

Compilers [Fall 2016] Programming Assignment VI

Objectives

1. To learn basic rules for generating RISC-style machine code equivalent to high-level abstract-syntax trees (ASTs).
2. To implement a code generator that converts Diminished Java (DJ) ASTs into Diminished Instruction Set Machine (DISM) programs.
3. To test a complete `dj2dism` compiler.

Due Date: Friday, December 2, 2016 (at 11:59pm).

Machine Details

Complete this assignment by yourself on the following CSEE network computers: `c4lab01`, `c4lab02`, ..., `c4lab20`. These machines are physically located in ENB 220. You can connect to these machines from home using SSH. (Example: Host name: `c4lab01.csee.usf.edu` Login ID and Password: <your NetID username and password>) You are responsible for ensuring that your programs compile and execute properly on these machines.

Assignment Description

This assignment asks you to complete your `dj2dism` compiler by implementing the code-generation phase of compilation.

Begin by downloading and studying the header file for the code-generation module. This file, `codegen.h`, is posted at <http://www.csee.usf.edu/~ligatti/compilers-16/as6>. Then implement the `generateDISM` method, which is declared in `codegen.h`, in a new file called `codegen.c`.

If you are completing this assignment by extending your Assignment V implementation (as is recommended), you will also have to make a few minor modifications to your `dj.y`:

1. The `codegen.h` header file needs to be included.
2. The main method in `dj.y`, which is the main method of the entire compiler, should now enforce that the compiler is invoked with the command “`dj2dism s.dj`”, where “`s.dj`” is the source-program filename and “`s`” can be any string.
3. After typechecking the input file “`s.dj`”, the main method of `dj.y` should open the file “`s.dism`”, into which the target program will be output. (If “`s.dism`” already exists, your compiler should overwrite the existing “`s.dism`” file.)
4. Having opened and obtained a file pointer for “`s.dism`”, the main method of `dj.y` should invoke the code generator by calling the `generateDISM` method.
5. Before exiting, the main method of `dj.y` should close the “`s.dism`” target-code file.

Notes on Code Generation in `dj2dism`

1. Your generated DISM-code file must contain exactly one instruction per line. This will make it easy to find the precise instruction at which your program halts. (The `sim-dism` virtual machine reports the instruction number at which the DISM program halted; if `sim-dism` reports that the program halted at instruction n and the code file

contains exactly one instruction per line, then we can find the halting instruction at code-file line number $n+1$.)

2. In order to satisfy the constraint that every line of the generated “.disasm” file must contain exactly one DISM instruction, you may sometimes find it convenient to write instructions of the form “#LABEL: mov 0 0” or “mov 0 0 ; comment”. In such instructions, the “mov 0 0” is used as an empty operation (i.e., a *no-op*) so that we can output a DISM label and/or a comment on an otherwise empty line of code.
3. Your code generator must implement the disjunction operator (`||`) in the standard short-circuit style. File *good6.dj* illustrates this operator’s dynamic semantics.
4. Your code generator must check and halt before dereferencing null-valued objects. Files *good20.dj*, *good21.dj*, and *good22.dj* illustrate the required behavior.

