eax loc_8048402: mov eax, dword [edx] test eax, eax jnz loc_8048400 mov eax, dword [edx+4] mov dword [ebp+arg_0], eax leave jmp __thunk_.close</pre>	$\forall F. (\exists \tau. C) \Rightarrow F \text{ where } C =$ $F.in_{stack0} \sqsubseteq \tau$ $\tau.load.\sigma32@0 \sqsubseteq \tau$ $\tau.load.\sigma32@4 \sqsubseteq \text{int} \wedge \# \text{FileDescriptor}$ $\text{int} \vee \# \text{SuccessZ} \sqsubseteq F.out_{eax}$ <pre>typedef struct { Struct_0 * field_0; int // #FileDescriptor field_4; } Struct_0; int // #SuccessZ close_last(const Struct_0 *);</pre>
---	---	--

Figure 2. Example C code (compiled with gcc 4.5.4 on Linux with flags `-m32 -O2`), disassembly, type scheme inferred from the machine code, and reconstructed C type. The tags `#FileDescriptor` and `#SuccessZ` encode inferred higher-level purposes.

that use user-defined allocators and wrappers to `malloc`. All of these functions would also need to be accurately identified and duplicated for each callsite.

A similar problem exists for functions like `free`, which is polymorphic in its lone parameter. Even more complex are functions like `memcpy`, which is polymorphic in its first two parameters and its return type, though the three types are not independent of each other. Furthermore, the polymorphic type signatures

```
malloc :  $\forall \tau. \text{size\_t} \rightarrow \tau^*$ 
free :  $\forall \tau. \tau^* \rightarrow \text{void}$ 
memcpy :  $\forall \alpha, \beta. (\beta \sqsubseteq \alpha) \Rightarrow (\alpha^* \times \beta^* \times \text{size\_t}) \rightarrow \alpha^*$ 
```

are all strictly more informative to the reverse engineer than the standard C signatures. How else could one know that the `void*` returned by `malloc` is not meant to be an opaque handle, but rather should be cast to some other pointer type?

In compiled C++ binaries, polymorphic functions are even more common. For example, a class member function must potentially accept both `base_t*` and `derived_t*` as types for this.

Foster et al. [11] noted that using bounded polymorphic type schemes for libc functions increased the precision of type-qualifier inference, at the level of source code. To advance the state of the art in machine-code type recovery, we believe it is important to also embrace polymorphic functions as a natural and common feature of machine code. Significant improvements to static type reconstruction—even for monomorphic types—will require the capability to infer polymorphic types of some nontrivial complexity.

2.3 Recursive Types

The relevance of recursive types for decompilation was recently discussed by Schwartz et al. [30], where lack of a recursive type system for machine code was cited as

an important source of imprecision. Since recursive data structures are relatively common, it is desirable that a type-inference scheme for machine code be able to represent and infer recursive types natively.

2.4 Offset and Reinterpreted Pointers

Unlike in source code, there is no syntactic distinction in machine code between a pointer-to-struct and a pointer-to-first-member-of-struct. For example, if X has type `struct { char*, FILE*, size_t }*` on a 32-bit platform, then it should be possible to infer that $X + 4$ can be safely passed to `fclose`; conversely, if $X + 4$ is passed to `fclose` we may need to infer that X points to a structure that, at offset 4, contains a `FILE*`. This affects the typing of local structures, as well: a structure on the stack may be manipulated using a pointer to its starting address or by manipulating the members directly, e.g., through the frame pointer.

These idioms, along with casts from `derived*` to `base*`, fall under the general class of *physical* [31] or *non-structural* [24] subtyping. In Retypd, we model these forms of subtyping using type scheme specialization (§ 3.5). Additional hints about the extent of local variables are found using data-delineation analysis [13].

2.5 Disassembly Failures

The problem of producing correct disassembly for stripped binaries is equivalent to the halting problem. As a result, we can never assume that our reconstructed program representation will be perfectly correct. Even sound analyses built on top of an unsound program representation may exhibit inconsistencies and quirks.

Thus, we must be careful that incorrect disassembly or analysis results from one part of the binary will not influence the correct type results we may have gathered for the rest of the binary. Type systems that model value assignments as type unifications are vulnerable to over-unification issues caused

by bad IR. Since unification is non-local, bad constraints in one part of the binary can degrade *all* type results.

Another instance of this problem arises from the use of register parameters. Although the x86 cdecl calling convention uses the stack for parameter passing, most optimized binaries will include many functions that pass parameters in registers for speed. Often, these functions do not conform to any standard calling convention. Although we work hard to ensure that only true register parameters are reported, conservativeness demands the occasional false positive.

Type-reconstruction methods that are based on unification are generally sensitive to precision loss due to false-positive register parameters. A common case is the “push ecx” idiom that reserves space for a single local variable in the stack frame of a function f . If ecx is incorrectly viewed as a register parameter of f in a unification-based scheme, whatever type variables are bound to ecx at each callsite to f will be mistakenly unified. In our early experiments, we found these overunifications to be a persistent and hard-to-diagnose source of imprecision.

In our early unification-based experiments, mitigation heuristics against overunification quickly ballooned into a disproportionately large and unprincipled component of type analysis. We designed Retypd’s subtype-based constraint system to avoid the need for such ad-hoc prophylactics against overunification.

2.6 Cross-casting and Bit Twiddling

Even at the level of source code, there are already many type-unsafe idioms in common use. Most of these idioms operate by directly manipulating the bit representation of a value, either to encode additional information or to perform computations that are not possible using the type’s usual interface. Some common examples include

- hashing values by treating them as untyped bit blocks [1],
- stealing unused bits of a pointer for tag information, such as whether a thunk has been evaluated [20],
- reducing the storage requirements of a doubly-linked list by XOR-combining the next and prev pointers, and
- directly manipulating the bit representation of another type, as in the quake3 inverse square root trick [29].

Because of these type-unsafe idioms, it is important that a type-inference scheme continues to produce useful results even in the presence of apparently contradictory constraints. We handle this situation in three ways:

1. separating the phases of constraint entailment, solving, and consistency checking,
2. modeling types with sketches (§ 3.5) that carry more information than C types, and
3. using unions to combine types with otherwise incompatible capabilities (e.g., τ is both int-like and pointer-like).

Table 1. Example field labels (type capabilities) in Σ .

Label	Variance	Capability
.in_L	$\ominus$	Function with input in location L .
.out_L	$\oplus$	Function with output in location L .
.load	$\oplus$	Readable pointer.
.store	$\ominus$	Writable pointer.
$\text{.}\sigma N@k$	$\oplus$	Has N -bit field at offset k .

2.7 Incomplete Points-to Information

Degradation of points-to accuracy on large programs has been identified as a source of type-precision loss in other systems [10]. Our algorithm can provide high-quality types even in the absence of points-to information. Precision can be further improved by increasing points-to knowledge via machine-code analyses such as VSA [6], but good results are already attained with no points-to analysis beyond the simpler problem of tracking the stack pointer.

2.8 Ad-hoc Subtyping

Programs may define an ad-hoc type hierarchy via typedefs. This idiom appears in the Windows API, where a variety of handle types are all defined as typedefs of void*. Some of the handle types are to be used as subtypes of other handles; for example, a GDI handle (HGDIOBJ) is a generic handle used to represent any one of the more specific HBRUSH, HPEN, etc. In other cases, a typedef may indicate a *supertype*, as in LPARAM or DWORD; although these are typedefs of int, they have the intended semantics of a generic 32-bit type, which in different contexts may be used as a pointer, an integer, a flag set, and so on.

To accurately track ad-hoc hierarchies requires a type system based around subtyping rather than unification. Models for common API type hierarchies are useful; still better is the ability for the end user to define or adjust the initial type hierarchy at run time. We support this feature by parameterizing the main type representation by an uninterpreted lattice Λ , as described in §3.5.

3. The Type System

The type system used by Retypd is based around the inference of *recursively constrained type schemes* (§3.1). Solutions to constraint sets are modeled by *sketches* (§3.5); the sketch associated to a value consists of a record of capabilities which that value holds, such as whether it can be stored to, called, or accessed at a certain offset. Sketches also include markings drawn from a customizable lattice $(\Lambda, \vee, \wedge, <:)$, used to propagate high-level information such as typedef names and domain-specific purposes during type inference.

