

Chapter 7

Intermediate-Code Generation

7.1 Introduction

The final goal of a compiler is to get programs written in a high-level language to run on a computer. This means that, eventually, the program will have to be expressed as machine code which can run on the computer. This does not mean that we need to translate directly from the high-level abstract syntax to machine code. Many compilers use a medium-level language as a stepping-stone between the high-level language and the very low-level machine code. Such stepping-stone languages are called *intermediate code*.

Apart from structuring the compiler into smaller jobs, using an intermediate language has other advantages:

- If the compiler needs to generate code for several different machine-architectures, only one translation to intermediate code is needed. Only the translation from intermediate code to machine language (*i.e.*, the *back-end*) needs to be written in several versions.
- If several high-level languages need to be compiled, only the translation to intermediate code need to be written for each language. They can all share the back-end, *i.e.*, the translation from intermediate code to machine code.
- Instead of translating the intermediate language to machine code, it can be *interpreted* by a small program written in machine code or a language for which a compiler or interpreter already exists.

The advantage of using an intermediate language is most obvious if many languages are to be compiled to many machines. If translation is done directly, the number of compilers is equal to the product of the number of languages and the number of machines. If a common intermediate language is used, one front-end (*i.e.*, compiler

to intermediate code) is needed for every language and one back-end (interpreter or code generator) is needed for each machine, making the total number of front-ends and back-ends equal to the sum of the number of languages and the number of machines.

If an interpreter for an intermediate language is written in a language for which there already exist implementations for the target machines, the same interpreter can be interpreted or compiled for each machine. This way, there is no need to write a separate back-end for each machine. The advantages of this approach are:

- No actual back-end needs to be written for each new machine, as long as the machine is equipped with an interpreter or compiler for the implementation language of the interpreter for the intermediate language.
- A compiled program can be distributed in a single intermediate form for all machines, as opposed to shipping separate binaries for each machine.
- The intermediate form may be more compact than machine code. This saves space both in distribution and on the machine that executes the programs (though the latter is somewhat offset by requiring the interpreter to be kept in memory during execution).

The disadvantage is speed: Interpreting the intermediate form will in most cases be a lot slower than executing translated code directly. Nevertheless, the approach has seen some success, *e.g.*, with Java.

Some of the speed penalty can be eliminated by translating the intermediate code to machine code immediately before or during execution of the program. This hybrid form is called *just-in-time compilation* and is often used for executing the intermediate code for Java.

We will in this book, however, focus mainly on using the intermediate code for traditional compilation, where the intermediate form will be translated to machine code by the back-end of the compiler.

7.2 Choosing an intermediate language

An intermediate language should, ideally, have the following properties:

- It should be easy to translate from a high-level language to the intermediate language. This should be the case for a wide range of different source languages.
- It should be easy to translate from the intermediate language to machine code. This should be true for a wide range of different target architectures.

- The intermediate format should be suitable for optimisations.

The first two of these properties can be somewhat hard to reconcile. A language that is intended as target for translation from a high-level language should be fairly close to this. However, this may be hard to achieve for more than a small number of similar languages. Furthermore, a high-level intermediate language puts more burden on the back-ends. A low-level intermediate language may make it easy to write back-ends, but puts more burden on the front-ends. A low-level intermediate language, also, may not fit all machines equally well, though this is usually less of a problem than the similar problem for front-ends, as machines typically are more similar than high-level languages.

A solution that may reduce the translation burden, though it does not address the other problems, is to have two intermediate levels: One, which is fairly high-level, is used for the front-ends and the other, which is fairly low-level, is used for the back-ends. A single shared translator is then used to translate between these two intermediate formats.

When the intermediate format is shared between many compilers, it makes sense to do as many optimisations as possible on the intermediate format. This way, the (often substantial) effort of writing good optimisations is done only once instead of in every compiler.

Another thing to consider when choosing an intermediate language is the “granularity”: Should an operation in the intermediate language correspond to a large amount of work or to a small amount of work?

The first of these approaches is often used when the intermediate language is interpreted, as the overhead of decoding instructions is amortised over more actual work, but it can also be used for compiling. In this case, each intermediate-code operation is typically translated into a sequence of machine-code instructions. When coarse-grained intermediate code is used, there is typically a fairly large number of different intermediate-code operations.

The opposite approach is to let each intermediate-code operation be as small as possible. This means that each intermediate-code operation is typically translated into a single machine-code instruction or that several intermediate-code operations can be combined into one machine-code operation. The latter can, to some degree, be automated as each machine-code instruction can be described as a sequence of intermediate-code instructions. When intermediate-code is translated to machine-code, the code generator can look for sequences that match machine-code operations. By assigning cost to each machine-code operation, this can be turned into a combinatorial optimisation problem, where the least-cost solution is found. We will return to this in chapter 8.

<i>Program</i>	→	[<i>Instructions</i>]
<i>Instructions</i>	→	<i>Instruction</i>
<i>Instructions</i>	→	<i>Instruction</i> , <i>Instructions</i>
<i>Instruction</i>	→	LABEL labelid
<i>Instruction</i>	→	id := <i>Atom</i>
<i>Instruction</i>	→	id := unop <i>Atom</i>
<i>Instruction</i>	→	id := id binop <i>Atom</i>
<i>Instruction</i>	→	id := <i>M</i> [<i>Atom</i>]
<i>Instruction</i>	→	<i>M</i> [<i>Atom</i>] := id
<i>Instruction</i>	→	GOTO labelid
<i>Instruction</i>	→	IF id relop <i>Atom</i> THEN labelid ELSE labelid
<i>Instruction</i>	→	id := CALL functionid (<i>Args</i>)
<i>Atom</i>	→	id
<i>Atom</i>	→	num
<i>Args</i>	→	id
<i>Args</i>	→	id , <i>Args</i>

Grammar 7.1: The intermediate language

7.3 The intermediate language

In this chapter we have chosen a fairly low-level fine-grained intermediate language, as it is best suited to convey the techniques we want to cover.

We will not treat translation of function calls until chapter 10, so a “program” in our intermediate language will, for the time being, correspond to the body of a function or procedure in a real program. For the same reason, function calls are initially treated as primitive operations in the intermediate language.

The grammar for the intermediate language is shown in grammar 7.1. A program is a sequence of instructions. The instructions are:

- A label. This has no effect but serves only to mark the position in the program as a target for jumps.
- An assignment of an atomic expression (constant or variable) to a variable.
- A unary operator applied to an atomic expression, with the result stored in a variable.

- A binary operator applied to a variable and an atomic expression, with the result stored in a variable.
- A transfer from memory to a variable. The memory location is an atomic expression.
- A transfer from a variable to memory. The memory location is an atomic expression.
- A jump to a label.
- A conditional selection between jumps to two labels. The condition is found by comparing a variable with an atomic expression by using a relational operator ($=$, \neq , $<$, $>$, \leq or \geq).
- A function call. The arguments to the function call are variables and the result is assigned to a variable. This instruction is used even if there is no actual result (*i.e.* if a procedure is called instead of a function), in which case the result variable is a dummy variable.

