

Decaf Language Reference

Mike Lam, James Madison University

Fall 2026

1 Introduction

Decaf is an imperative language similar to Java or C, but is greatly simplified compared to those languages. It will allow us to implement a compiler for the entire language in a single semester while still exploring all of the basic facets of a modern compiler.

This project is originally based on a project from course 6.035 at the Massachusetts Institute of Technology. However, significant changes have been made to the language. This document is the only authoritative reference to the version of Decaf used in CS 432 and CS 630 at James Madison University.

Here is an example program in Decaf:

```
// add.decaf - simple addition example

def int add(int x, int y)
{
    return x + y;          // add the two parameters
}

def int main()
{
    int a;
    a = 3;
    return add(a, 2);
}
```

Here is the output from running the reference compiler on the above program:

```
$ ./decaf add.decaf
RETURN VALUE = 5
```

Here are some important points of difference between Decaf and more general-purpose procedural languages like C and Java:

- Decaf is not object-oriented.
- Function declarations must begin with the `def` keyword.
- Variable declarations may not include initializations.
- Variable declarations must all occur together at the top of their scope.
- Many syntactic sugar operators (such as `+=` or `++`) are not included.
- Boolean expression evaluation is never short-circuited; both operands of `&&` and `||` are always evaluated.

2 Lexical Considerations

There are six token classes in Decaf:

Name	Description
ID	identifiers
KEY	keywords
DEC	decimal literals
HEX	hexadecimal literals
STR	string literals
SYM	symbols

Identifiers and keywords must begin with an alphabetic character and may contain alphanumeric characters and the underscore (`'_'`). A keyword can be thought of a special identifier that is “reserved” and cannot be used for actual variables or functions. Both keywords and identifiers are case-sensitive, and all keywords are lowercase. For example, `if` is a keyword, but `IF` is a variable name; `foo` and `Foo` are two different names referring to two distinct variables. The following are all of the supported keywords in Decaf:

```
def if else while return break continue int bool void true false
```

In addition, all of the following words are reserved; they are not currently keywords but may be used in future versions of the language and thus should not be used as identifier names:

```
for callout class interface extends implements new this string float double null
```

Keywords and identifiers must be separated by white space, or a token that is neither a keyword nor an identifier. For example, `iftrue` is a single identifier, not two distinct keywords. If a sequence begins with an alphabetic character, then it and the longest sequence of alphanumeric characters and underscores following it forms a token (either a keyword or an identifier).

Symbols can be split into three types: 1) grouping operators: parentheses, brackets, curly braces, assignment (equals), commas, and semicolons, 2) binary operators (BINOP) and 3) unary operators (UNOP). There are a variety of binary and unary operators in Decaf, including both arithmetic operators (e.g., plus and minus) and boolean operators (e.g., less-than and boolean OR). The following are all supported symbols in Decaf:

```
( { [ ] } ) , ; = + - * / % < > <= >= == != && || !
```

Integer literals must be representable in 64 bits (a longer literal is a parse error), and may be written in either decimal or hexadecimal form. The latter will always begin with the sequence `0x`, and may contain either lower- or upper-case letter digits. Integer literals may not be zero-padded; a leading zero ends the literal (so `007` is three tokens), and such input is caught during parsing, not lexing. Note that in Decaf integer literals are unsigned but integers themselves are signed. Thus, a negative “literal” in Decaf code is really an unsigned literal with the negation prefix operator (`-`) applied.

String literals must be enclosed in double quote marks and may contain four kinds of escaped characters: newlines (`'\n'`), tabs (`'\t'`), quotes (`'\"'`), and backslashes (`'\\'`). Unescaped newlines and carriage returns are not allowed in string literals. UTF-8 is the only supported character set.

Comments are started by `“//”` and are terminated by the end of the line. “Whitespace” may appear between any lexical tokens, and consists of one or more spaces, tabs, newlines, or carriage returns. Comments and whitespace have no effect on a Decaf program’s execution, and both should be discarded during the lexical analysis phase of compilation.