Retypd also supports recursively constrained type schemes that abstract over the set of types subject to a constraint set $\mathcal{C}$. The language of type constraints used by Retypd is weak enough that for any constraint set $\mathcal{C}$, satisfiability of $\mathcal{C}$ can be

Derived Type Variable Formation		Subtyping	
$\frac{\alpha \sqsubseteq \beta}{\text{VAR } \alpha} \text{ (T-LEFT)}$	$\frac{\alpha \sqsubseteq \beta, \text{VAR } \alpha.\ell}{\text{VAR } \beta.\ell} \text{ (T-INHERITL)}$	$\frac{\text{VAR } \alpha}{\alpha \sqsubseteq \alpha} \text{ (S-REFL)}$	$\frac{\alpha \sqsubseteq \beta, \text{VAR } \beta.\ell, \langle \ell \rangle = \oplus}{\alpha.\ell \sqsubseteq \beta.\ell} \text{ (S-FIELD}_{\oplus})$
$\frac{\alpha \sqsubseteq \beta}{\text{VAR } \beta} \text{ (T-RIGHT)}$	$\frac{\alpha \sqsubseteq \beta, \text{VAR } \beta.\ell}{\text{VAR } \alpha.\ell} \text{ (T-INHERITR)}$	$\frac{\alpha \sqsubseteq \beta, \beta \sqsubseteq \gamma}{\alpha \sqsubseteq \gamma} \text{ (S-TRANS)}$	$\frac{\alpha \sqsubseteq \beta, \text{VAR } \beta.\ell, \langle \ell \rangle = \ominus}{\beta.\ell \sqsubseteq \alpha.\ell} \text{ (S-FIELD}_{\ominus})$
$\frac{\text{VAR } \alpha.\ell}{\text{VAR } \alpha} \text{ (T-PREFIX)}$		$\frac{\text{VAR } \alpha.\text{load}, \text{VAR } \alpha.\text{store}}{\alpha.\text{store} \sqsubseteq \alpha.\text{load}} \text{ (S-POINTER)}$	

Figure 3. Deduction rules for the type system. α, β, γ represent derived type variables; ℓ represents a label in Σ .

reduced (in cubic time) to checking a set of scalar constraints $\kappa_1 <: \kappa_2$, where κ_i are constants belonging to Λ .

Thanks to the reduction of constraint satisfiability to scalar constraint checking, we can omit expensive satisfiability checks during type inference. Instead, we delay the check until the final stage when internal types are converted to C types for display, providing a natural place to instantiate union types that resolve any inconsistencies. Since compiler optimizations and type-unsafe idioms in the original source frequently lead to program fragments with unsatisfiable type constraints (§2.5, §2.6), this trait is particularly desirable.

3.1 Syntax: the Constraint Type System

Throughout this section, we fix a set $\mathcal{V}$ of *type variables*, an alphabet Σ of *field labels*, and a function $\langle \cdot \rangle : \Sigma \rightarrow \{\oplus, \ominus\}$ denoting the *variance* (Definition 3.2) of each label. We do not require the set Σ to be finite. Retypd makes use of a large set of labels; for simplicity, we will focus on those in Table 1.

Within $\mathcal{V}$, we assume there is a distinguished set of *type constants*. These type constants are symbolic representations $\bar{\kappa}$ of elements κ belonging to some lattice, but are otherwise uninterpreted. It is usually sufficient to think of the type constants as type names or semantic tags.

Definition 3.1. A *derived type variable* is an expression of the form αw with $\alpha \in \mathcal{V}$ and $w \in \Sigma^*$.

Definition 3.2. The variance of a label ℓ encodes the subtype relationship between $\alpha.\ell$ and $\beta.\ell$ when α is a subtype of β , formalized in rules S-FIELD $_{\oplus}$ and S-FIELD $_{\ominus}$ of Figure 3. The variance function $\langle \cdot \rangle$ can be extended to Σ^* by defining $\langle \varepsilon \rangle = \oplus$ and $\langle xw \rangle = \langle x \rangle \cdot \langle w \rangle$, where $\{\oplus, \ominus\}$ is the sign monoid with $\oplus \cdot \oplus = \ominus \cdot \ominus = \oplus$ and $\oplus \cdot \ominus = \ominus \cdot \oplus = \ominus$. A word $w \in \Sigma^*$ is called *covariant* if $\langle w \rangle = \oplus$, or *contravariant* if $\langle w \rangle = \ominus$.

Definition 3.3. Let $\mathcal{V} = \{\alpha_i\}$ be a set of base type variables. A *constraint* is an expression of the form $\text{VAR } X$ (“existence of the derived type variable X ”) or $X \sqsubseteq Y$ (“ X is a subtype of Y ”), where X and Y are derived type variables. A *constraint set over $\mathcal{V}$* is a finite collection $\mathcal{C}$ of constraints, where the type variables in each constraint are either type constants or members of $\mathcal{V}$. We will say that $\mathcal{C}$ *entails* c , denoted $\mathcal{C} \vdash c$, if c can be derived from the constraints in $\mathcal{C}$ using the de-

duction rules of Figure 3. We also allow projections: given a constraint set $\mathcal{C}$ with free variable τ , the projection $\exists \tau. \mathcal{C}$ binds τ as an “internal” variable in the constraint set. See τ in Figure 2 for an example or, for a more in-depth treatment of constraint projection, see Su et al. [34].

The field labels used to form derived type variables are meant to represent *capabilities* of a type. For example, the constraint $\text{VAR } \alpha.\text{load}$ means α is a readable pointer, and the derived type variable $\alpha.\text{load}$ represents the type of the memory region obtained by loading from α .

Let us briefly see how operations in the original program translate to type constraints, using C-like pseudocode for clarity. The full conversion from disassembly to type constraints is described in Appendix A of [22].

Value copies: When a value is moved between program variables in an assignment like $x := y$, we make the conservative assumption that the type of x may be upcast to a supertype of y . We will generate a constraint of the form $Y \sqsubseteq X$.

Loads and stores: Suppose that p is a pointer to a 32-bit type, and a value is loaded into x by the assignment $x := *p$. Then we will generate a constraint of the form $P.\text{load}.\sigma 32 @ 0 \sqsubseteq X$. Similarly, a store $*q := y$ results in the constraint $Y \sqsubseteq Q.\text{store}.\sigma 32 @ 0$.

In some of the pointer-based examples in this paper we omit the final $\sigma N @ k$ access after a `.load` or `.store` to simplify the presentation.

Function calls: Suppose the function f is invoked by $y := f(x)$. We will generate the constraints $X \sqsubseteq F.\text{in}$ and $F.\text{out} \sqsubseteq Y$, reflecting the flow of actuals to and from formals. Note that if we define $A.\text{in} = X$ and $A.\text{out} = Y$ then the two constraints are equivalent to $F \sqsubseteq A$ by the rules of Figure 3. This encodes the fact that the called function’s type must be at least as specific as the type used at the callsite.

One of the primary goals of our type-inference engine is to associate to each procedure a most-general type scheme.

Definition 3.4. A *type scheme* is an expression of the form $\forall \bar{\alpha}. C \Rightarrow \alpha_1$ where $\forall \bar{\alpha} = \forall \alpha_1 \dots \forall \alpha_n$ is quantification over a set of type variables, and C is a constraint set over $\{\alpha_i\}_{i=1..n}$.

Type schemes provide a way of encoding the pre- and post-conditions that a function places on the types in its calling

context. Without the constraint sets, we would only be able to represent conditions of the form “the input must be a subtype of X ” and “the output must be a supertype of Y ”. The constraint set $\mathcal{C}$ can be used to encode more interesting type relations between inputs and outputs, as in the case of `memcpy` (§2.2). For example, a function that returns the second 4-byte element from a `struct*` may have the type scheme $\forall \tau. (\tau.\text{in}.\text{load}.\sigma 32 @ 4 \sqsubseteq \tau.\text{out}) \Rightarrow \tau$.