An atomic expression is either a variable or a constant.

We have not specified the set of unary and binary operations, but we expect these to include normal integer arithmetic and bitwise logical operations.

We assume that all values are integers. Adding floating-point numbers and other primitive types is not difficult, though.

7.4 Syntax-directed translation

We will generate code using translation functions for each syntactic category, similarly to the functions we used for interpretation and type checking. We generate code for a syntactic construct independently of the constructs around it, except that the parameters of a translation function may hold information about the context (such as symbol tables) and the result of a translation function may (in addition to the generated code) hold information about how the generated code interfaces with its context (such as which variables it uses). Since the translation closely follows the syntactic structure of the program, it is called *syntax-directed translation*.

Given that translation of a syntactic construct is mostly independent of the surrounding and enclosed syntactic constructs, we might miss opportunities to exploit synergies between these and, hence, generate less than optimal code. We will try to remedy this in later chapters by using various optimisation techniques.

$$\begin{aligned}
 \textit{Exp} &\rightarrow \mathbf{num} \\
 \textit{Exp} &\rightarrow \mathbf{id} \\
 \textit{Exp} &\rightarrow \mathbf{unop} \textit{Exp} \\
 \textit{Exp} &\rightarrow \textit{Exp} \mathbf{binop} \textit{Exp} \\
 \textit{Exp} &\rightarrow \mathbf{id}(\textit{Exps}) \\
 \\
 \textit{Exps} &\rightarrow \textit{Exp} \\
 \textit{Exps} &\rightarrow \textit{Exp} , \textit{Exps}
 \end{aligned}$$

Grammar 7.2: A simple expression language

7.5 Generating code from expressions

Grammar 7.2 shows a simple language of expressions, which we will use as our initial example for translation. Again, we have let the set of unary and binary operators be unspecified but assume that the intermediate language includes all those used by the expression language. We assume that there is a function *transop* that translates the name of an operator in the expression language into the name of the corresponding operator in the intermediate language. The tokens **unop** and **binop** have the names of the actual operators as attributes, accessed by the function *getopname*.

When writing a compiler, we must decide what needs to be done at compile-time and what needs to be done at run-time. Ideally, as much as possible should be done at compile-time, but some things need to be postponed until run-time, as they need the actual values of variables, *etc.*, which are not known at compile-time. When we, below, explain the workings of the translation functions, we might use phrasing like “the expression is evaluated and the result stored in the variable”. This describes actions that are performed at run-time by the code that is generated at compile-time. At times, the textual description may not be 100% clear as to what happens at which time, but the notation used in the translation functions make this clear: Intermediate-language code is executed at run-time, the rest is done at compile time. Intermediate-language instructions may refer to values (constants and register names) that are generated at compile time. When instructions have operands that are written in *italics*, these operands are variables in the compiler that contain compile-time values that are inserted into the generated code. For example, if *place* holds the variable name $\tau 14$ and *v* holds the value 42, then the code template [*place* := *v*] will generate the code [$\tau 14 := 42$].

When we want to translate the expression language to the intermediate language, the main complication is that the expression language is tree-structured

while the intermediate language is flat, requiring the result of every operation to be stored in a variable and every (non-constant) argument to be in one. We use a function *newvar* to generate new variable names in the intermediate language. Whenever *newvar* is called, it returns a previously unused variable name.

We will describe translation of expressions by a translation function using a notation similar to the notation we used for type-checking functions in chapter 6.

Some attributes for the translation function are obvious: It must return the code as a synthesised attribute. Furthermore, it must translate variables and functions used in the expression language to the names these correspond to in the intermediate language. This can be done by symbol tables *vtable* and *ftable* that bind variable and function names in the expression language into the corresponding names in the intermediate language. The symbol tables are passed as inherited attributes to the translation function. In addition to these attributes, the translation function must use attributes to decide where to put the values of sub-expressions. This can be done in two ways:

- 1) The location of the values of a sub-expression can be passed up as a synthesised attribute to the parent expression, which decides on a position for its own value.
- 2) The parent expression can decide where it wants to find the values of its sub-expressions and pass this information down to these as inherited attributes.

Neither of these is obviously superior to the other. Method 1 has a slight advantage when generating code for a variable access, as it does not have to generate any code, but can simply return the name of the variable that holds the value. This, however, only works under the assumption that the variable is not updated before the value is used by the parent expression. If expressions can have side effects, this is not always the case, as the C expression “ $x+(x=3)$ ” shows. Our expression language does not have assignment, but it does have function calls, which may have side effects.

Method 2 does not have this problem: Since the value of the expression is created immediately before the assignment is executed, there is no risk of other side effects between these two points in time. Method 2 also has a slight advantage when we later extend the language to have assignment statements, as we can then generate code that calculates the expression result directly into the desired variable instead of having to copy it from a temporary variable.

Hence, we will choose method 2 for our translation function $Trans_{Exp}$, which is shown in figure 7.3.

The inherited attribute *place* is the intermediate-language variable that the result of the expression must be stored in.

$Trans_{Exp}(Exp, vtable, ftable, place) = \text{case } Exp \text{ of}$	
num	$v = \text{getvalue}(\mathbf{num})$ [$place := v$]
id	$x = \text{lookup}(vtable, \text{getname}(\mathbf{id}))$ [$place := x$]
unop Exp_1	$place_1 = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp_1, vtable, ftable, place_1)$ $op = \text{transop}(\text{getopname}(\mathbf{unop}))$ $code_1 ++ [place := op \ place_1]$
Exp_1 binop Exp_2	$place_1 = \text{newvar}()$ $place_2 = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp_1, vtable, ftable, place_1)$ $code_2 = Trans_{Exp}(Exp_2, vtable, ftable, place_2)$ $op = \text{transop}(\text{getopname}(\mathbf{binop}))$ $code_1 ++ code_2 ++ [place := place_1 \ op \ place_2]$
id ($Exps$)	$(code_1, [a_1, \dots, a_n])$ $= Trans_{Exps}(Exps, vtable, ftable)$ $fname = \text{lookup}(ftable, \text{getname}(\mathbf{id}))$ $code_1 ++ [place := \text{CALL } fname(a_1, \dots, a_n)]$

$Trans_{Exps}(Exps, vtable, ftable) = \text{case } Exps \text{ of}$	
Exp	$place = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp, vtable, ftable, place)$ $(code_1, [place])$
$Exp, Exps$	$place = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp, vtable, ftable, place)$ $(code_2, args) = Trans_{Exps}(Exps, vtable, ftable)$ $code_3 = code_1 ++ code_2$ $args_1 = place :: args$ $(code_3, args_1)$

Figure 7.3: Translating an expression

If the expression is just a number, the value of that number is stored in the *place*.

If the expression is a variable, the intermediate-language equivalent of this variable is found in *vtable* and an assignment copies it into the intended *place*.

A unary operation is translated by first generating a new intermediate-language variable to hold the value of the argument of the operation. Then the argument is translated using the newly generated variable for the *place* attribute. We then use an **unop** operation in the intermediate language to assign the result to the inherited *place*. The operator ++ concatenates two lists of instructions.

