

Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science

6.110, Spring 2024

Handout — Decaf Language

Monday, Feb 5

The project for the course is to write a compiler for a language called Decaf. Decaf is a simple imperative language similar to C or Pascal.

This handout describes the language syntax (i.e., the set of legal Decaf programs) and the language semantics (ascribing a meaning to each legal program). You may make extensions that expand the set of legal programs, provided that you include a clear description of the new programs that are accepted, and of their (reasonable) semantics; these extensions may not include syntactically standard Decaf programs explicitly ruled out by the semantic rules described in Section 4. A reasonable extension could include allowing new datatypes, or new language constructs. You may not make any extensions that change the semantics of legal programs or that rule out any legal programs.

1 Lexical Considerations

All Decaf keywords are lowercase. Keywords and identifiers are case-sensitive. For example, `if` is a keyword, but `IF` is an identifier; `foo` and `Foo` are two different identifiers referring to two distinct variables or methods.

The reserved words are:

```
bool  break  const  import  continue  else  false  for  while  if  int  return
len   true   void
```

Comments are started by `//` and are terminated by the end of the line. Block comments are started by `/*` and are terminated by a matching `*/`. You do not have to handle nested block comments (e.g., `/* /* */ */`).

White space may appear between any lexical tokens (but not within a single token). White space is defined as one or more spaces, tabs, line-break characters (carriage return, line feed, form feed), and comments.

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

String literals are composed of `<char>`s enclosed in double quotes. A character literal consists of a `<char>` enclosed in single quotes.

If a sequence begins with `0x`, then these first two characters and the longest following sequence of characters drawn from `[0-9a-fA-F]` form a hexadecimal integer literal. If a sequence begins with a decimal digit (but not `0x`), then the longest following sequence of decimal digits forms a decimal integer literal. A long sequence of digits (e.g., `123456789123456789123`) is scanned as a single token.

A `<char>` is any printable ASCII character (ASCII values between decimal value 32 and 126 inclusive) other than quote (`"`), single quote (`'`), or backslash (`\`), plus the 2-character sequences `\"` to

denote quote, `\'` to denote single quote, `\\` to denote backslash, `\t` to denote a literal tab, or `\n` to denote newline.

2 Reference Grammar

Notation	Meaning
$\langle \text{foo} \rangle$	means foo is a nonterminal.
foo	(in bold font) means that foo is a terminal i.e., a token or a part of a token.
$[x]$	means zero or one occurrence of x , i.e., x is optional; note that brackets in quotes <code>'[']'</code> are terminals.
x^*	means zero or more occurrences of x .
$x^+,$	a comma-separated list of one or more x 's. note that there is no comma following the last of x .
$\{ \}$	large braces are used for grouping; note that braces in quotes <code>'{ '}'</code> are terminals.
$ $	separates alternatives.

