Certifying Compilation and Run-time Code Generation

Luke Hornof (hornof@cis.upenn.edu) and Trevor Jim (tjim@cis.upenn.edu)
Computer and Information Science Department
University of Pennsylvania
Philadelphia, PA 19103

April 13, 1999

Abstract. A certifying compiler takes a source language program and produces object code, as well as a "certificate" that can be used to verify that the object code satisfies desirable properties, such as type safety and memory safety. Certifying compilation helps to increase both compiler robustness and program safety. Compiler robustness is improved since some compiler errors can be caught by checking the object code against the certificate immediately after compilation. Program safety is improved because the object code and certificate alone are sufficient to establish safety: even if the object code and certificate are produced on an unknown machine by an unknown compiler and sent over an untrusted network, safe execution is guaranteed as long as the code and certificate pass the verifier.

Existing work in certifying compilation has addressed statically generated code. In this paper, we extend this to code generated at run time. Our goal is to combine certifying compilation with run-time code generation to produce programs that are both verifiably safe and extremely fast. To achieve this goal, we present two new languages with explicit run-time code generation constructs: Cyclone, a type safe dialect of C, and TAL/T, a type safe assembly language. We have designed and implemented a system that translates a safe C program into Cyclone, which is then compiled to TAL/T, and finally assembled into executable object code. This paper focuses on our overall approach and the front end of our system; details about TAL/T will appear in a subsequent paper.

Keywords: program specialization, partial evaluation, run-time code generation, certifying compilation, proof-carrying code, typed assembly language

1. Introduction

1.1. Run-time specialization

Specialization is a program transformation that optimizes a program with respect to invariants. This technique has been shown to give dramatic speedups on a wide range of applications, including aircraft crew planning programs, image shaders, and operating systems (Augustsson, 1997; Guenter et al., 1995; Muller et al., 1998). *Run-time* specialization exploits invariants that become available during the execution of a program, generating optimized code on the fly. Opportunities for run-time specialization occur when dynamically changing values remain invariant for a period of time. For example, networking software can be specialized to a particular TCP connection or multicast tree.

Run-time code generation is tricky. It is hard to correctly write and reason about code that generates code; it is not obvious how to optimize or debug a program that

© 1999 *Kluwer Academic Publishers. Printed in the Netherlands.*

has yet to be generated. Early examples of run-time code generation include self-modifying code, and ad hoc code generators written by hand with a specific function in mind. These approaches proved complicated and error prone (Keppel et al., 1991).

More recent work has applied advanced programming language techniques to the problem. New source languages have been designed to facilitate run-time code generation by providing the programmer with high-level constructs and having the compiler implement the low-level details (Leone and Lee, 1994; Poletto et al., ; Taha and Sheard, 1997). Program transformations based on static analyses are now capable of automatically translating a normal program into a run-time code generating program (Consel et al., 1996; Grant et al., 1997; Hornof, 1997). And type systems can check run-time code generating programs at compile time, ensuring that certain bugs will not occur at run time (provided the compiler is correct) (Taha and Sheard, 1997; Wickline et al., 1998).

These techniques make it easier for programmers to use run-time code generation, but they do not address the concerns of the compiler writer or end user. The compiler writer still needs to implement a correct compiler—not easy even for a language without run-time code generation. The end user would like some assurance that executables will not crash their machine, even if the programs generate code and jump to it—behavior that usually provokes suspicion in security-concious users. We will address both of these concerns through another programming language technique, *certifying compilation*.

1.2. Certifying compilation

A certifying compiler takes a source language program and produces object code and a "certificate" that may help to show that the object code satisfies certain desirable properties (Morrisett et al., 1998; Necula and Lee, 1998). A separate component called the *verifier* examines the object code and certificate and determines whether the object code actually satisfies the properties. A wide range of properties can be verified, including memory safety (unallocated portions of memory are not accessed), control safety (code is entered only at valid entry points), and various security properties (e.g., highly classified data does appear on low security channels). Often, these properties are corollaries of type safety in an appropriate type system for the object code.

In this paper we will describe a certifying compiler for *Cyclone*, a high-level language that supports run-time code generation. Cyclone is compiled into *TAL/T*, an assembly language that supports run-time code generation. Cyclone and TAL/T are both type safe; the certificates of our system are the type annotations of the TAL/T output, and the verifier is the TAL/T type checker.

As compiler writers, we were motivated to implement Cyclone as a certifying compiler because we believe the approach enhances compiler correctness. For example, we were forced to develop a type system and operational semantics for TAL/T. This provides a formal framework for reasoning about object code that generates object code at run time. Eventually, we hope to prove that the compiler transforms type

correct source programs into type correct object programs, an important step towards proving correctness for the compiler. In the meantime, we use the verifier to type check the output of the compiler, so that we get immediate feedback when our compiler introduces type errors. As others have noted (Tarditi, 1997; Tarditi et al., 1996), this helps to identify and correct compiler bugs quickly.

We also wanted a certifying compiler to address the safety concerns of end users. In our system, type safety only depends on the certificate and the object code, and not on the method by which they are produced. Thus the end user does not have to rely on the programmer or the Cyclone compiler to ensure safety. This makes our system usable as the basis of security-critical applications like active networks and mobile code systems.

1.3. The Cyclone compiler

The Cyclone compiler is built on two existing systems, the Tempo specializer (Noël et al., 1998) and the Popcorn certifying compiler (Morrisett et al., 1998). It has three phases, shown in Fig. 1.

The first phase transforms a type safe C program into a Cyclone program that uses run-time code generation. It starts by applying the static analyses of the Tempo system to a C program and context information that specifies which function arguments are invariant. The Tempo front end produces an action-annotated program. We added an additional pass to translate the action-annotated program into a Cyclone run-time specializer.

The second phase verifies that the Cyclone program is type safe, and then compiles it into TAL/T. To do this, we modified the Popcorn compiler of Morrisett et al.; Popcorn compiles a type safe dialect of C into TAL, a typed assembly language. We extended the front end of Popcorn to handle Cyclone programs, and modified its back end so that it outputs TAL/T. TAL/T is TAL extended with instructions for manipulating *templates*, code fragments parameterized by holes, and their corresponding types. This compilation phase not only transforms high-level Cyclone constructs into low-level assembly instructions, but also transforms Cyclone types into TAL/T types.

The third phase first verifies the type safety of the TAL/T program. The type system of TAL/T ensures that the templates are combined correctly and that holes are filled in correctly. This paper describes our overall approach and the front end in detail, but the details of TAL/T will appear in a subsequent paper. Finally, the TAL/T program is assembled and linked into an executable.

This three phase design offers a very flexible user interface since it allows programs to be written in C, Cyclone, or TAL/T. In the simplest case, the user can simply write a C program (or reuse an existing program) and allow the system to handle the rest. If the user desires more explicit control over the code generation process, he may write (or modify) a Cyclone program. If very fine-grain control is desired, the user can fine-tune a TAL/T program produced by Cyclone, or can write one by hand. Note that,

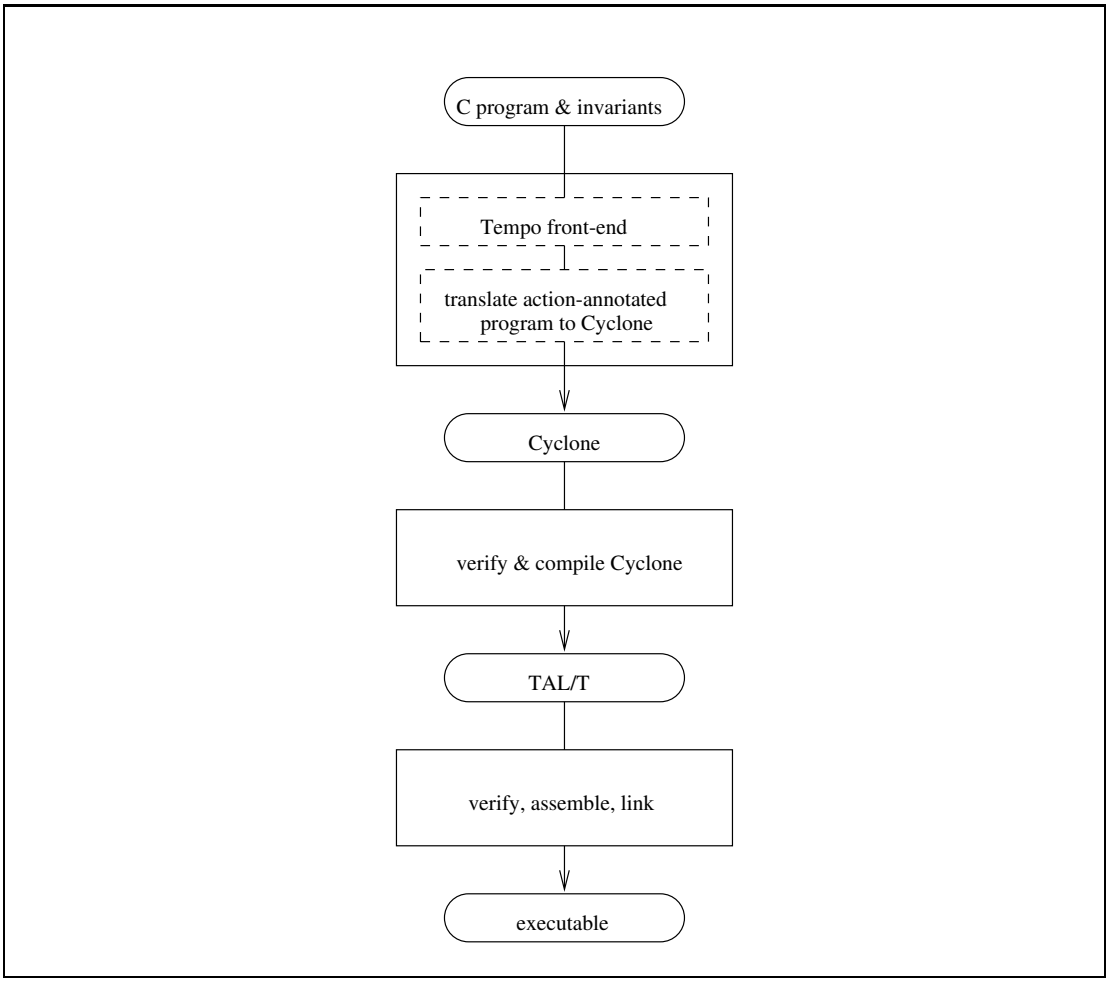

Figure 1. Overview of the Cyclone compiler

since verification is performed at the TAL/T level, the same program safety properties are guaranteed in all three of these cases.

1.4. Example

We now present an example that illustrates run-time code generation and the phases of our Cyclone compiler. Fig. 2 shows a modular exponentiation function, `mexp`, written in standard C. Its arguments are a base value, an exponent, and a modulus. Modular exponentiation is often used in cryptography; when the same key is used to encrypt or decrypt several messages, the function is called repeatedly with the same exponent and modulus. Thus `mexp` can benefit from specialization.

