

Examen

vendredi 20 mai 2019

Tout document papier est autorisé. Les ordinateurs, les téléphones portables, comme tout autre moyen de communication vers l'extérieur, doivent être éteints. Le temps à disposition est de **2h**.

On recommande de *bien lire* l'énoncé d'un exercice avant de commencer à le résoudre.

Exercice 1 (Compilation à la main de OCaml). Donner un bytecode de OCamlrun équivalent à l'expression de OCaml suivante :

```
let a = 2 in let f g = g a in f (fun x -> x+1)
```

Exercice 2. Pour chacune des listes d'instructions suivantes, deviner l'expression de OCaml qui l'a générée :

| | |
|----------------|---------------|
| (i). | (ii). |
| | closure L2, 0 |
| | push |
| const 3 | const 3 |
| push | push |
| const 2 | closure L1, 0 |
| mulint | push |
| push | acc 2 |
| const 2 | appterm 2, 4 |
| push | L1: acc 0 |
| acc 1 | offsetint 1 |
| modint | return 1 |
| push | restart |
| const 0 | L2: grab 1 |
| eqint | acc 1 |
| branchifnot L1 | push |
| const 3 | acc 2 |
| return 2 | push |
| L1: const 4 | acc 2 |
| return 2 | apply 1 |
| | addint |
| | return 2 |

Exercice 3. Considerons une methode `f` dans une classe `MaClasse` qui est de la forme

```
class MaClasse{
  public static int f(int x, int y){
    ....
    ....
  }
}
```

Pour chacune des portions de code octet JVM suivantes :

1. détailler le fonctionnement de JVM, notamment l'évolution des variables et de la pile de calcul des blocs d'activation lorsque la méthode est appelée avec les paramètres 2 et 3.
2. Trouver des instructions JAVA pour f qui génèrent le byte-code.

1.

```
public static int f(int, int);
  descriptor: (II)I
  flags: ACC_PUBLIC, ACC_STATIC
  Code:
    stack=2, locals=2, args_size=2
       0: iload_0
       1: iload_1
       2: if_icmpge      9
       5: iload_1
       6: iload_0
       7: isub
       8: ireturn
       9: iload_0
      10: iload_1
      11: isub
      12: ireturn
```

2.

```
public static int f(int, int);
  descriptor: (II)I
  flags: ACC_PUBLIC, ACC_STATIC
  Code:
    stack=2, locals=3, args_size=2
       0: iload_1
       1: istore_2
       2: iload_0
       3: iconst_1
       4: if_icmple     17
       7: iinc          0, -1
      10: iload_2
      11: iload_2
      12: imul
      13: istore_2
      14: goto          2
      17: iload_2
      18: ireturn
```

3.

```
public static int f(int, int);
  descriptor: (II)I
```

```

flags: ACC_PUBLIC, ACC_STATIC
Code:
  stack=3, locals=2, args_size=2
    0: iload_0
    1: iconst_1
    2: if_icmpgt      7
    5: iload_1
    6: ireturn
    7: iload_0
    8: iinc           0, -1
   11: iload_1
   12: iload_1
   13: imul
   14: invokestatic  #2           // Method f:(II)I
   17: ireturn

```

Exercice 4 (Comparaison OCamlrun et JVM). Comparer la pile d'exécution de JVM avec celle de OCamlrun. Détailler en particulier les points en communs et les différences. (Répondre en max 5 lignes).

Annexe OCamlrun

acc n Peeks the $n+1$ -th element of the stack and puts it into the accumulator.

envacc n Sets the accumulator to the field of index n of the environment.

apply n Sets extraArgs to $n-1$. Sets pc to the code value of the accumulator. Then sets the environment to the value of the accumulator.

appterm n, s Slides the n top elements from the stack towards bottom of $s - n$ positions. Then sets pc to the code value of the accumulator, the environment to the accumulator, and increases extraArgs by $n-1$.

return n Pops n elements from the stack. If extraArgs is strictly positive then it is decremented, pc is set to the code value of the accumulator, and the environment is set to the value of the accumulator. Otherwise, three values are popped from the stack and assigned to pc, environment and extraArgs.

restart Computes n , the number of arguments, as the size of the environment minus 2. Then pushes elements of the environment from index $n - 1$ to 2 onto the stack. Environment is set to the element of index 1 of the environment and extraArgs is increased by n .

grab n If extraArgs is greater than or equal to n , then extraArgs is decreased by n . Otherwise, creates a closure of extraArgs+3 elements in the accumulator. Code of this closure is set to pc - 3, element of index 1 is set to the environment and other elements are set to values popped from the stack. Then pc, environment, and extraArgs are popped from the stack.

closure ofs, n If n is greater than zero then the accumulator is pushed onto the stack. A closure of $n + 1$ elements is created into the accumulator. The code value of the closure is set to pc + ofs. Then, the other elements of the closure are set to values popped from the stack.

branchifnot ofs Performs an conditional jump by adding ofs to pc if the accumulator is zero.

eqint Sets the accumulator to a non-zero value or to zero whether the accumulator is equal to the value popped from the stack or not.

const n Sets the accumulator to n .

addint Sets the accumulator to the sum of the accumulator and the value popped from the stack.

mulint Sets the accumulator to the product of the accumulator by the value popped from the stack.

modint Sets the accumulator to the modulo of the accumulator by the value popped from the stack.

gtint Sets the accumulator to a non-zero value or to zero whether the accumulator is greater than the value popped from the stack or not.

offsetint ofs Adds ofs to the accumulator.

pop n Pops n elements from the stack.

push Pushes the accumulator onto the stack.

