Regular Expressions and Automata using Miranda

Simon Thompson
Computing Laboratory
Univerisity of Kent at Canterbury

May 1995

Contents

1. Introduction

In these notes Miranda is used as a vehicle to introduce regular expressions, pattern matching, and their implementations by means of non-deterministic and deterministic automata.

As part of the material, we give an implementation of the ideas, contained in a set of files. References to this material are scattered through the text. The files can be obtained by following the instructions in

`http://www.ukc.ac.uk/computer_science/Miranda_craft/regExp.html`

This material is based on the treatment of the subject in [Aho *et. al.*], but provides full implementations rather than their pseudo-code versions of the algorithms.

The material gives an illustration of many of the features of Miranda, including polymorphism (the states of an NFA can be represented by objects of any type); modularisation (the system is split across a number of modules); higher-order functions (used in finding limits of processes, for example) and other features. A tutorial introduction to Miranda can be found in [Thompson].

The paper begins with definitions of regular expressions, and how strings are matched to them; this also gives our first Miranda treatment also. After describing the abstract data type of sets we define non-deterministic finite automata, and their implementation in Miranda. We then show how to build an NFA corresponding to each regular expression, and how such a machine can be optimised, first by transforming it into a deterministic machine, and then by minimising the state space of the DFA. We conclude with a discussion of regular definitions, and show how recognisers for strings matching regular definitions can be built.

2. Regular Expressions

Regular expressions are patterns which can be used to describe sets of strings of characters of various kinds, such as

- the identifiers of a programming language – strings of alphanumeric characters which begin with an alphabetic character;
- the numbers – integer or real – given in a programming language; and so on.

There are five sorts of pattern, or regular expression:

ε	This is the Greek character *epsilon*, which matches the empty string.
`x`	`x` is any character. This matches the character itself.
$(r_1 \mid r_2)$	r_1 and r_2 are regular expressions.
$(r_1 r_2)$	r_1 and r_2 are regular expressions.
`(r)*`	`r` is a regular expression.

Examples of regular expressions include `(a|(ba))`, `((ba)|(ε|(a)*))` and `hello`.

In order to give a more readable version of these, it is assumed that `*` binds more tightly than juxtaposition (*i.e.* $(r_1 r_2)$), and that juxtaposition binds more tightly than $(r_1 | r_2)$. This means that $r_1 r_2$`*` will mean $(r_1 (r_2)$`*`$)$, *not* $((r_1 r_2))$`*`, and that $r_1 | r_2 r_3$ will mean $r_1 | (r_2 r_3)$, *not* $(r_1 | r_2) r_3$.

A Miranda algebraic type representing regular expressions is given by

```
reg ::= Epsilon |
        Literal char |
        Or reg reg |
        Then reg reg |
        Star reg
```

This definition and those which follow can be found in the file `regexp.m`. The Miranda representations of `(a|(ab))` and `((ba)|(ε|(a)*))` are

```
Or (Literal 'a') (Then (Literal 'a') (Literal 'b'))
Or (Then (Literal 'b') (Literal 'a'))
   (Or Epsilon (Star (Literal 'a')))
```

respectively. In order to shorten these definitions we will usually define constant literals such as

```
a = Literal 'a'
b = Literal 'b'
```

so that the expressions above become

```
Or a (Then a b)        Or (Then b a) (Or Epsilon (Star a))
```

If we use the infix forms of `Or` and `Then`, `$Or` and `$Then`, they read

```
a $Or (a $Then b)
(a $Then b) $Or (Epsilon $Or (Star a))
```

Functions over the type of regular expressions are defined by recursion over the structure of the expression. Examples include

```
literals :: reg -> [char]

literals Epsilon = []
literals (Literal ch) = [ch]
literals (Or r1 r2) = literals r1 ++ literals r2
literals (Then r1 r2) = literals r1 ++ literals r2
literals (Star r) = literals r
```

which prints a list of the literals appearing in a regular expression, and

```
printRE :: reg -> [char]

printRE Epsilon = "@"
printRE (Literal ch) = [ch]
printRE (Or r1 r2)
  = "(" ++ printRE r1 ++ "|" ++ printRE r2 ++ ")"
printRE (Then r1 r2)
  = "(" ++ printRE r1 ++ printRE r2 ++ ")"
printRE (Star r) = "(" ++ printRE r ++")*"
```

which gives a printable form of a regular expression. Note that `'@'` is used to represent epsilon in ASCII.