C code (invariant arguments in *italics*)

```
int mexp(int base, int exp, int mod)
{
  int s, t, u;

  s = 1; t = base; u = exp;

  while (u != 0) {
    if ((u&1) != 0)
      s = (s*t) % mod;
    t = (t*t) % mod;
    u >>= 1;
  }
  return(s);
}
```

Action-annotated code (*italicized* constructs can be evaluated)

```
int mexp(int base, int exp, int mod)
{
  int s, t, u;

  s = 1; t = base; u = exp;

  while (u != 0) {
    if ((u & 1) != 0)
      s = (s*t) % mod;
    t = (t*t) % mod;
    u >>= 1;
  }
  return(s);
}
```

Specialized source code (`exp = 2`, `mod = 1234`)

```
int mexp_sp(int base)
{
  int s, t;

  s = 1; t = base;

  t = (t*t) % 1234;
  s = (s*t) % 1234;
  t = (t*t) % 1234;

  return(s);
}
```

Figure 2. Specialization at the source level

To specialize the function with respect to a given exponent and modulus, the user indicates that the two arguments are *invariant*: the function will be called repeatedly with the same values for the invariant arguments. In Fig. 2, invariant arguments are shown in *italics*. A static analysis propagates this information throughout the program, producing an *action-annotated* program. Actions describe how each language construct will be treated during specialization. Constructs that depend only on invariants can be evaluated during specialization; these constructs are displayed in *italics* in the second part of the figure.

To understand how run-time specialization works, it is helpful to first consider how specialization could be achieved entirely within the source language. In our example, the specialized function `mexp_sp` of Fig. 2 is obtained from the action-annotated `mexp` when the exponent is 2 and the modulus is 1234. Italicized constructs of `mexp`, like the while loop, can be evaluated (note that the loop test depends only on the known arguments). Non-italicized constructs of `mexp` show up in the source code of `mexp_sp`. These constructs can only be evaluated when `mexp_sp` is called, because they depend on the unknown argument.

We can think of `mexp_sp` as being constructed by cutting and pasting together fragments of the source code of `mexp`. These fragments, or *templates*, are a central idea we used in designing Cyclone. Cyclone is a type safe dialect of C extended with four constructs that manipulate templates: `codegen`, `cut`, `splice`, and `fill`. Using these constructs, it is possible to write a Cyclone function that generates a specialized version of `mexp` at run time.

In fact, our system can automatically generate a Cyclone run-time specializer from an action-annotated program. Fig. 3 shows the Cyclone specializer produced from the action-annotated modular exponentiation function of Fig. 2. The function `mexp_gen` takes the two invariant arguments of the original `mexp` function and returns the function `mexp_sp`, a version of `mexp` specialized to those arguments. In the figure we have italicized code that will be evaluated when `mexp_gen` is called. Non-italicized code is template code that will be manipulated by `mexp_gen` to produce the specialized function. The template code will only be evaluated when the specialized function is itself called.

In our example, the `codegen` expression begins the code generation process by allocating a region in memory for the new function `mexp_sp`, and copying the first template into the region. This template includes the declarations of the function, its argument `base`, and local variables `s` and `t`, and also the initial assignments to `s` and `t`. Recall that this template code is *not* evaluated during the code generation process, but merely manipulated.

The `cut` statement marks the end of the template and introduces code (italicized) that will be evaluated during code generation: namely, the while loop. The while test and body, including the conditional statement, `splice` statements, and shift/assignment statement, will all be evaluated. After the while loop finishes, the template following the `cut` statement (containing `return(s)`) will be added to the code generation region.

```
int (int) mexp_gen(int exp, int mod)
{
  int u;

  u = exp;

  return codegen(
    int mexp_sp(int base) {
    int s, t;

    s = 1; t = base;

    cut
      while (u != 0) {
        if ((u & 1) != 0)
          splice  s = (s*t) % fill(mod);
        splice  t = (t*t) % fill(mod);
        u >>= 1;
      }
    return(s);
    });
}
```

Figure 3. A run-time specializer written in Cyclone

Evaluating a `splice` statement causes a template to be appended to the code generation region. In our example, each time the first `splice` statement is executed, an assignment to `s` is appended. Similarly, each time the second `splice` statement is executed, an assignment to `t` is appended. The effect of the while loop is thus to add some number of assignment statements to the code of `mexp_sp`; exactly how many, and which ones, is determined by the arguments of `mexp_gen`.

A `fill` expression can be used within a template, and it marks a *hole* in the template. When `fill(`e`)` is encountered in a template, e is evaluated at code generation time to a value, which is then used to fill the hole in the template. In our example, `fill` is used to insert the known modulus value into the assignment statements.

After code generation is complete, the newly generated function `mexp_sp` is returned as the result of `codegen`. It takes the one remaining argument of `mexp` to compute its result.

Cyclone programs can be evaluated symbolically to produce specialized source programs, like the one in Fig. 2; this is the basis of the formal operational semantics we give in section 2.4. In our implementation, however, we compile Cyclone source code

to object code, and we compile source templates into object templates. The Cyclone object code then manipulates object templates directly.

```
_mexp_gen:
    PUSH    0                   ; int u = 0
    MOV     EAX,[ESP+8]         ; int u = exp;
    MOV     [ESP+0],EAX
    CGSTART                     ; codegen(...
    CGDUMP  ECX,cdgn_beg$18     ; (dump 1st template)
    JMP     whiletest$21        ; while ...
whilebody$20:
    :                           ; if ((u & 1) != 0)
    CMP     EAX,ECX
    JE      ifend$24
    CGDUMP  ECX,splc_beg$25     ; (dump 2nd template)
    MOV     EAX,[ESP+12]        ; mod
    CGFILL  ECX,splc_beg$25,hole$29,EAX
ifend$24:
    CGDUMP  ECX,splc_beg$31     ; (dump 3rd template)
    MOV     EAX,[ESP+12]        ; mod
    CGFILL  ECX,splc_beg$26,hole$34,EAX
    MOV     ECX,1               ; u = u >> 1;
    MOV     EAX,[ESP+0]
    SAR     EAX,CL
    MOV     [ESP+0],EAX
whiletest$21:
    :
    CMP     EAX,ECX             ; ... (u != 0)
    JNE     whilebody$20
whileend$22:
    CGDUMP  ECX,cut_beg$33      ; (dump 4th template)
    CGEND   EAX
    ADD     ESP,4               ; (return spec. fun.)
    RETN
```

Figure 4. TAL/T code

Our object code, TAL/T, is an extension of TAL with instructions for manipulating object templates. Most of the TAL/T instructions are x86 machine instructions; the new template instructions are `CGSTART`, `CGDUMP`, `CGFILL`, `CGHOLE`, `TEMPLATE_START`, and `TEMPLATE_END`. For example, the Cyclone program in Fig. 3 is compiled into the TAL/T program shown in Figures 4 and 5. (We omitted some instructions to save space, and added source code fragments in comments to aid readability.)

```
TEMPLATE_START cdgn_beg$18,cdgn_beg$19
cdgn_beg$18                     ; (1st template)
    PUSH    0                   ; int s = 0;
    :
    MOV     EAX,[ESP+12]        ; t = base;
    MOV     [ESP+0],EAX
TEMPLATE_END cdgn_beg$19

TEMPLATE_START splc_beg$25,splc_end$26
splc_beg$25:                            ; (2nd template)
    CGHOLE  EAX,splc_beg$25,hole$29     ; fill(...)
    :
    MOV     [ESP+4],EAX                 ; s = mod(s*t,...));
TEMPLATE_END splc_end$26

TEMPLATE_START splc_beg$31,splc_end$32
splc_beg$31:                            ; (3rd template)
    CGHOLE  EAX,splc_beg$31,hole$34     ; fill(...)
    :
    MOV     [ESP+0],EAX                 ; t = mod(t*t,...));
TEMPLATE_END splc_end$32

TEMPLATE_START cut_beg$36,cut_end$37
cut_beg$36:                             ; (4th template)
    MOV     EAX,[ESP+4]                 ; return s;
    ADD     ESP,8
    RETN
TEMPLATE_END cut_end$37
```

Figure 5. TAL/T templates

Fig. 4 gives the code for `_mexp_gen`. It begins with x86 instructions for adding the local variable `u` to the stack and assigning it the value of the argument `exp`. Next, `CGSTART` is used to dynamically allocate a code generation region, and the first template is dumped (copied) into the region with the `CGDUMP` instruction. Next, the body of the loop is unrolled. Each Cyclone `splice` statement is compiled into a `CGDUMP` instruction, followed by instructions for computing hole values and a `CGFILL` instruction for filling in the hole. At the end of the `_mexp_gen` function, a final `CGDUMP` instruction outputs code for the last template.

Fig. 5 gives the code for each of the four templates. The first template allocates stack space for local variables `s` and `t` and assigns values to them. The second and third templates come from the statements contained within the Cyclone `splice` instructions,

i.e., the multiplications, mods, and assignments. The final template contains the code for `return(s)`. Each `CGHOLE` instruction introduces a placeholder inside a template, filled in during specialization as described above.

1.5. Summary

We designed a system for performing type safe run-time code generation. It has the following parts:

- C to action-annotated program translation
- Action-annotated program to Cyclone translation
- Cyclone language design
- Cyclone type system
- Cyclone verifier
- Cyclone to TAL/T compiler
- TAL/T language design
- TAL/T type system
- TAL/T verifier
- TAL/T to assembly translation
- Assembler/Linker

For some parts, we were able to reuse existing software. Specifically, we used Tempo for action-annotated program generation, Microsoft MASM for assembling, and Microsoft Visual C++ for linking. Other parts extend existing work. This was the case for the Cyclone language, type system, verifier, compiler, and the TAL/T language. Some components needed to be written from scratch, including the translation from an action-annotated program into a Cyclone program, and the definition of the new TAL/T instructions in terms of x86 instructions.

We've organized the rest of the paper as follows. In Section 2, we present the Cyclone language and its type system. In Section 3, we give a brief description of TAL/T; due to limited space we defer a full description to a later paper. We give implementation details and initial impressions about performance in Section 4. We discuss related work in Section 5, and future work in Section 6. Our final remarks are in Section 7.

2. Cyclone

2.1. Design decisions

Cyclone's `codegen`, `cut`, `splice`, and `fill` constructs were designed to express a template-based style of run-time code generation cleanly and concisely. We made some other design decisions based on Cyclone's relationship to the C programming language, and on implementation concerns.

First, because a run-time specializer is a function that returns a function as its result, we need higher order types in Cyclone. In C, higher order types can be written using pointer types, but Cyclone does not have pointers. Therefore, we introduce new notation for higher order types in Cyclone. For example:

```
int (float,int) f(int x) { ... }
```

This is a Cyclone function `f` that takes an int argument `x`, and returns a function taking a float and an int and returning an int. When `f` is declared and not defined, we use

```
int (int) (float,int) f;
```

Note that the type of the first argument appears to the left of the remaining arguments. This is consistent with the order the arguments would appear in C, using pointer types.

A second design decision concerns the extent to which we should support nested `codegen`'s. Consider the following example.

```
int (float) (int) f(int x) {
  return(codegen(
    int (int) g(float y) {
      return(codegen(
        int h(int z) {
        ... body of h ...
        }));
    }));
}
```

Here `f` is a function that generates a function `g` using `codegen` when called at run time. In turn, `g` will generate a function `h` each time it is called. Nested `codegen`'s are thus used to generate code that generates code; this is sometimes called incremental specialization.

We decided to support incremental specialization in Cyclone because it adds little complication to our type system and implementation, and it seems to have promising applications (Consel et al., 1993; Consel et al., 1999). Currently, we do not automatically generate nested `codegen`'s from C since the Tempo front end we are using is based on a two-level binding time analysis. However, a programmer can write nested

`codegen`'s explicitly in Cyclone, and it should be possible to use a multi-level binding time analysis (Glück and Jørgensen, 1996) to generate them automatically.