A binary operation is translated in a similar way. Two new intermediate-language variables are generated to hold the values of the arguments, then the arguments are translated and finally a binary operation in the intermediate language assigns the final result to the inherited *place*.

A function call is translated by first translating the arguments, using the auxiliary function *Trans_{Exps}*. Then a function call is generated using the argument variables returned by *Trans_{Exps}*, with the result assigned to the inherited *place*. The name of the function is looked-up in *f_{table}* to find the corresponding intermediate-language name.

Trans_{Exps} generates code for each argument expression, storing the results into new variables. These variables are returned along with the code, so they can be put into the argument list of the call instruction.

7.5.1 Examples of translation

Translation of expressions is always relative to symbol tables and a place for storing the result. In the examples below, we assume a variable symbol table that binds *x*, *y* and *z* to *v0*, *v1* and *v2*, respectively and a function table that binds *f* to *_f*. The place for the result is *t0* and we assume that calls to *newvar()* return, in sequence, the variables *t1*, *t2*, *t3*, ...

We start by the simple expression *x-3*. This is a binop-expression, so the first we do is to call *newvar()* twice, giving *place₁* the value *t1* and *place₂* the value *t2*. We then call *Trans_{Exp}* recursively with the expression *x*. When translating this, we first look up *x* in the variable symbol table, yielding *v0*, and then return the code [*t1* := *v0*]. Back in the translation of the subtraction expression, we assign this code to *code₁* and once more call *Trans_{Exp}* recursively, this time with the expression *3*. This is translated to the code [*t2* := *3*], which we assign to *code₂*. The final result is produced by *code₁*++*code₂*+[*t0* := *t1* - *t2*] which yields [*t1* := *v0*, *t2* := *3*, *t0* := *t1* - *t2*]. We have translated the source-language operator - to the intermediate-language operator -.

The resulting code looks quite suboptimal, and could, indeed, be shortened to [*t0* := *v0* - *3*]. When we generate intermediate code, we want, for simplicity, to

$$\begin{aligned}
Stat &\rightarrow Stat ; Stat \\
Stat &\rightarrow \mathbf{id} := Exp \\
Stat &\rightarrow \mathbf{if} Cond \mathbf{then} Stat \\
Stat &\rightarrow \mathbf{if} Cond \mathbf{then} Stat \mathbf{else} Stat \\
Stat &\rightarrow \mathbf{while} Cond \mathbf{do} Stat \\
Stat &\rightarrow \mathbf{repeat} Stat \mathbf{until} Cond \\
\\
Cond &\rightarrow Exp \mathbf{relop} Exp
\end{aligned}$$

Grammar 7.4: Statement language

treat each subexpression independently of its context. This may lead to superfluous assignments. We will look at ways of getting rid of these when we treat machine code generation and register allocation in chapters 8 and 9.

A more complex expression is $3+f(x-y,z)$. Using the same assumptions as above, this yields the code

```

t1 := 3
  t4 := v0
  t5 := v1
  t3 := t4 - t5
  t6 := v2
  t2 := CALL _f(t3,t6)
t0 := t1 + t2

```

We have, for readability, laid the code out on separate lines rather than using a comma-separated list. The indentation indicates the depth of calls to $Trans_{Exp}$ that produced the code in each line.

Suggested exercises: 7.1.

7.6 Translating statements

We now extend the expression language in figure 7.2 with statements. The extensions are shown in grammar 7.4.

When translating statements, we will need the symbol table for variables (for translating assignment), and since statements contain expressions, we also need *f*table so we can pass it on to $Trans_{Exp}$.

Just like we use *newvar* to generate new unused variables, we use a similar function *newlabel* to generate new unused labels. The translation function for statements is shown in figure 7.5. It uses an auxiliary translation function for conditions shown in figure 7.6.

A sequence of two statements are translated by putting the code for these in sequence.

An assignment is translated by translating the right-hand-side expression using the left-hand-side variable as target location (*place*).

When translating statements that use conditions, we use an auxiliary function *TransCond*. *TransCond* translates the arguments to the condition and generates an IF-THEN-ELSE instruction using the same relational operator as the condition. The target labels of this instruction are inherited attributes to *TransCond*.

An if-then statement is translated by first generating two labels: One for the then-branch and one for the code following the if-then statement. The condition is translated by *TransCond*, which is given the two labels as attributes. When (at run-time) the condition is true, the first of these are selected, and when false, the second is chosen. Hence, when the condition is true, the then-branch is executed followed by the code after the if-then statement. When the condition is false, we jump directly to the code following the if-then statement, hence bypassing the then-branch.

An if-then-else statement is treated similarly, but now the condition must choose between jumping to the then-branch or the else-branch. At the end of the then-branch, a jump bypasses the code for the else-branch by jumping to the label at the end. Hence, there is need for three labels: One for the then-branch, one for the else-branch and one for the code following the if-then-else statement.

If the condition in a while-do loop is true, the body must be executed, otherwise the body is by-passed and the code after the loop is executed. Hence, the condition is translated with attributes that provide the label for the start of the body and the label for the code after the loop. When the body of the loop has been executed, the condition must be re-tested for further passes through the loop. Hence, a jump is made to the start of the code for the condition. A total of three labels are thus required: One for the start of the loop, one for the loop body and one for the end of the loop.

A repeat-until loop is slightly simpler. The body precedes the condition, so there is always at least one pass through the loop. If the condition is true, the loop is terminated and we continue with the code after the loop. If the condition is false, we jump to the start of the loop. Hence, only two labels are needed: One for the start of the loop and one for the code after the loop.

Suggested exercises: 7.2.

<i>Trans_{Stat}(Stat, vtable, ftable) = case Stat of</i>	
<i>Stat₁ ; Stat₂</i>	<i>code₁ = Trans_{Stat}(Stat₁, vtable, ftable) code₂ = Trans_{Stat}(Stat₂, vtable, ftable) code₁++code₂</i>
<i>id := Exp</i>	<i>place = lookup(vtable, getname(id)) Trans_{Exp}(Exp, vtable, ftable, place)</i>
<i>if Cond then Stat₁</i>	<i>label₁ = newlabel() label₂ = newlabel() code₁ = Trans_{Cond}(Cond, label₁, label₂, vtable, ftable) code₂ = Trans_{Stat}(Stat₁, vtable, ftable) code₁++[LABEL label₁]++code₂ ++[LABEL label₂]</i>
<i>if Cond then Stat₁ else Stat₂</i>	<i>label₁ = newlabel() label₂ = newlabel() label₃ = newlabel() code₁ = Trans_{Cond}(Cond, label₁, label₂, vtable, ftable) code₂ = Trans_{Stat}(Stat₁, vtable, ftable) code₃ = Trans_{Stat}(Stat₂, vtable, ftable) code₁++[LABEL label₁]++code₂ ++[GOTO label₃, LABEL label₂] ++code₃++[LABEL label₃]</i>
<i>while Cond do Stat₁</i>	<i>label₁ = newlabel() label₂ = newlabel() label₃ = newlabel() code₁ = Trans_{Cond}(Cond, label₂, label₃, vtable, ftable) code₂ = Trans_{Stat}(Stat₁, vtable, ftable) [LABEL label₁]++code₁ ++[LABEL label₂]++code₂ ++[GOTO label₁, LABEL label₃]</i>
<i>repeat Stat₁ until Cond</i>	<i>label₁ = newlabel() label₂ = newlabel() code₁ = Trans_{Stat}(Stat₁, vtable, ftable) code₂ = Trans_{Cond}(Cond, label₂, label₁, vtable, ftable) [LABEL label₁]++code₁ ++code₂++[LABEL label₂]</i>

