Universitetet i Oslo
Institutt for Informatikk

Reliable Systems
Martin Steffen

INF 5110: Compiler construction

Spring 2022 — **Series 2** — 2. 2. 2022

Topic: Context free grammars (Exercises with hints for solution)

Issued: 2. 2. 2022

This exercise set covers more than one lecture. It's about grammars, and partly for the lectures about *parsing*. We might not be able to cover it within 2 hours.

Exercise 1 (First- and follow sets) Compute the *First* and *Follow*-sets for the grammar Figure 1.

$$
\begin{array}{rcl}
\mathit{exp} & \to & \mathit{term}\ \mathit{exp}' \\
\mathit{exp}' & \to & \mathit{addop}\ \mathit{term}\ \mathit{exp}' \mid \boldsymbol{\epsilon} \\
\mathit{addop} & \to & + \mid - \\
\mathit{term} & \to & \mathit{factor}\ \mathit{term}' \\
\mathit{term}' & \to & \mathit{mulop}\ \mathit{factor}\ \mathit{term}' \mid \boldsymbol{\epsilon} \\
\mathit{mulop} & \to & * \\
\mathit{factor} & \to & (\ \mathit{exp}\ ) \mid \textbf{number}
\end{array}
$$

Figure 1: Expression grammar (left-recursion removed)

Solution:

Listing 1: First sets

```
for all X ∈ Σ_T ∪ {ϵ} do
    First[X] := {X}
end;

for all non-terminals A do
  First[A] := {}
end
while there are changes to any First[A] do
  for each production A → X_1 ... X_n do
    k := 1;
    continue := true
    while continue = true and k ≤ n do
      First[A] := First[A] ∪ First[X_k] \ {ϵ}
      if ϵ ∉ First[X_k] then continue := false
      k := k + 1
    end;
    if   continue = true
```

```
18      then First[A] := First[A] ∪ {ϵ}
19    end;
20  end
```

We have learnt algorithms to do that. They are repeated in this exercise for easy reference in Listing 1 and 2.

Listing 2: Follow sets

```
1  Follow[S] := {$}
2  for all non-terminals A ≠ S do
3    Follow[A] := {}
4  end
5  while there are changes to any Follow-set do
6    for each production A → X1 ... Xn do
7      for each Xi which is a non-terminal do
8        Follow[Xi] := Follow[Xi] ∪ (First(Xi+1 ... Xn) \ {ϵ})
9        if ϵ ∈ First(Xi+1 Xi+2 ... Xn)
10       then Follow[Xi] := Follow[Xi] ∪ Follow[A]
11     end
12   end
13 end
```

But one can also informally try to figure it out (at the danger that one forgets some symbols, especially when dealing with nullable symbols and ϵ-productions, one has to watch out for those). In any case, the first-sets are simpler. Furthermore, the definition of the follow sets *depends* on determining the first-sets. Therefore, in such an exercise, one *always* starts with the first-sets, and only then attempts the follow-sets.

For the first sets: the main complication are *nullable symbols,* and the grammar has those, so we need to watch out.

In the "recursive definition" of the algorithm for the first sets one should notice that the definition defines the first set not just for non-terminals, but for terminals and ϵ, as well! That is different (a bit) from the tables the book goes through (of course only marginally). Actually I also think the algo in [1, Figure 4.6] does not work, as it ignores those cases. In that algo, never any terminal symbol is added. Let's therefore ignore this version.

Here's a "run" for the first-sets (the non-terminals are abbreviated in the obvious manner to safe space). Pass 4 is not shown, nothing would change there compared to pass 3 and the algo finishes.

non-term	pass 1	pass 2	pass 3
e			$($, **number**
e'		$\epsilon, +, -$	
e'	$\{\epsilon\}$	$\epsilon, +, -$	
a	$+$		
a	$+, -$		
t		$($, **number**	
t'		$\epsilon, *$	
t'	ϵ	$\epsilon, *$	
m	$*$		
f	$($		
f	$($, **number**		

Here's the "collapse" of the result, i.e., that's the end result.

