

Programming Language Technology

Exam, 08 April 2021 at 08.30 – 12.30 on Canvas

Course codes: Chalmers DAT151, GU DIT231. As re-exam, also DAT150 and DIT230.
Exam supervision: Andreas Abel. Questions may be asked in Zoom breakout room, by email (<mailto:andreas.abel@gu.se>, subject: **PLT exam**) or telephone (+46 31 772 1731).

Exam review: Modalities will be announced later.

Allowed aids:

- All exam questions have to be solved *individually*.
- *No communication* of any form is permitted during the exam, including conversation, telephone, email, chat, asking questions in internet fora etc.
- All course materials can be used, including the book, lecture notes, previous exam solutions, own lab solution, etc. Any material copied verbatim should be marked as *quotation with reference* to the source.
- Publicly available *documentation* on the internet may be consulted freely to prepare the solution. *Small* portions of code and text from publicly available resources may be reused in the solution if *clearly marked* as quotation and *properly referencing* the source.

Any violation of the above rules and further common sense rules applicable to an examination, including *plagiarism* or *sharing solutions* with others, will lead to immediate failure of the exam (grade U), and may be subject to further persecution.

Grading scale: VG = 5, G = 4/3, U.

To pass, you need to deliver complete answers to two out of questions 1-3. (Typos, bugs, and minor omissions are not a problem as long as your answer demonstrates good understanding of the subject matter.) For a Chalmers grade 4 you need complete answers to all of the questions 1-3. A VG/5 requires excellent answers on questions 1-3.

Submission instructions:

- Please answer the questions in English.
- The solutions need to be submitted as one **.zip** archive, named according to schema **FirstName LastName Personnummer.zip**. Checklist:
 - Lovelace.cf
 - Sum.adb
 - Question2.{txt|md|pdf|...}
 - Question3.{txt|md|pdf|...}
 - (other relevant files)

In the following, a fragment *Lovelace* of the Ada programming language is described, in its syntax and semantics. Two example programs, **Primes.adb** and **Factorial.adb** are included to clarify the specification. In the exam, you are asked to describe a grammar, a type checker, and a compiler for Lovelace.

1. A *program* consists of:
 - (a) imports, in Lovelace fixed to the two lines

```
with Ada.Integer_Text_IO;  
use  Ada.Integer_Text_IO;
```
 - (b) main header: **procedure** *identifier* **is**
 - (c) a list of *function definitions*,
 - (d) a list of main *variable declarations*,
 - (e) a main *block*.

Running a program will execute the variable declarations and the statements of the block (from which functions can be called).

2. A *variable declaration* is a non-empty comma-separated list of identifiers a colon, a *type*, colon-equals, an initializing expression, and a semicolon. The scope of the initializing expression are the functions and variables declared before, *not* including the variable(s) we are just initializing. Note: If we declare several variables, the initializing expression is evaluated again *for each* variable.
3. A *type* is **Integer** or **Boolean**.
4. A *function definition* consists of:
 - (a) header: **function** *identifier* *parenthesized-parameters* **return type is**,
 - (b) a list of local *variable declarations*,
 - (c) body: a *block* for the function, terminated by a semicolon.

The *parameters* are a non-empty semicolon-separated list of *parameter declarations* each of which consists of: *identifier* colon *type*.

A function needs to be called (see *function call* expression) with the correct number of arguments of the correct type. The call will execute the block with parameters initialized to their respective argument value and local variables initialized to their value (see ??). The execution of the function ends when a **return** statement is encountered.

The joint list of parameters and local variables may not have any duplicates.

5. A *block* for a function or procedure with name *identifier* is started by keyword **begin** and ended by **end identifier** semicolon. In between is a non-empty list of *statements*, each terminated by a semicolon.
6. A statement can be one of the following. The typing and execution of the statements is like in C/C++/Java unless noted otherwise.