3.2 Deduction Rules

The deduction rules for our type system appear in Figure 3. Most of the rules are self-evident under the interpretation in Definition 3.3, but a few require some additional motivation.

S-FIELD $_{\oplus}$ and **S-FIELD $_{\ominus}$** : These rules ensure that field labels act as co- or contra-variant type operators, generating subtype relations between derived type variables from subtype relations between the original variables.

T-INHERITL and **T-INHERITR**: The rule T-INHERITL should be uncontroversial, since a subtype should have all capabilities of its supertype. The rule T-INHERITR is more unusual since it moves capabilities in the other direction; taken together, these rules require that two types in a subtype relation must have exactly the same set of capabilities. This is a form of structural typing, ensuring that comparable types have the same shape.

Structural typing appears to be at odds with the need to cast more capable objects to less capable ones, as described in §2.4. Indeed, T-INHERITR eliminates the possibility of forgetting capabilities during value assignments. But we still maintain this capability at procedure invocations due to our use of polymorphic type schemes. An explanation of how type-scheme instantiation enables us to forget fields of an object appears in §3.4, with more details in §E.1.2 of [22].

These rules ensure that Retypd can perform “iterative variable recovery”; lack of iterative variable recovery was cited by the creators of the Phoenix decompiler [30] as a common cause of incorrect decompilation when using TIE [17] for type recovery.

S-POINTER: This rule is a consistency condition ensuring that the type that can be loaded from a pointer is a supertype of the type that can be stored to a pointer. Without this rule, pointers would provide a channel for subverting the type system. An example of how this rule is used in practice appears in §3.3.

The deduction rules of Figure 3 are simple enough that each proof may be reduced to a normal form (see Theorem B.1 in [22]). An encoding of the normal forms as transition sequences in a modified pushdown system is used to provide a compact representation of the entailment closure $\bar{\mathcal{C}} = \{c \mid \mathcal{C} \vdash c\}$. The pushdown system modeling $\bar{\mathcal{C}}$ is queried and manipulated to provide most of the interesting type-inference functionality. An outline of this functionality appears in §5.2.

<pre>f() { p = q; *p = x; y = *q; }</pre>	<pre>g() { p = q; *q = x; y = *p; }</pre>
---	---

Figure 4. Two programs, each mediating a copy from x to y through a pair of aliased pointers.

3.3 Modeling Pointers

To model pointers soundly in the presence of subtyping, we found that our initial naïve approach suffered from unexpected difficulties when combined with subtyping. Following the C type system, it seemed natural to model pointers by introducing an injective unary type constructor `Ptr`, so that `Ptr( $T$ )` is the type of pointers to T . In a unification-based type system, this approach works as expected.

In the presence of subtyping, a new issue arises. Consider the two programs in Figure 4. Since the type variables P and Q associated to `p`, `q` can be seen to be pointers, we can begin by writing $P = \text{Ptr}(\alpha)$, $Q = \text{Ptr}(\beta)$. The first program will generate the constraint set $\mathcal{C}_1 = \{\text{Ptr}(\beta) \sqsubseteq \text{Ptr}(\alpha), X \sqsubseteq \alpha, \beta \sqsubseteq Y\}$ while the second generates $\mathcal{C}_2 = \{\text{Ptr}(\beta) \sqsubseteq \text{Ptr}(\alpha), X \sqsubseteq \beta, \alpha \sqsubseteq Y\}$. Since each program has the effect of copying the value in x to y , both constraint sets should satisfy $\mathcal{C}_i \vdash X \sqsubseteq Y$. To do this, the pointer subtype constraint must entail some constraint on α and β , but which one?

If we assume that `Ptr` is covariant, then $\text{Ptr}(\beta) \sqsubseteq \text{Ptr}(\alpha)$ entails $\beta \sqsubseteq \alpha$ and so $\mathcal{C}_2 \vdash X \sqsubseteq Y$, but $\mathcal{C}_1 \not\vdash X \sqsubseteq Y$. On the other hand, if we make `Ptr` contravariant then $\mathcal{C}_1 \vdash X \sqsubseteq Y$ but $\mathcal{C}_2 \not\vdash X \sqsubseteq Y$.

It seems that our only recourse is to make subtyping degenerate to type equality under `Ptr`: we are forced to declare that $\text{Ptr}(\beta) \sqsubseteq \text{Ptr}(\alpha) \vdash \alpha = \beta$, which of course means that $\text{Ptr}(\beta) = \text{Ptr}(\alpha)$ already. This is a catastrophe for subtyping as used in machine code, since many natural subtype relations are mediated through pointers. For example, the unary `Ptr` constructor cannot handle the simplest kind of C++ class subtyping, where a derived class physically extends a base class by appending new member variables.

The root cause of the difficulty seems to be in conflating two capabilities that (most) pointers have: the ability to be written through and the ability to be read through. In Retypd, these two capabilities are modeled using different field labels `.store` and `.load`. The `.store` label is contravariant, while the `.load` label is covariant.

To see how the separation of pointer capabilities avoids the loss of precision suffered by `Ptr`, we revisit the two example programs. The first generates the constraint set

$$\mathcal{C}'_1 = \{Q \sqsubseteq P, \quad X \sqsubseteq P.\text{store}, \quad Q.\text{load} \sqsubseteq Y\}$$

By rule T-INHERITR we may conclude that Q also has a field of type `.store`. By S-POINTER we can infer that $Q.\text{store} \sqsubseteq Q.\text{load}$. Finally, since `.store` is contravariant and

$Q \sqsubseteq P$, $\text{S-FIELD}_\ominus$ says we also have $P.\text{store} \sqsubseteq Q.\text{store}$. Putting these parts together gives the subtype chain

$$X \sqsubseteq P.\text{store} \sqsubseteq Q.\text{store} \sqsubseteq Q.\text{load} \sqsubseteq Y$$

The second program generates the constraint set

$$C'_2 = \{Q \sqsubseteq P, \quad X \sqsubseteq Q.\text{store}, \quad P.\text{load} \sqsubseteq Y\}$$

Since $Q \sqsubseteq P$ and P has a field `.load`, we conclude that Q has a `.load` field as well. Next, S-POINTER requires that $Q.\text{store} \sqsubseteq Q.\text{load}$. Since `.load` is covariant, $Q \sqsubseteq P$ implies that $Q.\text{load} \sqsubseteq P.\text{load}$. This gives the subtype chain

$$X \sqsubseteq Q.\text{store} \sqsubseteq Q.\text{load} \sqsubseteq P.\text{load} \sqsubseteq Y$$

By splitting out the read- and write-capabilities of a pointer, we can achieve a sound account of pointer subtyping that does not degenerate to type equality. Note the importance of the consistency condition S-POINTER : this rule ensures that writing through a pointer and reading the result cannot subvert the type system.

The need for separate handling of read- and write-capabilities in a mutable reference has been rediscovered multiple times. A well-known instance is the covariance of the array type constructor in Java and C#, which can cause runtime type errors if the array is mutated; in these languages, the read capabilities are soundly modeled only by sacrificing soundness for the write capabilities.

3.4 Non-structural Subtyping and T-INHERITR

It was noted in §3.2 that the rule T-INHERITR leads to a system with a form of structural typing: any two types in a subtype relation must have the same capabilities. Superficially, this seems problematic for modeling typecasts that forget about fields, such as a cast from `derived*` to `base*` when `derived*` has additional fields (§2.4).

The missing piece that allows us to effectively forget capabilities is instantiation of callee type schemes at a callsite. To demonstrate how polymorphism enables forgetfulness, consider the example type scheme $\forall F. (\exists \tau. C) \Rightarrow F$ from Figure 2. The function `close_last` can be invoked by providing any actual-in type α , such that $\alpha \sqsubseteq F.\text{in_stack0}$; in particular, α can have *more* fields than those required by C . We simply select a more capable type for the existentially-quantified type variable τ in C . In effect, we have used *specialization* of polymorphic types to model *non-structural* subtyping idioms, while *subtyping* is used only to model *structural* subtyping idioms. This restricts our introduction of non-structural subtypes to points where a type scheme is instantiated, such as at a call site.