$$\begin{aligned}
\langle \text{program} \rangle &\rightarrow \langle \text{import_decl} \rangle^* \langle \text{field_decl} \rangle^* \langle \text{method_decl} \rangle^* \\
\langle \text{import_decl} \rangle &\rightarrow \textbf{import} \langle \text{id} \rangle \textbf{;} \\
\langle \text{field_decl} \rangle &\rightarrow [\textbf{const}] \langle \text{type} \rangle \{ \{ \langle \text{id} \rangle \mid \langle \text{id} \rangle \textbf{'[']' } [\langle \text{int_literal} \rangle] \textbf{' '}] \} [\textbf{=} \langle \text{initializer} \rangle] \}^+, \textbf{;} \\
\langle \text{method_decl} \rangle &\rightarrow \{ \langle \text{type} \rangle \mid \textbf{void} \} \langle \text{id} \rangle ([\{ \langle \text{type} \rangle \langle \text{id} \rangle \}^+,]) \langle \text{block} \rangle \\
\langle \text{block} \rangle &\rightarrow \textbf{'{' } } \langle \text{field_decl} \rangle^* \langle \text{statement} \rangle^* \textbf{' ' } \\
\langle \text{type} \rangle &\rightarrow \textbf{int} \mid \textbf{bool} \\
\langle \text{initializer} \rangle &\rightarrow \langle \text{literal} \rangle \mid \langle \text{array_literal} \rangle \\
\langle \text{statement} \rangle &\rightarrow \langle \text{location} \rangle \langle \text{assign_expr} \rangle \textbf{;} \\
&\mid \langle \text{method_call} \rangle \textbf{;} \\
&\mid \textbf{if} (\langle \text{expr} \rangle) \langle \text{block} \rangle [\textbf{else} \langle \text{block} \rangle] \\
&\mid \textbf{for} (\langle \text{id} \rangle = \langle \text{expr} \rangle \textbf{;} \langle \text{expr} \rangle \textbf{;} \langle \text{for_update} \rangle) \langle \text{block} \rangle \\
&\mid \textbf{while} (\langle \text{expr} \rangle) \langle \text{block} \rangle \\
&\mid \textbf{return} [\langle \text{expr} \rangle] \textbf{;} \\
&\mid \textbf{break} \textbf{;} \\
&\mid \textbf{continue} \textbf{;} \\
\langle \text{for_update} \rangle &\rightarrow \{ \langle \text{location} \rangle \langle \text{assign_expr} \rangle \} \mid \langle \text{method_call} \rangle \\
\langle \text{assign_expr} \rangle &\rightarrow \langle \text{assign_op} \rangle \langle \text{expr} \rangle \mid \langle \text{increment} \rangle \\
\langle \text{assign_op} \rangle &\rightarrow \textbf{=} \mid \textbf{+=} \mid \textbf{-=} \mid \textbf{*=} \mid \textbf{/=} \mid \textbf{\%=} \\
\langle \text{increment} \rangle &\rightarrow \textbf{++} \mid \textbf{--}
\end{aligned}$$

$\langle \text{method_call} \rangle \rightarrow \langle \text{method_name} \rangle ([\langle \text{expr} \rangle^+,])$
 $\quad \quad \quad | \quad \langle \text{method_name} \rangle ([\langle \text{import_arg} \rangle^+,])$

$\langle \text{method_name} \rangle \rightarrow \langle \text{id} \rangle$

$\langle \text{location} \rangle \rightarrow \langle \text{id} \rangle$
 $\quad \quad \quad | \quad \langle \text{id} \rangle \text{ '[' } \langle \text{expr} \rangle \text{ ']'}$

$\langle \text{expr} \rangle \rightarrow \langle \text{location} \rangle$
 $\quad \quad \quad | \quad \langle \text{method_call} \rangle$
 $\quad \quad \quad | \quad \langle \text{literal} \rangle$
 $\quad \quad \quad | \quad \text{len } (\langle \text{id} \rangle)$
 $\quad \quad \quad | \quad \langle \text{expr} \rangle \langle \text{bin_op} \rangle \langle \text{expr} \rangle$
 $\quad \quad \quad | \quad - \langle \text{expr} \rangle$
 $\quad \quad \quad | \quad ! \langle \text{expr} \rangle$
 $\quad \quad \quad | \quad (\langle \text{expr} \rangle)$

$\langle \text{import_arg} \rangle \rightarrow \langle \text{expr} \rangle | \langle \text{string_literal} \rangle$

$\langle \text{bin_op} \rangle \rightarrow \langle \text{arith_op} \rangle | \langle \text{rel_op} \rangle | \langle \text{eq_op} \rangle | \langle \text{cond_op} \rangle$

$\langle \text{arith_op} \rangle \rightarrow + | - | * | / | \%$

$\langle \text{rel_op} \rangle \rightarrow < | > | <= | >=$

$\langle \text{eq_op} \rangle \rightarrow == | !=$

$\langle \text{cond_op} \rangle \rightarrow \&\& | ||$

$\langle \text{literal} \rangle \rightarrow [-] \langle \text{int_literal} \rangle | \langle \text{char_literal} \rangle | \langle \text{bool_literal} \rangle$