Exercises

1. Write a more readable form of the expression `((((a|b)|c)((a)*|(b)*))(c|d))`.
2. What is the unabbreviated form of `((x?)*(y?)*)+`?

3. Matching regular expressions

Regular expressions are patterns. We should ask which strings match each regular expression.

ε The empty string matches epsilon.

`x` The character `x` matches the pattern `x`, for any character `x`.

$(r_1|r_2)$ The string `st` will match $(r_1|r_2)$ if `st` matches either r_1 or r_2 (or both).

(r_1r_2) The string `st` will match (r_1r_2) if `st` can be split into two substrings st_1 and st_2, $st = st_1{+}{+}st_2$, so that st_1 matches r_1 and st_2 matches r_2.

$(r)*$ The string `st` will match $(r)*$ if `st` can be split into zero or more substrings, $st = st_1{+}{+}st_2{+}{+}...{+}{+}st_n$, each of which matches `r`. The zero case implies that the empty string will match $(r)*$ for *any* regular expression `r`.

This can be implemented in Miranda, in the file `matches.m`. The first three cases are a simple transliteration of the definitions above.

```
matches :: reg -> string -> bool

matches Epsilon st = (st = "")
matches (Literal ch) st = (st = [ch])
matches (Or r1 r2) st
  = matches r1 st \/ matches r2 st
```

In the case of juxtaposition, we need an auxiliary function which gives the list containing all the possible ways of splitting up a list.

```
splits :: [*] -> [ ([*],[*]) ]

splits st = [ (take n st,drop n st) | n <- [0..#st] ]
```

For example, `splits [2,3]` is `[([],[2,3]),([2],[3]),([2,3],[])]`. A string will match `(Then r1 r2)` if at least one of the splits gives strings which match `r1` and `r2`.

```
matches (Then r1 r2) st
  = or [matches r1 s1 & matches r2 s2 | (s1,s2)<-splits st]
```

The final case is that of `Star`. We can explain `a*` as either ε or as `a` followed by `a*`.

We can use this to implement the check for the match, but it is problematic when `a` can be matched by ε. When this happens, the match is tested recursively on the same string, giving an infinite loop. This is avoided by disallowing an epsilon match on `a` – the first match on `a` has to be non-trivial.

```
matches (Star r) st
  = matches Epsilon st \/
      or [ matches r s1 & matches (Star r) s2 |
                     (s1,s2) <- frontSplits st ]
```

`frontSplits` is defined like `splits` but so as to exclude the split `([],st)`.

Exercises

3. Argue that the string ε matches `(a|(bc)*)*` and that the string `abba` matches `a((b|a)*(ba)*)`.

4. Why does the string `bab` not match `a((b|a)*(ba)*)`?

5. Give informal descriptions of the sets of strings matching the following regular expressions.

`(a|b)*a(a|b)*a(a|b)`   `(a|b)*a(a|b)(a|b)`

ε`|a|b|ba|b?(ab)+a?`

6. Give regular expressions describing the following sets of strings

- All strings of `a`s and `b`s containing at most two `a`s.
- All strings of `a`s and `b`s containing exactly two `a`s.
- All strings of `a`s and `b`s of length at most three.
- All strings of `a`s and `b`s which contain no repeated adjacent characters, that is no substring of the form `aa` or `bb`.

4. Sets

A set is a collection of elements of a particular type, which is both like and unlike a list. Lists are familiar from Miranda, and examples include

```
[Joe,Sue,Ben]      [Ben,Sue,Joe]
[Joe,Sue,Sue,Ben]  [Joe,Sue,Ben,Sue]
```

Each of these lists is different – not only do the elements of a list matter, but also the *order* in which they occur, and their *multiplicity* (the number of times each element occurs).

In many situations, order and multiplicity are irrelevant. If we want to talk about the collection of people coming to our birthday party, we just want the names – we cannot invite someone more than once, so multiplicity is not important; the order we might list them in is also of no interest. In other words, all we want to know is the *set* of people coming. In the example above, this is the set containing `Joe`, `Sue` and `Ben`.

Sets can be implemented in a number of ways in Miranda, and the precise form is not important for the user. It is sensible to declare the type as an abstract data type, or `abstype`, illustrated in Figure 1. The `abstype` and its implementation are found in the file `sets.m`.