3.5 Semantics: the Poset of Sketches

The simple type system defined by the deduction rules of Figure 3 defines the *syntax* of legal derivations in our type system. The constraint solver of §5.2 is designed to find a simple representation for all conclusions that can be derived

Figure 5. A sketch instantiating the type scheme in Figure 2.

from a set of type constraints. Yet there is no notion of what a type *is* inherent to the deduction rules of Figure 3. We have defined the rules of the game, but not the equipment with which it should be played.

We found that introducing C-like entities at the level of constraints or types resulted in too much loss of precision when working with the challenging examples described in §2. Consequently we developed the notion of a *sketch*, a kind of regular tree labeled with elements of an auxiliary lattice Λ . Sketches are related to the recursive types studied by Amadio and Cardelli [3] and Kozen et al. [16], but do not depend on *a priori* knowledge of the ranked alphabet of type constructors.

Definition 3.5. A *sketch* is a (possibly infinite) tree T with edges labeled by elements of Σ and nodes marked with elements of a lattice Λ , such that T only has finitely many subtrees up to labeled isomorphism. By collapsing isomorphic subtrees, we can represent sketches as deterministic finite state automata with each state labeled by an element of Λ . The set of sketches admits a lattice structure, with operations described by Figure 18 in [22].

The lattice of sketches serves as the model in which we interpret type constraints. The interpretation of the constraint $\text{VAR } \alpha.u$ is “the sketch S_α admits a path from the root with label sequence u ”, and $\alpha.u \sqsubseteq \beta.v$ is interpreted as “the sketch obtained from S_α by traversing the label sequence u is a subsketch (in the lattice order) of the sketch obtained from S_β by traversing the sequence v .”

The main utility of sketches is that they are nearly a free tree model [25] of the constraint language. Any constraint set C is satisfiable over the lattice of sketches, as long as C cannot prove an impossible subtype relation in the auxiliary lattice Λ . In particular, we can always solve the fragment of C that does not reference constants in Λ . Stated operationally, we can always recover the tree structure of sketches that potentially solve C . This observation is formalized by the following theorem:

Theorem 3.1. Suppose that C is a constraint set over the variables $\{\tau_i\}_{i \in I}$. Then there exist sketches $\{S_i\}_{i \in I}$, such that $w \in S_i$ if and only if $C \vdash \text{VAR } \tau_i.w$.

Proof. The idea is to symmetrize $\sqsubseteq$ using an algorithm that is similar in spirit to Steensgaard’s method of almost-linear-time pointer analysis [33]. Begin by forming a graph with one node $n(\alpha)$ for each derived type variable appearing

in $\mathcal{C}$, along with each of its prefixes. Add a labeled edge $n(\alpha) \xrightarrow{\ell} n(\alpha.\ell)$ for each derived type variable $\alpha.\ell$ to form a graph G . Now quotient G by the equivalence relation $\sim$ defined by $n(\alpha) \sim n(\beta)$ if $\alpha \sqsubseteq \beta \in \mathcal{C}$, and $n(\alpha') \sim n(\beta')$ whenever there are edges $n(\alpha) \xrightarrow{\ell} n(\alpha')$ and $n(\beta) \xrightarrow{\ell'} n(\beta')$ in G with $n(\alpha) \sim n(\beta)$ where either $\ell = \ell'$ or $\ell = \text{.load}$ and $\ell' = \text{.store}$.

By construction, there exists a path through $G/\sim$ with label sequence u starting at the equivalence class of τ_i if and only if $\mathcal{C} \vdash \text{VAR } \tau_i.u$; the (regular) set of all such paths yields the tree structure of S_i . $\square$

Working out the lattice elements that should label S_i is a trickier problem; the basic idea is to use the same automaton Q constructed during constraint simplification (Theorem 5.1 in [22]) to answer queries about which type constants are upper and lower bounds on a given derived type variable. The full algorithm is listed in §D.4 of [22].

In Retypd, we use a large auxiliary lattice Λ containing hundreds of elements that includes a collection of standard C type names, common typedefs for popular APIs, and user-specified semantic classes such as `#FileDescriptor` in Figure 2. This lattice helps model ad-hoc subtyping and preserve high-level semantic type names, as discussed in §2.8.

Note. Sketches are just one of many possible models for the deduction rules that could be proposed. A general approach is to fix a poset $(\mathcal{T}, <:)$ of types, interpret $\sqsubseteq$ as $<:$, and interpret co- and contra-variant field labels as monotone (resp. antimonotone) functions $\mathcal{T} \rightarrow \mathcal{T}$.

The separation of syntax from semantics allows for a simple way to parameterize the type-inference engine by a model of types. By choosing a model $(\mathcal{T}, \equiv)$ with a symmetric relation $\equiv \subseteq \mathcal{T} \times \mathcal{T}$, a unification-based type system similar to SecondWrite [10] is generated. On the other hand, by forming a lattice of type intervals and interval inclusion, we would obtain a type system similar to TIE [17] that outputs upper and lower bounds on each type variable.

4. Analysis Framework

4.1 IR Reconstruction

Retypd is built on top of GrammaTech’s machine-code-analysis tool CodeSurfer for Binaries. CodeSurfer carries out common program analyses on binaries for multiple CPU architectures, including x86, x86-64, and ARM. CodeSurfer is used to recover a high-level IR from the raw machine code; type constraints are generated directly from this IR, and resolved types are applied back to the IR and become visible to the GUI and later analysis phases.

CodeSurfer achieves platform independence through TSL [18], a language for defining a processor’s concrete semantics in terms of concrete numeric types and mapping types that model flag, register, and memory banks. Interpreters for a given abstract domain are automatically created from

the concrete semantics simply by specifying the abstract domain $\mathcal{A}$ and an interpretation of the concrete numeric and mapping types. Retypd uses CodeSurfer’s recovered IR to determine the number and location of inputs and outputs to each procedure, as well as the program’s call graph and per-procedure control-flow graphs. An abstract interpreter then generates sets of type constraints from the concrete TSL instruction semantics. A detailed account of the abstract semantics for constraint generation appears in Appendix A of [22].

4.2 Approach to Type Resolution

After the initial IR is recovered, type inference proceeds in two stages: first, type-constraint sets are generated in a bottom-up fashion over the strongly-connected components of the callgraph. Pre-computed type schemes for externally linked functions may be inserted at this stage. Each constraint set is simplified by eliminating type variables that do not belong to the SCC’s interface; the simplification algorithm is outlined in §5. Once type schemes are available, the callgraph is traversed bottom-up, assigning sketches to type variables as outlined in §3.5. During this stage, type schemes are specialized based on the calling contexts of each function. Appendix F of [22] lists the full algorithms for constraint simplification (§F.1) and solving (§F.2).

4.3 Translation to C Types

The final phase of type resolution converts the inferred sketches to C types for presentation to the user. Since C types and sketches are not directly comparable, this resolution phase necessarily involves the application of heuristic conversion policies. Restricting the heuristic policies to a single post-inference phase provides us with the flexibility to generate high-quality, human-readable C types while maintaining soundness and generality during type reconstruction.

Example 4.1. A simple example involves the generation of const annotations on pointers. We decided on a policy that only introduced const annotations on function parameters, by annotating the parameter at location L when the constraint set $\mathcal{C}$ for procedure p satisfies $\mathcal{C} \vdash \text{VAR } p.\text{in}_L.\text{load}$ and $\mathcal{C} \not\vdash \text{VAR } p.\text{in}_L.\text{store}$. Retypd appears to be the first machine-code type-inference system to infer const annotations; a comparison of our recovered annotations to the original source code appears in §6.4.

Example 4.2. A more complex policy is used to decide between union types and generic types when incompatible scalar constraints must be resolved. Retypd merges comparable scalar constraints to form antichains in Λ ; the elements of these antichains are then used for the resulting C union type.