A final design decision concerns the extent to which Cyclone should support lexically scoped bindings. In the last example, the function `h` is nested inside of two other functions, `f` and `g`. In a language with true lexical scoping, the arguments and local variables of these outer functions would be visible within the inner function: `f`, `x`, `g`, and `y` could be used in the body of `h`.

We decided that we would *not* support full lexical scoping in Cyclone. Our scoping rule is that in the body of a function, only the function itself, its arguments and local variables, and top-level variables are visible. This is in keeping with C's character as a low-level, machine- and systems-oriented language: the operators in the language are close to those provided by the machine, and the cost of executing a program is not hidden by high-level abstractions. We felt that closures and lambda lifting, the standard techniques for supporting lexical scoping, would stray too far from this. If lexical scoping is desired, the programmer can introduce explicit closures. Or, lexical scoping can be achieved using the Cyclone features, for example, if `y` is needed in the body of `h`, it can be accessed using `fill(y)`.

2.2. Syntax

Now we formalize a core calculus of Cyclone. Full Cyclone has, in addition, structures, unions, arrays, void, break and continue, and for and do loops.

Figure 6 gives the grammar of Cyclone. We use c to range over constants and b to range over base types. There is an implicit signature assigning types to constants, so that we can speak of "the type of c." We use x to range over variable names and f to range over function names; function names are distinct from variable names. In particular, *values*, v, are constants or function names, but not variable names. We use z to range over both kinds of names.

Heaps, H, are sequences of definitions, for both variables and functions. The variables and functions defined in a heap must be distinct; in our semantics we implicitly use alpha conversion when necessary to achieve this. If a heap H contains definitions for variables and functions $z_1, \ldots, z_n$ with types $t_1, \ldots, t_n$, then $\widetilde{H}$ is defined to be the sequence $t_1\ z_1, \ldots, t_n\ z_n$. We sometimes consider a sequence $\widetilde{H} = t_1\ z_1, \ldots, t_n\ z_n$ to be a finite function from variables to types: $\widetilde{H}(z_i) = t_i$ if $1 \leq i \leq n$. This assumes that the z_i are distinct. We achieve this by alpha conversion when necessary, and by imposing some standard syntactic restrictions on Cyclone programs (the names of a function and its formal parameters must be distinct, and global names must be distinct).

We write D for a declaration sequence that contains declarations for variables only, and not for functions. Such sequences are used to declare the arguments and local variables of functions. We use $\cdot$ for the empty declaration sequence.

If $t = b\ m_1\ \cdots\ m_n$, then we define $t \bullet m$ to be the type $b\ m\ m_1\ \cdots\ m_n$. Thus $t \bullet m$ is the type of a function from m to t; in the arrow notation used in functional languages, this would be written $m \to t = m \to m_1 \to \cdots \to m_n \to b$. We use this new

Constants	c		
Base Types	b		
Variable Names	x		
Function Names	f		
Values	v	::=	$c \mid f$
Names	z	::=	$x \mid f$
Modifiers	m	::=	$(t_1, \ldots, t_n)$
Types	t	::=	$b\ m_1 \cdots m_n$
Heaps	H	::=	$\cdot$
		$\mid$	$H,\ t\ x = v$
		$\mid$	$H,\ t\ f(D)\ s$
Heap decls.	$\widetilde{H}$	::=	$t_1\ z_1, \ldots, t_n\ z_n$
Declarations	d	::=	$t\ x$
Decl. sequences	D	::=	$d_1, \ldots, d_n$
Statements	s	::=	$e;$
		$\mid$	$d = e;$
		$\mid$	$\{\ s_1 \cdots s_n\ \}$
		$\mid$	`if (`e`)` s_1 `else` s_2
		$\mid$	`while (`e`)` s
		$\mid$	`return` e`;`
		$\mid$	`splice` s
		$\mid$	`cut` s
Expressions	e	::=	x
		$\mid$	c
		$\mid$	f
		$\mid$	$e_0(e_1, \ldots, e_n)$
		$\mid$	x `=` e
		$\mid$	`codegen(`$t\ x(D)\ s$`)`
		$\mid$	`fill(`e`)`

Figure 6. The grammar of core Cyclone

notation instead of arrow notation because we want to be consistent with C's syntax for types (see section 2.1).

We introduce additional notation that is helpful in giving the type of a function from its definition. If $D = t_1\ x_1, \ldots, t_k\ x_n$, then $\overline{D}$ is defined to be the modifier $(t_1, \ldots, t_n)$, so that a function definition $t\ x(D)\ s$ declares x to be of type $t \bullet \overline{D}$.

For example, the function `mexp_gen()` from Section 1 has type

$$\begin{aligned}(\texttt{int (int)}) \bullet (\overline{\texttt{int exp, int mod}}) &= (\texttt{int (int)}) \bullet \texttt{(int, int)} \\ &= \texttt{int (int,int) (int)}.\end{aligned}$$

2.3. Typing rules

We define type environments E to support Cyclone's scoping rules:

$$\begin{aligned} E ::= \ & \texttt{start(}t\ f\texttt{(}D_{\text{params}}\texttt{);}\ D_{\text{local}}\texttt{);}\widetilde{H}_{\text{global}} \\ \mid \ & \texttt{frame(}t\ x\texttt{(}D_{\text{params}}\texttt{);}\ D_{\text{local}}\texttt{);}E \\ \mid \ & \texttt{hidden(}t\ x\texttt{(}D_{\text{params}}\texttt{);}\ D_{\text{local}}\texttt{);}E \end{aligned}$$

A type environment is a stack-like sequence of *frames*. Each frame gives the type of a function, as well as the types of its parameters and local variables. The stack-like structure of environments corresponds to the lexical nesting of functions in the Cyclone program. The rightmost frame of a Cyclone environment is a `start` frame, and it corresponds to an outermost, or top-level, function. All other frames correspond to nested functions introduced by `codegen`. In our typing rules, when a `codegen` expression is encountered, a new `frame` is introduced; it appears to the left of all the other frames in the environment, and it hides the declarations of frames to the right, which correspond to outer functions. When a `cut` statement or `fill` expression is encountered within the `codegen` expression, the `frame` temporarily becomes a `hidden` frame. The idea is that the declarations of a `hidden` frame are not accessible, even though it appears to the left of other frames in the environment; the declarations of a frame further to the right will become visible instead. The typing of a `suspend` statement reverses this, turning `hidden` back into frame.

The *visible declarations* $\text{vis}(E)$ of an environment E are defined in Figure 7. Informally, $\text{vis}(E)$ says that the declarations of the first (leftmost) non-hidden frame and the global declarations are visible, and all other declarations are not visible. Note that $\text{vis}(E)$ is a sequence of declarations, so we may write $\text{vis}(E)(z)$ for the type of z in E. Global declarations are visible in every environment; $\text{global}(E)$ is used to recursively examine an environment and return the global declarations.

Figure 7 also defines two other functions on environments: $\text{rtype}(E)$ is the return type for the function of the first non-hidden frame, and $E + d$ is the environment obtained by adding declaration d to the local declarations of the first non-hidden frame.

The example in Figure 8 demonstrates how Cyclone type environments are used. The figure shows a heap H, environments associated with different program points, and the names visible in each environment. The global declarations $\widetilde{H}$ include declarations for the variable `a` and function `f`, and are visible everywhere in the program. The initial environment used to type the function `f` includes the globals and a `start` frame for `f`. This is immediately extended with a declaration for the local variable `b`. Inside the

$$\text{vis}(E) = \begin{cases} D', D, t \bullet \overline{D}\ f, \widetilde{H} & \text{if } E = \texttt{start(}t\ f\texttt{(}D\texttt{);}\ D'\texttt{);}\widetilde{H} \\ D', D, t \bullet \overline{D}\ x, \text{global}(E') & \text{if } E = \texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \\ \text{vis}(E') & \text{if } E = \texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \end{cases}$$

$$\text{global}(E) = \begin{cases} \widetilde{H} & \text{if } E = \texttt{start(}t\ f\texttt{(}D\texttt{);}\ D'\texttt{);}\widetilde{H} \\ \text{global}(E') & \text{if } E = \texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \\ \text{global}(E') & \text{if } E = \texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \end{cases}$$

$$\text{rtype}(E) = \begin{cases} t & \text{if } E = \texttt{start(}t\ f\texttt{(}D\texttt{);}\ D'\texttt{);}\widetilde{H} \\ t & \text{if } E = \texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \\ \text{rtype}(E') & \text{if } E = \texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \end{cases}$$

$$E + d = \begin{cases} \texttt{start(}t\ f\texttt{(}D\texttt{);}\ d\texttt{,}D'\texttt{);}\widetilde{H} & \text{if } E = \texttt{start(}t\ f\texttt{(}D\texttt{);}\ D'\texttt{);}\widetilde{H} \\ \texttt{frame(}t\ x\texttt{(}D\texttt{);}\ d\texttt{,}D'\texttt{);}E' & \text{if } E = \texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \\ \texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' + d & \text{if } E = \texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \end{cases}$$

Figure 7. Cyclone environment functions

`codegen` expression, the definition of function `g` adds a new `frame` to the environment. The set of visible names now includes `g` and its formal parameter `c`, while the local variables of function `f` are no longer visible. The `cut` statement changes the `frame` environment to a `hidden` environment. This changes the set of visible variables, but retains the `frame` information so that it can be restored at the `splice` statement and after the `cut`. At the end of `g`'s definition, the `frame` is removed.

The typing rules of Cyclone programs are given in Figure 9. We assume that there is a distinguished function name, `main`, and for any $\widetilde{H}$ we define $\text{main}(\widetilde{H})$ to be the environment `start(int main(); `$\cdot$`);`$\widetilde{H}$. The function name `main` is not permitted anywhere else. The first typing rules states that a program, i.e., a heap H and the body of the main function s, is well typed if the heap is well typed with respect to its declarations, $\widetilde{H}$, and the body is well typed with respect to $\text{main}(\widetilde{H})$.

The next set of rules takes a list of declarations $\widetilde{H}_1$ and a heap H_2 and checks that each definition in the heap is well typed. A variable definition $t\ x = v$ is well typed if v has type t in the environment $\text{main}(\widetilde{H}_1)$. A function definition $t\ \ f\texttt{(}D'\texttt{)}\ \ s$ is well typed if the body s is well typed in the environment consisting of the global declarations and the frame for the outermost function f.

The typing rules shown in Figure 10 for statements and expressions are mostly standard. For example, a conditional statement is well typed if its test is an integer and both of its branches are well typed. A function call is well typed if the call has the correct number of actual parameters with the right types.