The implementation we have given represents a set as an *ordered list of elements without repetitions*. The individual functions are described and implemented as follows.

`sing a` is the singleton set, consisting of the single element `a`

```
sing a = [a]
```

`union,inter,diff` give the union, intersection and difference of two sets. The union consists of the elements occurring in either set (or both), the intersection of those elements in both sets and the difference of those elements in the first but not the second set; their definitions are given in Figure 2.

The `empty` set is the empty list

```
empty = []
```

`memset x a` tests whether `a` is a member of the set `x`. Note that this is an optimisation of the function `member` over lists; since the list is ordered, we need look no further once we have found an element greater than the one we seek.

```
memset [] b     = False
memset (a:x) b = memset x b      , if a<b
               = True            , if a=b
               = False           , otherwise
```

```
abstype          set *
with
 sing                 :: * -> set *
 union,inter,diff     :: set * -> set * -> set *
 empty                :: set *
 memset               :: set * -> * -> bool
 subset               :: set * -> set * -> bool
 eqset                :: set * -> set * -> bool
 mapset               :: (* -> **) -> set * -> set **
 filterset,separate   :: (*->bool) -> set * -> set *
 foldset              :: (* -> * -> *) -> * -> set * -> *
 makeset              :: [*] -> set *
 showset              :: (*->[char]) -> set * -> [char]
 card                 :: set * -> num
 flatten              :: set * -> [*]
 setlimit             :: (set * -> set *) -> set * -> set *
```

Figure 1: The set abstype

```
union [] y         = y
union x []         = x
union (a:x) (b:y)  = a : union x (b:y)    , if a<b
                   = a : union x y        , if a=b
                   = b : union (a:x) y    , otherwise
inter [] y = []
inter x [] = []
inter (a:x) (b:y)  = inter x (b:y)        , if a<b
                   = a : inter x y        , if a=b
                   = inter (a:x) y        , otherwise
diff [] y = []
diff x [] = x
diff (a:x) (b:y)   = a : diff x (b:y)     , if a<b
                   = diff x y             , if a=b
                   = diff (a:x) y         , otherwise
```

Figure 2: Set operations

`subset x y` tests whether `x` is a subset of `y`; that is whether every element of `x` is an element of `y`.

```
subset [] y = True
subset x [] = False
subset (a:x) (b:y) = False                , if a<b
                   = subset x y           , if a=b
                   = subset (a:x) y       , if a>b
```

`eqset x y` tests whether two sets are equal.

```
eqset = (=)
```

`mapset`, `filterset` and `foldset` behave like `map`, `filter` and `foldr` except that they operate over sets. `separate` is a synonym for `filterset`.

```
mapset f = makeset . (map f)
filterset = filter
separate  = filterset
foldset = foldr
```

`makeset` turns a list into a set

```
makeset = remdups . sort
          where
          remdups [] = []
          remdups [a] = [a]
          remdups (a:x) = a : remdups x , if a < b
                        = remdups x     , otherwise
                          where
                          b = hd x
```

`showset f` gives a printable version of a set, one item per line, using the function `f` to give a printable version of each element.

```
showset f = concat . (map ((++"\n") . f))
```

`card x` gives the number of elements in `x`

```
card = (#)
```

`flatten x` turns a set `x` into an ordered list of the elements of the set

```
flatten = id
```

`setlimit f x` gives the 'limit' of the sequence

```
x , f x , f (f x) , f (f (f x)) , ...
```

that is the first element in the sequence whose successor is equal, as a set, to the element itself. In other words, keep applying `f` until a fixed point or limit is reached.

```
setlimit f x = x                , if eqset x next
             = setlimit f next  , otherwise
               where
               next = f x
```

Exercises

7. How is the function `powerSet :: set * -> set (set *)` which returns the set of all subsets of a set defined?

8. How would you define the functions

```
setUnion :: set (set *) -> set *
setInter :: set (set *) -> set *
```

which return the union and intersection of a set of sets?

9. Can infinite sets (of numbers, for instance) be adequately represented by ordered lists? Can you tell if two infinite lists are equal, for instance?

10. The abstype `set *` can be represented in a number of different ways. Alternatives include: arbitrary lists (rather than ordered lists without repetitions), and boolean valued functions, that is elements of the type `* -> bool`. Give implementations of the type using these two representations.