Hints on Organizing the Code Generator

You may find it helpful to organize your *codegen.c* such that it contains the following (in addition to the definition of the `generateDISM` method declared in *codegen.h*):

```
#define MAX_DISM_ADDR 65535

/* Global for the DISM output file */
FILE *fout;

/* Global to remember the next unique label number to use */
unsigned int labelNumber = 0;

/* Declare mutually recursive functions (defs and docs appear below) */
void codeGenExpr(ASTree *t, int classNumber, int methodNumber);
void codeGenExprs(ASTree *expList, int classNumber, int methodNumber);

/* Print a message and exit under an exceptional condition */
void internalCGerror(char *msg)

/* Using the global classesST, calculate the total number of fields,
   including inherited fields, in an object of the given type */
int getNumObjectFields(int type)

/* Generate code that increments the stack pointer */
void incSP()

/* Generate code that decrements the stack pointer */
void decSP()

/* Output code to check for a null value at the top of the stack.
   If the top stack value (at M[SP+1]) is null (0), the DISM code
   output will halt. */
void checkNullDereference()

/* Generate DISM code for the given single expression, which appears
   in the given class and method (or main block).
   If classNumber < 0 then methodNumber may be anything and we assume
   we are generating code for the program's main block. */
void codeGenExpr(ASTree *t, int classNumber, int methodNumber)
```

```

/* Generate DISM code for an expression list, which appears in
   the given class and method (or main block).
   If classNumber < 0 then methodNumber may be anything and we assume
   we are generating code for the program's main block. */
void codeGenExprs(ASTree *expList, int classNumber, int methodNumber)

/* Generate DISM code as the prologue to the given method or main
   block. If classNumber < 0 then methodNumber may be anything and we
   assume we are generating code for the program's main block. */
void genPrologue(int classNumber, int methodNumber)

/* Generate DISM code as the epilogue to the given method or main
   block. If classNumber < 0 then methodNumber may be anything and we
   assume we are generating code for the program's main block. */
void genEpilogue(int classNumber, int methodNumber)

/* Generate DISM code for the given method or main block.
   If classNumber < 0 then methodNumber may be anything and we assume
   we are generating code for the program's main block. */
void genBody(int classNumber, int methodNumber)

/* Map a given (1) static class number, (2) a method number defined
   in that class, and (3) a dynamic object's type to:
   (a) the dynamic class number and (b) the dynamic method number that
   actually get called when an object of type (3) dynamically invokes
   method (2).
   This method assumes that dynamicType is a subtype of staticClass. */
void getDynamicMethodInfo(int staticClass, int staticMethod,
    int dynamicType, int *dynamicClassToCall, int *dynamicMethodToCall)

/* Emit code for the program's vtable, beginning at label #VTABLE.
   The vtable jumps (i.e., dispatches) to code based on
   (1) the dynamic calling object's address (at M[SP+4]),
   (2) the calling object's static type (at M[SP+3]), and
   (3) the static method number (at M[SP+2]). */
void genVTable()

```

More Hints

My *codegen.c* file is 618 lines of code (98 of which are comments/whitespace). As with Assignment V, this is a relatively challenging assignment, requiring you to write several hundred lines of code. Please budget your time accordingly.

Many examples of valid DJ programs, on which you can test your compiler, are posted at <http://www.cse.usf.edu/~ligatti/compilers-16/as1/dj/examples/good/>. A complete solution to this assignment will correctly compile all of the valid DJ programs into equivalent, valid DISM code that can be executed in the sim-dism virtual machine. As usual, though, we will grade your compiler's performance on DJ programs that have not been distributed to the class.

Grading

Your grade on this assignment is determined by the level to which you implement the desired functionality:

- *Level I:* generateDISM correctly generates DISM code for all DJ programs that (1) declare no classes and (2) declare no local variables. Expect to write about 250-300 lines of code for this level (or about 200-250 without counting comments and whitespace).
- *Level II:* generateDISM correctly generates DISM code for all DJ programs that declare no classes. For this level, expect to write about 70 lines of code beyond that of Level I.
- *Level III:* generateDISM correctly generates DISM code for all DJ programs.

Undergraduate students will earn: 80% credit for reaching Level I, 100% credit for reaching Level II, and 110% credit for reaching Level III.

Graduate students will earn: 60% credit for reaching Level I, 75% credit for reaching Level II, and 100% credit for reaching Level III.

All students will earn +15% extra credit for having implemented a complete Level-III dj2dism compiler from scratch (i.e., without using any of the provided object-code files) during this semester.

Compilation of the Compiler

Use the following commands to compile the complete dj2dism compiler from scratch:

```
> flex dj.l
> bison -v dj.y
> gcc dj.tab.c ast.c symtbl.c typecheck.c codegen.c -o dj2dism
```

Alternatively, you can download working versions of the lexer, parser, AST, symbol-table, and type-checker modules as object-code files *dj.tab.o*, *ast.o*, *symtbl.o*, and *typecheck.o* from <http://www.cse.usf.edu/~ligatti/compilers-16/as6/>. With these object-code files, and the header files *ast.h*, *symtbl.h*, *typecheck.h*, and *codegen.h*, you can compile dj2dism with:

```
> gcc dj.tab.o ast.o symtbl.o typecheck.o codegen.c -o dj2dism
```

Example Executions

Your compiler should continue to report lexical, syntactic, and semantic errors in the source program before exiting. As with previous assignments, you need not report multiple errors. If there are no lexical, syntactic, or semantic errors, your compiler generates DISM code.