Decaf programs should be stored in files with the `“.decaf”` extension, and a single Decaf program may not span multiple files.

3 Syntax

The following is the Decaf grammar (in EBNF form):

$$\begin{aligned}
 \textit{Program} &\rightarrow (\textit{VarDecl} \mid \textit{FuncDecl})^* \\
 \textit{VarDecl} &\rightarrow \textit{Type} \textit{ID} ([' \texttt{DEC} '])? ';' \\
 \textit{Type} &\rightarrow \texttt{int} \mid \texttt{bool} \mid \texttt{void} \\
 \textit{FuncDecl} &\rightarrow \texttt{def} \textit{Type} \textit{ID} '(' \textit{Params}? ')' \textit{Block} \\
 \textit{Params} &\rightarrow \textit{Type} \textit{ID} (', \textit{Type} \textit{ID})^* \\
 \textit{Block} &\rightarrow \{' \textit{VarDecl}^* \textit{Stmt}^* \}' \\
 \textit{Stmt} &\rightarrow \textit{Loc} '=' \textit{Expr} ';' \\
 &\mid \textit{FuncCall} ';' \\
 &\mid \texttt{if} '(' \textit{Expr} ')' \textit{Block} (\texttt{else} \textit{Block})? \\
 &\mid \texttt{while} '(' \textit{Expr} ')' \textit{Block} \\
 &\mid \texttt{return} \textit{Expr}? ';' \\
 &\mid \texttt{break} ';' \\
 &\mid \texttt{continue} ';' \\
 \textit{Expr} &\rightarrow \textit{BinExpr} \\
 \textit{BinExpr} &\rightarrow \textit{BinExpr} \texttt{BINOP} \textit{BinExpr} \\
 &\mid \textit{UnaryExpr} \\
 \textit{UnaryExpr} &\rightarrow \texttt{UNOP} \textit{BaseExpr} \\
 &\mid \textit{BaseExpr} \\
 \textit{BaseExpr} &\rightarrow '(' \textit{Expr} ')' \\
 &\mid \textit{Loc} \\
 &\mid \textit{FuncCall} \\
 &\mid \textit{Lit} \\
 \textit{Loc} &\rightarrow \textit{ID} ([' \textit{Expr} '])? \\
 \textit{FuncCall} &\rightarrow \textit{ID} '(' \textit{Args}? ')' \\
 \textit{Args} &\rightarrow \textit{Expr} (', \textit{Expr})^* \\
 \textit{Lit} &\rightarrow \texttt{DEC} \mid \texttt{HEX} \mid \texttt{STR} \mid \texttt{true} \mid \texttt{false}
 \end{aligned}$$

The following is a key to some of the meta-notation used above:

Example	Description
<i>Foo</i>	non-terminal
foo	keyword
F00	token class (see Section 2)
'a'	symbol consisting of character 'a'
<i>x</i> ?	zero or one occurrences of <i>x</i> (i.e., optional)
<i>x</i> *	zero or more occurrences of <i>x</i>
()	used for grouping
	designates alternatives

Although this is not reflected in the reference grammar on the previous page, all binary operations in Decaf are left-associative. In addition, there are seven levels of operator precedence, shown in the table below from highest to lowest:

<i>Operators</i>	<i>Comments</i>
- !	unary negation (integer and boolean)
* / %	integer multiplication, division, and remainder
+ -	integer addition and subtraction
< <= >= >	boolean ordinal relation
== !=	boolean equality
&&	boolean conjunction (AND)
	boolean disjunction (OR)

As written above, this grammar is neither LL(1) nor LR(1); however, it can be transformed (using standard CFG transformations) into a grammar that is both. The process of performing these transformations is a useful exercise for the reader and a vital component of building a parser for the language.

Note that variables may be declared `void`; such declarations are allowed by the grammar but will be flagged during the static analysis phase.