5. Non-deterministic Finite Automata

A Non-deterministic Finite Automaton or NFA is a simple machine which can be used to recognise regular expressions. It consists of four components

- A finite set of states, S.
- A finite set of moves.
- A start state (in S).
- A set of terminal or final states (a subset of S).

In Miranda (see file `nfa_types.m`) this is written

```
nfa * ::= NFA (set *)
              (set (move *))
              *
              (set *)
```

This has been represented by an algebraic type rather than a 4-tuple simply for readability. The type of states can be different in different applications, and indeed in the following we use both numbers and sets of numbers as states.

A move is between two states, and is either given by a character, or an ε.

```
move * ::= Move * char * |
           Emove * *
```

The first example of an NFA, called `M`, follows.

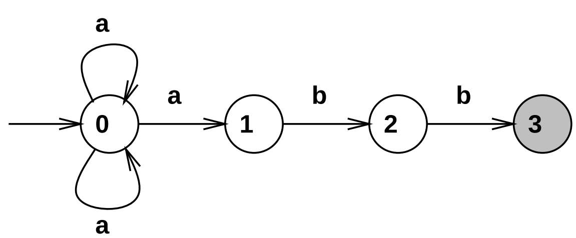

The states are `0,1,2,3`, with the start state `0` indicated by an incoming arrow, and the final states indicated by shaded circles. In this case there is a single final state, `3`. The moves are indicated by the arrows, marked with characters `a` and `b` in this case. From state `0` there are two possible moves on symbol `a`, to `1` and to remain at `0`. This is one source of the non-determinism in the machine.

The Miranda representation of the machine is

```
NFA
(makeset [0..3])
(makeset [ Move 0 'a' 0 ,
           Move 0 'a' 1 ,
           Move 0 'b' 0 ,
           Move 1 'b' 2 ,
           Move 2 'b' 3 ])
0
(sing 3)
```

A second example, called `N`, is illustrated below.

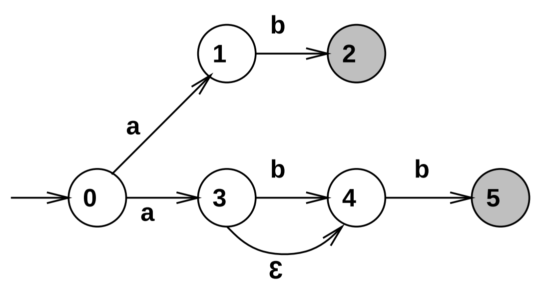

The Miranda representation of this machine is

```
NFA
(makeset [0..5])
(makeset [ Move 0 'a' 1,
           Move 1 'b' 2,
           Move 0 'a' 3,
           Move 3 'b' 4,
           Emove 3 4,
           Move 4 'b' 5 ])
0
(makeset [2,5])
```

This machine contains two kinds of non-determinism. The first is at state `0`, from which it is possible to move to either `1` or `3` on reading `a`. The second occurs at state `3`: it is possible to move 'invisibly' from state `3` to state `4` on the epsilon move, `Emove 3 4`.

The Miranda code for these machines together with a function `print_nfa` to print an nfa whose states are numbered can be found in the file `nfa_misc.m`.

How do these machines recognise strings? A move can be made from one state `s` to another `t` either if the machine contains `Emove s t` or if the next symbol to be read is, say, `a` and the machine contains a move `Move s a t`. A string will be *accepted* by a machine if there is a sequence of moves through states of the machine starting at the start state and terminating at one of the terminal states – this is called an *accepting path*. For instance, the path

$$\mathtt{0} \xrightarrow{\mathtt{a}} \mathtt{1} \xrightarrow{\mathtt{b}} \mathtt{2} \xrightarrow{\mathtt{b}} \mathtt{3}$$

is an accepting path through `M` for the string `abb`. This means that the machine `M` accepts this string. Note that other paths through the machine are possible for this string, an example being

$$\mathtt{0} \xrightarrow{\mathtt{a}} \mathtt{0} \xrightarrow{\mathtt{b}} \mathtt{0} \xrightarrow{\mathtt{b}} \mathtt{0}$$