The typing rules for `codegen`, `cut`, `splice`, and `fill` are the interesting ones. `Codegen` starts the process of run time code generation. To type `codegen(`$t\ \ x\texttt{(}D\texttt{)}\ \ s$`)`

H	Environment (note, $\widetilde{H}$ = `int () f,int a`)	Visible names
int a = 0;		
int (int) f()	`start(int f(); ·);`$\widetilde{H}$	{a, f}
{		
int b = 0;	`start(int f(); int b);`$\widetilde{H}$	{a, f, b}
return codegen(		
`int g(int c) {`	`frame(int g(int c); ·);`	
	`start(int f(); int b);`$\widetilde{H}$	{a, f, g, c}
cut {	`hidden(int g(int c); ·);`	
	`start(int f(); int b);`$\widetilde{H}$	{a, f, b}
splice `c++;`	`frame(int g(int c); );`	
	`start(int f(); int b);`$\widetilde{H}$	{a, f, g, c}
}		
`return(c);`	`frame(int g(int c); ·);`	
	`start(int f(); int b);`$\widetilde{H}$	{a, f, g, c}
});		
}	`start(int f(); int b);`$\widetilde{H}$	{a, f, b}

Figure 8. Example of Cyclone environments

$\boxed{\vdash H, s}$ (H, s is a well-formed program)

$$\frac{\widetilde{H} \vdash H, \quad \text{main}(\widetilde{H}) \vdash s}{\vdash H, s}$$

$\boxed{\widetilde{H_1} \vdash H_2}$ (H_2 is a well-formed heap)

$$\widetilde{H_1} \vdash \cdot$$

$$\frac{\widetilde{H_1} \vdash H_2, \quad \text{main}(\widetilde{H_1}) \vdash v : t}{\widetilde{H_1} \vdash H_2,\ t\ x = v} \quad \widetilde{H_1} \vdash (x) = t$$

$$\frac{\widetilde{H_1} \vdash H_2, \quad \texttt{start(}t\ f\texttt{(}D\texttt{);}\ \cdot\ \texttt{);}\widetilde{H_1} \vdash s}{\widetilde{H_1} \vdash H_2,\ t\ f\texttt{(}D\texttt{)}\ s} \quad \widetilde{H_1}(f) = t \bullet \overline{D}$$

Figure 9. Typing rules of Cyclone programs

$\boxed{E \vdash s}$ (s is a well-formed statement)

$$\frac{E \vdash e : t}{E \vdash e;}$$

$$\frac{E \vdash e : t}{E \vdash t\ x = e;}$$

$$\frac{}{E \vdash \{\ \}}$$

$$\frac{E \vdash d = e; \quad E + d \vdash \{\ s_1\ \cdots\ s_n\ \}}{E \vdash \{\ d = e;\ s_1\ \cdots\ s_n\ \}}$$

$$\frac{E \vdash s_0,\ E \vdash \{\ s_1\ \cdots\ s_n\ \}}{E \vdash \{\ s_0\ s_1\ \cdots\ s_n\ \}} \quad s_0 \neq d = e;$$

$$\frac{E \vdash e : \texttt{int}, \quad E \vdash s_1, \quad E \vdash s_2}{E \vdash \texttt{if (}e\texttt{)}\ s_1\ \texttt{else}\ s_2}$$

$$\frac{E \vdash e : \texttt{int}, \quad E \vdash s}{E \vdash \texttt{while (}e\texttt{)}\ s}$$

$$\frac{E \vdash e : t}{E \vdash \texttt{return}\ e;} \quad \text{rtype}(E) = t$$

$$\frac{\texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{)};E \vdash s}{\texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{)};E \vdash \texttt{splice}\ s}$$

$$\frac{\texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{)};E \vdash s}{\texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{)};E \vdash \texttt{cut}\ s}$$

$\boxed{E \vdash e : t}$ (e has type t)

$E \vdash x : t$ if $\text{vis}(E)(x) = t$

$E \vdash c : t$ where t is the type of the constant c

$$\frac{E \vdash e_0 : t \bullet (t_1, \ \ldots, \ t_n),\ E \vdash e_1 : t_1,\ \ldots,\ E \vdash e_n : t_n}{E \vdash e_0(e_1, \ \ldots, e_n) : t}$$

$$\frac{E \vdash x : t,\ E \vdash e : t}{E \vdash x\ \texttt{=}\ e : t}$$

$$\frac{\texttt{frame(}t\ x\texttt{(}D\texttt{);}\ \cdot\texttt{)};E \vdash s}{E \vdash \texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ s\texttt{)} : t \bullet \widetilde{D}}$$

$$\frac{\texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{)};E \vdash e : t}{\texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{)};E \vdash \texttt{fill(}e\texttt{)} : t}$$

$E \vdash f : t$ if $\text{vis}(E)(f) = t$

Figure 10. Typing rules of Cyclone statements and expressions

in an environment E, we type the body s of the function in an environment `frame(`t x`(`D`); ·); `E. This makes the function x and its parameters D visible in the body s, while any enclosing function, parameters, and local variables will be hidden.

An expression `fill(`e`)` should only appear within a template. Our typing rule ensures this by looking at the environment: it must have the form `frame(`t x`(`D`); `D'`); `E. If so, the expression `fill(`e`)` is typed if e is typed in the environment `hidden(`t x`(`D`); `D'`); `E. That is, the function being generated with `codegen`, as well as its parameters and local variables, are hidden when computing the value that will fill the hole. This is necessary because the parameters and local variables will not become available until the function is called; they will not be available when the hole is filled.

The rules for `cut` and `splice` are similar. Like `fill`, `cut` can only be invoked within a template, so it changes `frame` to `hidden`. `Splice` is the dual of `cut`; it changes a frame `hidden` by `cut` back into a visible `frame`. Thus `splice` introduces a template, and `cut` interrupts a template.

2.4. Operational semantics of Cyclone

Here we define a small-step operational semantics for Cyclone.

To accommodate the semantics, we extend Cyclone expressions with a new form:

$$e \ ::= \ \cdots \mid \langle\!\langle s \rangle\!\rangle_t$$

The new expression form will be used in the semantics of function calls; $\langle\!\langle s \rangle\!\rangle_t$ should evaluate to the return value of a function call. The statement s comes from the body of the function, and it should finish computing by evaluating a `return` statement with a value of type t.

During the evaluation of $\langle\!\langle s \rangle\!\rangle_t$, s could make another function call. This would produce a nested expression $\langle\!\langle s' \rangle\!\rangle_{t'}$. These nested $\langle\!\langle \cdot \rangle\!\rangle_t$ expressions essentially form the call stack of the program. In order to type the call stack, we need to modify our environment to keep track of the type that should be returned for the topmost function call on the stack. Therefore, we extend our type environments E with a new kind of frame:

$$E \ ::= \ \cdots \mid \texttt{return}\ t;\ E$$

A `return` frame corresponds to an element of the call stack. The environment function rtype(E) is redefined in the obvious way:

$$\text{rtype}(E) \ = \ \begin{cases} t & \text{if } E = \texttt{start(}t\ f\texttt{(}D\texttt{);}\ D'\texttt{);}\widetilde{H} \\ t & \text{if } E = \texttt{frame(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \\ \text{rtype}(E') & \text{if } E = \texttt{hidden(}t\ x\texttt{(}D\texttt{);}\ D'\texttt{);}E' \\ t & \text{if } E = \texttt{return}\ t\texttt{;}E' \end{cases}$$

For a `return` frame, rtype returns the type stored in the frame. Otherwise, rtype is defined exactly as before.

Finally, we require a new typing rule, presented below:

$\boxed{E \vdash e : t}$ (e has type t)

$$\frac{\texttt{return}\ t; E \vdash s}{E \vdash \langle\!\langle s \rangle\!\rangle_t : t}$$

Each time a function call is evaluated, an expression $\langle\!\langle s \rangle\!\rangle_t$ is introduced. This typing rule says that the body s is well typed if the environment E, extended with a return frame containing type t, is well typed. When the `return` statement within the body is type-checked, the function rtype() will extract the type t from the environment and compare it with the type of the value returned.

2.4.1. *Redexes*
Our rewriting semantics is defined by identifying *redexes*, which are the syntactic subterms where rewriting will take place. We have two kinds of redexes: expression redexes and statement redexes. The redexes are summarized in the following table.

Expression redexes	Statement redexes
x	v;
$v_0(v_1,\ldots,v_n)$	$t\ x = v$;
$x = v$	{ {} $s_0 \cdots s_n$ }
$\langle\!\langle \{\} \rangle\!\rangle_t$	`if` (v) s_1 s_2
`codegen`($t\ x(D)\ \overline{s}$)	`while` (e) s
`codegen`($t\ x(D)$ TSS[`cut` {}])	`return` v;
`codegen`($t\ x(D)$ TSE[`fill`(v)])	`splice` s
`fill`(e)	`cut` s

2.4.2. *Evaluation contexts*
A program may contain many redexes ready to be evaluated. We use evaluation contexts to identify which of these redexes should be evaluated first. An evaluation context is a statement or an expressions with a hole, $[\cdot]$, in place of a subexpression or substatement. The hole will be used to specify the redex to evaluate. Our evaluation contexts are given in Figure 11.

We need about four times the number of contexts as usual, because Cyclone distinguishes statements from expressions, and because we have both normal evaluation and evaluation under templates.

The contexts SS[·], SE[·], ES[·], and EE[·] are used to specify normal evaluation, i.e., evaluation that does not generate code. The evaluation strategy is the standard,