Figure 7.5: Translation of statements

$Trans_{Cond}(Cond, label_t, label_f, vtable, ftable) = \text{case } Cond \text{ of}$	
$Exp_1 \text{ relop } Exp_2$	$t_1 = \text{newvar}()$ $t_2 = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp_1, vtable, ftable, t_1)$ $code_2 = Trans_{Exp}(Exp_2, vtable, ftable, t_2)$ $op = \text{transop}(\text{getopname}(\text{relop}))$ $code_1 ++ code_2 ++ [\text{IF } t_1 \text{ opt}_2 \text{ THEN } label_t \text{ ELSE } label_f]$

Figure 7.6: Translation of simple conditions

7.7 Logical operators

Logical conjunction, disjunction and negation are often available for conditions, so we can write, *e.g.*, $x = y$ **or** $y = z$, where **or** is a logical disjunction operator. There are typically two ways to treat logical operators in programming languages:

- 1) Logical operators are similar to arithmetic operators: The arguments are evaluated and the operator is applied to find the result.
- 2) The second operand of a logical operator is not evaluated if the first operand is sufficient to determine the result. This means that a logical **and** will not evaluate its second operand if the first evaluates to **false**, and a logical **or** will not evaluate the second operand if the first is **true**.

The first variant is typically implemented by using bitwise logical operators and uses 0 to represent **false** and a nonzero value (typically 1 or -1) to represent **true**. In C, there is no separate boolean type. The integer 1 is used for logical truth¹ and 0 for falsehood. Bitwise logical operators **&** (bitwise **and**) and **|** (bitwise **or**) are used to implement the corresponding logical operations. Logical negation is *not* handled by bitwise negation, as the bitwise negation of 1 is not 0. Instead, a special logical negation operator **!** is used that maps any non-zero value to 0 and 0 to 1. We assume an equivalent operator is available in the intermediate language.

The second variant is called *sequential logical operators*. In C, these are called **&&** (logical **and**) and **||** (logical **or**).

Adding non-sequential logical operators to our language is not too difficult. Since we have not said exactly which binary and unary operators exist in the intermediate language, we can simply assume these include relational operators, bitwise logical operations and logical negation. We can now simply allow any expression² as a condition by adding the production

¹Actually, any non-zero value is treated as logical truth.

²If it is of boolean type, which we assume has been verified by the type checker.

$$Cond \rightarrow Exp$$

to grammar 7.4. We then extend the translation function for conditions as follows:

$Trans_{Cond}(Cond, label_t, label_f, vtable, ftable) = \text{case } Cond \text{ of}$	
$Exp_1 \mathbf{relop} Exp_2$	$t_1 = \text{newvar}()$ $t_2 = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp_1, vtable, ftable, t_1)$ $code_2 = Trans_{Exp}(Exp_2, vtable, ftable, t_2)$ $op = \text{transop}(\text{getopname}(\mathbf{relop}))$ $code_1 ++ code_2 ++ [\text{IF } t_1 \text{ op } t_2 \text{ THEN } label_t \text{ ELSE } label_f]$
Exp	$t = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp, vtable, ftable, t)$ $code_1 ++ [\text{IF } t \neq 0 \text{ THEN } label_t \text{ ELSE } label_f]$

We need to convert the numerical value returned by $Trans_{Exp}$ into a choice between two labels, so we generate an IF instruction that does just that.

The rule for relational operators is now actually superfluous, as the case it handles is covered by the second rule (since relational operators are assumed to be included in the set of binary arithmetic operators). However, we can consider it an optimisation, as the code it generates is shorter than the equivalent code generated by the second rule. It will also be natural to keep it separate when we add sequential logical operators.

7.7.1 Sequential logical operators

We will use the same names for sequential logical operators as C, *i.e.*, $\&\&$ for logical **and**, $\|\|$ for logical **or** and $!$ for logical negation. The extended language is shown in figure 7.7. Note that we allow an expression to be a condition as well as a condition to be an expression. This grammar is highly ambiguous (not least because **binop** overlaps **relop**). As before, we assume such ambiguity to be resolved by the parser before code generation. We also assume that the last productions of Exp and $Cond$ are used as little as possible, as this will yield the best code.

The revised translation functions for Exp and $Cond$ are shown in figure 7.8. Only the new cases for Exp are shown.

As expressions, **true** and **false** are the numbers 1 and 0.

A condition $Cond$ is translated into code that chooses between two labels. When we want to use a condition as an expression, we must convert this choice into a number. We do this by first assuming that the condition is false and hence assign 0 to the target location. We then, if the condition is true, jump to code that assigns 1 to the target location. If the condition is false, we jump around this code, so

$Exp \rightarrow \mathbf{num}$
 $Exp \rightarrow \mathbf{id}$
 $Exp \rightarrow \mathbf{unop} \ Exp$
 $Exp \rightarrow \ Exp \ \mathbf{binop} \ Exp$
 $Exp \rightarrow \mathbf{id}(Exps)$
 $Exp \rightarrow \mathbf{true}$
 $Exp \rightarrow \mathbf{false}$
 $Exp \rightarrow \ Cond$

$Exps \rightarrow \ Exp$
 $Exps \rightarrow \ Exp \ , \ Exps$

$Cond \rightarrow \ Exp \ \mathbf{relop} \ Exp$
 $Cond \rightarrow \mathbf{true}$
 $Cond \rightarrow \mathbf{false}$
 $Cond \rightarrow \ ! \ Cond$
 $Cond \rightarrow \ Cond \ \&\& \ Cond$
 $Cond \rightarrow \ Cond \ || \ Cond$
 $Cond \rightarrow \ Exp$