All that is needed for the machine to accept is *one* accepting path; it does not affect acceptance if there are other non-accepting (or indeed accepting) paths. More than one accepting path can exist. Machine `N` accepts the string `ab` by both

$$\mathtt{0} \xrightarrow{\mathtt{a}} \mathtt{1} \xrightarrow{\mathtt{b}} \mathtt{2}$$

and

$$\mathtt{0} \xrightarrow{\mathtt{a}} \mathtt{3} \xrightarrow{\varepsilon} \mathtt{4} \xrightarrow{\mathtt{b}} \mathtt{5}$$

A machine will *reject* a string only when there is no accepting path. Machine `N` rejects the string `a`, since the two paths through the machine labelled by `a` fail to terminate in a final state:

$$\mathtt{0} \xrightarrow{\mathtt{a}} \mathtt{1} \qquad \mathtt{0} \xrightarrow{\mathtt{a}} \mathtt{3}$$

Machine `N` rejects the string `aa` since there is *no* path through the machine labelled by `aa`: after reading `a` the machine can be in state `1`, `3` or `4`, from none of these can an `a` move be made.

6. Simulating an NFA

As was explained in the last section, a string `st` is accepted by a machine `M` when there is at least one accepting path labelled by `st` through `M`, and is rejected by `M` when no such path exists.

The key to implementation is to explore simultaneously *all* possible paths through the machine labelled by a particular string. Take as an informal example the string `ab` and the machine `N`. After reading no input, the machine can only be in state `0`. On reading an `a` there are moves to states `1` and `3`; however this is not the whole story. From state `3` it is possible to make an ε-move to state `4`, so after reading `a` the machine can be in any of the states {`1,3,4`}.

On reading a `b`, we have to look for all the possible `b` moves from each of the states {`1,3,4`}. From `1` we can move to `2`, from `3` to `4` and from `4` to `5` – no ε-moves are possible from the states {`2,4,5`}, and so the states accessible after reading the string `ab` are {`2,4,5`}. Is this string to be accepted by `N`? We accept it exactly if the set contains a final state – it contains both `2` and `5`, so it is accepted. Note that the states accessible after reading `a` are {`1,3,4`}; this set contains no final state, and so the machine `N` *rejects* the string `a`.

There is a general pattern to this process, which consists of a repetition of

- Take a set of states, such as {`1,3,4`}, and find the set of states accessible by a move on a particular symbol, e.g. `b`. In this case it is the set {`2,4,5`}. This is called `onemove` in `nfa_lib.m`.

- Take a set of states, like {`1,3`}, and find the set of states accessible from the states by *zero* or more ε-moves. In this example, it is the set {`1,3,4`}. This is the ε-closure of the original set, and is called `closure` in `nfa_lib.m`.

The functions `onemove` and `closure` are composed in the function `onetrans`, and this function is iterated along the string by the `trans` function of `implement_nfa.m`.

Implementation in Miranda

We discuss the development of the function

```
trans :: nfa * -> string -> set *
```

top-down. Iteration along a string is given by `foldl`

```
foldl :: (set * -> char -> set *)
               -> set * -> string -> set *

foldl f r []     = r
foldl f r (c:cs) = foldl f (f r c) cs
```

The first argument, `f`, is the step function, taking a set and a character to the states accessible from the set on the character. The second argument, `r`, is the starting state, and the final argument is the string along which to iterate.

How does the function operate? If given an empty string, the start state is the result. If given a string `(c:cs)`, the function is called again, with the tail of the string, `cs`, and with a new starting state, `(f r c)`, which is the result of applying the step function to the starting set of states and the first character of the string. Now to develop `trans`.

```
trans mach str
    = foldl step startset str
      where
      step set ch = onetrans mach ch set
      startset = closure mach (sing (startstate mach))
```

`step` is derived from `onetrans` simply by suppling its machine argument `mach`, similarly `startset` is derived from the machine `mach`, using the functions `closure` and `startstate`. All these functions are defined in `nfa_lib.m`. We discuss their definitions now.

```
onetrans :: nfa * -> char -> set * -> set *

onetrans mach c x = closure mach (onemove mach c x)
```

Next, we examine `onemove`,

```
onemove :: nfa * -> char -> set * -> set *

onemove (NFA states moves start term) c x
  = makeset [ s | t <- flatten x ;
                  Move z d s <- flatten moves ;
                  z=t ; c=d ]
```

The essential idea here is to run through the elements `t` of the set `x` and the set of moves, `moves` looking for all `c`-moves originating at `t`. For each of these, the result of the move, `s`, goes into the resulting set.