Because the provided grammar is not yet suitable for directly building a parser, there are many possible correct LL(1) and LR(1) grammars, and thus many correct parse trees for any given Decaf program. Because of this, it is important to design a clearly-defined *abstract syntax tree* format that is independent of any specific parsing implementation. Here is a suggested list of Decaf AST objects:

```

Program                (contains lists of VarDecls and FuncDecls)
VarDecl
FuncDecl              (contains a Block)
Block                 (contains a list of VarDecls and a list of statements)

Statements:
  Assignment          (contains a Location and an expression)
  Conditional          (contains an expression and either one or two Blocks)
  WhileLoop           (contains an expression and a Block)
  Return              (contains an optional expression)
  Break
  Continue

Expressions:
  BinaryOp            (contains two child expressions)
  UnaryOp             (contains one child expression)
  Location
  FuncCall            (contains a list of expressions)
  Literal

```

The exact implementation of your AST is up to you, but it should be designed to be easy to traverse and manipulate.

4 Semantics

A Decaf program consists of a series of variable and function declarations.

4.1 Variables

All Decaf variables are statically typed, and there are only two valid basic types: signed integers that are 64 bits wide (`int`) and booleans (`bool`). Variables may be declared outside a function; such variables are considered “global” variables and are accessible from all functions. Variables declared inside a function or block are only visible inside the corresponding lexical scope (see Section 4.3 for details).

Decaf supports simple, fixed-length, one-dimensional, static arrays. Arrays must be declared global and their size is specified at their declaration time. Arrays are indexed from 0 to $N - 1$, where N is the size of the array (which must be greater than zero). The usual bracket notation is used to index arrays; an index must be present for all array accesses and the index expression must evaluate to an integer. Because arrays have a compile-time fixed size and cannot be declared as parameters (or local variables), there is no facility for querying the length of an array variable in Decaf. Arrays may be of either basic type (integers or booleans).

Assignment is only permitted for scalar values. Decaf uses value-copy semantics, and the assignment “`<loc> = <expr>`” copies the value resulting from the evaluation of `<expr>` into `<loc>`. If a variable is referenced before it is assigned, the result is undefined. For assignments to an array element, the index is evaluated before the value.

4.2 Functions

In Decaf, functions must be declared using the “`def`” keyword. Functions must either be “`void`” or have a single return value; they may take an unlimited number of parameters. Each parameter must have a formally-declared type. Arrays may NOT be passed as function parameters.

Every Decaf program must contain a function called `main` that takes no parameters and returns an `int`. Execution of the program begins at the `main` function. Functions may be called before their definition in the source code.

Decaf does not provide support for variadic parameters (e.g., “`printf`” in C).

4.3 Scope

Decaf uses simple static scoping rules. There are at least three valid scopes for any statement in a Decaf program: the global scope, the function declaration scope, and a function body block/local scope.

The global scope consists of names of variables and functions introduced at the top level of the source code. The function declaration scope consists of the formal parameters introduced in a function declaration. Additional block/local scopes exist within each block in the code, including the block that represents the body of a function declaration; these can also be part of `if` or `while` statements.

An identifier introduced in a function or block scope can *shadow* an identifier from the global scope. In this case, the identifier may only be referenced until the variable leaves scope. Similarly, identifiers introduced in local scopes can shadow identifiers in less deeply nested scopes, the function scope, and the global scope.

No identifier may be defined more than once in the same scope. Thus, variable and function names must all be distinct in the global scope, formal parameter names must be distinct for each function, and local variable names must be distinct in each local block scope.

4.4 Function Calls

Function invocation involves (1) passing argument values from the caller to the callee, (2) executing the body of the callee, and (3) returning to the caller, possibly with a result.

Arguments are passed by value and the semantics can be thought of in terms of assignment: the formal arguments of a function are like local variables of the function and are initialized by assignment to the values resulting from the evaluation of the argument expressions. The arguments are evaluated from left to right.

The body of the callee is then executed by executing the statements of its function body in sequence.