- (a) A return statement: **return** *expression*. Returns from the current function with the value of the *expression*.
 - (b) An assignment: *identifier* colon-equals *expression*.
 - (c) A conditional: **if** *expression* **then** *statements*, optionally followed by **else** *statements*, terminated by **end if**.
 - (d) A while-loop: **while** *expression* **loop** *statements* **end loop**.
 - (e) A for-loop: **for** *identifier* **in** *expression* dot-dot *expression* **loop** *statements* **end loop**. The for-loop declares a new variable *identifier* of type **Integer**, the so-called loop variable. This variable is only in scope in the *statements* and it may shadow other variables. The first expression denotes the initial value of the loop variable and the second expression the final value. Both values are integers and computed before the loop starts. If the final value is below the initial value, the loop is not executed. Otherwise, the loop variable is set to the initial value. The statement is executed, and the loop variable is incremented by one. The actions of the previous sentence are repeated as long as the loop variable is not larger than the final value.
 - (f) A print statement: **put** followed by a parenthesized expression of type **Integer**. Prints the value and a newline character to the standard output.
7. An expression can be one of the following. Typing and interpretation of expressions is like in C/C++/Java unless noted otherwise.
- (a) A variable: *identifier*.
 - (b) A boolean constant **true** or **false**.
 - (c)
 - (d) A *function call*: *identifier* followed by a parenthesized non-empty comma-separated list of *expressions*.
 - (e) A parenthesized expression.
 - (f) A infix binary operation: *expression operator expression*. All operators are left associative. Operators come in four binding strengths:
 - i. Multiplicative operators, bind strongest:
 - integer multiplication *****,
 - integer division **div**,
 - integer remainder **mod**.
 - ii. Additive operators, next in binding strength:
 - integer addition **+**,
 - integer subtraction **-**.
 - iii. Relational operators, but-last in strength: Equality operators **=** (equal) and **/=** (not equal) and integer comparison operators **<**, **<=**, **>**, and **>=** with the usual meaning.
 - iv. Short-circuiting logical operators, least in binding strength:
 - v. boolean conjunction **and** **then**,
 - vi. boolean disjunction **or** **else**.

Operators are always applied to two expressions of the same type. Equality operators apply to booleans and to integers. Like in C/C++/Java, boolean conjunction and disjunction are short-circuiting, i.e., if the left operand determines the value of the operation, the right operand is not evaluated.

8. An *identifier* starts with a letter, followed by a possibly empty sequence of letters, digits, and underscores. (Note: this is different from BNFC's **Ident** token type.)
9. An *integer literal* is a non-empty sequence of digits.

Comments start with double-dash (--) and last until the end of the line.

An identifier is *never* in scope before its declaration. The detailed scoping rules are:

1. Functions are in scope *after* their declaration: in their own body, in functions defined later, and in the main block. There is no mutual recursion. All functions must have distinct names.
2. The parameters and local variables of a function must be distinct. They are only in scope in the corresponding initializing expressions (see above) and the function body. They may shadow function identifiers.
3. The main variables (as well as all functions) are in scope in the main block. The names of the main variables must be distinct from each other and from the functions.

```
-- Factorial.adb
```

```
with Ada.Integer_Text_IO;  
use  Ada.Integer_Text_IO;
```

```
procedure Factorial is
```

```
    function factorial (n : Integer) return Integer is  
    begin  
        if n < 2 then  
            return 1;  
        else  
            return n * factorial(n - 1);  
        end if;  
    end factorial;
```

```
    n : Integer := 7;  
begin  
    put(factorial(n));  
end Factorial;
```

```

-- Primes.adb

with Ada.Integer_Text_IO;
use  Ada.Integer_Text_IO;

procedure Primes is

    function prime (n : Integer) return Boolean is
        i : Integer := 3;
    begin
        if n <= 2          then return (n = 2); end if;
        if n mod 2 = 0    then return false;  end if;
        while i * i <= n loop
            if n mod i = 0 then return false;  end if;
            i := i + 2;
        end loop;
        return true;
    end prime;

    -- Test 100 numbers for primality, starting with 1.
    lower : Integer := 1;
    upper : Integer := 100 + lower - 1;

begin
    for n in lower .. upper loop
        if prime(n) then
            put(n);
        end if;
    end loop;
end Primes;

```

Question 1 (Grammar)

1. Write an Lovelace program **Sum.adb** that computes and prints the sum of the integers from 1 to 100. This program should contain a function **sum** with two integer parameters determining the range (e.g. “from 1 to 100”), and the main block should call this function with arguments 1 and 100.
2. Write a labelled BNF grammar for Lovelace in a file **Lovelace.cf** and create a parser from this grammar using BNFC. For the best evaluation, the parser should be free of conflicts (shift/reduce and reduce/reduce).
3. Recommended: Test your parser on **Primes.adb**, **Factorial.adb** and **Sum.adb**.

Deliverables: files **Lovelace.cf** and **Sum.adb**.