Example 4.3. The initial type-simplification stage results in types that are as general as possible. Often, this means that types are found to be more general than is strictly helpful to a (human) observer. A policy is applied that specializes type schemes to the most *specific* scheme that is compatible

with all statically-discovered uses. For example, a C++ object may include a getter function with a highly polymorphic type scheme, since it could operate equally well on any structure with a field of the correct type at the correct offset. But we expect that in every calling context, the getter will be called on a specific object type (or perhaps its derived types). We can specialize the getter’s type by choosing the least polymorphic specialization that is compatible with the observed uses. By specializing the function signature before presenting a final C type to the user, we trade some generality for types that are more likely to match the original source.

5. The Simplification Algorithm

In this section, we sketch an outline of the simplification algorithm at the core of the constraint solver. The complete algorithm appears in Appendix D of [22].

5.1 Inferring a Type Scheme

The goal of the simplification algorithm is to take an inferred type scheme $\forall \bar{\alpha}. C \Rightarrow \tau$ for a procedure and create a smaller constraint set C' , such that any constraint on τ implied by C is also implied by C' .

Let C denote the constraint set generated by abstract interpretation of the procedure being analyzed, and let $\bar{\alpha}$ be the set of free type variables in C . We could already use $\forall \bar{\alpha}. C \Rightarrow \tau$ as the constraint set in the procedure’s type scheme, since the input and output types used in a valid invocation of f are tautologically those that satisfy C . Yet, as a practical matter, we cannot use the constraint set directly, since this would result in constraint sets with many useless free variables and a high growth rate over nested procedures.

Instead, we seek to generate a *simplified constraint set* C' , such that if c is an “interesting” constraint and $C \vdash c$ then $C' \vdash c$ as well. But what makes a constraint interesting?

Definition 5.1. For a type variable τ , a constraint is called *interesting* if it has one of the following forms:

- A capability constraint of the form $\text{VAR } \tau.u$
- A recursive subtype constraint of the form $\tau.u \sqsubseteq \tau.v$
- A subtype constraint of the form $\tau.u \sqsubseteq \bar{\kappa}$ or $\bar{\kappa} \sqsubseteq \tau.u$ where $\bar{\kappa}$ is a type constant.

We will call a constraint set C' a *simplification* of C if $C' \vdash c$ for every interesting constraint c , such that $C \vdash c$. Since both C and C' entail the same set of constraints on τ , it is valid to replace C with C' in any valid type scheme for τ .

Simplification heuristics for set-constraint systems were studied by Fähndrich and Aiken [12]; our simplification algorithm encompasses all of these heuristics.

5.2 Unconstrained Pushdown Systems

The constraint-simplification algorithm works on a constraint set C by building a pushdown system $\mathcal{P}_C$ whose transition sequences represent valid derivations of subtyping judgements. We briefly review pushdown systems and some necessary generalizations here.

Definition 5.2. An *unconstrained pushdown system* is a triple $\mathcal{P} = (\mathcal{V}, \Sigma, \Delta)$ where $\mathcal{V}$ is the set of *control locations*, Σ is the set of *stack symbols*, and $\Delta \subseteq (\mathcal{V} \times \Sigma^*)^2$ is a (possibly infinite) set of *transition rules*. We will denote a transition rule by $\langle X; u \rangle \hookrightarrow \langle Y; v \rangle$ where $X, Y \in \mathcal{V}$ and $u, v \in \Sigma^*$. We define the set of *configurations* to be $\mathcal{V} \times \Sigma^*$. In a configuration (p, w) , p is called the *control state* and w the *stack state*.

Note that we require neither the set of stack symbols, nor the set of transition rules, to be finite. This freedom is required to model the derivation S-POINTER of Figure 3, which corresponds to an infinite set of transition rules.

Definition 5.3. An unconstrained pushdown system $\mathcal{P}$ determines a *transition relation* $\longrightarrow$ on the set of configurations: $(X, w) \longrightarrow (Y, w')$ if there is a suffix s and a rule $\langle X; u \rangle \hookrightarrow \langle Y; v \rangle$, such that $w = us$ and $w' = vs$. The transitive closure of $\longrightarrow$ is denoted $\xrightarrow{*}$.

With this definition, we can state the primary theorem behind our simplification algorithm.

Theorem 5.1. Let C be a constraint set and $\mathcal{V}$ a set of base type variables. Define a subset S_C of $(\mathcal{V} \cup \Sigma)^* \times (\mathcal{V} \cup \Sigma)^*$ by $(Xu, Yv) \in S_C$ if and only if $C \vdash Xu \sqsubseteq Yv$. Then S_C is a regular set, and an automaton Q to recognize S_C can be constructed in $O(|C|^3)$ time.

Proof. The basic idea is to treat each $Xu \sqsubseteq Yv \in C$ as a transition rule $\langle X; u \rangle \hookrightarrow \langle Y; v \rangle$ in the pushdown system $\mathcal{P}$. In addition, we add control states $\# \text{START}$, $\# \text{END}$ with transitions $\langle \# \text{START}; X \rangle \hookrightarrow \langle X; \varepsilon \rangle$ and $\langle X; \varepsilon \rangle \hookrightarrow \langle \# \text{END}; X \rangle$ for each $X \in \mathcal{V}$. For the moment, assume that (1) all labels are covariant, and (2) the rule S-POINTER is ignored. By construction, $(\# \text{START}, Xu) \xrightarrow{*} (\# \text{END}, Yv)$ in $\mathcal{P}$ if and only if $C \vdash Xu \sqsubseteq Yv$. A theorem of Büchi [27] ensures that for any two control states A and B in a standard (not unconstrained) pushdown system, the set of all pairs (u, v) with $(A, u) \xrightarrow{*} (B, v)$ is a regular language; Caucal [8] gives a saturation algorithm that constructs an automaton to recognize this language.

In the full proof, we add two novelties: first, we support contravariant stack symbols by encoding variance data into the control states and transition rules. The second novelty involves the rule S-POINTER; this rule is problematic since the natural encoding would result in infinitely many transition rules. We extend Caucal’s construction to lazily instantiate all necessary applications of S-POINTER during saturation. For details, see Appendix D of [22]. $\square$

Since C will usually entail an infinite number of constraints, this theorem is particularly useful: it tells us that the full set of constraints entailed by C has a finite encoding by an automaton Q . Further manipulations on the constraint closure, such as efficient minimization, can be carried out on Q . By restricting the transitions to and from $\# \text{START}$ and

#END, the same algorithm is used to eliminate type variables, producing the desired constraint simplifications.

5.3 Overall Complexity of Inference

The saturation algorithm used to perform constraint-set simplification and type-scheme construction is, in the worst case, cubic in the number of subtype constraints to simplify. Since some well-known pointer analysis methods also have cubic complexity (such as Andersen [4]), it is reasonable to wonder if Retypd’s “points-to free” analysis really offers a benefit over a type-inference system built on top of points-to analysis data.

To understand where Retypd’s efficiencies are found, first consider the n in $O(n^3)$. Retypd’s core saturation algorithm is cubic in the number of subtype constraints; due to the simplicity of machine-code instructions, there is roughly one subtype constraint generated per instruction. Furthermore, Retypd applies constraint simplification on each procedure in isolation to eliminate the procedure-local type variables, resulting in constraint sets that only relate procedure formal-ins, formal-outs, globals, and type constants. In practice, these simplified constraint sets are small.

Since each procedure’s constraint set is simplified independently, the n^3 factor is controlled by the largest procedure size, not the overall size of the binary. By contrast, source-code points-to analysis such as Andersen’s are generally cubic in the overall number of pointer variables, with exponential duplication of variables depending on the call-string depth used for context sensitivity. The situation is even more difficult for machine-code points-to analyses such as VSA, since there is no syntactic difference between a scalar and a pointer in machine code. In effect, every program variable must be treated as a potential pointer.