$\langle \text{array_literal} \rangle \rightarrow \{ ' \langle \text{literal} \rangle^+, ' \}$

$\langle \text{id} \rangle \rightarrow \langle \text{alpha} \rangle \langle \text{alpha_num} \rangle^*$

$\langle \text{alpha_num} \rangle \rightarrow \langle \text{alpha} \rangle | \langle \text{digit} \rangle$

$\langle \text{alpha} \rangle \rightarrow \mathbf{a} | \mathbf{b} | \dots | \mathbf{z} | \mathbf{A} | \mathbf{B} | \dots | \mathbf{Z} | _$

$\langle \text{digit} \rangle \rightarrow \mathbf{0} | \mathbf{1} | \mathbf{2} | \dots | \mathbf{9}$

$\langle \text{hex_digit} \rangle \rightarrow \langle \text{digit} \rangle | \mathbf{a} | \mathbf{b} | \mathbf{c} | \mathbf{d} | \mathbf{e} | \mathbf{f} | \mathbf{A} | \mathbf{B} | \mathbf{C} | \mathbf{D} | \mathbf{E} | \mathbf{F}$

$\langle \text{int_literal} \rangle \rightarrow \langle \text{decimal_literal} \rangle | \langle \text{hex_literal} \rangle$

$\langle \text{decimal_literal} \rangle \rightarrow \langle \text{digit} \rangle \langle \text{digit} \rangle^*$

$\langle \text{hex_literal} \rangle \rightarrow \mathbf{0x} \langle \text{hex_digit} \rangle \langle \text{hex_digit} \rangle^*$

$\langle \text{bool_literal} \rangle \rightarrow \mathbf{true} | \mathbf{false}$

$\langle \text{char_literal} \rangle \rightarrow ' \langle \text{char} \rangle '$

$\langle \text{string_literal} \rangle \rightarrow " \langle \text{char} \rangle^* "$

3 Semantics

A Decaf program consists of a single file. A program consists of import declarations, field declarations and method declarations. Import declarations introduce methods from other libraries to the Decaf file. Field declarations introduce variables that can be accessed globally by all methods in the program. Method declarations introduce functions/procedures. The program must contain a declaration for a method called `main` that has type `void` and takes no parameters. Execution of a Decaf program starts at method `main`.

3.1 Types

There are two basic types in Decaf: `int` and `bool`. In addition, there are single-dimensional arrays of integers (`int[]`) and booleans (`bool[]`).

An `int` is a 64-bit signed integer. A `bool` is either `true` or `false`, and should be passed as a `1` or `0` respectively when invoking external methods.

Arrays may be declared in any scope. An array declaration must contain either a size $[N]$ or an initializer, but not both. All arrays are one-dimensional and have a compile-time fixed size. Arrays are indexed from 0 to $N - 1$, where $N > 0$ is the size of the array. Arrays are indexed by the usual bracket notation `a[i]`.

3.2 Scope Rules

Decaf has simple and quite restrictive scope rules. All identifiers must be defined (textually) before use. For example:

- A variable must be declared before it is used.
- A method can be called only by code appearing after its header. Note that recursive methods are allowed.

There are at least two valid scopes at any point in a Decaf program: the global scope and the method scope. The global scope consists of imports, fields, and methods introduced in the top level of the program. The method scope consists of method parameters introduced in a method declaration. Additional local scopes exist within each `{block}` in the code.

Identifiers introduced within a scope are visible in that scope and in all nested scopes. Identifiers within a more deeply nested scope shadow identifiers in a less deeply nested scope. For example, a variable declared in a method shadows a method of the same name in the global scope.