non-term	*First*
exp	(, **number**
exp'	ϵ, +, −
addop	+, −
term	(, **number**
$term'$	ϵ, ∗
mulop	∗
factor	(, **number**

As a side remark: In the result, the primed non-terminals $term'$ and exp' are the ones containing ϵ in the first set, indicating that they are nullable. That's not a coincidence. The grammar we are dealing with is (one variant of) the expression grammar, namely one on which the algo for left recursion removal has been applied. The algo introduces new terminals with appropriate rules. By convention and to be a bit systematic, for a non-terminals A with left recursion, a new non-terminal A' is added decorated with a prime, and the corresponding massaged rules contain an ϵ-production. In the current example, the new symbols are $term'$ and exp'.

So far for the first sets. The algo for the *follow sets* is kind of similar in spirit to the first-set calculation, it's likewise a of "saturation" algorithm. But the "table" used to simulate the run is organized slightly differently, because the slots now correspond to the *right-hand sides* of the production, not the left-hand sides alone. The *initialization* of the algo start by having $Follow(exp) = \{\$\}$, all the rest is the empty set. As a general rule, we never add any ϵ to the follow-set, they simply don't belong there. However, we need to check if ϵ is in the considered *First*-set, as that plays a role.

We should also remember that we need in the construction not directly $First(X)$ but (corresponding) definition for sequences of symbols $First(X_1 \dots X_k)$. By "corresponding" it is meant that "in spirit", it's the same definition, only applied to sequences. The lecture provides the definition of the first set of a word (but not in pseudo code form). Basically it's an iterated application of the definition of the first-set for *one* symbol (and taking into account nullable symbols in the analogous way that the definition/code of first-sets of a symbol treats that already (when dealing with the right-hand side of a production). So it's nothing really new.

In the table, **F** stands for the follow set. I write also here += for the operation increasing the current value of a given set (representing `Follow[X] := Follow[X]` $\cup\{\dots\}$).

In the productions on the right-most column, I indicate the longest "post-fix" of the right-hand side which is nullable (that can be seen from the result of the first-sets). That's helpful, because those situations require special treatment by the algorithm. For example, in the situation

$$exp' \to addop\ term\ exp' \ ,$$

exp' is nullable, but the longer $term\ exp'$ is not.[1] To treat the non-terminal *addop* in the corresponding case is easier as the postfix after it is not nullable. To treat the non-terminal *term* is slightly more complex, as the exp' is nullable. One should not forget: to tread exp' leads to a post-fix of ϵ, which also counts as nullable, therefore the "speacial treatment" applies to exp', as well.

[1] Neither is $addop\ term\ exp'$, but that's not crucial.

production	init	pass 1	pass 2
$exp \to term\ exp'$			
$exp' \to addop\ term\ exp'$			
$exp' \to \epsilon$			
$addop \to +$			
$addop \to -$			
$term \to factor\ term'$			
$term' \to mulop\ factor\ term'$			
$term' \to \epsilon$			
$mulop \to *$			
$factor \to (\ exp\ )$			
$factor \to \mathbf{number}$			

production	init	pass 1	pass 2
$exp \to term\ exp'$	$\$$	$\mathbf{F}(term) = \{+,-\} \cup \{\$\}$ $\mathbf{F}(exp') = \{\$\}$	$\mathbf{F}(term)$ += $\{)\}$ $\mathbf{F}(exp')$ += $\{)\}$
$exp' \to addop\ term\ exp'$		$\mathbf{F}(addop) = \{(, \mathbf{number}\}$ $\mathbf{F}(term) = \{+,-,\$\} \cup \{\$\}$ $\mathbf{F}(exp') = \mathbf{F}(exp')$	
$exp' \to \epsilon$		$\mathbf{F}(\epsilon)$ = not a case	
$addop \to +$		+ is a terminal, nothing's done	
$addop \to -$		same	
$term \to factor\ term'$		$\mathbf{F}(factor) = \{*\} \cup \mathbf{F}(term)$ $\mathbf{F}(term') = \mathbf{F}(term) = \{\$,+,-\}$	$\mathbf{F}(factor)$ += $\{)\}$ $\mathbf{F}(term')$ += $\{)\}$
$term' \to mulop\ factor\ term'$		$\mathbf{F}(mulop) = \{(, \mathbf{number}\}$ $\mathbf{F}(factor) = \{*\} \cup \mathbf{F}(term') = \{*,+,-,\$\}$ $\mathbf{F}(term') =$ -	
$term' \to \epsilon$		not a case	
$mulop \to *$		* is a terminal	
$factor \to (\ exp\ )$		$\mathbf{F}(exp) = \{,)\}$	
$factor \to \mathbf{number}$		terminal	