On our benchmark suite of real-world programs, we found that execution time for Retypd scales slightly below $O(N^{1.08})$, where N is the number of program instructions. The following back-of-the-envelope calculation can heuristically explain much of the disparity between the $O(N^3)$ theoretical complexity and the $O(N^{1.08})$ measured complexity. On our benchmark suite, the maximum procedure size n grew roughly like $n \approx N^{2/5}$. We could then expect that a per-procedure analysis would perform worst when the program is partitioned into $N^{3/5}$ procedures of size $N^{2/5}$. On such a program, a per-procedure $O(n^k)$ analysis may be expected to behave more like an $O(N^{3/5} \cdot (N^{2/5})^k) = O(N^{(3+2k)/5})$ analysis overall. In particular, a per-procedure cubic analysis like Retypd could be expected to scale like a global $O(N^{1.8})$ analysis. The remaining differences in observed versus theoretical execution time can be explained by the facts that real-world constraint graphs do not tend to exercise the simplification algorithm’s worst-case behavior, and that the distribution of procedure sizes is heavily weighted towards small procedures.

```

INT_PTR Proto_EnumAccounts(WPARAM wParam,
                           LPARAM lParam)
{
    *(int*)wParam = accounts.getCount();
    *(PROTOACCOUNT**)lParam =
        accounts.getArray();
    return 0;
}

```

Figure 6. Ground-truth types declared in the original source do not necessarily reflect program semantics. Example from `miranda32`.

6. Evaluation

6.1 Implementation

Retypd is implemented as a module within CodeSurfer for Binaries. By leveraging the multi-platform disassembly capabilities of CodeSurfer, it can operate on x86, x86-64, and ARM code. We performed the evaluation using minimal analysis settings, disabling value-set analysis (VSA) but computing affine relations between the stack and frame pointers. Enabling additional CodeSurfer phases such as VSA can greatly improve the reconstructed IR, at the expense of increased analysis time.

Existing type-inference algorithms such as TIE [17] and SecondWrite [10] require some modified form of VSA to resolve points-to data. Our approach shows that high-quality types can be recovered in the absence of points-to information, allowing type inference to proceed even when computing points-to data is too unreliable or expensive.

6.2 Evaluation Setup

Our benchmark suite consists of 160 32-bit x86 binaries for both Linux and Windows, compiled with a variety of gcc and Microsoft Visual C/C++ versions. The benchmark suite includes a mix of executables, static libraries, and DLLs. The suite includes the same coreutils and SPEC2006 benchmarks used to evaluate REWARDS, TIE, and SecondWrite [10, 17, 19]; additional benchmarks came from a standard suite of real-world programs used for precision and performance testing of CodeSurfer for Binaries. All binaries were built with optimizations enabled and debug information disabled.

Ground truth is provided by separate copies of the binaries that have been built with the same settings, but with debug information included (DWARF on Linux, PDB on Windows). We used IdaPro [15] to read the debug information, which allowed us to use the same scripts for collecting ground-truth types from both DWARF and PDB data.

All benchmarks were evaluated on a 2.6 GHz Intel Xeon E5-2670 CPU, running on a single logical core. RAM utilization by CodeSurfer and Retypd combined was capped at 10GB.

Benchmark	Description	Instructions
<i>CodeSurfer benchmarks</i>		
libidn	Domain name translator	7K
Tutorial00	Direct3D tutorial	9K
zlib	Compression library	14K
ogg	Multimedia library	20K
distributor	UltraVNC repeater	22K
libbz2	BZIP library, as a DLL	37K
glut	The glut32.dll library	40K
pngtest	A test of libpng	42K
freelut	The freelut.dll library	77K
miranda	IRC client	100K
XMail	Email server	137K
yasm	Modular assembler	190K
python21	Python 2.1	202K
quake3	Quake 3	281K
TinyCad	Computed-aided design	544K
Shareaza	Peer-to-peer file sharing	842K
<i>SPEC2006 benchmarks</i>		
470.lbm	Lattice Boltzmann Method	3K
429.mcf	Vehicle scheduling	3K
462.libquantum	Quantum computation	11K
401.bzip2	Compression	13K
458.sjeng	Chess AI	25K
433.milc	Quantum field theory	28K
482.sphinx3	Speech recognition	43K
456.hmmer	Protein sequence analysis	71K
464.h264ref	Video compression	113K
445.gobmk	GNU Go AI	203K
400.perlbench	Perl core	261K
403.gcc	C/C++/Fortran compiler	751K

Figure 7. Benchmarks used for evaluation. All binaries were compiled from source using optimized release configurations. The SPEC2006 benchmarks were chosen to match the benchmarks used to evaluate SecondWrite [10].

Our benchmark suite includes the individual binaries in Figure 7 as well as the collections of related binaries shown in Figure 10. We found that programs from a single collection tended to share a large amount of common code, leading to highly correlated benchmark results. For example, even though the coreutils benchmarks include many tools with very disparate purposes, all of the tools make use of a large, common set of statically linked utility routines. Over 80% of the .text section in tail consists of such routines; for yes, the number is over 99%. The common code and specific idioms appearing in coreutils make it a particularly low-variance benchmark suite.

In order to avoid over-representing these program collections in our results, we treated these collections as clusters in the data set. For each cluster, we computed the average of each metric over the cluster, then inserted the average as a

single data point to the final data set. Because Retypd performs well on many clusters, this averaging procedure tends to reduce our overall precision and conservativeness measurements. Still, we believe that it gives a less biased depiction of the algorithm’s expected real-world behavior than does an average over all benchmarks.

6.3 Sources of Imprecision

Although Retypd is built around a sound core of constraint simplification and solving, there are several ways that imprecision can occur. As described in §2.5, disassembly failures can lead to unsound constraint generation. Second, the heuristics for converting from sketches to C types are lossy by necessity. Finally, we treat the source types as ground truth, leading to “failures” whenever Retypd recovers an accurate type that does not match the original program—a common situation with type-unsafe source code.

A representative example of this last source of imprecision appears in Figure 6. This source code belongs to the miranda32 IRC client, which uses a plugin-based architecture; most of miranda32’s functionality is implemented by “service functions” with the fixed signature `int ServiceProc(WPARAM, LPARAM)`. The types `WPARAM` and `LPARAM` are used in certain Windows APIs for generic 16- and 32-bit values. The two parameters are immediately cast to other types in the body of the service functions, as in Figure 6.

6.4 const Correctness

As a side-effect of separately modeling `.load` and `.store` capabilities, Retypd is easily able to recover information about how pointer parameters are used for input and/or output. We take this into account when converting sketches to C types; if a function’s sketch includes `.inL.load` but not `.inL.store` then we annotate the parameter at `L` with `const`, as in Figure 5 and Figure 2. Retypd appears to be the first machine-code type-inference system to infer `const` annotations directly.

On our benchmark suite, we found that 98% of parameter `const` annotations in the original source code were recovered by Retypd. Furthermore, Retypd inferred `const` annotations on many other parameters; unfortunately, since most C and C++ code does not use `const` in every possible situation, we do not have a straightforward way to detect how many of Retypd’s additional `const` annotations are correct.

Manual inspection of the missed `const` annotations shows that most instances are due to imprecision when analyzing one or two common statically linked library functions. This imprecision then propagates outward to callers, leading to decreased `const` correctness overall. Still, we believe the 98% recovery rate shows that Retypd offers a useful approach to `const` inference.

Figure 8. Distance to ground-truth types and size of the interval between inferred upper and lower bounds. Smaller distances represent more accurate types; smaller interval sizes represent increased confidence.

6.5 Comparisons to Other Tools

We gathered results over several metrics that have been used to evaluate SecondWrite, TIE, and REWARDS. These metrics were defined by Lee et al. [17] and are briefly reviewed here.

TIE infers upper and lower bounds on each type variable, with the bounds belonging to a lattice of C-like types. The lattice is naturally stratified into levels, with the distance between two comparable types roughly being the difference between their levels in the lattice, with a maximum distance of 4. A recursive formula for computing distances between pointer and structural types is also used. TIE also determines a policy that selects between the upper and lower bounds on a type variable for the final displayed type.

TIE considers three metrics based on this lattice: the conservativeness rate, the interval size, and the distance. A type interval is *conservative* if the interval bounds overapproximate the declared type of a variable. The *interval size* is the lattice distance from the upper to the lower bound on a type variable. The *distance* measures the lattice distance from the final displayed type to the ground-truth type. REWARDS and SecondWrite both use unification-based algorithms, and have been evaluated using the same TIE metrics. The evaluation of REWARDS using TIE metrics appears in Lee et al. [17].