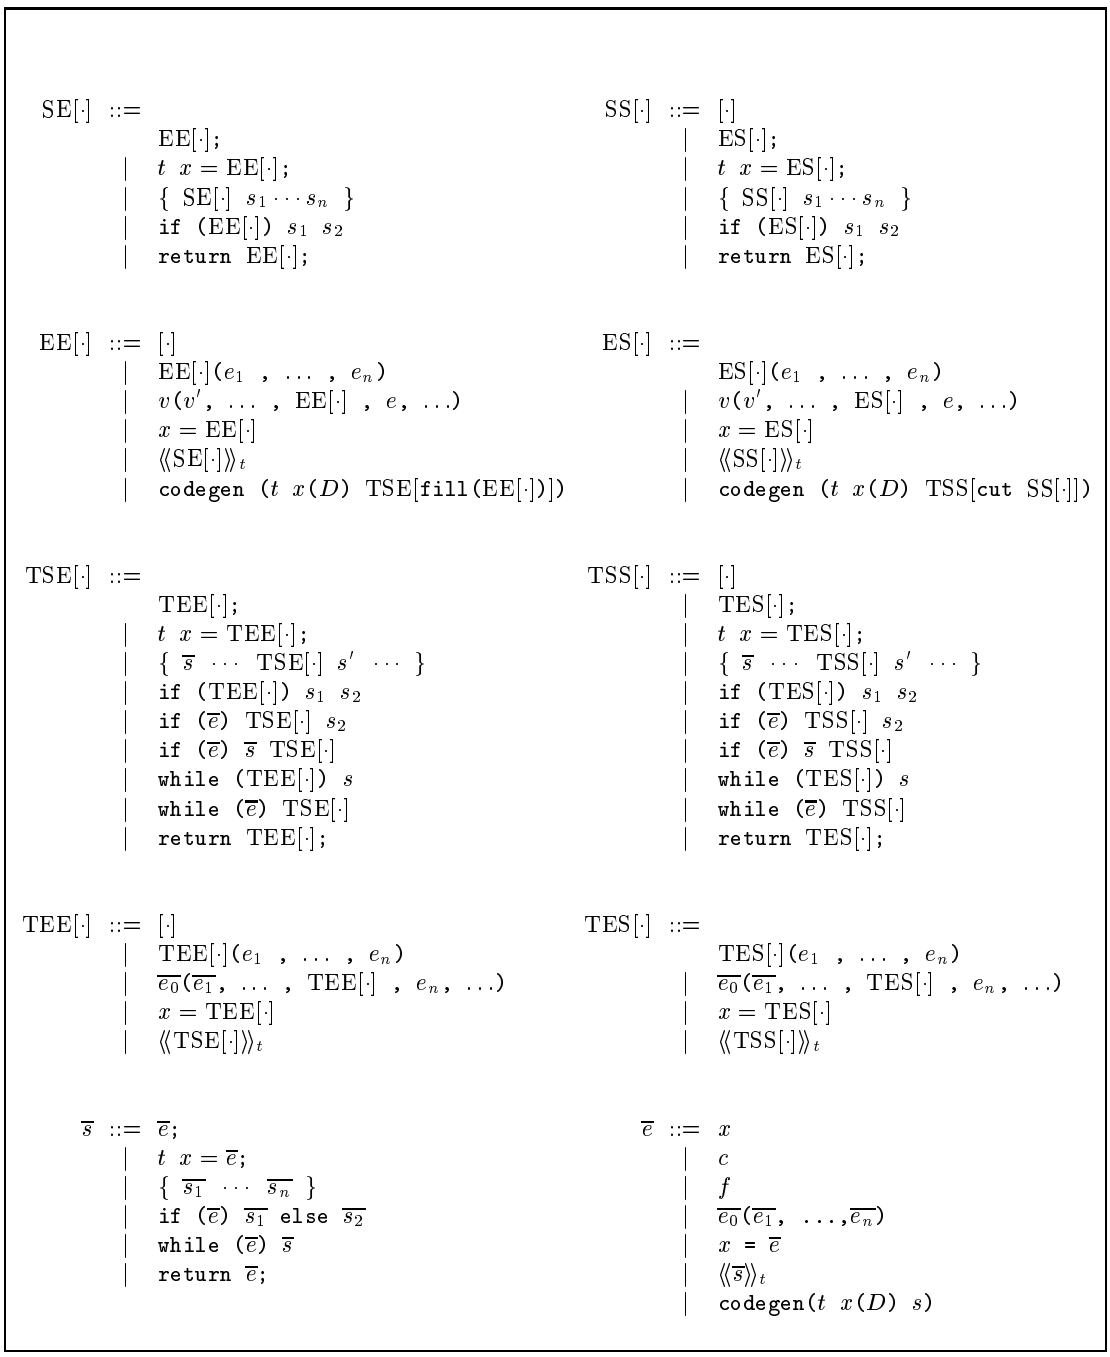

Figure 11. Evaluation contexts for Cyclone's operational semantics

left-to-right strategy. The contexts are named mnemonically, for example, a context SE[·] becomes a statement when its hole is filled by an expression, and a context SS[·] becomes a statement when its hole is filled by a statement.

The contexts TSS[·], TSE[·], TES[·], and TEE[·] are used to specify evaluation within templates. In templates, evaluation proceeds in a left-to-right fashion, but only under a `cut` statement or `fill` expression. The template contexts are defined using $\overline{s}$ and $\overline{e}$, which are statements and expressions that do not contain `cut`'s, `splice`'s, or `fill`'s to evaluate. `Splice` must be handled specially in template code: if it occurs in a template without a corresponding `cut`, it causes an error.

The following lemma shows that our evaluation contexts and redexes exactly capture all of the cases we will need to consider to define a complete, deterministic evaluation relation.

Lemma (Parsing). For any statement s, exactly one of the following cases holds.

1. There is a unique context SS[·] such that $s = \text{SS}[s']$, where s' is a statement redex;
2. there is a unique context SE[·] such that $s = \text{SE}[e]$, where e is an expression redex; or
3. $s = \{\}$.

Proof. The lemma is proved by induction, with a strengthened induction hypothesis. We prove the conjunction of the lemma and the following proposition:

For any expression e, exactly one of the following cases holds.

1. There is a unique context SS[·] such that $e = \text{ES}[s]$, where s is a statement redex;
2. there is a unique context SE[·] such that $e = \text{EE}[e']$, where e' is an expression redex; or
3. e is a value.

The proof is a simultaneous induction on s and e. It is useful to prove two similar lemmas separately:

Lemma A. For any statement s, exactly one of the following cases holds.

1. There is a unique context TSS[·] such that $s = \text{TSS}[\texttt{cut}\ s']$ or $s = \text{TSS}[\texttt{splice}\ s']$, where s' is some statement;
2. there is a unique context TSE[·] such that $s = \text{TSE}[\texttt{fill(}e\texttt{)}]$, where e is some expression; or
3. s has the form $\overline{s}$.

Lemma B. For any expression e, exactly one of the following cases holds.

1. There is a unique context TES[·] such that $e = \text{TES}[\texttt{cut}\ s]$ or $e = \text{TES}[\texttt{splice}\ s]$, where s is some statement;
2. there is a unique context TEE[·] such that $e = \text{TEE}[\texttt{fill}(e')]$, where e' is some expression; or
3. e has the form $\overline{e}$.

These lemmas are easy to prove by simultaneous induction.

Now we show some representative cases of the proof of the Parsing Lemma.

– $s = e;$. Then by induction exactly one of the following cases holds.

- e is a value v. Then $v;$ is a statement redex and $s = \text{SS}[v;]$ where $\text{SS}[\cdot] = [\cdot]$. Note, SS[·] and $v;$ are unique: if not, there is some other $\text{SS}'[\cdot]$ and statement redex s' such that
$$e; = \text{SS}'[s'].$$
By the grammar of SS[·], we must have ES[·] such that $\text{SS}'[\cdot] = \text{ES}[\cdot];$ and $e = \text{ES}[s']$. However, by induction, this is not possible: e is a value.
Furthermore, there is no SE[·] and expression redex e' such that
$$s = e; = \text{SE}[e'].$$
If there were, by the grammar of SE[·], there would be EE[·] such that $\text{SE}[\cdot] = \text{EE}[\cdot];$ and $e = \text{EE}[e']$. Once again, the induction hypothesis says that since e is a value, it cannot be parsed this way.
Finally, s is not $\{\}$.
- $e = \text{ES}[s']$ for a unique ES[·] and statement redex s'. Then $s = \text{SS}[s']$ where $\text{SS}[\cdot] = \text{ES}[\cdot];$.
Uniqueness is argued as above.
- $e = \text{EE}[e']$ for a unique EE[·] and expression redex e'. Then $s = \text{SS}[e']$ where $\text{SS}[\cdot] = \text{ES}[\cdot];$.
Uniqueness is argued as above.

– $e = \texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ s\texttt{)}$. Then by Lemma A exactly one of the following cases holds.

- s has the form $\overline{s}$. Then e is an expression redex; define $\text{EE}[\cdot] = [\cdot]$. Then $e = \text{EE}[e]$ where e is an expression redex.
Furthermore, EE[·] and e are unique: if not, by the grammar of EE[·] there would have to be TSE[·] and e' such that
$$e = \texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \text{TSE}[\texttt{fill(}e'\texttt{)}]\texttt{)}$$

and this is not possible by the definition of $\overline{s}$.

Similarly, there is no ES[·] and statement redex s' such that $e = \mathrm{ES}[s']$: there would have to be TSS[·] and s'' such that

$$e = \texttt{codegen}(t\ \ x(D)\ \ \mathrm{TSS}[\texttt{cut}\ \ s''])$$

and this is not possible by the definition of $\overline{s}$.

Finally, note that e is not a value.

- $s = \mathrm{TSS}[\texttt{cut}\ \ s']$. Then by induction exactly one of the following cases holds.
  - $s' = \{\}$. Then e is an expression redex; define EE[·] = [·]. Then $e = \mathrm{EE}[e]$ where e is an expression redex.
    Furthermore, EE[·] and e are unique: if not, by the grammar of EE[·] there must be TSE[·] and e' such that

    $$e = \texttt{codegen}(t\ \ x(D)\ \ \mathrm{TSE}[\texttt{fill}(e')]),$$

    that is, $s = \mathrm{TSE}[\texttt{fill}(e')]$. This is impossible by Lemma A.
    Similarly, there is no ES[·] and statement redex s' such that $e = \mathrm{ES}[s']$.
    Finally, note that e is not a value.
  - $s' = \mathrm{SS}[s'']$ for unique SS[·] and statement redex s''. Then $e = \mathrm{ES}'[s'']$ where $\mathrm{ES}'[\cdot] = \texttt{codegen}(t\ \ x(D)\ \ \mathrm{TSS}[\texttt{cut}\ \ \mathrm{SS}[\cdot]])$ and s'' is a statement redex.
    Uniqueness is shown as in the previous case.
  - $s' = \mathrm{SE}[e']$ for unique SE[·] and expression redex e'. Then $e = \mathrm{EE}[e']$ where $\mathrm{EE}[\cdot] = \texttt{codegen}(t\ \ x(D)\ \ \mathrm{TSS}[\texttt{cut}\ \ \mathrm{SE}[\cdot]])$ and e' is an expression redex.
    Uniqueness is shown as in the previous case.
- $s = \mathrm{TSS}[\texttt{fill}(e')]$. This case is proved just as the last case.
- $s = \mathrm{TSS}[\texttt{splice}\ \ s']$. Define $\mathrm{ES}[\cdot] = \texttt{codegen}(t\ \ x(D)\ \ \mathrm{TSS}[\cdot])$. Then $e = \mathrm{ES}[\texttt{splice}\ \ s']$, and $\texttt{splice}\ \ s'$ is a statement redex.
  Furthermore, ES[·] and $\texttt{splice}\ \ s'$ are unique: if not, by the grammar of ES[·] there must be TSS′[·], SS[·], and statement redex s'' such that

  $$e = \texttt{codegen}(t\ \ x(D)\ \ \mathrm{TSS}'[\texttt{cut}\ \ \mathrm{SS}[s'']]),$$

  that is, $s = \mathrm{TSS}'[\texttt{cut}\ \ \mathrm{SS}[s'']]$. This is impossible by Lemma A.
  Also, there is no EE[·] and expression redex e' such that $e = \mathrm{EE}[e']$. We cannot have EE[·] = [·] and $e = e'$ because e is not an expression redex. And we cannot have $\mathrm{EE}[\cdot] = \texttt{codegen}(t\ \ x(D)\ \ \mathrm{TSE}[\cdot])$, because then $\mathrm{TSE}[e'] = \mathrm{TSS}[\texttt{splice}\ \ s']$, contrary to Lemma A.
  Finally, e is not a value.

– $e =$ `fill(`e'`)`. Then e is an expression redex; define EE[$\cdot$] = [$\cdot$]. Then $e =$ EE[e] where e is an expression redex.

  Furthermore, EE[$\cdot$] and e are unique, since the grammar of EE[$\cdot$] does not define a context starting with `fill`.

  Also, there is no ES[$\cdot$] and statement redex s such that $e =$ ES[s], since the grammar of ES[$\cdot$] does define a context starting with `fill`.

  Finally, e is not a value.

– The remaining cases are proved similarly.

End proof.

2.4.3. *Context-dependent redexes*
The `return` and `splice` redexes are not self-contained; their rewrite rules are context-dependent. For example, a `splice` statement that does not appear within a `codegen` expression and `cut` statement should result in an error. The following definition makes these context dependencies precise.

Definition.

– A context SS[$\cdot$] is in *return mode* if its hole does not occur within $\langle\langle\cdot\rangle\rangle_t$.

– A context SS[$\cdot$] is in *splice mode* if its hole does not occur within $\langle\langle\cdot\rangle\rangle_t$ or `cut`.

Nested function calls produce nested $\langle\langle\cdot\rangle\rangle_t$ expressions. Return mode is used to match up a `return` statement with the innermost $\langle\langle\cdot\rangle\rangle_t$ expression in which it is located. If a context is in return mode, then a `return` statement in the hole of the context returns to a $\langle\langle\cdot\rangle\rangle_t$ expression outside of the context.