The "run" of the algo resp. the table used to represent the run shows the difference in the structure of the follow-calculation. In contrast to the first-algorithm, there there are 3 loops. The outer loop, which correspond to the colums (the "passes"), The second loop going through all grammar productions, and finally, for each production, the inner loop going to the symbols on the right-hand side one by one; the latter one is missing in the first-algo.

non-term	**F**
exp	$\$,)$
exp'	$\$,)$
addop	$(, \mathbf{number}$
term	$\$,), +, -$
$term'$	$\$,), +, -$
mulop	$(, \mathbf{number}$
factor	$\$,), +, -, *$

Remark 2021 on traversals: The presented run of the first-set algorithms followed blindly the algo as presented in the lecture. In particular it made the following choice: it takes the productions of the grammar *in the presented order* and iterates them from the first to the last in one pass, then goes to the next pass and does it again (in the same order), etc. That's the way the hints here fill out the corresponding table.

In the live exercises (in ZOOM), the soluton was presented differently, not following the order of the rules as given. I don't 100% remember, perhaps it went through the rules in *reverse* order. That actually is a smarter way treating things, insofar that the solution stabilizes faster. If course, it depends on the grammar, what is the best traversal strategy is, the algorithm will

in general not know in advance. Indeed, the "algo" as presented does prescribe a particular order, it just states "for each rule, do the following". The solution here goes through them in order of appearance, the zoom-presentation choose a different, probably smarter order. I have also written some general remark about traversal orders and worklist algorithms in blog-form, but the text is not pensum, □

Exercise 2 (Nullable) Describe an algorithm that finds all nullable non-terminals without first finding the first-sets.

Solution: It should be clear that there *is* an algo, already from the fact that in the lecture we discussed one (in different representations). Apart from that: taking the original definition of nullability: from that one it's not immediately clear, because the definition states a condition like

$$A \Rightarrow^* \boldsymbol{\epsilon} \ .$$

However, given a grammar, one can intuitively make a "graph search", taking a non-terminal, and then look at "chains of productions" (which corresponds to a graph traversal) and see if one hits on $\boldsymbol{\epsilon}$ without passing through a terminal. Doing that for all non-terminals would answer the question.

However, that's not the smartest way to do it and that's not the way we handled the first- and follow-sets algorithmically (perhaps also compare the ϵ-closure). Very generally, one is better off to calculate "simultaneously" the first- and follow-sets, resp. in the task *here*, calculate the question of nullability simultaneously for all non-terminals (not one by one). That may (very roughly) be compared to, for instance the idea behind Dijsktra's shortest-path algo (covered for instance in INF2220): Also that algo works calculating shortest path *for all pairs of nodes* not simply for the one particular pair one might be interested in (and iterate that for all pairs).

Anyhow: a good solution to the task here is to do better than the nullability-one-by-one solution.

One way of doing it is: we arrange for a data structure nullable, which is a set of non-terminals (perhaps concretely some collection data structure), containing all non-terms of the given grammar which are nullable. To be more precise: that set contains all nullable non-terms *when the algo terminates.* During the run of the algo, it contains "the current knowledge or estimation" which of the non-terms are already known to be nullable. It's a crucial characteristic of this kind of algorithm (like the follow/first calculations as well, and many others) that the "estimation grows only in one direction", meaning that during the run of the algo, the current knowledge of nullable symbols, i.e, the corresponding set data structure, only *grows.* In each stage (or "pass" or iteration), the algo potentially *adds* further symbols, when detecting nullability of further symbols so far not (yet) known to be nullable. The algo terminates, if now *new* symbols are detected meaning basically that the set of nullable symbols does not *grow* any more (it *stabilizes*).