No identifier may be defined more than once in the same scope. Thus field and method names must all be distinct in the global scope, method parameters must all be distinct in a method scope, and local variables must all be distinct in a given block scope. Similarly, no identifier may be defined in the method parameters and the top level block of a method.

3.3 Locations

Locations in Decaf are either global (i.e., defined at the top level of the program) or local (i.e., defined as method parameters or local variables). Each location has a type. Locations of types `int` and `bool` contain scalar integer values and boolean values respectively. Locations of types `int[]` and `bool[]` denote arrays of integers and booleans respectively. Constant locations are declared with `const`.

If an initializer is present when the location is declared, the location is initialized to the value of the initializer. Otherwise, each location is initialized to a default value. Locations of type `int` have a default value of `0`, and values of type `bool` have a default value of `false`. Locations of type `int[]` and `bool[]` with no initializer must have a length parameter in the declaration, and have a default value of an array of that length with all elements initialized to their respective default values.

Global locations are initialized when the program starts. Local locations are initialized when the variable enters scope. If a variable enters scope multiple times (e.g., as in the body of a loop), it is initialized each time it enters scope.

Since arrays are statically sized in Decaf, global arrays may be allocated in the static data space of a program and need not be allocated on the heap. Local arrays may be dynamically allocated on the stack or statically allocated on the heap when appropriate.

3.4 Assignment

Assignment is only permitted for non-constant scalar locations (note that initialization is distinct from assignment). Decaf uses value-copy semantics for assignments.

The assignment `<location> = <expr>` copies the value resulting from the evaluation of `<expr>` into `<location>`, and is only valid when `<location>` and `<expr>` have the same type.

The compound assignment `<location> += <expr>` increments the value stored in `<location>` by `<expr>`, and is only valid for both `<location>` and `<expr>` of type `int`.

Other compound assignment operators `<location> -= <expr>`, `<location> *= <expr>`, `<location> /= <expr>`, and `<location> %= <expr>` work similarly to `<location> += <expr>`, but they perform subtraction, multiplication, division, and modulo, respectively.

The assignments `<location> ++` and `<location> --` increment and decrement the value stored in `<location>` by `1` respectively, and are only valid for `<location>` of type `int`.

In assignments containing `<expr>`s with side effects, the left hand side is evaluated before the right hand side. Both are evaluated exactly once, and are evaluated completely before the assignment is performed.

Note that it is legal to assign to a method parameter variable within a method body. Such assignments affect only the method scope.

3.5 Method Invocation and Return

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

Parameters to methods are passed by value (with the exception of arrays and strings passed to external functions, which is discussed in Section 3.8). The parameters of a method are considered to be like local variables of the method 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 method body in sequence.

A method that has no declared result type can only be called as a statement; it cannot be used in an expression. Such a method 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 method 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. It is illegal for control to reach the textual end of a method that returns a result.

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

3.6 Control Statements

3.6.1 `if`

The `if` statement has the usual semantics. First, the $\langle \text{expr} \rangle$ is evaluated. If the result is `true`, the first arm is executed. Otherwise, the `else` arm is executed, if it exists.

3.6.2 `while`

The `while` statement has the usual semantics. First, the $\langle \text{expr} \rangle$ is evaluated. If the result is `false`, control moves to the statement following the `while` statement. Otherwise, the loop body is executed. If control in a `while` statement reaches the end of the loop body, the `while` statement is executed again.

3.6.3 `for`

The `for` statement has the usual C-like semantics. The $\langle \text{id} \rangle$ is the loop index variable, and must have been declared in scope as an integer variable. Because it must be an identifier, this means that array index locations are not valid loop index variables. Before entering the loop body, it is assigned the value of the first $\langle \text{expr} \rangle$. This expression is evaluated once, just prior to reaching the loop for the first time.

The second $\langle \text{expr} \rangle$ is the ending condition of the loop, which must evaluate to type `bool`. It is evaluated before each iteration of the loop body. If the result is `false`, control moves to the statement following the `for` statement. Otherwise, the loop body is executed. Note that $\langle \text{expr} \rangle$ may be an expression with side effects (e.g., a method call).