A function that is declared with a `void` return type can only be called as a statement, i.e., it cannot be used in an expression. Such a function returns control to the caller when `return` is called (no result expression is allowed) or when the textual end of the callee is reached.

A function that returns a result may be called as part of an expression, in which case the result of the call is the result of evaluating the expression in the `return` statement when this statement is reached. If control reaches the textual end of a function that returns a result, the result is undefined.

A function that returns a result may also be called as a statement. In this case, the result is ignored.

4.5 Control Structures

The `if` statement has the same semantics as in standard procedural programming languages. First, the conditional expression is evaluated. If the result is true, the “true” block is executed. Otherwise, the “else” block is executed, if it exists. Since Decaf requires that the “true” and “else” blocks be enclosed in braces, there is no ambiguity in matching an “else” block with its corresponding `if` statement.

The `while` statement also has the same semantics as in standard procedural languages. First, the guard expression is evaluated. If the result is true, the loop body is executed. After the loop body executes one iteration, the guard expression is re-evaluated. If the guard expression evaluates to false at any point at which it is evaluated, the loop terminates and execution resumes at the statement following the loop.

Decaf provides support for the standard `continue` and `break` statements, which immediately transfer control to the beginning or end of the loop (respectively). In the case of the former, the guard expression should be re-checked before the loop body is executed again. These statements may not occur outside of a `while` loop.

4.6 Expressions

Expressions follow the normal rules for evaluation. In the absence of other constraints, operators with the same precedence are evaluated from left to right, and operands are also evaluated left to right. Parentheses may be used to override normal precedence.

A location expression evaluates to the value contained by the location. The semantics of this differ depending on whether the location is an l-value or an r-value. Array operations are discussed in Section 4.1.

Method invocation expressions are discussed in Section 4.4.

Integer literals evaluate to their integer value. Boolean literals are evaluated to the integer values 1 (true) and 0 (false). String literals do not evaluate to anything; since variables cannot store strings, their only valid use is as an argument to predefined functions (e.g., the `print_str` function; see Section 4.7).

The arithmetic operators (arith op and unary minus) have their usual precedence and meaning, as do the relational operators (rel op). `%` computes the remainder of dividing its operands.

Relational operators are used to compare integer expressions and may not be used on boolean variables. The equality operators (`==` and `!=`) are valid for both `int` and `boolean` types, and can be used to compare any two expressions having the same type. The result of a relational operator or equality operator has type `boolean`. Boolean expression evaluation is never short-circuited; both operands of `&&` and `||` are always evaluated.

4.7 Input and Output

Decaf supports only very primitive I/O. There is currently no support for input; all input values must be hard-coded in the source. Output is supported using the following predefined library functions:

```
void print_str(string value)
void print_int(int value)
void print_bool(bool value)
```

The first two operate as expected. The last one (`print_bool`) prints the internal representation of the boolean (i.e., 1 for true and 0 for false). None of these functions add a newline at the end of the input; if newlines are desired they must be printed explicitly using the `print_str` function.

These functions are not actually implemented in Decaf (they are provided by the Decaf standard library) and thus must be manually added to the global symbol table during static analysis to ensure that calls to them are properly type-checked. Here are some examples of their use:

```
def int main()
{
    print_str("Hello!");           // result: "Hello!"
    print_int(2+3);                // result: "5"
    print_bool(5 < 9);             // result: "1"
    print_bool(5 < 2);             // result: "0"
    return 0;
}
```

5 Interpreter/Architecture Details

The reference compiler for Decaf includes an interpreter that simulates running a Decaf program compiled to ILOC on an architecture that can directly execute ILOC. This architecture has limitations that are relevant to Decaf programmers although they are not part of the Decaf language itself.