The *input* of the algo is a grammar in BNF. Output is the set of nullable non-terminal symbols, kept in, say *nullable*. The general structure is (as for the first and follow algos,and similar ones):

```
1  initialize nullable
2  while not stabilized
3    incease nullable (looking at the grammar, production after production)
4  end
5  return nullable
```

It is important to realize: going through the grammar "production after production" (or "terminal after terminal" or whatever) does *not* mean, that, after having done one sweep through all the productions, one is done. That is typically not the case. One stops if, going *repeatedly*

through the productions, it turns out that no new info can be learnt ("stabilization", "saturation"), *then* we terminate. See also how we filled in the tables with the "rounds" or "passes" when calculating the first- and follow-sets.

1. Initialization: add all non-terminals to *nullable*.
2. Repeat, until no more elements are added: if there's a non-terminal A with

$$A \to \ldots B_1 B_2 \ldots B_n$$

where *all* B_i are already (at the current stage) members of *nullable*, then add A to *nullable*.

Alternatively, more in line with the data structure of the first-algo: we could arrange the data structure in such a way that one has an boolean *array*, indexed by the non-terminals. It's of course the same. □

Exercise 3 (Associativity and precedence) Take the binary ops $+$, $-$, $*$, $/$ and $\uparrow$. Let's agree also on the following precedences and associativity

op	precedence	associativity
$+, -$	low	left assoc.
$*, /$	higher	left. assoc.
$\uparrow$	highest	right. assoc

Write an *unambiguous* grammar that captures the given precedences and associativies (of course, directly with a BNF grammar, without allowing yourself specifying those requirements as extra side-conditions).

Solution: We have learned in the lecture how it works, at least for the operators *except* the exponentiation.

The "flat grammar", simple, elegant, but with utter disrespect for associativity and precedence could look as follows (it's not required for the task):

$$\begin{array}{rcl} \mathit{exp} & \to & \textbf{number} \mid (\ \mathit{exp}\ ) \mid \mathit{exp}\ \mathit{op}\ \ \mathit{exp} \\ \mathit{op} & \to & + \mid - \mid * \mid \uparrow \end{array}$$

The ok grammar (without exponentiation) looked as follows

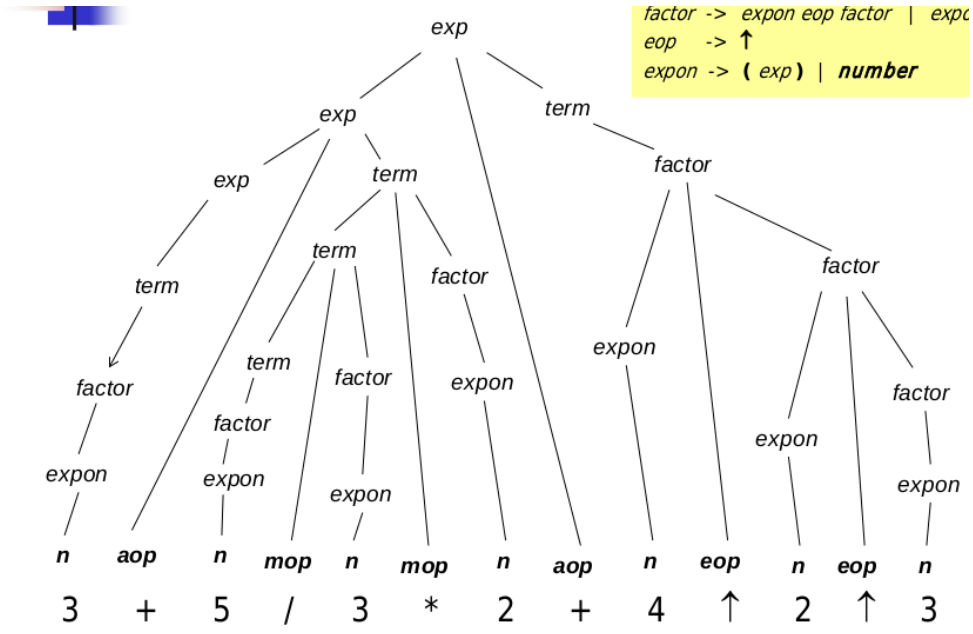
$$\begin{array}{rcl} \mathit{exp} & \to & \mathit{exp}\ \ \mathit{addop}\ \ \mathit{term} \mid \mathit{term} \\ \mathit{addop} & \to & + \mid - \\ \mathit{term} & \to & \mathit{term}\ \ \mathit{mulop}\ \ \mathit{factor} \mid \mathit{factor} \\ \mathit{mulop} & \to & * \\ \mathit{factor} & \to & (\ \mathit{exp}\ ) \mid \textbf{number} \end{array}$$