Similarly, splice mode is used to match up a `splice` statement with its corresponding `cut` statement. The condition that the hole of SS[$\cdot$] not occur within $\langle\langle\cdot\rangle\rangle_t$ is used to rule out cases like the following:

```
codegen(int g() {
        cut ⟪ splice {} ⟫t;
        return 3;
        })
```

Here, `splice` occurs within the $\langle\langle\cdot\rangle\rangle_t$ expression. The expression is not well-typed, because $\langle\langle\cdot\rangle\rangle_t$ causes a `return` frame to be pushed onto the environment, and `splice` requires a `hidden` frame on the environment. Intuitively, any code within $\langle\langle\cdot\rangle\rangle_t$ must have come from a function call, so this is an example of a `splice` across a function call, which is not allowed in Cyclone. Splice mode is defined to detect this case so that the expression evaluates to error in the semantics.

$$
\begin{array}{rcl}
H, \mathrm{SE}[x] & \rightarrow & \begin{cases} H, \mathrm{SE}[v] \\ \quad \text{if } H = H_1, t\ x = v, H_2 \\ \text{error} \\ \quad \text{otherwise} \end{cases} \\[2ex]
H, \mathrm{SE}[v_0\texttt{(}v_1\texttt{, ..., }v_n\texttt{)}] & \rightarrow & \begin{cases} H, \mathrm{SE}[\langle\!\langle \{d_1 = v_1\texttt{; } \cdots\texttt{; } d_n = v_n\texttt{; } s\} \rangle\!\rangle_t] \\ \quad \text{if } v_0 = f \text{ and } H = H_1, t\ f\texttt{(}d_1, \ldots, d_n\texttt{)}\ s, H_2 \\ \text{error} \\ \quad \text{otherwise} \end{cases} \\[2ex]
H, \mathrm{SE}[x = v] & \rightarrow & \begin{cases} H_1, t\ x = v, H_2, \mathrm{SE}[v] \\ \quad \text{if } H = H_1, t\ x = v', H_2 \\ \text{error} \\ \quad \text{otherwise} \end{cases} \\[2ex]
H, \mathrm{SE}[\langle\!\langle \{\} \rangle\!\rangle_t] & \rightarrow & \mathrm{error}_{\mathrm{return}} \\[2ex]
H, \mathrm{SS}[v\texttt{;}] & \rightarrow & H, \mathrm{SS}[\{\}] \\[2ex]
H, \mathrm{SS}[t\ x = v\texttt{;}] & \rightarrow & H, t\ x = v, \mathrm{SS}[\{\}] \\
 & & \text{where } x \text{ does not appear in } H \\[2ex]
H, \mathrm{SS}[\texttt{\{ \{\} }\ s_0 \cdots s_n\ \texttt{\}}] & \rightarrow & H, \mathrm{SS}[\texttt{\{ }\ s_0 \cdots s_n\ \texttt{\}}] \\[2ex]
H, \mathrm{SS}[\texttt{if (}v\texttt{) }\ s_1\ s_2] & \rightarrow & \begin{cases} H, \mathrm{SS}[s_2] \\ \quad \text{if } v = 0 \\ H, \mathrm{SS}[s_1] \\ \quad \text{otherwise} \end{cases} \\[2ex]
H, \mathrm{SS}[\texttt{while (}e\texttt{) }\ s] & \rightarrow & H, \mathrm{SS}[\texttt{if (}e\texttt{) \{ }\ s\ \texttt{while (}e\texttt{) }\ s\ \texttt{\} \{\}}] \\[2ex]
H, \mathrm{SS}[\texttt{return }\ v\texttt{;}] & \rightarrow & \begin{cases} H, \mathrm{SE}[v] \\ \quad \text{if } \mathrm{SS}[\cdot] = \mathrm{SE}[\langle\!\langle \mathrm{SS}'[\cdot] \rangle\!\rangle_t], \\ \quad \text{where } \mathrm{SS}'[\cdot] \text{ is in return mode} \\ \mathrm{error}_{\mathrm{return}} \\ \quad \text{if } \mathrm{SS}[\cdot] \text{ is in return mode} \end{cases}
\end{array}
$$

Figure 12. Rewrite rules for Cyclone's standard operators

2.4.4. *Rewrite rules*

Finally, the rewriting relation $\rightarrow$ is defined by the rewrite rules in Figures 12 and 13. The relation rewrites a heap and statement H, s. The standard rules are given in Figure 12. For example, a variable expression is rewritten in terms of its definition in the heap, and a `while` statement is rewritten by unrolling the loop one time. Notice that the rewrite rule for a function call replaces the call with a $\langle\!\langle \cdot \rangle\!\rangle_t$ expression containing the body of the function (as well as variable definitions corresponding to

$$H, \mathrm{SE}[\texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \overline{s}\texttt{)}] \rightarrow H, t\ f\texttt{(}D\texttt{)}\ \overline{s}[x := f], \mathrm{SE}[f] \quad \text{where } f \text{ is a fresh function name}$$

$$H, \mathrm{SE}[\texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \mathrm{TSS}[\texttt{cut}\ \{\}]\texttt{)}] \rightarrow \begin{cases} H, \mathrm{SE}[\texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \mathrm{TSS}[\{\}]\texttt{)}] \\ \quad \text{if TSS}[\cdot] \text{ is in return mode} \\ \text{error} \\ \quad \text{if TSS}[\cdot] \text{ is not in return mode} \end{cases}$$

$$H, \mathrm{SE}[\texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \mathrm{TSE}[\texttt{fill(}v\texttt{)}]\texttt{)}] \rightarrow \begin{cases} H, \mathrm{SE}[\texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \mathrm{TSE}[v]\texttt{)}] \\ \quad \text{if TSE}[\cdot] \text{ is in return mode} \\ \text{error} \\ \quad \text{if TSE}[\cdot] \text{ is not in return mode} \end{cases}$$

$$H, \mathrm{SS}[\texttt{cut}\ s] \rightarrow \text{error}$$

$$H, \mathrm{SE}[\texttt{fill(}e\texttt{)}] \rightarrow \text{error}$$

$$H, \mathrm{SS}[\texttt{splice}\ s] \rightarrow \begin{cases} H, \mathrm{SE}[\texttt{codegen(}t\ x\ \texttt{(}D\texttt{)}\ \mathrm{TSS}[s\ \texttt{cut}\ \mathrm{SS}'[\{\}]]\texttt{)}] \\ \quad \text{if SS}[\cdot] = \\ \qquad \mathrm{SE}[\texttt{codegen(}t\ x\ \texttt{(}D\texttt{)} \\ \qquad\qquad \mathrm{TSS}[\texttt{cut}\ \mathrm{SS}'[\texttt{splice}\ s]]\texttt{)}], \\ \quad \text{where SS}'[\cdot] \text{ is in splice mode} \\ \text{error} \\ \quad \text{otherwise} \end{cases}$$

Figure 13. Rewrite rules for Cyclone's code generation operators

the parameter bindings). The corresponding rule for the `return` statement returns its value v to the innermost enclosing $\langle\!\langle \cdot \rangle\!\rangle_t$ expression. If there is no $\langle\!\langle \cdot \rangle\!\rangle_t$ expression, then a function call is not in progress, and the program evaluates to error. Function calls are only defined for functions in the heap; an expression such as `c(3,4)` always evaluates to error—for simplicity, we did not support primitive functions in the semantics, though this would not be difficult to add.

Figure 13 gives the rules for Cyclone's code generation constructs. The behavior of `codegen` is specified by the first three rules. In the first rule, the body of the function being generated contains no `cut`'s or `fill`'s to evaluate, so code generation is complete, and the new function is added to the heap. The second rule shows the interaction between `codegen` and `cut`. As a statement `cut` s evaluates under `codegen`, the body s is rewritten until it becomes the empty statement, $\{\}$. Then `cut` $\{\}$ is simply replaced by $\{\}$, which becomes part of the generated function. However, if there is an expression $\langle\!\langle \cdot \rangle\!\rangle_t$ between the `codegen` and `cut`, an error is produced: we prohibit `cut`'s across function calls, just as we do `splice`'s.

The third rule gives the interaction between `codegen` and `fill`. As an expression `fill(`e`)` evaluates under `codegen`, e is rewritten until it becomes a value v. Then v

replaces the `fill` expression, becoming part of the generated function. Note that v is either a function f defined in the heap or a constant c, so it is in scope whether it appears inside or outside of template code.

Two additional rules handle another error case for `cut` and `fill`: if they are evaluated outside of template code, they result in an error. The grammars of SS[·] and SE[·] are defined intentionally so that their holes do not appear in template code.

The last rule, for `splice`, is the most complicated. Essentially it handles all of the cases that required two rules for `cut` and `fill`. The reason that we use one rule for `splice` has to do with the Parsing Lemma. Consider

$$s_0 = \texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \text{TSE}[\texttt{cut}\ \text{SS}[\texttt{splice}\ s]]\texttt{);}$$

Naively, we could make this a statement redex, with a rule something like the first rule for `cut`. To handle the error where `splice` is used outside of `codegen`, we would also have to make `splice` s a statement redex. But, this would break the Parsing Lemma: s_0 would have two parses. Namely,

$$\begin{aligned} s_0 &= \text{SS}_1[s_0], \quad \text{where } \text{SS}_1[\cdot] = [\cdot] \\ s_0 &= \text{SS}_2[\texttt{splice}\ s], \quad \text{where } \text{SS}_2[\cdot] = \texttt{codegen(}t\ x\texttt{(}D\texttt{)}\ \text{TSE}[\texttt{cut}\ \text{SS}[\cdot]]\texttt{);} \end{aligned}$$

We avoid this parsing ambiguity by having only one redex for `splice`, and we use splice mode to disambiguate the cases.

2.4.5. *Safety*

By the Parsing Lemma and a case analysis of the rewrite rules, we can see that a program H, s can rewrite in four disjoint ways:

- $H, s \to H', s'$ for some heap and statement H', s';
- $H, s \to \text{error}$;
- $H, s \to \text{error}_{\text{return}}$; or
- s is the statement $\{\}$, in which case evaluation is halted (there is no H', s' such that $H, s \to H', s'$, and H, s does not rewrite to error or $\text{error}_{\text{return}}$).

There are two kinds of error. Most errors are prevented by the type system; these are simply lumped together under "error." These errors include applying a function to the wrong number of arguments, trying to apply a non-function, using undefined variables, etc. There is also an error, "$\text{error}_{\text{return}}$," that is not prevented by our type system. It occurs when a function completes without executing a `return` statement, or there is a `return` without a corresponding function call. In our implementation, when this happens execution halts and the error is reported to the user. The full implementation has a few more errors of this sort, including array bound errors, stack overflow, and out of memory errors.

The safety theorem is stated as follows.

Theorem (Safety). If $\vdash H, s$, then there is no H', s' such that $H, s \rightarrow^* H', s' \rightarrow$ error.

The theorem is proved by showing that types are preserved by each rule rewriting $H, s \rightarrow H', s'$, and by showing that if a rule rewrites $H, s \rightarrow$ error, then H, s is not well-typed.

3. TAL/T

The output of the Cyclone compiler is a program in TAL/T, an extension of the Typed Assembly Language (TAL) of Morrisett et al. (Morrisett et al., 1998). In designing TAL/T, our primary concern was to retain the low-level, assembly language character of TAL. Most TAL instructions are x86 machine instructions, possibly annotated with type information. The exceptions are a few macros, such as `malloc`, that would be difficult to type in their expanded form; each macro expands to a short sequence of x86 instructions. Since each instruction is simple, the trusted components of the system—the typing rules, the verifier, and the macros—are also simple. This gives us a high degree of confidence in the correctness and safety of the system.