Grammar 7.7: Example language with logical operators

$Trans_{Exp}(Exp, vtable, ftable, place) = \text{case } Exp \text{ of}$	
⋮	
true	$[place := 1]$
false	$[place := 0]$
Cond	$label_1 = \text{newlabel}()$ $label_2 = \text{newlabel}()$ $code_1 = Trans_{Cond}(Cond, label_1, label_2, vtable, ftable)$ $[place := 0] ++ code_1$ $++ [LABEL label_1, place := 1]$ $++ [LABEL label_2]$
$Trans_{Cond}(Cond, label_t, label_f, vtable, ftable) = \text{case } Cond \text{ of}$	
$Exp_1 \mathbf{relop} Exp_2$	$t_1 = \text{newvar}()$ $t_2 = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp_1, vtable, ftable, t_1)$ $code_2 = Trans_{Exp}(Exp_2, vtable, ftable, t_2)$ $op = \text{transop}(\text{getopname}(\mathbf{relop}))$ $code_1 ++ code_2 ++ [IF t_1 op t_2 THEN label_t ELSE label_f]$
true	$[GOTO label_t]$
false	$[GOTO label_f]$
$! Cond_1$	$Trans_{Cond}(Cond_1, label_f, label_t, vtable, ftable)$
$Cond_1 \ \&\& \ Cond_2$	$arg_2 = \text{newlabel}()$ $code_1 = Trans_{Cond}(Cond_1, arg_2, label_f, vtable, ftable)$ $code_2 = Trans_{Cond}(Cond_2, label_t, label_f, vtable, ftable)$ $code_1 ++ [LABEL arg_2] ++ code_2$
$Cond_1 \ \ Cond_2$	$arg_2 = \text{newlabel}()$ $code_1 = Trans_{Cond}(Cond_1, label_t, arg_2, vtable, ftable)$ $code_2 = Trans_{Cond}(Cond_2, label_t, label_f, vtable, ftable)$ $code_1 ++ [LABEL arg_2] ++ code_2$
Exp	$t = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp, vtable, ftable, t)$ $code_1 ++ [IF t \neq 0 THEN label_t ELSE label_f]$

Figure 7.8: Translation of sequential logical operators

the value remains 0. We could equally well have done things the other way around, *i.e.*, first assign 1 to the target location and modify this to 0 when the condition is false.

It gets a bit more interesting in $Trans_{Cond}$, where we translate conditions. We have already seen how comparisons and expressions are translated, so we move directly to the new cases.

The constant `true` condition just generates a jump to the label for true conditions, and, similarly, `false` generates a jump to the label for false conditions.

Logical negation generates no code by itself, it just swaps the attribute-labels for true and false when translating its argument. This negates the effect of the argument condition.

Sequential logical **and** is translated as follows: The code for the first operand is translated such that if it is false, the second condition is not tested. This is done by jumping straight to the label for false conditions when the first operand is false. If the first operand is true, a jump to the code for the second operand is made. This is handled by using the appropriate labels as arguments to the call to $Trans_{Cond}$. The call to $Trans_{Cond}$ for the second operand uses the original labels for true and false. Hence, both conditions have to be true for the combined condition to be true.

Sequential **or** is similar: If the first operand is true, we jump directly to the label for true conditions without testing the second operand, but if it is false, we jump to the code for the second operand. Again, the second operand uses the original labels for true and false.

Note that the translation functions now work even if **binop** and **unop** do not contain relational operators or logical negation, as we can just choose the last rule for expressions whenever the **binop** rules do not match. However, we can not in the same way omit non-sequential (*e.g.*, bitwise) **and** and **or**, as these have a different effect (*i.e.*, they always evaluate both arguments).

We have, in the above, used two different nonterminals for conditions and expressions, with some overlap between these and consequently ambiguity in the grammar. It is possible to resolve this ambiguity by rewriting the grammar and get two non-overlapping syntactic categories in the abstract syntax. Another solution is to join the two nonterminals into one, *e.g.*, Exp and use two different translation functions for this: Whenever an expression is translated, the translation function most appropriate for the context is chosen. For example, `if-then-else` will choose a translation function similar to $Trans_{Cond}$ while assignment will choose a one similar to the current $Trans_{Exp}$.

Suggested exercises: 7.3.

7.8 Advanced control statements

We have, so far, shown translation of simple conditional statements and loops, but some languages have more advanced control features. We will briefly discuss how such can be implemented.

Goto and labels. Labels are stored in a symbol table that binds each to a corresponding label in the intermediate language. A jump to a label will generate a GOTO statement to the corresponding intermediate-language label. Unless labels are declared before use, an extra pass may be needed to build the symbol table before the actual translation. Alternatively, an intermediate-language label can be chosen and an entry in the symbol table be created at the first occurrence of the label even if it is in a jump rather than a declaration. Subsequent jumps or declarations of that label will use the intermediate-language label that was chosen at the first occurrence. By setting a mark in the symbol-table entry when the label is declared, it can be checked that all labels are declared exactly once.

The scope of labels can be controlled by the symbol table, so labels can be local to a procedure or block.

Break/exit. Some languages allow exiting loops from the middle of the loop-body by a `break` or `exit` statement. To handle these, the translation function for statements must have an extra inherited parameter which is the label that a `break` or `exit` statement must jump to. This attribute is changed whenever a new loop is entered. Before the first loop is entered, this attribute is undefined. The translation function should check for this, so it can report an error if a `break` or `exit` occurs outside loops. This should, rightly, be done during type-checking (see chapter 6), though.

C's `continue` statement, which jumps to the start of the current loop, can be handled similarly.

Case-statements. A case-statement evaluates an expression and selects one of several branches (statements) based on the value of the expression. In most languages, the case-statement will be exited at the end of each of these statements. In this case, the case-statement can be translated as an assignment that stores the value of the expression followed by a nested `if-then-else` statement, where each branch of the case-statement becomes a `then-branch` of one of the `if-then-else` statements (or, in case of the default branch, the final `else-branch`).

In C, the default is that *all* case-branches following the selected branch are executed unless the case-expression (called `switch` in C) is explicitly terminated with a `break` statement (see above) at the end of the branch. In this case, the case-statement can still be translated to a nested `if-then-else`, but the branches of

these are now GOTO's to the code for each case-branch. The code for the branches is placed in sequence after the nested `if-then-else`, with `break` handled by GOTO's as described above. Hence, if no explicit jump is made, one branch will fall through to the next.

7.9 Translating structured data

So far, the only values we have used are integers and booleans. However, most programming languages provide floating-point numbers and structured values like arrays, records (structs), unions, lists or tree-structures. We will now look at how these can be translated. We will first look at floats, then at one-dimensional arrays, multi-dimensional arrays and finally other data structures.

7.9.1 Floating-point values

Floating-point values are, in a computer, typically stored in a different set of registers than integers. Apart from this, they are treated the same way we treat integer values: We use temporary variables to store intermediate expression results and assume the intermediate language has binary operators for floating-point numbers. The register allocator will have to make sure that the temporary variables used for floating-point values are mapped to floating-point registers. For this reason, it may be a good idea to let the intermediate code indicate which temporary variables hold floats. This can be done by giving them special names or by using a symbol table to hold type information.

7.9.2 Arrays

We extend our example language with one-dimensional arrays by adding the following productions:

$$\begin{aligned} Exp &\rightarrow Index \\ Stat &\rightarrow Index := Exp \\ Index &\rightarrow \mathbf{id}[Exp] \end{aligned}$$

Index is an array element, which can be used the same way as a variable, either as an expression or as the left part of an assignment statement.

We will initially assume that arrays are zero-based (*i.e.* the lowest index is 0).

Arrays can be allocated statically, *i.e.*, at compile-time, or *dynamically*, *i.e.*, at run-time. In the first case, the *base address* of the array (the address at which index 0 is stored) is a compile-time constant. In the latter case, a variable will contain the base address of the array. In either case, we assume that the symbol table for variables binds an array name to the constant or variable that holds its base address.