Distance and interval size: Retypd shows substantial improvements over other approaches in the distance and interval-size metrics, indicating that it generates more accurate types with less uncertainty. The mean distance to the ground-truth type was 0.54 for Retypd, compared to 1.15 for dynamic TIE, 1.53 for REWARDS, 1.58 for static TIE, and 1.70 for SecondWrite. The mean interval size shrunk to 1.2 with Retypd, compared to 1.7 for SecondWrite and 2.0 for TIE.

Figure 9. Conservativeness and pointer accuracy metric. Perfect type reconstruction would be 100% conservative and match on 100% of pointer levels. Note that the y axis begins at 70%.

Multi-level pointer accuracy: ElWazeer et al. [10] also introduced a multi-level pointer-accuracy rate that attempts to quantify how many “levels” of pointers were correctly inferred. On SecondWrite’s benchmark suite, Retypd attained a mean multi-level pointer accuracy of 91%, compared with SecondWrite’s reported 73%. Across all benchmarks, Retypd averages 88% pointer accuracy.

Conservativeness: The best type system would have a high conservativeness rate (few unsound decisions) coupled with a low interval size (tightly specified results) and low distance (inferred types are close to ground-truth types). In each of these metrics, Retypd performs about as well or better than existing approaches. Retypd’s mean conservativeness rate is 95%, compared to 94% for TIE. But note that TIE was evaluated only on `coreutils`; on that cluster, Retypd’s conservativeness was 98%. SecondWrite’s overall conservativeness is 96%, measured on a subset of the SPEC2006 benchmarks; Retypd attained a slightly lower 94% on this subset.

It is interesting to note that Retypd’s conservativeness rate on `coreutils` is comparable to that of REWARDS, even though REWARDS’ use of dynamic execution traces suggests it would be more conservative than a static analysis by virtue of only generating feasible type constraints.

6.6 Performance

Although the core simplification algorithm of Retypd has cubic worst-case complexity, it only needs to be applied on a per-procedure basis. This suggests that the real-world scaling behavior will depend on the distribution of procedure sizes, not on the whole-program size.

Cluster	Count	Description	Instructions	Distance	Interval	Conserv.	Ptr. Acc.	Const
freeglut-demos	3	freeglut samples	2K	0.66	1.49	97%	83%	100%
coreutils	107	GNU coreutils 8.23	10K	0.51	1.19	98%	82%	96%
vpx-d	8	VPx decoders	36K	0.63	1.68	98%	92%	100%
vpx-e	6	VPx encoders	78K	0.63	1.53	96%	90%	100%
sphinx2	4	Speech recognition	83K	0.42	1.09	94%	91%	99%
putty	4	SSH utilities	97K	0.51	1.05	94%	86%	99%
Retypd, as reported				0.54	1.20	95%	88%	98%
Retypd, without clustering				0.53	1.22	97%	84%	97%

Figure 10. Clusters in the benchmark suite. For each metric, the average over the cluster is given. If a cluster average is worse than Retypd’s overall average for a certain metric, a box is drawn around the entry.

Figure 11. Type-inference time on benchmarks (25 trials). The line indicates the best-fit exponential $t = 0.000804 \cdot N^{1.078}$ modeling the average time, demonstrating slightly superlinear scaling. The coefficient of determination is $R^2 = 0.962$.

In practice, Figure 11 suggests that Retypd gives nearly linear performance over the benchmark suite, which ranges in size from 2K to 840K instructions. To measure Retypd’s performance, we used numerical regression to find the best-fit model $T = \alpha N^\beta$ relating execution time T to program size N . This results in the relation $T = 0.000804 \cdot N^{1.078}$ with coefficient of determination $R^2 = 0.962$, suggesting that over 96% of the variation in performance data can be explained by this model. In other words, on real-world programs Retypd demonstrates nearly linear scaling of execution time. The cubic worst-case per-procedure behavior of the constraint solver does not translate to cubic behavior overall. Similarly, we found that the sub-linear model $m = 0.023 \cdot N^{0.882}$ explains 98.7% of the memory usage in Retypd.

Note. The regressions above were performed by numerically fitting exponential models in (N, T) and (N, m) space, rather than analytically fitting linear models in log-log space. Our

Figure 12. Average type-inference memory usage on benchmarks. The line indicates the best-fit exponential $m = 0.023 \cdot N^{0.882}$, with coefficient of determination $R^2 = 0.987$. Over 25 trials, the average standard deviation in memory usage was 1.4%.

models then minimize the error in the predicted values of T and m , rather than minimizing errors in $\log T$ or $\log m$. Linear regression in log-log space results in the less-predictive models $T = 0.0002 \cdot N^{1.14}$ with $R^2 = 0.9091$ and $m = 0.07 \cdot N^{0.78}$ with $R^2 = 0.9094$.

The constraint-simplification workload of Retypd would be straightforward to parallelize over the directed acyclic graph of strongly-connected components in the callgraph, further reducing the scaling constant. If a reduction in memory usage is required, Retypd could swap constraint sets to disk; the current implementation keeps all working sets in RAM.

7. Related Work

Machine-code type recovery: Hex-Ray’s reverse engineering tool IdaPro [15] is an early example of type reconstruction via static analysis. The exact algorithm is proprietary, but it appears that IdaPro propagates types through unification from

library functions of known signature, halting the propagation when a type conflict appears. *IdaPro*’s reconstructed IR is relatively sparse, so the type propagation fails to produce useful information in many common cases, falling back to the default `int` type. However, the analysis is very fast.

SecondWrite [10] is an interesting approach to static IR reconstruction with a particular emphasis on scalability. The authors combine a best-effort VSA variant for points-to analysis with a unification-based type-inference engine. Accurate types in *SecondWrite* depend on high-quality points-to data; the authors note that this can cause type accuracy to suffer on larger programs. In contrast, *Retypd* is not dependent on points-to data for type recovery and makes use of subtyping rather than unification for increased precision.

TIE [17] is a static type-reconstruction tool used as part of Carnegie Mellon University’s binary-analysis platform (BAP). *TIE* was the first machine-code type-inference system to track subtype constraints and explicitly maintain upper and lower bounds on each type variable. As an abstraction of the C type system, *TIE*’s type lattice is relatively simple; missing features, such as recursive types, were later identified by the authors as an important target for future research [30].

HOWARD [32] and *REWARDS* [19] both take a dynamic approach, generating type constraints from execution traces. Through a comparison with *HOWARD*, the creators of *TIE* showed that static type analysis can produce higher-precision types than dynamic type analysis, though a small penalty must be paid in conservativeness of constraint-set generation. *TIE* also showed that type systems designed for static analysis can be easily modified to work on dynamic traces; we expect the same is true for *Retypd*, though we have not yet performed these experiments.

Most previous work on machine-code type recovery, including *TIE* and *SecondWrite*, either disallows recursive types or only supports recursive types by combining type-inference results with a points-to oracle. For example, to infer that x has a the type `struct S { struct S *, ... }*` in a unification-based approach like *SecondWrite*, first we must have resolved that x points to some memory region M , that M admits a 4-byte abstract location α at offset 0, and that the type of α should be unified with the type of x . If pointer analysis has failed to compute an explicit memory region pointed to by x , it will not be possible to determine the type of x correctly. The complex interplay between type inference, points-to analysis, and abstract-location delineation leads to a relatively fragile method for inferring recursive types. In contrast, our type system can infer recursive types even when points-to facts are completely absent.

Robbins et al. [28] developed an SMT solver equipped with a theory of rational trees and applied it to type reconstruction. Although this allows for recursive types, the lack of subtyping and the performance of the SMT solver make it difficult to scale this approach to real-world binaries. Except for test cases on the order of 500 instructions, precision of the recovered types was not assessed.

Related type systems: The type system used by *Retypd* is related to the recursively constrained types (rc types) of Eifrig, Smith, and Trifonov [9]. *Retypd* generalizes the rc type system by building up all types using flexible records; even the function-type constructor $\rightarrow$, taken as fundamental in the rc type system, is decomposed into a record with in and out fields. This allows *Retypd* to operate without the knowledge of a fixed signature from which type constructors are drawn, which is essential for analysis of stripped machine code.