TAL already has instructions that are powerful enough to generate code at run time: `malloc` and move are sufficient. The problem with this approach is in the types. If we `malloc` a region for code, what is its type? Clearly, by the end of the code generation process, it should have the type of TAL code that can be jumped to. But at the start of code generation, when it is not safe to jump to, it must have a different type. Moreover, the type of the region should change as we move instructions into it. The TAL type system is not powerful enough to show that a sequence of `malloc` and move instructions results in a TAL program that can safely be jumped to.

Our solution, TAL/T, is an extension of TAL with some types and macros for manipulating templates. Since this paper focuses on Cyclone and the front end of system, we will only sketch the ideas of TAL/T here. Full details will appear in a subsequent paper.

In TAL, a procedure is just the label or address of a sequence of TAL instructions. A procedure is called by jumping to the label or address. The type of a procedure is a precondition that says that on entry, the x86 registers should contain values of particular types. For example, if a procedure is to return it will have a precondition saying that a return address should be accessible through the stack pointer when it is jumped to.

In TAL/T, a template is also the label of a sequence of instructions. Unlike a TAL procedure, however, a template is not meant to be jumped to. For example, it might need to be concatenated with another template to form a TAL procedure. Thus the type of a template includes a postcondition as well as a precondition. Our typing rules for the template instructions of TAL/T will ensure that before a template is dumped into a code generation region, its precondition matches the postcondition of the previous template dumped. Also, a template may have holes that need to be filled; the types of these holes are also given in the type of the template.

```
int f(int x) {
  return(codegen(
    int g(int y) {
      return y+1;
    })(x));
}

int h(int x)(int) {
  return(codegen(
    int k(int y) {
      return(fill(f(x)))
    }));
}
```

Figure 14. An example showing that two `codegen` expressions can be executing at once. When called, `h` starts generating `k`, but stops in the middle to call `f` which generates `g`.

The type of a code generation region is very similar to that of a template: it includes types for the holes that remain to be filled in the region, the precondition of the first template that was dumped, and the postcondition of the last template that was dumped. When all holes have been filled and a template with no postcondition is dumped, the region will have a type consisting of just a precondition, i.e., the type of a TAL procedure. At this point code generation is finished and the result can be jumped to.

Now we give a brief description of the new TAL/T macros. This is intended to be an informal description showing that each macro does not go beyond what is already in TAL—the macros are low level, and remain close to machine code.

The macros manipulate an implicit stack of code generation regions. Each region in the stack is used for a function being generated by a `codegen`. The stack is needed because it is possible to have two `codegen` expressions executing at once (for an example, see Figure 14).

- `cgstart` initiates run-time code generation by allocating a new code generation region. This new region is pushed onto the stack of code generation regions and becomes the "current" region. The `cgstart` macro is about as complicated as `malloc`.

- `cgdump` r`,` L copies the template at label L into the current code generation region. After execution, the register r points to the copy of the template, and can be used to fill holes in the copy. `Cgdump` is our most complicated macro: its core is a simple string-copy loop, but it must also check that the current code

```
int f() {
  return(codegen(
    int g(int i) {
      cut { return 4; }
      return(i+1);
    })(7));
}
```

Figure 15. An example that shows the need for `cgabort`. When called, the function `f` starts generating function `g` but aborts in the middle (it returns 4).

generation region has enough room for a copy of the template. If there is not enough room, `cgdump` allocates a new region twice the size of the old region, copies the contents of the old region plus the new template to the new region, and replaces the old region with the new on the region stack. This is the most complex TAL/T instruction, consisting of roughly twenty x86 instructions.

- `cghole` r, L_{template}, L_{hole} is a move instruction containing a hole. It should be used in a template with label L_{template}, and declares the hole L_{hole}.
- `cgfill` r_1, L_{template}, L_{hole}, r_2 fills the hole of a template; it is a simple move instruction. Register r_1 should point to a copy of the template at label L_{template}, which should have a hole with label L_{hole}. Register r_2 contains the value to put in the hole.
- `cgfillrel` fills the hole of a template with a pointer into a second template; like `cgfill` it expands to a simple move instruction. It is needed for jumps between templates.
- `cgabort` aborts a code generation; it pops the top region off the region stack. It is needed when the run-time code generation of a function stops in the middle, as in the example of Figure 15.
- `cgend` r finalizes the code generation process: the current region is popped off the region stack and put into register r. TAL can then jump to location r.

4. Implementation Status

We now describe some key aspects of our implementation. As previously mentioned, some components were written from scratch, while others were realized by modifying existing software.

4.1. Action-annotated program to Cyclone

We translate Tempo action-annotated programs into run-time specializers written in Cyclone. Using the Tempo front end, this lets us automatically generate a Cyclone program from a C program.

An action-annotated program distinguishes two kinds of code: normal code that will be executed during specialization, indicated in italics in Fig. 2; and template code that will emitted during specialization (non-italicized code). The annotated C program is translated into a Cyclone program that uses `codegen`, `cut`, `splice`, and `fill`. Since italicized constructs will be executed during code generation, they will occur outside `codegen`, or within a `cut` statement or a `fill` expression. Non-italicized constructs will be placed within a `codegen` expression or `splice` statement.

Our algorithm operates in two modes: "normal" mode translates constructs that should be executed at code generation time and "template" mode translates constructs that will be part of a template. The algorithm performs a recursive descent of the action-annotated abstract syntax, keeping track of which mode it is in. It starts off in "normal" mode and produces Cyclone code for the beginning of the run-time specializer: its arguments (the invariants) and any local variables and initial statements that are annotated with italics. When the first non-italic construct is encountered, a `codegen` expression is issued, putting the translation into "template" mode. The rest of the program is translated as follows.

An italic statement or expression must be translated in "normal" mode. Therefore, if the translation is in "template" mode, we insert `cut` (if we are processing a statement) or `fill` (if we are processing an expression) and switch into "normal" mode. Similarly, a non-italic statement should be translated in "template" mode; here we insert `splice` and switch modes if necessary. It isn't possible to encounter a non-italic expression within an italic expression.

Another step needs to be taken during this translation since specialization is *speculative*, i.e., both branches of a conditional statement can be optimistically specialized when the conditional test itself cannot be evaluated. This means that during specialization, the store needs to be saved prior to specializing one branch and restored before specializing the other branch. Therefore, we must introduce Cyclone statements to save and restore the store when translating such a conditional statement. This is the same solution used by Tempo (Consel et al., 1996).

4.2. Cyclone to TAL/T

To compile Cyclone to TAL/T, we extended an existing compiler, the Popcorn compiler of Morrisett et al. Popcorn is written in Caml, and it compiles a type safe dialect of C into TAL, a typed assembly language (Morrisett et al., 1998). Currently, Popcorn is a very simple, stack based compiler, though it is being extended with register allocation and more sophisticated optimizations.

The Popcorn compiler works by performing a traversal of the abstract syntax tree, emitting TAL code as it goes. It uses an environment data structure of the following form:

```
type env = { local_env: (id * int) list;
             args_on_stack: int }
```

The environment maintains the execution state of each function as it is compiled. The field `local_env` contains each variable identifier and its corresponding stack offset. Arguments are pushed onto the stack prior to entry to the function body; the field `args_on_stack` records the number of arguments, so they can be popped off the stack upon exiting the function.

To compile Cyclone we needed to extend the environment datatype: first, because Cyclone switches between generating normal code and template code, and second, because Cyclone has nested functions. Therefore, we use environments with the same structure as the environments used in Cyclone's typing rules:

```
type cyclone_env =
    Start of env * (id list)
  | Frame of env * cyclone_env
  | Hidden of env * cyclone_env
```

That is, environments are sequences of type frames for functions. A frame can either be start, normal, or hidden. Once we have this type of environment, visible bindings are defined as they are for vis(E) in Section 2.

An `Start` frame contains the local environment for a top-level function as well as the global identifiers. A `Frame` is used when compiling template code. A new `Frame` environment is created each time `codegen` is encountered. A `Frame` becomes `Hidden` to switch back to "normal" mode when a `cut` or `fill` is encountered.

Popcorn programs are compiled by traversing the abstract syntax tree and translating each Popcorn construct into the appropriate TAL instructions; the resulting sequence of TAL instructions is the compiled program. Compiling a Cyclone program, however, is more complicated; it is performed in two phases. The first phase alternates between generating normal and template TAL/T instructions and a second phase rearranges the instructions to put them in their proper place. In order for the instructions to be rearranged in the second phase, the first phase interleaves special *markers* with the TAL/T instructions:

```
type marker =
    M_TemplateBeg of id * id
  | M_TemplateEnd
  | M_Fill of id * exp
```

These markers are used to indicate which instructions are normal, which belong within a template, and which are used to fill holes. `M_TemplateBeg` takes two arguments, the beginning and ending label of a template, and is issued at the beginning

of a template (when `codegen` or `splice` is encountered, or `cut` ends). Similarly, `M_TemplateEnd` is issued at the end of a template (at the end of a `codegen` or `splice`, or the beginning of a `cut`). Note that between corresponding `M_TemplateBeg` and `M_TemplateEnd` markers, other templates may begin and end. Therefore, these markers can be nested. When a hole is encountered, a `M_Fill` marker is issued. The first argument of `M_Fill` is a label for the hole inside the template. The second argument is the Cyclone source code expression that should fill the hole.

The following example shows how the `cut` statement is compiled.

```
fun compile_stmt stmt cyclone_env =
 match stmt of
  Cut s ->
   match cyclone_env with
     Start _ -> raise Error
   | Frame(env,cyclone_env') ->
       cg_fill_holes (Hidden(env,cyclone_env'));
       compile_stmt s (Hidden(env,cyclone_env'));
       emit_mark(M_TemplateBeg(id_new "a", id_new "b"));
   | Hidden _ -> raise Error
```

The function `compile_stmt` takes a Cyclone statement and an environment, and emits TAL/T instructions as a side-effect. The first thing to notice is that a cut can only occur when the compiler is in "template" mode, in which case the environment begins with `Frame`. A `cut` statement ends a template. Therefore, `cg_fill_holes` is called, which emits a `M_TemplateEnd` marker, and emits TAL/T code to dump the template and fill its holes. Filling holes must be done using a "normal" environment, and therefore the first frame becomes `Hidden`. Next, `compile_stmt` is called recursively to compile the statement `s` within the cut. Since the statement `s` should also be compiled in normal mode, it also keeps the first frame `Hidden`. Finally, a `M_TemplateBeg` marker is emitted so that the compilation of any constructs following the cut will occur within a new template.

The second phase of the code generation uses the markers to rearrange the code. The TAL/T instructions issued between a `M_TemplateBeg` and a `M_TemplateEnd` marker are extracted and made into a template. The remaining, normal instructions are concatenated to make one function; hole filling instructions are inserted after the instruction which dumps the template that contains the hole. The example in Fig. 4 shows a TAL/T program after the second phase is completed; the normal code includes instructions to dump templates and fill holes, and is followed by the templates.

4.3. TAL/T to executable

TAL is translated into assembly code by expanding each TAL macro into a sequence of x86 instructions. Similarly, the new TAL/T macros expand into a sequence of x86 and TAL instructions. A description of each TAL/T macro is given in Section 3. The

resulting x86 assembly language program is assembled with Microsoft MASM and linked with the Microsoft Visual C++ linker.

4.4. Initial Impressions

We have implemented our system and have started testing it on programs to assess its strengths and weaknesses. Since there is currently a lot of interest in specializing interpreters, we decided to explore this type of application program. A state-of-the-art program specializer such as Tempo typically achieves a speedup between 2 and 20, depending on the interpreter and program interpreted. To see how our system compares, we took a bytecode interpreter available in the Tempo distribution and ran it through our system.