Now we have to include the exponentiation. In the same way we did for the terms and factors: we need a new non-terminal, here say *expon*, and furthermore, we need to get it right-associative. Therefore, the rule for *factor* is formulated using *right-recursion:* the *factor* occurs on the right, not on the left (as for *term* and *exp*)

$$\begin{array}{rcl} \mathit{exp} & \to & \mathit{exp}\ \mathit{addop}\ \mathit{term} \mid \mathit{term} \\ \mathit{addop} & \to & + \mid - \\ \mathit{term} & \to & \mathit{term}\ \mathit{mulop}\ \mathit{term} \mid \mathit{factor} \\ \mathit{mulop} & \to & * \\ \mathit{factor} & \to & \mathit{expon}\ \mathit{eop}\ \mathit{factor} \mid \mathit{expon} \\ \mathit{eop} & \to & \uparrow \\ \mathit{expon} & \to & (\ \mathit{exp}\ ) \mid \textbf{number} \end{array}$$

Here's an example CST (concrete syntax tree) for the expression.

$$3 + 5/3 * 2 + 4 \uparrow 2 \uparrow 3$$



□

Exercise 4 (Tiny grammar) For the grammar given answer the following questions:

- Is the grammar *unambiguious?*
- How can we change the grammar, so that TINY allows empty statements?
- How can we arrange it that semicolons are required *in between* statements, not *after* statements?
- What's the precedence and associativity of the different operators?

program	→	*stmts*
stmts	→	*stmts* **;** *stmt* | *stmt*
stmt	→	*if-stmt* | *repeat-stmt* | *assign-stmt*
	|	*read-stmt* | *write-stmt*
if-stmt	→	**if** *expr* **then** *stmt* **end**
	|	**if** *expr* **then** *stmt* **else** *stmt* **end**
repeat-stmt	→	**repeat** *stmts* **until** *expr*
assign-stmt	→	**identifier** **:=** *expr*
read-stmt	→	**read identifier**
write-stmt	→	**write** *expr*
expr	→	*simple-expr comparison-op simple-expr* | *simple-expr*
comparison-op	→	$<$ | $=$
simple-expr	→	*simple-expr addop term* | *term*
addop	→	$+$ | $-$
term	→	*term mulop factor* | *factor*
mulop	→	$*$ | $/$
factor	→	**(** *expr* **)** | **number** | **identifier**

Solution: [of Exercise 4]

ambiguity: If one is "smart" one might be aware that many properties of context-free grammar (not to mention context-free languages ...) are "hard", even undecidable. Whether or not a grammar is unambiguous or not is one of them: that is in general undecidable (and that immediately makes also the question if a context free language has an unambiguous grammar undecidable). Of course, for this particular grammar the question can be expected to have a definite answer (if only for the reason that otherwise the question would be meaninglessly hard ...). Still, those remarks about *general* undecidability could warn the reader that ambiguity of grammars is not a trivial matter (and in particular there's no algorithm that definitely can check the issue). That also might imply that showing that this particular grammar *is* unambiguous (should that be the case) is not easy: we would have to find an argument, that each word of terminals of the TINY-language has a unique parse tree, and there are infinitely many. That sounds hard (especially since we have not covered that question in the lecture). It seems more easy, should that be the case, to prove that the grammar is *ambiguous.* All we need is *one* word with 2 different parse trees. However, this general "smart" meta-consideration will lead us down the wrong path right here, the answer will be that the grammar seems unambiguous (but we won't make an attempt in proving it).