The use of CFL reachability to perform polymorphic subtyping first appeared in Rehof and Fähndrich [26], extending previous work relating simpler type systems to graph reachability [2, 23]. *Retypd* continues by adding type-safe handling of pointers and a simplification algorithm that allows us to compactly represent the type scheme for each function.

CFL reachability has also been used to extend the type system of Java [14] and C++ [11] with support for additional type qualifiers. Our reconstructed `const` annotations can be seen as an instance of this idea, although our qualifier inference is not separated from type inference.

To the best of our knowledge, no prior work has applied polymorphic type systems with subtyping to machine code.

8. Future Work

One interesting avenue for future research could come from the application of dependent and higher-rank type systems to machine-code type inference, although we rapidly approach the frontier where type inference is undecidable. A natural example of dependent types appearing in machine code is `malloc`, which could be typed as $\text{malloc} : (n : \text{size_t}) \rightarrow \tau_n$ where τ_n denotes the common supertype of all n -byte types. The key feature is that the *value* of a parameter determines the *type* of the result.

Higher-rank types are needed to properly model functions that accept pointers to polymorphic functions as parameters. Such functions are not entirely uncommon; for example, any function that is parameterized by a custom polymorphic allocator will have rank ≥ 2 .

Retypd was implemented as an inference phase that runs after *CodeSurfer*’s main analysis loop. We expect that by moving *Retypd* into *CodeSurfer*’s analysis loop, there will be an opportunity for interesting interactions between IR generation and type reconstruction.

9. Conclusion

By examining a diverse corpus of optimized binaries, we have identified a number of common idioms that are stumbling blocks for machine-code type inference. For each of these idioms, we identified a type-system feature that could enable the difficult code to be properly typed. We gathered these features into a type system and implemented the inference algorithm in the tool *Retypd*. Despite removing the requirement for points-to data, *Retypd* is able to accurately and conservatively type a wide variety of real-world binaries.

We assert that Retypd demonstrates the utility of high-level type systems for reverse engineering and binary analysis.

Acknowledgments

The authors would like to thank Vineeth Kashyap and the anonymous reviewers for their many useful comments on this manuscript, and John Phillips, David Ciarletta, and Tim Clark for their help with test automation.

References

- [1] ISO/IEC TR 19768:2007: Technical report on C++ library extensions, 2007.
- [2] O. Agesen. Constraint-based type inference and parametric polymorphism. In *Static Analysis Symposium*, pages 78–100, 1994.
- [3] R. M. Amadio and L. Cardelli. Subtyping recursive types. *ACM Transactions on Programming Languages and Systems*, 15(4):575–631, 1993.
- [4] L. O. Andersen. *Program analysis and specialization for the C programming language*. PhD thesis, University of Copenhagen, 1994.
- [5] G. Balakrishnan, R. Gruian, T. Reps, and T. Teitelbaum. CodeSurfer/x86 – a platform for analyzing x86 executables. In *Compiler Construction*, pages 250–254, 2005.
- [6] G. Balakrishnan and T. Reps. Analyzing Memory Accesses in x86 Executables, in *Compiler Construction*, pages 5–23, 2004.
- [7] A. Carayol and M. Hague. Saturation algorithms for model-checking pushdown systems. In *International Conference on Automata and Formal Languages*, pages 1–24, 2014.
- [8] D. Caucal. On the regular structure of prefix rewriting. *Theoretical Computer Science*, 106(1):61–86, 1992.
- [9] J. Eifrig, S. Smith, and V. Trifonov. Sound polymorphic type inference for objects. In *Object-Oriented Programming, Systems, Languages, and Applications*, pages 169–184, 1995.
- [10] K. ElWazeer, K. Anand, A. Kotha, M. Smithson, and R. Barua. Scalable variable and data type detection in a binary rewriter. In *Programming Language Design and Implementation*, pages 51–60, 2013.
- [11] J. S. Foster, R. Johnson, J. Kodumal, and A. Aiken. Flow-insensitive type qualifiers. *ACM Transactions on Programming Languages and Systems*, 28(6):1035–1087, 2006.
- [12] M. Fähndrich and A. Aiken. Making set-constraint program analyses scale. In *Workshop on Set Constraints*, 1996.
- [13] D. Gopan, E. Driscoll, D. Nguyen, D. Naydich, A. Loginov, and D. Melski. Data-delineation in software binaries and its application to buffer-overflow discovery. In *International Conference on Software Engineering*, pages 145–155, 2015.
- [14] D. Greenfieldboyce and J. S. Foster. Type qualifier inference for Java. In *Object-Oriented Programming, Systems, Languages, and Applications*, pages 321–336, 2007.
- [15] Hex-Rays. Hex-Rays IdaPro. <http://www.hex-rays.com/products/ida/>, 2015.
- [16] D. Kozen, J. Palsberg, and M. I. Schwartzbach. Efficient recursive subtyping. *Mathematical Structures in Computer Science*, 5(01):113–125, 1995.
- [17] J. Lee, T. Avgerinos, and D. Brumley. TIE: Principled reverse engineering of types in binary programs. In *Network and Distributed System Security Symposium*, pages 251–268, 2011.
- [18] J. Lim and T. Reps. TSL: A system for generating abstract interpreters and its application to machine-code analysis. *ACM Transactions on Programming Languages and Systems*, 35(1):4, 2013.
- [19] Z. Lin, X. Zhang, and D. Xu. Automatic reverse engineering of data structures from binary execution. In *Network and Distributed System Security Symposium*, 2010.
- [20] S. Marlow, A. R. Yakushev, and S. Peyton Jones. Faster laziness using dynamic pointer tagging. In *International Conference on Functional Programming*, pages 277–288, 2007.
- [21] L. Mauborgne and X. Rival. Trace partitioning in abstract interpretation based static analyzers. In *Programming Languages and Systems*, pages 5–20, 2005.
- [22] M. Noonan, A. Loginov, and D. Cok. Polymorphic type inference for machine code (extended version). URL <http://arxiv.org/abs/1603.05495>.
- [23] J. Palsberg and P. O’Keefe. A type system equivalent to flow analysis. *ACM Transactions on Programming Languages and Systems*, 17(4):576–599, 1995.
- [24] J. Palsberg, M. Wand, and P. O’Keefe. Type inference with non-structural subtyping. *Formal Aspects of Computing*, 9(1):49–67, 1997.
- [25] F. Pottier and D. Rémy. The essence of ML type inference. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 10. MIT Press, 2005.
- [26] J. Rehof and M. Fähndrich. Type-based flow analysis: From polymorphic subtyping to cfl-reachability. In *Principles of Programming Languages*, pages 54–66, 2001.
- [27] J. Richard Büchi. Regular canonical systems. *Archive for Mathematical Logic*, 6(3):91–111, 1964.
- [28] E. Robbins, J. M. Howe, and A. King. Theory propagation and rational-trees. In *Principles and Practice of Declarative Programming*, pages 193–204, 2013.
- [29] M. Robertson. *A Brief History of InvSqrt*. PhD thesis, University of New Brunswick, 2012.
- [30] E. J. Schwartz, J. Lee, M. Woo, and D. Brumley. Native x86 decompilation using semantics-preserving structural analysis and iterative control-flow structuring. In *USENIX Security Symposium*, pages 353–368, 2013.
- [31] M. Siff, S. Chandra, T. Ball, K. Kunchithapadam, and T. Reps. Coping with type casts in C. In *Software Engineering—ESEC/FSE’99*, pages 180–198, 1999.
- [32] A. Slowinska, T. Stancescu, and H. Bos. Howard: A dynamic excavator for reverse engineering data structures. In *Network and Distributed System Security Symposium*, 2011.
- [33] B. Steensgaard. Points-to analysis in almost linear time. In *Principles of Programming Languages*, pages 32–41, 1996.
- [34] Z. Su, A. Aiken, J. Niehren, T. Priesnitz, and R. Treinen. The first-order theory of subtyping constraints. In *Principles of Programming Languages*, pages 203–216, 2002.