Preliminary results show that Cyclone achieves a speedup of over 3. This is encouraging, since this is roughly the speedup Tempo achieves on similar programs. A more precise comparison of the two systems still needs to be done, however. On the other hand, in our initial implementation, the cost of generating code is higher than in Tempo. One possible reason is that for safety, we allocate our code generation regions at run time, and perform bounds checks as we dump templates. The approach taken by Tempo, choosing a maximum buffer size at compile time and allocating a buffer of that size, is faster but not safe.

5. Related Work

Propagating types through all stages of a compiler, from the front end to the back end, has been shown to aid robust compiler construction: checking type safety after each stage quickly identifies compiler bugs (Tarditi, 1997; Tarditi et al., 1996). Additionally, Necula and Lee have shown that proving properties at the assembly language level is useful for safe execution of untrusted mobile code (Necula and Lee, 1998). So far, this approach has been taken only for statically generated code. Our system is intended to achieve these same goals for dynamically generated code.

Many of the ideas in Cyclone were derived from the Tempo run-time specializer (Consel and Noël, 1996; Hornof, 1997; Hornof and Noyé, 1997). We designed Cyclone and TAL/T with a template-based approach in mind, and we use the Tempo front end for automatic template identification. Another run-time specializer, DyC, shares some of the same features, such as static analyses and a template-like back end (Auslander et al., 1996; Grant et al., ; Grant et al., 1997). There are, however, some important differences between Cyclone and these systems. We have tried to make our compiler more robust than Tempo and DyC, by making Cyclone type safe, and by using types to verify the safety of compiled code. Like Tempo and DyC, Cyclone can automatically construct specializers, but in addition, Cyclone also gives the programmer explicit control over run-time code generation, via the `codegen`, `cut`, `splice`, and `fill` constructs. It is even possible for us to hand-tweak the specializers produced by the Tempo

front end with complete type safety. Like DyC, we can perform optimizations such as inter-template code motion, since we are writing our own compiler. Tempo's strategy of using an unmodified, existing compiler limits the optimizations that it can perform.

ML-box, Meta-ML, and 'C are all systems that add explicit code generation constructs to existing languages. ML-box and Meta-ML are type safe dialects of ML (Leone and Lee, 1994; Wickline et al., 1998; Taha and Sheard, 1997), while 'C is an unsafe dialect of C (Poletto et al., 1997; Poletto et al.,). All three systems have features for combining code fragments that go beyond what we provide in Cyclone. For example, in 'C it is possible to generate functions that have n arguments, where n is a value computed at run time; this is not possible in Cyclone, ML-box, or Meta-ML. On the other hand, 'C cannot generate a function that generates a function; this can be done in Cyclone (using nested `codegen`s), and also in ML-box and Meta-ML. An advantage we gain from not having sophisticated features for manipulating code fragments is simplicity: for example, the Cyclone type system does not need a new type for code fragments. The most fundamental difference, however, is that the overall system we present will provide type safety not only at the source level, but also at the object level. This makes our system more robust and makes it usable in a proof carrying code system.

6. Future Work

In this paper we presented a framework for performing safe and robust run-time code generation. Our compiler is based on a simple, stack-based, certifying compiler written by Morrisett et al. They are extending the compiler with register allocation and other standard optimizations, and we expect to merge Cyclone with their improvements.

We are interested in studying template-specific optimizations. For example, because templates appear explicitly in TAL/T, we plan to study inter-template optimizations, such as code motion between templates. Performing inter-template optimizations is more difficult in a system, like Tempo, based on an existing compiler that is not aware of templates.

We are also interested in analyses that could statically bound the size of the dynamic code generation region. This would let us allocate exactly the right amount of space when we begin generating code for a function, and would let us eliminate bounds checks during template dumps.

We would like to extend the front end of Tempo so that it takes Cyclone, and not just C, as input. This would mean extending the analyses of Tempo to handle Cyclone, which is an n-level language like ML-box. Additionally, we may implement the analysis of Glück and Jørgensen (Glück and Jørgensen, 1996) to produce n-level Cyclone from C or Cyclone.

7. Conclusion

We have designed a programming language and compiler that combines dynamic code generation with certified compilation. Our system, Cyclone, has the following features.

7.0.0.1. *Robust dynamic code generation* Existing dynamic code generation systems only prove safety at the source level. Our approach extends this to object code. This means that bugs in the compiler that produce unsafe run-time specializers can be caught at compile time, before the specializer itself is run. This is extremely helpful because of the complexity of the analyses and transformations involved in dynamic code generation.

7.0.0.2. *Flexibility and Safety* Cyclone produces dynamic code generators that exploit run-time invariants to produce optimized programs. The user interface is flexible, since the final executable can be generated from a C program, a Cyclone program, or TAL/T assembly code. Type safety is statically verified in all three cases.

7.0.0.3. *Safe execution of untrusted, dynamic, mobile code generators* This approach can be used to extend a proof-carrying code system to include dynamic code generation. Since verification occurs prior to run time, there is no run-time cost incurred for the safety guarantees. Sophisticated optimization techniques can be employed in the certifying compiler. The resulting system could produce mobile code that is not only safe, but potentially extremely fast.

Acknowledgements. We were able to implement Cyclone quickly because we worked from the existing Tempo and TAL implementations. We'd like to thank Charles Consel and the Tempo group, and Greg Morrisett and the TAL group, for making this possible. The paper was improved by feedback from Julia Lawall and the anonymous referees.

References

ACM: 1996, 'Proceedings of the ACM SIGPLAN '96 Conference on Programming Language Design and Implementation'. Philadelphia, Pennsylvania:. *SIGPLAN Notices* 31(5), May 1996.

ACM: 1997, 'Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation'. Amsterdam, The Netherlands:.

ACM: 1998, 'Proceedings of the ACM SIGPLAN'98 Conference on Programming Language Design and Implementation (PLDI)'. Montreal, Canada:. *SIGPLAN Notices* 33(5), May 1998.

Augustsson, L.: 1997, 'Partial Evaluation in Aircraft Crew Planning'. in (ACM, 1997), pp. 127–136.

Auslander, J., M. Philipose, C. Chambers, S. J. Eggers, and B. N. Bershad: 1996, 'Fast, Effective Dynamic Compilation'. in (ACM, 1996), pp. 149–159. *SIGPLAN Notices* 31(5), May 1996.

Consel, C., L. Hornof, F. Noël, J. Noyé, and E. Volanschi: 1996, 'A Uniform Approach for Compile-Time and Run-Time Specialization'. In: O. Danvy, R. Glück, and P. Thiemann (eds.): *Partial Evaluation, International Seminar, Dagstuhl Castle*. pp. 54–72.

Consel, C., R. Marlet, and P. Boinot: 1999, 'Efficient Incremental Run-Time Specialization for Free'. In: *Proceedings of the ACM SIGPLAN'99 Conference on Programming Language Design and Implementation (PLDI)*. Atlanta, Georgia.

Consel, C. and F. Noël: 1996, 'A General Approach for Run-Time Specialization and its Application to C'. In: *Conference Record of POPL '96: The* 23rd *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. St. Petersburg Beach, Florida, pp. 145–156.

Consel, C., C. Pu, and J. Walpole: 1993, 'Incremental Partial Evaluation: The Key to High Performance, Modularity and Portability in Operating Systems'. In: *Partial Evaluation and Semantics-Based Program Manipulation, Copenhagen, Denmark, June 1993*. pp. 44–46.

Glück, R. and J. Jørgensen: 1996, 'Fast Binding-Time Analysis for Multi-Level Specialization'. In: D. Bjørner, M. Broy, and I. V. Pottosin (eds.): *Perspectives of System Informatics. Proceedings*, Vol. 1181 of *Lecture Notes in Computer Science*. pp. 261–272.

Grant, B., M. Mock, M. Philipose, C. Chambers, and S. J. Eggers, 'DyC: An Expressive Annotation-Directed Dynamic Compiler for C'. *Theoretical Computer Science*. To appear.

Grant, B., M. Mock, M. Philipose, C. Chambers, and S. J. Eggers: 1997, 'Annotation-Directed Run-Time Specialization in C'. in (ACM, 1997), pp. 163–178.

Guenter, B., T. Knoblock, and E. Ruf: 1995, 'Specializing Shaders'. In: *Computer Graphics Proceedings*. pp. 343–350.

Hornof, L.: 1997, 'Static Analyses for the Effective Specialization of Realistic Applications'. Ph.D. thesis, Université de Rennes I.

Hornof, L. and J. Noyé: 1997, 'Accurate Binding-Time Analysis for Imperative Languages: Flow, Context, and Return Sensitivity'. in (ACM, 1997), pp. 63–73.

Keppel, D., S. Eggers, and R. Henry: 1991, 'A Case for Runtime Code Generation'. Technical Report 91-11-04, Department of Computer Science, University of Washington, Seattle, WA.

Leone, M. and P. Lee: 1994, 'Lightweight Run-Time Code Generation'. In: *ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation*. Orlando, Florida, pp. 97–106. University of Melbourne, Australia, Department of Computer Science, Technical Report 94/9.

Morrisett, G., D. Walker, K. Crary, and N. Glew: 1998, 'From System F to Typed Assembly Language'. In: *Conference Record of POPL '98: The 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. San Diego, California, pp. 85–97.

Muller, G., R. Marlet, E. Volanschi, C. Consel, C. Pu, and A. Goel: 1998, 'Fast, Optimized Sun RPC Using Automatic Program Specialization'. In: *Proceedings of the 18th International Conference on Distributed Computing Systems*. Amsterdam, The Netherlands.

Necula, G. C. and P. Lee: 1998, 'The Design and Implementation of a Certifying Compiler'. in (ACM, 1998), pp. 333–344. *SIGPLAN Notices* 33(5), May 1998.

Noël, F., L. Hornof, C. Consel, and J. Lawall: 1998, 'Automatic, Template-Based Run-time Specialization : Implementation and Experimental Study'. In: *International Conference on Computer Languages*. Chicago, IL, pp. 132–142. Also available as IRISA report PI-1065.

Poletto, M., D. R. Engler, and M. F. Kaashoek: 1997, 'tcc: A System for Fast, Flexible, and High-level Dynamic Code Generation'. In: *Proceedings of the ACM SIGPLAN'97 Conference on Programming Language Design and Implementation (PLDI)*. Las Vegas, Nevada, pp. 109–121. *SIGPLAN Notices* 32(5), May 1997.

Poletto, M., W. C. Hsieh, D. R. Engler, and M. F. Kaashoek, ''C and tcc: A Language and Compiler for Dynamic Code Generation'. *ACM Transactions on Programming Languages and Systems*. To appear.

Taha, W. and T. Sheard: 1997, 'Multi-Stage Programming with Explicit Annotations'. in (ACM, 1997), pp. 203–217.

Tarditi, D.: 1997, 'Design and Implementation of Code Optimizations for a Type-Directed Compiler for Standard ML'. Ph.D. thesis, Carnegie Mellon University. CMU–CS–97–108.

Tarditi, D., G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee: 1996, 'TIL: A Type-Directed Optimizing Compiler for ML'. in (ACM, 1996), pp. 181–192. *SIGPLAN Notices* 31(5), May 1996.

Wickline, P., P. Lee, and F. Pfenning: 1998, 'Run-time Code Generation and Modal-ML'. in (ACM, 1998), pp. 224–235. *SIGPLAN Notices* 33(5), May 1998.