$Trans_{Exp}(Exp, vtable, ftable, place) = \text{case } Exp \text{ of}$	
$Index$	$(code_1, address) = Trans_{Index}(Index, vtable, ftable)$ $code_1 ++ [place := M[address]]$
$Trans_{Stat}(Stat, vtable, ftable) = \text{case } Stat \text{ of}$	
$Index := Exp$	$(code_1, address) = Trans_{Index}(Index, vtable, ftable)$ $t = \text{newvar}()$ $code_2 = Trans_{Exp}(Exp, vtable, ftable, t)$ $code_1 ++ code_2 ++ [M[address] := t]$
$Trans_{Index}(Index, vtable, ftable) = \text{case } Index \text{ of}$	
$id[Exp]$	$base = \text{lookup}(vtable, \text{getname}(id))$ $t = \text{newvar}()$ $code_1 = Trans_{Exp}(Exp, vtable, ftable, t)$ $code_2 = code_1 ++ [t := t * 4, t := t + base]$ $(code_2, t)$

Figure 7.9: Translation for one-dimensional arrays

Most modern computers are byte-addressed, while integers typically are 32 or 64 bits long. This means that the index used to access array elements must be multiplied by the size of the elements (measured in bytes), *e.g.*, 4 or 8, to find the actual offset from the base address. In the translation shown in figure 7.9, we use 4 for the size of integers. We show only the new parts of the translation functions for *Exp* and *Stat*.

We use a translation function $Trans_{Index}$ for array elements. This returns a pair consisting of the code that evaluates the address of the array element and the variable that holds this address. When an array element is used in an expression, the contents of the address is transferred to the target variable using a memory-load instruction. When an array element is used on the left-hand side of an assignment, the right-hand side is evaluated, and the value of this is stored at the address using a memory-store instruction.

The address of an array element is calculated by multiplying the index by the size of the elements and adding this to the base address of the array. Note that *base* can be either a variable or a constant (depending on how the array is allocated, see below), but since both are allowed as the second operator to a **binop** in the intermediate language, this is no problem.

Allocating arrays

So far, we have only hinted at how arrays are allocated. As mentioned, one possibility is static allocation, where the base-address and the size of the array are

known at compile-time. The compiler, typically, has a large address space where it can allocate statically allocated objects. When it does so, the new object is simply allocated after the end of the previously allocated objects.

Dynamic allocation can be done in several ways. One is allocation local to a procedure or function, such that the array is allocated when the function is entered and deallocated when it is exited. This typically means that the array is allocated on a stack and popped from the stack when the procedure is exited. If the sizes of locally allocated arrays are fixed at compile-time, their base addresses are constant offsets from the stack top (or from the *frame pointer*, see chapter 10) and can be calculated from this at every array-lookup. However, this does not work if the sizes of these arrays are given at run-time. In this case, we need to use a variable to hold the base address of each array. The address is calculated when the array is allocated and then stored in the corresponding variable. This can subsequently be used as described in *TransIndex* above. At compile-time, the array-name will in the symbol table be bound to the variable that at runtime will hold the base-address.

Dynamic allocation can also be done globally, so the array will survive until the end of the program or until it is explicitly deallocated. In this case, there must be a global address space available for run-time allocation. Often, this is handled by the operating system which handles memory-allocation requests from all programs that are running at any given time. Such allocation may fail due to lack of memory, in which case the program must terminate with an error or release memory enough elsewhere to make room. The allocation can also be controlled by the program itself, which initially asks the operating system for a large amount of memory and then administrates this itself. This can make allocation of arrays faster than if an operating system call is needed every time an array is allocated. Furthermore, it can allow the program to use *garbage collection* to automatically reclaim arrays that are no longer in use.

These different allocation techniques are described in more detail in chapter 12.

Multi-dimensional arrays

Multi-dimensional arrays can be laid out in memory in two ways: *row-major* and *column-major*. The difference is best illustrated by two-dimensional arrays, as shown in Figure 7.10. A two-dimensional array is addressed by two indices, *e.g.*, (using C-style notation) as $a[i][j]$. The first index, i , indicates the *row* of the element and the second index, j , indicates the *column*. The first row of the array is, hence, the elements $a[0][0]$, $a[0][1]$, $a[0][2]$, ... and the first column is $a[0][0]$, $a[1][0]$, $a[2][0]$, ...³

In row-major form, the array is laid out one row at a time and in column-major

³Note that the coordinate system, following computer-science tradition, is rotated 90° clockwise compared to mathematical tradition.

	1st column	2nd column	3rd column	...
1st row	a[0][0]	a[0][1]	a[0][2]	...
2nd row	a[1][0]	a[1][1]	a[1][2]	...
3rd row	a[2][0]	a[2][1]	a[2][2]	...
⋮	⋮	⋮	⋮	⋮

Figure 7.10: A two-dimensional array

form it is laid out one column at a time. In a 3×2 array, the ordering for row-major is

$$a[0][0], a[0][1], a[1][0], a[1][1], a[2][0], a[2][1]$$

For column-major the ordering is

$$a[0][0], a[1][0], a[2][0], a[0][1], a[1][1], a[2][1]$$

If the size of an element is *size* and the sizes of the dimensions in an *n*-dimensional array are $dim_0, dim_1, \dots, dim_{n-2}, dim_{n-1}$, then in row-major format an element at index $[i_0][i_1] \dots [i_{n-2}][i_{n-1}]$ has the address

$$base + ((\dots (i_0 * dim_1 + i_1) * dim_2 \dots + i_{n-2}) * dim_{n-1} + i_{n-1}) * size$$

In column-major format the address is

$$base + ((\dots (i_{n-1} * dim_{n-2} + i_{n-2}) * dim_{n-3} \dots + i_1) * dim_0 + i_0) * size$$

Note that column-major format corresponds to reversing the order of the indices of a row-major array. *i.e.*, replacing i_0 and dim_0 by i_{n-1} and dim_{n-1} , i_1 and dim_1 by i_{n-2} and dim_{n-2} , and so on.

We extend the grammar for array-elements to accommodate multi-dimensional arrays:

$$\begin{aligned} Index &\rightarrow \mathbf{id}[Exp] \\ Index &\rightarrow Index[Exp] \end{aligned}$$

and extend the translation functions as shown in figure 7.11. This translation is for row-major arrays. We leave column-major arrays as an exercise.