Your compiler should enforce that the given source-program filename ends in “.dj”:

```
> ./dj2dism
Usage: dj2dism file.dj
> ./dj2dism good99.dj
ERROR: could not open file good99.dj
> ./dj2dism good1.dism
Error: Input filename does not match *.dj
```

The compiler produces no output (besides the target-code file) when compiling a valid DJ program:

```
> ./dj2dism good1.dj
>
```

After compiling *good1.dj*, we can find the dj2dism-generated code in file *good1.dism*:

```
mov 7 65535 ; initialize FP
mov 6 65535 ; initialize SP
mov 5 1 ; initialize HP
mov 0 0 ; ALLOCATE STACK SPACE FOR MAIN LOCALS
mov 0 0 ; BEGIN METHOD/MAIN-BLOCK BODY
mov 1 0
str 6 0 1 ; M[SP] <- R[r1] (a nat literal)
mov 1 1
sub 6 6 1 ; SP--
bgt 6 5 #0 ; branch if SP>HP
mov 1 77 ; error code 77 => out of stack memory
hlt 1 ; out of stack memory! (SP <= HP)
#0: mov 0 0
hlt 0 ; NORMAL TERMINATION AT END OF MAIN BLOCK
```

(Of course, the DISM files your compiler generates must only be functionally equivalent—rather than identical to—the DISM files my compiler generates. The one exception is that your compiler may produce DISM code that halts with whatever codes you wish. Mine uses the code shown above, plus error code 66 to indicate that insufficient heap memory exists for a new object, 88 to indicate a null-pointer dereference, and 99 to indicate that a vtable entry wasn't found).

Debugging Hint: When trying to understand and correct the behavior of the DISM code your compiler generates, you may find it helpful to insert *ptn* instructions into the generated DISM code. For example, to see the address to which the SP refers at some point, you could insert *lod 4 6 0* and *ptn 4* at that point. You may also find it helpful to execute some of your DISM programs in *sim-dism* with the debug-mode turned on.

As expected, we can execute file *good1.dism* in the *sim-dism* virtual machine:

```
> ./sim-dism good1.dism
Simulation completed with code 0 at PC=13.
```

Of all the provided *good*.dj* files, my dj2dism generates the most DISM code for *good9*.

```
> ./dj2dism good9.dj
> ./sim-dism good9.dism
1
2
4
5
6
1
2
4
4
0
33
Simulation completed with code 0 at PC=2070.
```

We've waited all semester to compile and run our *sub.dj* programs from Assignment I.

```
> ./dj2dism sub.dj
> ./sim-dism sub.dism
Enter a natural number: 1
Enter a natural number: 2
Enter a natural number: 3
Enter a natural number: 4
Enter a natural number: 0
Enter a natural number: 1
Enter a natural number: 4
Enter a natural number: 0
1
0
Enter a natural number: 2
Enter a natural number: 3
Enter a natural number: 0
1
1
Enter a natural number: 3
Enter a natural number: 2
Enter a natural number: 0
0
0
Enter a natural number: 0
Simulation completed with code 0 at PC=665.
```

Notice how much smaller—and more efficient—our hand-written *sub.disms* were! My hand-written *sub.dism* is 53 instructions, but my compiled *sub.dism* is 1422 instructions.

Submission Notes

- Type the following pledge as an initial comment in your *codegen.c* file: “I pledge my Honor that I have not cheated, and will not cheat, on this assignment.” Type your name after the pledge. Not including this pledge will lower your grade 50%.
- Upload and submit your *codegen.c* file in Canvas, or if you've implemented the entire Level-III compiler by yourself, upload the full compiler as a zip or tar file.
- You may submit your assignment in Canvas as many times as you like; we'll grade your latest submission.
- When submitting your code in Canvas, include with your submission a comment indicating at which level we should grade your assignment: I, II, or III. If you don't indicate a level, we'll grade your submission as a Level-III submission.
- For every day that your assignment is late (up to 3 days), your grade reduces 10%.
- To make it easier for our teaching assistant to read and evaluate your code, use spaces rather than tabs in your code and avoid long lines of code (I try to limit lines to 80 characters in width).
- Your programs will be graded on both correctness and style, so include good comments, well-chosen variable names, etc. in your programs.
- *codegen.c*—and the DISM code it outputs—must be well commented and formatted.