Features of the simulator architecture:

- 64-bit words (for both integers and booleans)
- 64-bit signed integer virtual registers (up to 2048 of them; see below)
- Four special-purpose 64-bit integer registers:
 1. IP: instruction pointer
 2. SP: stack pointer
 3. BP: base pointer
 4. RET: function return value
- 64Kb address space with stack and static data regions
- Fixed-size, read-only code region indexed by instruction

The most important limitation to keep in mind is the 64Kb address space, which limits the size of the stack and therefore the depth of recursive functions. Two further limits constrain program size rather than run-time behavior: the simulator holds at most 4096 instructions, and the virtual register bank has 2048 registers.

The virtual registers form one bank shared by all functions, numbered consecutively across the entire program, and are *not* preserved across a call, so a value needed afterward must sit in a register the callee never reuses or in memory. For non-recursive calls the global numbering supplies that property automatically, since caller and callee never share a register number. It does *not* hold for recursion: a recursive callee reuses its own register

numbers, so unallocated ILOC does not support recursive functions. Register allocation is what fixes this, by spilling live registers at call sites via the precall/postreturn sequence of Section 9.

In addition, the output functions listed in the previous section must be implemented using the `PRINT` instruction in ILOC, and as such must be handled as special cases during code generation.

6 Implementation-Defined Behavior

The following behaviors are not dictated by this specification. They are recorded here because the reference implementation enforces each choice consistently, and a program that depends on one is depending on an implementation detail rather than on Decaf.

- Signed overflow wraps. Arithmetic is two’s-complement modulo 2^{64} ; `9223372036854775807 + 1` evaluates to `-9223372036854775808`. No overflow is diagnosed.
- Integer division truncates toward zero, so `-17 / 5` is `-3`, not `-4`.
- The sign of `%` follows the dividend, like division: `-17 % 5` is `-2` and `17 % -5` is `2`.
- Division or modulo by zero is a run-time error. The program stops with a diagnostic; the result is not defined and execution does not continue.
- Integer literals in $[2^{63}, 2^{64})$ are reinterpreted as their (negative) two’s-complement value. A literal at or above 2^{64} is a compile-time error.
- Array indices are not bounds-checked. An out-of-range index reads or writes whatever occupies that address, and is a run-time error only when it falls outside the address space entirely.

Several limits also constrain what the implementation will accept. They are not part of the language, but a program that exceeds one is rejected by the reference compiler or ILOC simulator:

- Source files: 65535 bytes.
- Individual tokens (identifiers, literals, string literals): 255 characters.
- Expression nesting: 256 levels.
- Global variables: 8160 words in total (between the static data region’s start and its end).
- Virtual registers: 2048, and instructions: 4096 — see Section 5. Both scale with total program size.

Simulated memory is word-addressed: unaligned addresses are run-time errors.

7 Type Checking

Here is an incomplete list of basic type inference and type checking rules for Decaf. Each rule consists of a series of antecedents or premises (above the line) and a conclusion (below the line). In these rules, $\Gamma \vdash e : \tau$ means that “ e has type τ in environment Γ ”. If τ is omitted, the statement simply means that “ e is well-typed” (i.e., it has no type errors). Here, Γ (called the *typing context* or *type environment*) refers to the local symbol table (with the ability to look up symbols recursively through its parent table). The absence of Γ indicates that no symbol table is necessary to type-check e .

Expressions:

$$\begin{array}{c}
\text{TDec} \frac{}{\vdash \text{DEC} : \mathbf{int}} \quad \text{THex} \frac{}{\vdash \text{HEX} : \mathbf{int}} \quad \text{TStr} \frac{}{\vdash \text{STR} : \mathbf{str}} \\
\\
\text{TTrue} \frac{}{\vdash \mathbf{true} : \mathbf{bool}} \quad \text{TFalse} \frac{}{\vdash \mathbf{false} : \mathbf{bool}} \quad \text{TSubExpr} \frac{\Gamma \vdash e : \tau}{\Gamma \vdash \text{'(' e ' '}' : \tau} \\
\\
\text{TLoc} \frac{\text{ID} : \tau \in \Gamma}{\Gamma \vdash \text{ID} : \tau} \quad \text{TArrLoc} \frac{\text{ID} : \tau[] \in \Gamma \quad \Gamma \vdash e : \mathbf{int}}{\Gamma \vdash \text{ID} \text{'[' e ']' : \tau}} \\
\\
\text{TFuncCall} \frac{\text{ID} : (\tau_1, \tau_2, \dots, \tau_n) \rightarrow \tau_r \in \Gamma \quad \Gamma \vdash e_1 : \tau_1 \quad \Gamma \vdash e_2 : \tau_2 \quad \dots \quad \Gamma \vdash e_n : \tau_n}{\Gamma \vdash \text{ID} \text{'(' e_1, e_2, \dots, e_n '}' : \tau_r} \\
\\
\text{TNot} \frac{\Gamma \vdash e : \mathbf{bool}}{\Gamma \vdash \text{'!' e '}' : \mathbf{bool}} \quad \text{TNeg} \frac{\Gamma \vdash e : \mathbf{int}}{\Gamma \vdash \text{'-' e '}' : \mathbf{int}} \\
\\
\text{TAdd} \frac{\Gamma \vdash e_1 : \mathbf{int} \quad \Gamma \vdash e_2 : \mathbf{int}}{\Gamma \vdash e_1 \text{'+' e_2 '}' : \mathbf{int}} \quad (\text{similar for TSub (-), TMul (*), TDiv (/) and TMod (\%)}) \\
\\
\text{TLe} \frac{\Gamma \vdash e_1 : \mathbf{int} \quad \Gamma \vdash e_2 : \mathbf{int}}{\Gamma \vdash e_1 \text{'<' e_2 '}' : \mathbf{bool}} \quad (\text{similar for TLeq (<=), TGe (>), and TGeq (>=)}) \\
\\
\text{TEq} \frac{\Gamma \vdash e_1 : \tau \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 \text{'==' e_2 '}' : \mathbf{bool}} \quad (\text{similar for TNeq (!=)}) \\
\\
\text{TAnd} \frac{\Gamma \vdash e_1 : \mathbf{bool} \quad \Gamma \vdash e_2 : \mathbf{bool}}{\Gamma \vdash e_1 \text{'\&\&' e_2 '}' : \mathbf{bool}} \quad (\text{similar for TOr (||)})
\end{array}$$

Statements:

$$\begin{array}{c}
\text{TIf} \frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma \vdash b}{\Gamma \vdash \text{'if ' (' e ' '}' b}} \quad \text{TIfElse} \frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma \vdash b_1 \quad \Gamma \vdash b_2}{\Gamma \vdash \text{'if ' (' e ' '}' b_1 \text{' else ' } b_2}} \quad \text{TWhile} \frac{\Gamma \vdash e : \mathbf{bool} \quad \Gamma \vdash b}{\Gamma \vdash \text{'while ' (' e ' '}' b}} \\
\\
\text{TAssign} \frac{\Gamma \vdash \text{ID} : \tau \quad \Gamma \vdash e : \tau}{\Gamma \vdash \text{ID} \text{'=' e ';'}} \quad \text{TArrAssign} \frac{\Gamma \vdash \text{ID}[0] : \tau \quad \Gamma \vdash e_i : \mathbf{int} \quad \Gamma \vdash e_v : \tau}{\Gamma \vdash \text{ID} \text{'[' e_i ']' '=' e_v ';'}} \\
\\
\text{TBlock} \frac{\Gamma' \vdash s_1 \quad \Gamma' \vdash s_2 \quad \dots \quad \Gamma' \vdash s_n}{\Gamma \vdash \text{'{' Vars, s_1, s_2, \dots, s_n '}'}} \quad \text{where } \Gamma' = \Gamma \cup \{v_i : \tau_i \in \text{Vars}\}
\end{array}$$

8 Operational Semantics

Here is an incomplete list of rules for generating ILOC code from Decaf code. In this sense they also encode *operational semantics* for the Decaf language. Each rule consists of a series of antecedents (above the line) and a conclusion (below the line). In these rules, $s \rightarrow C$ means “statement s generates code C .” For expressions, $e \rightarrow \langle C, r \rangle$ means that “expression e generates code C with the result stored in virtual register r .” The semicolon operator (;) is used to concatenate code.