SOLUTION: Grammar (file Lovelace.cf):

```
-- BNFC Grammar of Lovelace, a fragment of Ada

Prg.      Program ::=
           "with" "Ada" "." "Integer_Text_IO" ";"
           "use"  "Ada" "." "Integer_Text_IO" ";"
           "procedure" Id "is" [Def] Body ";" ;
terminator Def      ";"";

Bdy.      Body ::= [VarDecl] Block;
terminator VarDecl ";"";

-- # Declarations

FunDef.   Def ::= "function" Id "(" [ParDecl] ")" "return" Type "is" Body;
separator nonempty ParDecl ";"";

ParDecl.  ParDecl ::= Id ":" Type;

-- ## Variable declarations

VarDecl.  VarDecl ::= [Id] ":" Type "!=" Exp;
separator nonempty Id ","";

-- # Types

TBool.    Type ::= "Boolean";
TInt.     Type ::= "Integer";

internal
  TVoid.   Type ::= "Void";

-- # Blocks

Blck.     Block ::= "begin" [Stm] "end" Id;
terminator nonempty Stm ";"";

-- # Statements

SReturn.  Stm ::= "return" Exp;
SAssign.  Stm ::= Id "!=" Exp;
SIf.      Stm ::= "if" Exp "then" [Stm] "end" "if";
SIfElse.  Stm ::= "if" Exp "then" [Stm] "else" [Stm] "end" "if";
SWhile.   Stm ::= "while" Exp "loop" [Stm] "end" "loop";
SFor.     Stm ::= "for" Id "in" Exp ".." Exp "loop" [Stm] "end" "loop";
SPut.     Stm ::= "put" "(" Exp ")";
```

```

-- # Expressions

EVar.      Exp4 ::= Id;
ETrue.     Exp4 ::= "true";
EFalse.    Exp4 ::= "false";
EInt.      Exp4 ::= Integer;
ECall.     Exp4 ::= Id "(" [Exp] ")";
EMul.      Exp3 ::= Exp3 MulOp Exp4;
EAdd.      Exp2 ::= Exp2 AddOp Exp3;
ECmp.      Exp1 ::= Exp1 CmpOp Exp2;
ELog.      Exp  ::= Exp  LogOp Exp1;

coercions  Exp 4;
separator  nonempty Exp ", ";

-- # Operators

OTimes.    MulOp ::= "*"  ;
ODiv.      MulOp ::= "div";
OMod.      MulOp ::= "mod";

OPlus.     AddOp ::= "+"  ;
OMinus.    AddOp ::= "-"  ;

OLt.       CmpOp ::= "<"  ;
OLtEq.     CmpOp ::= "<=" ;
OGt.       CmpOp ::= ">"  ;
OGtEq.     CmpOp ::= ">=" ;
OEq.       CmpOp ::= "="  ;
ONEq.      CmpOp ::= "/="  ;

-- ## Short-cutting logical operators

OAnd.      LogOp ::= "and" "then";
OOr.       LogOp ::= "or"  "else" ;

-- # Identifiers

token      Id letter (letter | digit | '_' )*;

-- # Comments

comment    "--";

```

Summation program (file `Sum.adb`):

```
-- Sum.adb

with Ada.Integer_Text_IO; -- put for Integer
use  Ada.Integer_Text_IO;

procedure Sum is

  function sum (lower : Integer; upper : Integer) return Integer is
    sum : Integer := 0;
  begin
    for i in lower .. upper loop
      sum := sum + i;
    end loop;
    return sum;
  end sum;

begin
  put(sum(1,100));
end Sum;
```


Question 2 (Type checker): Write a specification of a type checker for the Lovelace language of Question 1. The type checker receives an abstract syntax tree of a Lovelace program and shall throw an error if any of the scoping or typing rules are violated.

Deliverable: **submit a text document** with name **Question2** (plus file extension) that contains the specification. The text document can be a plain text file possibly using markup (like markdown) or a PDF.

The specification should have the following structure:

- A. State. Describe the components of the *state* of the type checker and how these components are implemented, i.e., which data structure (like list, map, integer...) is used for each component.
- B. Initialization and run: Describe how the state is initialized and how the type checker (??) is started (i.e., which arguments are given to the type checker).
- C. Syntax-directed traversal: Describe the type checker: Write an explanation how each relevant Lovelace construct (expression, statement, block, declaration, ...) is checked (or its type inferred). You may use judgements and rules or pseudo-code or *precise* language.
- D. API (optional): If you used helper functions to manipulate the state in item ??, describe them here.

The specification can use the names from your BNFC grammar.

The specification should be written in a high-level but self-contained way so that an *informed outsider* can implement the type checker easily following your specification. An informed outsider shall be a person who has very good programming skills and good familiarity with programming language technology in general, but no specific knowledge about the Lovelace language nor access to the course material.