Annexe JVM

aload_n The n must be an index into the local variable array of the current frame. The local variable at n must contain a reference. The objectref in the local variable at n is pushed onto the operand stack.

iaload The arrayref must be of type reference and must refer to an array whose components are of type long. The index must be of type int. Both arrayref and index are popped from the operand stack. The long value in the component of the array at index is retrieved and pushed onto the operand stack.

istore_n The n must be an index into the local variable array of the current frame. The value on the top of the operand stack must be of type int. It is popped from the operand stack, and the value of the local variable at n is set to value.

areturn The objectref must be of type reference and must refer to an object of a type that is assignment compatible with the type represented by the return descriptor of the current method. If the current method is a synchronized method, the monitor entered or reentered on invocation of the method is updated and possibly exited as if by execution of a monitorexit instruction in the current thread. If no exception is thrown, objectref is popped from the operand stack of the current frame and pushed onto the operand stack of the frame of the invoker. Any other values on the operand stack of the current method are discarded.

The interpreter then reinstates the frame of the invoker and returns control to the invoker.

goto branchbyte1 branchbyte2 The unsigned bytes branchbyte1 and branchbyte2 are used to construct a signed 16-bit branchoffset, where branchoffset is $(\text{branchbyte1} \ll 8) \mid \text{branchbyte2}$. Execution proceeds at that offset from the address of the opcode of this goto instruction. The target address must be that of an opcode of an instruction within the method that contains this goto instruction.

iadd Both value1 and value2 must be of type int. The values are popped from the operand stack. The int result is $\text{value1} + \text{value2}$. The result is pushed onto the operand stack.

iconst_n Push the int constant n (-1, 0, 1, 2, 3, 4 or 5) onto the operand stack.

idiv Both value1 and value2 must be of type int. The values are popped from the operand stack. The int result is the value of the Java programming language expression $\text{value1} / \text{value2}$. The result is pushed onto the operand stack.

isub Both value1 and value2 must be of type int. The values are popped from the operand stack. The int result is $\text{value1} - \text{value2}$. The result is pushed onto the operand stack.

if_icmp<cond> branchbyte1 branchbyte2 Both value1 and value2 must be of type int. They are both popped from the operand stack and compared. All comparisons are signed. The possible comparisons are : eq, ne, lt, le, gt, ge.

If the comparison succeeds, the unsigned branchbyte1 and branchbyte2 are used to construct a signed 16-bit offset, where the offset is calculated to be $(\text{branchbyte1} \ll 8) \mid \text{branchbyte2}$. Execution then proceeds at that offset from the address of the opcode of this if_icmp<cond> instruction. The target address must be that of an opcode of an instruction within the method that contains this if_icmp<cond> instruction.

Otherwise, execution proceeds at the address of the instruction following this `if_icmp<cond>` instruction.

if<cond> branchbyte1 branchbyte2 The value must be of type `int`. It is popped from the operand stack and compared against zero. Cfr `if_icmp<cond>`.

ifnonnull branchtype1 branchtype2 The value must be of type reference. It is popped from the operand stack. If value is not null, the unsigned `branchbyte1` and `branchbyte2` are used to construct a signed 16-bit offset, where the offset is calculated to be $(\text{branchbyte1} \ll 8) \mid \text{branchbyte2}$. Execution then proceeds at that offset from the address of the opcode of this `ifnonnull` instruction. The target address must be that of an opcode of an instruction within the method that contains this `ifnonnull` instruction.

Otherwise, execution proceeds at the address of the instruction following this `ifnonnull` instruction.

iinc index const The index is an unsigned byte that must be an index into the local variable array of the current frame. The `const` is an immediate signed byte. The local variable at `index` must contain an `int`. The value `const` is first sign-extended to an `int`, and then the local variable at `index` is incremented by that amount.

imul Both `value1` and `value2` must be of type `int`. The values are popped from the operand stack. The `int` result is `value1 * value2`. The result is pushed onto the operand stack.

invokestatic indexbyte1 indexbyte2 The unsigned `indexbyte1` and `indexbyte2` are used to construct an index into the run-time constant pool of the current class, where the value of the index is $(\text{indexbyte1} \ll 8) \mid \text{indexbyte2}$. The run-time constant pool item at that index must be a symbolic reference to a method, which gives the name and descriptor of the method as well as a symbolic reference to the class in which the method is to be found. The named method is resolved. The resolved method must not be an instance initialization method or the class or interface initialization method. It must be static, and therefore cannot be abstract.

On successful resolution of the method, the class that declared the resolved method is initialized if that class has not already been initialized.

The operand stack must contain `nargs` argument values, where the number, type, and order of the values must be consistent with the descriptor of the resolved method.

....

If the method is not native, the `nargs` argument values are popped from the operand stack. A new frame is created on the Java Virtual Machine stack for the method being invoked. The `nargs` argument values are consecutively made the values of local variables of the new frame, with `arg1` in local variable 0 (or, if `arg1` is of type `long` or `double`, in local variables 0 and 1) and so on. Any argument value that is of a floating-point type undergoes value set conversion prior to being stored in a local variable. The new frame is then made current, and the Java Virtual Machine `pc` is set to the opcode of the first instruction of the method to be invoked. Execution continues with the first instruction of the method.

ireturn cfr `areturn`.

iload_n cfr `aload_n`.