Expressions:

$$\begin{array}{c}
\text{STrue} \frac{}{\text{true} \rightarrow \langle \text{loadI } 1 \Rightarrow r, r \rangle} \quad \text{SFalse} \frac{}{\text{false} \rightarrow \langle \text{loadI } 0 \Rightarrow r, r \rangle} \\
\\
\text{SInt} \frac{}{\text{INT} \rightarrow \langle \text{loadI INT} \Rightarrow r, r \rangle} \quad \text{SLoc} \frac{r_b = \mathbf{base}(\text{ID}) \quad x_o = \mathbf{offset}(\text{ID})}{\text{ID} \rightarrow \langle \text{loadAI } [r_b + x_o] \Rightarrow r, r \rangle} \\
\\
\text{SArrLoc} \frac{e \rightarrow \langle C_e, r_e \rangle \quad x_s = \mathbf{size}(\text{ID}) \quad r_b = \mathbf{base}(\text{ID})}{\text{ID}[e] \rightarrow \langle C_e; \text{multI } r_e, x_s \Rightarrow r_o; \text{loadAO } [r_b + r_o] \Rightarrow r, r \rangle} \\
\\
\text{SFuncCall} \frac{e_1 \rightarrow \langle C_1, r_1 \rangle \quad e_2 \rightarrow \langle C_2, r_2 \rangle \quad \dots \quad e_n \rightarrow \langle C_n, r_n \rangle}{\text{ID}(e_1, e_2, \dots, e_n) \rightarrow \langle C_1; C_2; \dots C_n; (\text{push } r_n; \dots \text{push } r_2; \text{push } r_1; \text{call ID}; \text{addI SP, } r_n \Rightarrow \text{SP; i2i RET} \Rightarrow r_r), r_r \rangle} \\
\\
\text{SNot} \frac{e \rightarrow \langle C, r_1 \rangle}{\text{'!'} e \rightarrow \langle C; \text{not } r_1 \Rightarrow r_2, r_2 \rangle} \quad \text{SAdd} \frac{e_1 \rightarrow \langle C_1, r_1 \rangle \quad e_2 \rightarrow \langle C_2, r_2 \rangle}{e_1 \text{'+' } e_2 \rightarrow \langle C_1; C_2; \text{add } r_1, r_2 \Rightarrow r_3, r_3 \rangle} \\
\\
\text{(similar for SNeg (-))} \quad \text{(similar for SSub (-), SMul (*), SDiv (/), SAnd (\&\&), and SOr (||))} \\
\\
\text{(Note: SMod (\%) is purposefully omitted as an exercise for the reader)} \\
\\
\text{SLe} \frac{e_1 \rightarrow \langle C_1, r_1 \rangle \quad e_2 \rightarrow \langle C_2, r_2 \rangle}{e_1 \text{'<'} e_2 \rightarrow \langle C_1; C_2; \text{cmp_LT } r_1, r_2 \Rightarrow r_3, r_3 \rangle} \\
\\
\text{(similar for SLeq (<=), SGe (>), SGeq (>=), SNeq (!=), and SEq (==))}
\end{array}$$