The specification will be judged on clarity and correctness.

SOLUTION:

A. State

The state of the type checker has two following components:

- Signature **sig**: a finite map from identifiers to function types (**Def**).
- Environment **env**: a finite map from identifiers to types.
- Return type **returnType** that is **Nothing** when we are checking the main procedure.

A function type consists of a *list of types* holding the types of the arguments and a *return type*.

B. Initialization and run

The type checker receives a **Program** consisting of a list of function definitions, a list of variable declarations, and a main block.

1. State component **sig** is initialized to the empty map.
2. The functions are checked in order first to last.
3. Component **env** is reset to the empty map and **returnType** to **Nothing**.
4. The variables declarations are checked in order first to last.
5. The statements of the main block are checked.

C. Type Checker

The type checker is a collection of mutually recursive procedures and functions that check function and variable declarations and statements and infer or check types of expressions. At any point a type error can be thrown.

A function definition, consisting of a function name, a parameter list, a return type, a list of variable declarations, and a block with a list of statements is checked as follows:

1. If the function name is already a key in **sig**, throw error *duplicate function*. Otherwise, extend **sig** by a binding of the function name to the function type. The latter is a pair consisting of the list of the parameter types in left-to-right order, and the return type.
2. Set **returnType** to the return type.
3. Empty **env** and then add the parameters with their type one by one, from first to last. If a parameter is already in **env**, throw error *duplicate parameter*.
4. Check the variable declarations.
5. Check the statements.

A variable declaration, consisting of a list of variable names, a type, and an initializing expression is checked as follows:

1. Check the expression against the type.
2. If we are checking the main variable declarations, throw error *duplicate declaration* if one of the variable identifiers is bound in **sig**.
3. Bind each of the variables to the type in **env**. If the name of a variable is already present in **env**, throw error *duplicate variable*.

Statements are checked as follows:

- List of statements: in first-to-last order, check each statement.
- Return statement `return e`: If `returnType` is `Nothing`, throw error *illegal return*. Otherwise, check e against the `returnType`.
- Print statement `put(e)`: Check e against type `Integer`.
- Assignment `x := e`: Lookup the type t of x in `env`. If no type is found, throw error *unbound variable*. Otherwise, check e against t .
- Conditional `if e then s1 [else s2] end if`: Check e against type `Boolean`. Check statements s_1 and, if present, also statements s_2 .
- Loop statement `while e loop s end loop`: Check e against type `Boolean`. Check statements s .
- Loop `for x in e1 .. e2 loop s end loop`: Check e_1 and e_2 against type `Integer`. Save `env`. Bind x to `Integer` in map `env`. Check statements s . Restore the old `env`.

To check an expression e against a type t , infer the type t' of e . If t' differs from t , throw error *type mismatch*.

The type of an expression is inferred as follows:

- Literals `true` and `false` have type `Boolean`.
- Numeric literals have type `Integer`.
- A variable x has the type as stored in `env`. If the variable is not present in this map, throw error *unbound variable*.
- A call to function f with arguments e_1, \dots, e_n is inferred as follows:
 1. If f is bound in `env`, throw error *illegal call to variable*.
 2. If f is not bound in `sig`, throw error *unbound function*; otherwise, we have the types of its parameters and its return type which is the type of the call.
 3. If the number of arguments does not match the number of parameters, throw error *wrong number of arguments*.
 4. Check each argument against its respective parameter type.
- Arithmetic operations are of type `Integer`, provided the two expressions check against type `Integer`.
- Logical operations are of type `Boolean`, provided the two expressions check against type `Boolean`.
- Equality operators are of type `Boolean`, provided both arguments infer to the same type.
- The other comparison operators are of type `Boolean`, provided both arguments check against type `Integer`.

Question 3 (Compilation): Specify a compiler from Lovelace to JVM. The compiler takes a type-correct abstract syntax tree of a Lovelace program as input and translates this into Jasmin method definitions which are printed to the standard output.

Deliverable: **submit a text document** with name **Question3** (plus file extension) that contains the specification. Instructions analogous to Question 2 apply. In particular, follow the same structure: A. State, B. Initialization and run, C. Compilation schemes, D. API.

Restrictions of the task:

1. The compiler does not have to output a full Jasmin class file, only the methods corresponding to the defined Lovelace functions and a **main** method for the main block. (You may assume that no Lovelace function is called **main**.)
2. You need not output **.limit** pragmas (stack/locals).
3. You may simply use the Lovelace function identifiers for the corresponding Jasmin method names.
4. You can assume a Java method that can be called to output an integer.
5. You need not care about Java modifiers like **public** or **static**.
6. It is sufficient to treat one logical, one arithmetical, and one comparison operator.
7. Choose *one* of **if-then** or **if-then-else** or **while**.