After each iteration of the loop body, the $\langle \text{for_update} \rangle$ is executed. It may have side effects thus must be evaluated each time. The $\langle \text{location} \rangle$ updated does not need to be the same as the loop index variable. Control then returns to checking the ending condition.

3.6.4 `break` and `continue`

The `break` statement causes control to exit the innermost loop, moving control to the statement following the loop. The `continue` statement causes control to exit the current iteration of the innermost loop, moving control to the condition check at the beginning of the loop for `while` loops, and to the `<for_update>` for `for` loops.

3.6.5 `return`

The `return` statement causes control to exit the current method, returning to the caller.

3.7 Expressions

Expressions follow the normal rules for evaluation. All binary operators are left associative.

A location expression evaluates to the value contained by the location.

Method invocation expressions are discussed in *Method Invocation and Return*. Array operations are discussed in *Types*. I/O related expressions are discussed in *External Library*.

Integer literals evaluate to their integer value. Character literals evaluate to their integer ASCII values: e.g., `'A'` represents the integer 65. The type of a character literal is `int`.

The `len` operator evaluates the size of an array to an `int` value. Because array lengths are static and known at compile time, this is a constant expression which can be evaluated at compile time. Note that the `len` operator is not a method call.

The arithmetic operators (`<arith_op>`) and unary minus `-` have their usual meaning. The division operator `/` computes the integer part of the quotient of its operands (e.g., `5/2` evaluates to `2`). The mod operator `%` computes the remainder of dividing its operands (e.g., `5%2` evaluates to `1`). Division and mod by zero have undefined behavior, and will not be tested. The result of an arithmetic operator has type `int`.

Relational operators (`<rel_op>`) are used to compare integer expressions, and also have their usual meaning. The result of a relational operator or equality operator has type `bool`.

The equality operators, `==` (equal) and `!=` (not equal) are defined for `int` and `bool` types only, and can be used to compare any two expressions having the same type. The result of an equality operator has type `bool`.

The boolean connectives `&&` and `||` are interpreted using short circuit evaluation as in Java. That is, the right hand side is not executed (and therefore has no side effects if it contains a method call) if the result of the expression is already known (i.e., if the left hand side is `false` for `&&` or `true` for `||`).

The logical not operator `!` has the usual meaning, and is defined for `bool` expressions only.

Operator precedence, from highest to lowest:

<i>Operators</i>	<i>Comments</i>
-	unary minus
!	logical not
* / %	multiplication, division, remainder
+ -	addition, subtraction
< <= >= >	relational
== !=	equality
&&	conditional and
	conditional or

Note that this precedence is not reflected in the reference grammar.

3.8 External Library

Decaf includes a method for calling external functions, similar to the C language. An external function is put in scope with the `import` keyword at the top of the file. The syntax (as specified in the grammar) is:

```
import <id>    ;
```

All external functions are treated as if they return `int`. Once imports have been declared, they may be called similarly to any function. The one exception is that arguments to imports may contain string literals. Normal Decaf methods may not contain string literals as arguments.

3.8.1 Import Arguments

When calling an import, the following rules apply to the arguments:

Expressions of type `bool` and `int` are passed as integers. Expressions of type `int []` and `bool []` are passed as pointers to the first element of the array. A string literal is passed as a pointer to the first character of a null-terminated character array. The return value is passed back as an integer.

The user of an `import` is responsible for ensuring that the arguments given match the signature of the function, and that the return value is only used if the underlying library function actually returns a value of appropriate type. Arguments are passed to the function in the system's standard calling convention. The compiler is not responsible for verifying that imports have the correct number or type of arguments.

3.8.2 External I/O Functions

In addition to accessing the standard C library using `import`, an I/O function can be written in C (or any other language), compiled using standard tools, linked with the runtime system, and accessed by the `import` mechanism. This is the only way to access I/O functions in Decaf.