That makes it a priori plausible that the given grammar is ambiguous, we only are requested to nail down at least one "part" which causes ambiguity (there may be more than one reason, but we only need to find one).

Before doing so, and as a side remark: there are doable approaches which allows (indirectly) to prove unambiguity in some cases: for instance, if one could show that a given grammar is LL(k) or LR(k) or one of such classes, then we indirectly know the grammar is *unambiguous.*

Now: what's the reason for ambiguity? There might be the "usual suspects". Since the grammar is for a "real" (toy) language and not an arbitrary artificial grammar, we should concentrate on those usual suspects.

As we learned, a common source of ambiguity in grammars of programming languages is that, especially for binary constructs, the "grouping" is not fixed by the grammar (which then gives ambiuity for the parse tree). Typically that involves question of "associativity" or "precedence" (as discussed with the previous exercise). Another example we had, though not with plain binary operators but related nonetheless, is the phenomenon of "dangling else".

The dangling else is not a problem here.

Now, what about associativity. The grammar seem to make use of the techniques we used in the lecture, splitting up expressions into factors and terms. So, that part takes care of precedence and associativity of multiplication, addition, etc.

Now, remains the suspicious comparison expressions (comparing two simple expressions). At first sight that might look having problems with associativity. However: it would be ambiguous if we had a production for expressions like

$$\textit{exp} \to \textit{exp comparison-op exp} \tag{1}$$

or the simple expressions would be done like that:

$$\textit{simple-expr} \to \textit{simple-expr comparison-op simple-expr} \tag{2}$$

or similar arrangements. would lead to ambiguity.

So, since the grammar is not of any of the forms, it is very plausible that it's *unambiguous*.

Additional remarks: If one had an different version, for instance with a production of the form of equation (1) or (2), and thus the answer would be that the grammar is indeed ambiguous, in that case a good answer would give one example of an expression and give 2 different parse trees for that.

As a rather esoteric side remark: Technically, the question actually is not whether *parts* of the grammar are ambiguous (for instance the part starting with the expression non-terminal in the changed version), but if the *overall* grammar is ambiguous. So in principle, the grammar could be unambigous, even if expressions were parsed ambiguously, namely in the (weird) case when expressions can never be derived from the start symbol. Such grammars with unreachable parts or parts that can never be derived into a word of terminals are obviously "defective" in that there should never be "useless" symbols or productions in a grammar.

Anyway, the TINY grammar (and all decent grammars) does not suffer from such defects. Given an ambiguous grammar, pointing out and explaining the ambiguity (or lack thereof) for expressions is an answer good enough.

Empty statements: That's trivial, we simply add $stmt \to \epsilon$. This will make it possible to write statement sequences like `;;;;;`.

Semicolon as terminator, not separator: We could do

$$stmts \to stmts\ stmt\ \mathbf{;} \mid stmt\ \mathbf{;}$$

See also the language specification for the oblig, where the grammar is written in EBNF, not BNF.

Associativity and precedence:

op	precedence	associativity
$*, /$	highest	left
$+, -$	medium	left
$<, =$	lowest	non-associative

We can also think of semicolon and it's associativity. Except that it's not really counted among operators, typically.

Associativity and precedences for expressions and statements becomes more tricky if one deals with languages which "mixes" them. For instance, as is C-like languages, that every statement is *also* an expression. Then we have to think of

```
1    5 + x = 5
```

If for instance in Java

```
1  public class Stmtasexpr {
2      static int x = 23;
3      public static void main(String[] arg) {
4          //       x= 5 + x = 5;
5          x = 5 + (x = 5);
6      }
7  }
```

□

References

[1] K. Louden. *Compiler Construction, Principles and Practice.* PWS Publishing, 1997.