However, the compiler needs to output proper JVM instructions (not pseudo machine code).

Good luck!

SOLUTION:

A. State

The state of the compiler consists of the following components:

1. A finite map **context** from identifiers to natural numbers (local variable addresses).
2. A natural number **nextAddress** denoting the next free slot in the JVM variable store of the currently compiled method.
3. A stream **labels** of so far unused label names. Elements are taken from this stream whenever a new label name is needed.

Since we are directly printing the generated Jasmin to the standard output, we need not store any generated code. Further, since we pretend that we can use Lovelace function identifiers for the corresponding Jasmin methods, we need no function “signature”.

B. Initialization and run

Given a type- and scope-correct Lovelace program, compilation proceeds as follows:

1. The component **labels** is initialized to an infinite stream of distinct label names, e.g. **L0**, **L1**, **...**
2. Each function definition is **compiled** (see below)
3. Output: **.method main**
4. The **context** is reset to an empty map and **nextAddress** to 0.

5. Each variable declaration is **compiled** (see below).
6. The main block is **compiled** (see below).
7. Output: **return**
8. Output: **.end method**

C. Compiler

The compiler is specified as an overloaded procedure **compile** in pseudo-code acting on Lovelace abstract syntax from Question 1.

-- Definitions

```
compile (FunDef f pars t body):
  emit (.method f)
  context      <- empty
  nextAddress <- 0
  for (x:t in pars): addVar x
  compile body
  emit (ldc 0)
  emit (ireturn)
  emit (.end method)
```

```
compile (Bdy vardecls block):
  for (d in vardecls): compile d
  compile block
```

```
compile (VarDcl xs t e):
  for (x in xs):
    compile e
    a <- addVar x
    emit (istore a)
```

-- Statements

```
compile (Blck ss):
  for (s : ss) compile s
```

```
compile (SReturn e):
  compile e
  emit (ireturn)
```

```
compile (SAssign x e):
  a <- lookupVar x
  compile e
  emit (istore a)
```

```
compile (SPut e):
```

```

compile e
emit (invokestatic put)

compile (SIf e s1):
done <- newLabel
compile e
emit (ifeq done)
compile s1
emit (done:)

compile (SIfElse e s1 s2):
else, done <- newLabel
compile e
emit (ifeq else)
compile s1
emit (goto done)
emit (else:)
compile s2
emit (done:)

compile (SWhile e s):
start, done <- newLabel
emit (start:)
compile e
emit (ifeq done)
compile s
emit (goto start)
emit (done:)

compile (SFor x e1 e2 s):
start, done <- newLabel
compile e2
compile e1
a <- addVar x
emit (start:)
emit (dup2)           -- copy bounds for comparison
emit (if_icmplt done)
emit (istore a)
compile s
emit (iload a)
emit (ldc 1)
emit (iadd)          -- bounds are again on top of stack
emit (goto start)
emit (done:)
emit pop2
releaseVar x

```

-- Expressions

```
compile (EInt i): emit (ldc i)
compile (ETrue) : emit (ldc 1)
compile (EFalse): emit (ldc 0)
```

```
compile (EId x):
  a <- lookupVar x
  emit (iload a)
```

```
compile (ECall f es):
  for (e : es) compile e
  emit (invokestatic f)
```

```
compile (EAdd e1 OMinus e2):
  compile e1
  compile e2
  emit (isub)
```

```
compile (ECmp e1 OLt e2):
  true, done <- newLabel
  compile e1
  compile e2
  emit (if_icmplt true)
  emit (ldc 0)
  emit (goto done)
  emit (true:)
  emit (ldc 1)
  emit (done:)
```

```
compile (EMul e1 OAnd e2):
  false, done <- newLabel
  compile e1
  emit (dup)
  emit (ifeq done) -- short-circuit if e1 is false
  emit (pop)
  compile e2
  emit (done:)
```

D. API

- `emit (text)`: Write *text* and newline to standard output.
- `addVar x`: Set address of *x* to `nextAddress` in `context`. Increase `nextAddress` by one. Return address of *x*.
- `releaseVar x`: Set `nextAddress` to value of *x* in `context`. Drop *x* from this map.
- `lookupVar x`: Return the address of *x* in map `context`.
- `newLabel`: Extract the next element from stream `labels`.