4 Semantic Rules

These rules place additional constraints on the set of valid Decaf programs besides the constraints implied by the grammar. A program that is grammatically well-formed and does not violate any of the following rules is called a *legal* program. A robust compiler will explicitly check each of these rules, and will generate an error message describing each violation it is able to find. A robust compiler will generate at least one error message for each illegal program, but will generate no errors for a legal program.

1. No identifier is declared twice in the same scope. This include `import` identifiers, which exist in the global scope.
2. No identifier is used before it is declared.
3. The program contains a definition for a method called `main` that has type `void` and takes no parameters. Note that since execution starts at method `main`, any methods defined after `main` will never be executed.
4. All types of initializers must match the type of the variable being initialized.
5. Array initializers have either a declared length or an initializer list, but not both.
6. If present, the `<int_literal>` in an array declaration must be greater than 0.
7. The number and types of parameters in a method call (non-import) must be the same as the number and types of the declared parameters for the method.
8. If a method call is used as an expression, the method must return a result.
9. String literals and array variables may not be used as parameters to non-import methods. (Note: an expression like `a[0]` is not an array variable).
10. A `return` statement must not have a return value unless it appears in the body of a method that is declared to return a value.
11. The expression in a `return` statement must have the same type as the declared result type of the enclosing method definition.
12. An `<id>` used as a `<location>` must name a declared local/global variable or parameter.
13. The `<id>` in a method statement must be a declared method or import.
14. For all locations of the form `<id>[<expr>]:`
 - (a) `<id>` must be an array variable, and
 - (b) the type of `<expr>` must be `int`.
15. The argument of the `len` operator must be an array variable.
16. The `<expr>` in an `if` or `while` statement must have type `bool`, as well as the second `<expr>` of a `for` statement.
17. The operands of the unary minus, `<arith.op>`s and `<rel.op>`s must have type `int`.

18. The operands of $\langle \text{eq_op} \rangle$ s must have the same type, either `int` or `bool`.
19. The operands of $\langle \text{cond_op} \rangle$ s and the operand of logical not (`!`) must have type `bool`.
20. The $\langle \text{location} \rangle$ and the $\langle \text{expr} \rangle$ in an assignment, $\langle \text{location} \rangle = \langle \text{expr} \rangle$, must have the same type.
21. The $\langle \text{location} \rangle$ and the $\langle \text{expr} \rangle$ in a compound assignment, $\langle \text{location} \rangle += \langle \text{expr} \rangle$, $\langle \text{location} \rangle -= \langle \text{expr} \rangle$, $\langle \text{location} \rangle *= \langle \text{expr} \rangle$, $\langle \text{location} \rangle /= \langle \text{expr} \rangle$, and $\langle \text{location} \rangle \%= \langle \text{expr} \rangle$, must be of type `int`. The same is true of the $\langle \text{location} \rangle$ in `++` and `--` statements.
22. Declarations of `const` locations must have an initializer
23. Locations declared `const` must not be assigned to
24. All `break` and `continue` statements must be contained within the body of a `for` or a `while` statement.
25. All integer literals must be in the range $-9223372036854775808 \leq x \leq 9223372036854775807$ (64 bits).

5 Run Time Checking

In addition to the constraints described above, which are statically enforced by the compiler's semantic checker, the following constraints are enforced dynamically. The compiler's code generator must emit code to perform these checks; violations are discovered at run-time.

1. The subscript of an array must be in bounds.
2. Control must not fall off the end of a method that is declared to return a result.

When a run-time error occurs, an appropriate error message is output to the terminal and the program terminates. If the subscript of an array is found to be out of bounds, the program must terminate with exit value -1. If control falls off the end of a method that is declared to return a result, the program must terminate with exit value -2. The error messages output should be helpful to the programmer trying to find the problem in the source program.