Statements:

$$\begin{array}{c}
\text{SIf} \frac{e \rightarrow \langle C_e, r_e \rangle \quad b \rightarrow C_b}{\text{if '(' } e \text{')' } b \rightarrow C_e; \text{cbr } r_e \Rightarrow l_1, l_2; l_1::; C_b; l_2:} \\
\\
\text{SIfElse} \frac{e \rightarrow \langle C_e, r_e \rangle \quad b_1 \rightarrow C_1 \quad b_2 \rightarrow C_2}{\text{if '(' } e \text{')' } b_1 \text{ else } b_2 \rightarrow C_e; \text{cbr } r_e \Rightarrow l_1, l_2; l_1::; C_1; \text{jump } l_3; l_2::; C_2; l_3:} \\
\\
\text{SWhile} \frac{e \rightarrow \langle C_e, r_e \rangle \quad b \rightarrow C_b}{\text{while '(' } e \text{')' } b \rightarrow l_1::; C_e; \text{cbr } r_e \Rightarrow l_2, l_3; l_2::; C_b; \text{jump } l_1; l_3:} \\
\\
\text{SAssign} \frac{e \rightarrow \langle C_e, r_e \rangle \quad r_b = \mathbf{base}(\text{ID}) \quad x_o = \mathbf{offset}(\text{ID})}{\text{ID} = e \rightarrow C_e; \text{storeAI } r_e \Rightarrow [r_b + x_o]} \\
\\
\text{SArrAssign} \frac{e_i \rightarrow \langle C_i, r_i \rangle \quad e_e \rightarrow \langle C_e, r_e \rangle \quad x_s = \mathbf{size}(\text{ID}) \quad r_b = \mathbf{base}(\text{ID})}{\text{ID}[e_i] = e_e \rightarrow C_i; C_e; \text{multI } r_i, x_s \Rightarrow r_o; \text{storeAO } r_e \Rightarrow [r_b + r_o]} \\
\\
\text{SBlock} \frac{s_1 \rightarrow C_1 \quad s_2 \rightarrow C_2 \quad \dots \quad s_n \rightarrow C_n}{\text{'{' } s_1, s_2, \dots, s_n \text{'}} \rightarrow C_1; C_2; \dots C_n}
\end{array}$$

9 Application Binary Interface

The following calling conventions should be respected by all ILOC code generated from Decaf programs. In these conventions the *caller* is the procedure making the call and the *callee* is the procedure that is being called. Further, the *prologue* and *epilogue* are segments of code located at the beginning and end of every procedure, respectively. Finally, the *precall* and *postreturn* sequences are segments of code located immediately preceding and immediately following the actual `call` instruction.

- In the precall, the caller must perform the following tasks in order:
 1. Calculate the values of all actual parameters in order.
 2. Push all actual parameter values on the stack **in reverse order**.
 3. Save any live general-purpose registers (if applicable).
- In the prologue, the callee must perform the following tasks in order:
 1. Save the caller's base pointer by pushing it on the stack.
 2. Set the callee's base pointer using the stack pointer.
 3. Allocate space for local variables and/or spilled registers by decrementing the stack pointer. This operation must be emitted even if the number of bytes required is zero.
- The callee must access its parameters using positive static offsets from the base pointer.
- The callee must access its local variables using negative static offsets from the base pointer.
- In the epilogue, the callee must perform the following tasks in order:
 1. Save the return value in the `RET` register (if applicable).
 2. Restore the stack pointer from the base pointer.
 3. Restore the caller's base pointer by popping it from the stack.
- In the postreturn, the caller must de-allocate any space allocated for parameters that were pushed onto the stack.

Here is an example from the program in Section 1, showing the ILOC instructions that are executed in order:

```
...
// PRECALL
loadAI [BP-8] => r5           // calculate parameters (in order)
loadI 2 => r6
push r6                      // push parameters (reverse order)
push r5

call add                     // push return address and jump

// PROLOGUE
push BP                      // save caller base pointer
i2i SP => BP                  // save stack pointer as callee base pointer
addI SP, 0 => SP              // allocate space for local vars (0 bytes)

...

loadAI [BP+16] => r0          // parameters (positive offset from BP)
loadAI [BP+24] => r1

...

i2i r2 => RET                 // save return value

// EPILOGUE
i2i BP => SP                  // restore stack pointer
pop BP                       // restore caller's base pointer
return                       // pop return address and jump

// POSTRETURN
addI SP, 16 => SP             // deallocate parameters
...
```