With these extensions, the symbol table must return both the base-address of the array and a list of the sizes of the dimensions. Like the base-address, the dimension sizes can either be compile-time constants or variables that at run-time will hold the sizes. We use an auxiliary translation function $Calc_{Index}$ to calculate the position of

$Trans_{Exp}(Exp, vtable, ftable, place) = \text{case } Exp \text{ of}$	
$Index$	$(code_1, address) = Trans_{Index}(Index, vtable, ftable)$ $code_1 ++ [place := M[address]]$

$Trans_{Stat}(Stat, vtable, ftable) = \text{case } Stat \text{ of}$	
$Index := Exp$	$(code_1, address) = Trans_{Index}(Index, vtable, ftable)$ $t = \text{newvar}()$ $code_2 = Trans_{Exp}(Exp_2, vtable, ftable, t)$ $code_1 ++ code_2 ++ [M[address] := t]$

$Trans_{Index}(Index, vtable, ftable) =$	
$(code_1, t, base, []) = Calc_{Index}(Index, vtable, ftable)$ $code_2 = code_1 ++ [t := t * 4, t := t + base]$ $(code_2, t)$	

$Calc_{Index}(Index, vtable, ftable) = \text{case } Index \text{ of}$	
$id[Exp]$	$(base, dims) = \text{lookup}(vtable, \text{getname}(id))$ $t = \text{newvar}()$ $code = Trans_{Exp}(Exp, vtable, ftable, t)$ $(code, t, base, \text{tail}(dims))$
$Index[Exp]$	$(code_1, t_1, base, dims) = Calc_{Index}(Index, vtable, ftable)$ $dim_1 = \text{head}(dims)$ $t_2 = \text{newvar}()$ $code_2 = Trans_{Exp}(Exp, vtable, ftable, t_2)$ $code_3 = code_1 ++ code_2 ++ [t_1 := t_1 * dim_1, t_1 := t_1 + t_2]$ $(code_3, t_1, base, \text{tail}(dims))$

Figure 7.11: Translation of multi-dimensional arrays

an element. In $Trans_{Index}$ we multiply this position by the element size and add the base address. As before, we assume the size of elements is 4.

In some cases, the sizes of the dimensions of an array are not stored in separate variables, but in memory next to the space allocated for the elements of the array. This uses fewer variables (which may be an issue when these need to be allocated to registers, see chapter 9) and makes it easier to return an array as the result of an expression or function, as only the base-address needs to be returned. The size information is normally stored just before the base-address so, for example, the size of the first dimension can be at address $base - 4$, the size of the second dimension at $base - 8$ and so on. Hence, the base-address will always point to the first element of the array no matter how many dimensions the array has. If this strategy is used, the necessary dimension-sizes must be loaded into variables when an index is calculated. Since this adds several extra (somewhat costly) loads, optimising compilers often try to re-use the values of previous loads, *e.g.*, by doing the loading once outside a loop and referring to variables holding the values inside the loop.

Index checks

The translations shown so far do not test if an index is within the bounds of the array. Index checks are fairly easy to generate: Each index must be compared to the size of (the dimension of) the array and if the index is too big, a jump to some error-producing code is made. If the comparison is made on unsigned numbers, a negative index will look like a very large index. Hence, a single conditional jump is inserted at every index calculation.

This is still fairly expensive, but various methods can be used to eliminate some of these tests. For example, if the array-lookup occurs within a `for`-loop, the bounds of the loop-counter may guarantee that array accesses using this variable will be within bounds. In general, it is possible to make an analysis that finds cases where the index-check condition is subsumed by previous tests, such as the exit test for a loop, the test in an `if-then-else` statement or previous index checks. See section 11.4 for details.

Non-zero-based arrays

We have assumed our arrays to be zero-based, *i.e.*, that the indices start from 0. Some languages allow indices to be arbitrary intervals, *e.g.*, -10 to 10 or 10 to 20 . If such are used, the starting-index must be subtracted from each index when the address is calculated. In a one-dimensional array with known size and base-address, the starting-index can be subtracted (at compile-time) from base-address instead. In a multi-dimensional array with known dimensions, the starting-indices can be multiplied by the sizes of the dimensions and added together to form a

single constant that is subtracted from the base-address instead of subtracting each starting-index from each index.

7.9.3 Strings

Strings are usually implemented in a fashion similar to one-dimensional arrays. In some languages (*e.g.* C or pre-ISO-standard Pascal), strings *are* just arrays of characters.

However, strings often differ from arrays in that the length is not static, but can vary at run-time. This leads to an implementation similar to the kind of arrays where the length is stored in memory, as explained in section 7.9.2. Another difference is that the size of a character is typically one byte (unless 16-bit Unicode characters are used), so the index calculation does not multiply the index by the size (as this is 1).

Operations on strings, *e.g.*, concatenation and substring extraction, are typically implemented by calling library functions.

7.9.4 Records/structs and unions

Records (structs) have many properties in common with arrays. They are typically allocated in a similar way (with a similar choice of possible allocation strategies), and the fields of a record are typically accessed by adding an offset to the base-address of the record. The differences are:

- The types (and hence sizes) of the fields may be different.
- The field-selector is known at compile-time, so the offset from the base address can be calculated at this time.

The offset for a field is simply the sum of the sizes of all fields that occur before it. For a record-variable, the symbol table for variables must hold the base-address and the offsets for each field in the record. The symbol table for types must hold the offsets for every record type, such that these can be inserted into the symbol table for variables when a record of this type is declared.

In a union (sum) type, the fields are not consecutive, but are stored at the same address, *i.e.*, the base-address of the union. The size of an union is the maximum of the sizes of its fields.

In some languages, union types include a *tag*, which identifies which variant of the union is stored in the variable. This tag is stored as a separate field before the union-fields. Some languages (*e.g.* Standard ML) enforce that the tag is tested when the union is accessed, others (*e.g.* Pascal) leave this as an option to the programmer.

Suggested exercises: 7.8.

7.10 Translating declarations

In the translation functions used in this chapter, we have several times required that “The symbol table must contain ...”. It is the job of the compiler to ensure that the symbol tables contain the information necessary for translation. When a name (variable, label, type, *etc.*) is declared, the compiler must keep in the symbol-table entry for that name the information necessary for compiling any use of that name. For scalar variables (*e.g.*, integers), the required information is the intermediate-language variable that holds the value of the variable. For array variables, the information includes the base-address and dimensions of the array. For records, it is the offsets for each field and the total size. If a type is given a name, the symbol table must for that name provide a description of the type, such that variables that are declared to be that type can be given the information they need for their own symbol-table entries.

The exact nature of the information that is put into the symbol tables will depend on the translation functions that use these tables, so it is usually a good idea to write first the translation functions for *uses* of names and then translation functions for their declarations.

Translation of function declarations will be treated in chapter 10.

7.10.1 Example: Simple local declarations

We extend the statement language by the following productions:

$$\begin{aligned} Stat &\rightarrow Decl ; Stat \\ Decl &\rightarrow \text{int } \mathbf{id} \\ Decl &\rightarrow \text{int } \mathbf{id}[\mathbf{num}] \end{aligned}$$

We can, hence, declare integer variables and one-dimensional integer arrays for use in the following statement. An integer variable should be bound to a location in the symbol table, so this declaration should add such a binding to *vtable*. An array should be bound to a variable containing its base address. Furthermore, code must be generated for allocating space for the array. We assume arrays are heap allocated and that the intermediate-code variable *HP* points to the first free element of the (upwards growing) heap. Figure 7.12 shows the translation of these declarations. When allocating arrays, no check for heap overflow is done.

7.11 Further reading

A comprehensive discussion about intermediate languages can be found in [35].

$Trans_{Stat}(Stat, vtable, ftable) = \text{case } Stat \text{ of}$	
$Decl ; Stat_1$	$(code_1, vtable_1) = Trans_{Decl}(Decl, vtable)$ $code_2 = Trans_{Stat}(Stat_1, vtable_1, ftable)$ $code_1 ++ code_2$
$Trans_{Decl}(Decl, vtable) = \text{case } Decl \text{ of}$	
$\text{int } \mathbf{id}$	$t_1 = \text{newvar}()$ $vtable_1 = \text{bind}(vtable, \text{getname}(\mathbf{id}), t_1)$ $([], vtable_1)$
$\text{int } \mathbf{id}[\mathbf{num}]$	$t_1 = \text{newvar}()$ $vtable_1 = \text{bind}(vtable, \text{getname}(\mathbf{id}), t_1)$ $([t_1 := HP, HP := HP + (4 * \text{getvalue}(\mathbf{num}))], vtable_1)$

Figure 7.12: Translation of simple declarations

Functional and logic languages often use high-level intermediate languages, which are in many cases translated to lower-level intermediate code before emitting actual machine code. Examples of such intermediate languages can be found in [23], [8] and [6].

Another high-level intermediate language is the Java Virtual Machine [29]. This language has single instructions for such complex things as calling virtual methods and creating new objects. The high-level nature of JVM was chosen for several reasons:

- By letting common complex operations be done by single instructions, the code is smaller, which reduces transmission time when sending the code over the Internet.
- JVM was originally intended for interpretation, and the complex operations also helped reduce the overhead of interpretation.
- A program in JVM is *validated* (essentially type-checked) before interpretation or further translation. This is easier when the code is high-level.

Exercises

Exercise 7.1

Use the translation functions in figure 7.3 to generate code for the expression $2+g(x+y, x*y)$. Use a *vtable* that binds x to v_0 and y to v_1 and an *ftable* that binds g to $_g$. The result of the expression should be put in the intermediate-code variable r (so the *place* attribute in the initial call to $Trans_{Exp}$ is r).

Exercise 7.2

Use the translation functions in figures 7.5 and 7.6 to generate code for the statement

```
x:=2+y;
if x<y then x:=x+y;
repeat
  y:=y*2;
  while x>10 do x:=x/2
until x<y
```

use the same *vtable* as in exercise 7.1.

Exercise 7.3

Use the translation functions in figures 7.5 and 7.8 to translate the following statement

```
if x<=y && !(x=y || x=1)
then x:=3
else x:=5
```

use the same *vtable* as in exercise 7.1.

Exercise 7.4

De Morgan's law tells us that $!(p \ || \ q)$ is equivalent to $(!p) \ \&\& \ (!q)$. Show that these generate identical code when compiled with *Trans_{Cond}* from figure 7.8.

Exercise 7.5

Show that, in any code generated by the functions in figures 7.5 and 7.8, every IF-THEN-ELSE instruction will be followed by one of the target labels.

Exercise 7.6

Extend figure 7.5 to include a `break`-statement for exiting loops, as described in section 7.8, *i.e.*, extend the statement syntax by

$$Stat \rightarrow \text{break}$$

and add a rule for this to *Trans_{Stat}*. Add whatever extra attributes you may need to do this.

Exercise 7.7

We extend the statement language with the following statements:

$$\begin{aligned} Stat &\rightarrow \mathbf{labelid} : \\ Stat &\rightarrow \mathbf{goto labelid} \end{aligned}$$

for defining and jumping to labels.

Extend figure 7.5 to handle these as described in section 7.8. Labels have scope over the entire program (statement) and need not be defined before use. You can assume that there is exactly one definition for each used label.

Exercise 7.8

Show translation functions for multi-dimensional arrays in column-major format.

Hint: Starting from figure 7.11, it may be a good idea to rewrite the productions for *Index* so they are right-recursive instead of left-recursive, as the address formula for column-major arrays groups to the right. Similarly, it is a good idea to reverse the list of dimension sizes, so the size of the rightmost dimension comes first in the list.

Exercise 7.9

When statements are translated using the functions in figure 7.5, it will often be the case that the statement immediately following a label is a GOTO statement, *i.e.*, we have the following situation:

$$\begin{aligned} \text{LABEL } &label_1 \\ \text{GOTO } &label_2 \end{aligned}$$

It is clear that any jump to *label₁* can be replaced by a jump to *label₂*, and that this will result in faster code. Hence, it is desirable to do so. This is called jump-to-jump optimisation, and can be done after code-generation by a post-process that looks for these situations. However, it is also possible to avoid most of these situations by modifying the translation function.

This can be done by adding an extra inherited attribute *endlabel*, which holds the name of a label that can be used as the target of a jump to the end of the code that is being translated. If the code is immediately followed by a GOTO statement, *endlabel* will hold the target of this GOTO rather than a label immediately preceding this.

- a) Add the *endlabel* attribute to $Trans_{Stat}$ from figure 7.5 and modify the rules so *endlabel* is exploited for jump-to-jump optimisation. Remember to set *endlabel* correctly in recursive calls to $Trans_{Stat}$.

b) Use the modified $Trans_{Stat}$ to translate the following statement:

```
while x>0 do {
  x := x-1;
  if x>10 then x := x/2
}
```

The curly braces are used as disambiguators, though they are not part of grammar 7.4.

Use the same $vtable$ as exercise 7.1 and use `endlab` as the *endlabel* for the whole statement.

Exercise 7.10

In figure 7.5, `while` statements are translated in such a way that every iteration of the loop executes an unconditional jump (`GOTO` in addition to the conditional jumps in the loop condition).

Modify the translation so each iteration only executes the conditional jumps in the loop condition, *i.e.*, so an unconditional jump is saved in every iteration. You may have to add an unconditional jump outside the loop.

Exercise 7.11

Logical conjunction is associative: $p \wedge (q \wedge r) \Leftrightarrow (p \wedge q) \wedge r$.

Show that this also applies to the sequential conjunction operator `&&` when translated as in figure 7.8, *i.e.*, that $p \&\&(q \&\&r)$ generates the same code (up to renaming of labels) as $(p \&\&q) \&\&r$.

Exercise 7.12

Figure 7.11 shows translation of multi-dimensional arrays in row-major layout, where the address of each element is found through multiplication and addition. On machines with fast memory access but slow multiplication, an alternative implementation of multi-dimensional arrays is sometimes used: An array with dimensions $dim_0, dim_1, \dots, dim_n$ is implemented as a one-dimensional array of size dim_0 with pointers to dim_0 different arrays each of dimension dim_1, \dots, dim_n , which again are implemented in the same way (until the last dimension, which is implemented as a normal one-dimensional array of values). This takes up more room, as the pointer arrays need to be stored as well as the elements. But array-lookup can be done using only addition and memory accesses.

- a) Assuming pointers and array elements need four bytes each, what is the total number of bytes required to store an array of dimensions $dim_0, dim_1, \dots, dim_n$?
- b) Write translation functions for array-access in the style of figure 7.11 using this representation of arrays. Use addition to multiply numbers by 4 for scaling indices by the size of pointers and array elements.