The definition uses list comprehensions, so it is necessary first to `flatten` the sets x and moves into lists, and then to convert the list comprehension into a set by means of makeset.

```
closure :: nfa * -> set * -> set *

closure (NFA states moves start term)
  = setlimit add
    where
    add stateset = union stateset (makeset accessible)
                   where
                   accessible
                     = [ s | x <- flatten stateset ;
                             Emove y s <- flatten moves ;
                             y=x ]
```

The essence of closure is to take the limit of the function which adds to a set of states all those states which are accessible by a single ε-move; in the limit we get a set to which no further states can be added by ε-transitions. Adding the states got by single ε-moves is accomplished by the function add and the auxiliary definition accessible which resembles the construction of onemove.

7. Implementing an example

The machine P is illustrated by

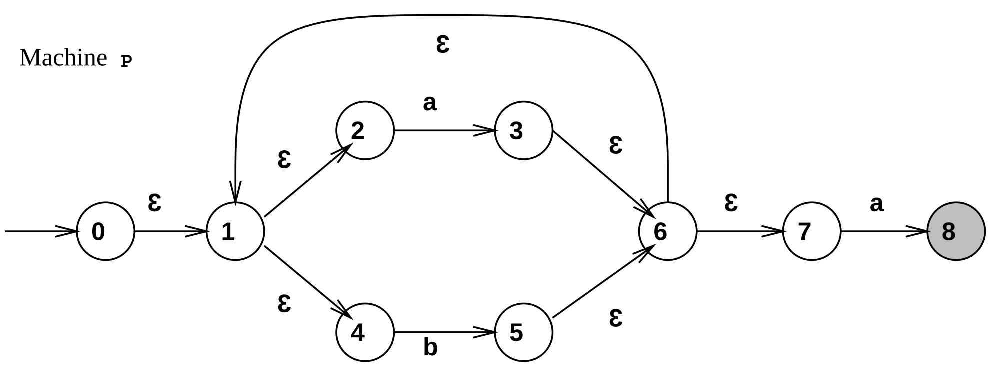

Exercise

11. Give the Miranda definition of the machine P.

The ε-closure of the set {0} is the set {0,1,2,4}. Looking at the definition of `closure` above, the first application of the function `add` to {0} gives the set {0,1}; applying `add` to this gives {0,1,2,4}. Applying `add` to this set gives the same set, hence this is the value of `setlimit` here. The set of states with which we start the simulation is therefore {0,1,2,4}. Suppose the first input is `a`; applying `onemove` reveals only one `a` move, from `2` to `3`. Taking the closure of the set {3} gives the set {1,2,3,4,6,7}. A `b` move from here is only from `4` to `5`; closing under ε-moves gives {1,2,4,5,6,7}. An `a` move from here is possible in two ways: from `2` to `3` and from `7` to `8`; closing up {3,8} gives {1,2,3,4,6,7,8}. Is the string `aba` therefore accepted by `P`? Yes, because `8` is a member of {1,2,3,4,6,7,8}. This sequence can be illustrated thus

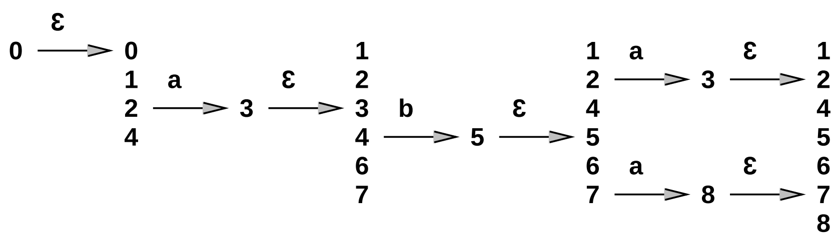

Exercise

12. Show that the string `abb` is not accepted by the machine `P`.

8. Building NFAs from regular expressions

For each regular expression it is possible to build an NFA which accepts exactly those strings matching the expression. The machines are illustrated in Figure 3.

The construction is by induction over the structure of the regular expression: the machines for an character and for ε are given outright, and for complex expressions, the machines are built from the machines representing the parts. It is straightforward to justify the construction.

`(e|f)` Any path through `M(e|f)` must be either a path through `M(e)` or a path through `M(f)` (with ε at the start and end.

`ef` Any path through `M(ef)` will be a path through `M(e)` followed by a path through `M(f)`.

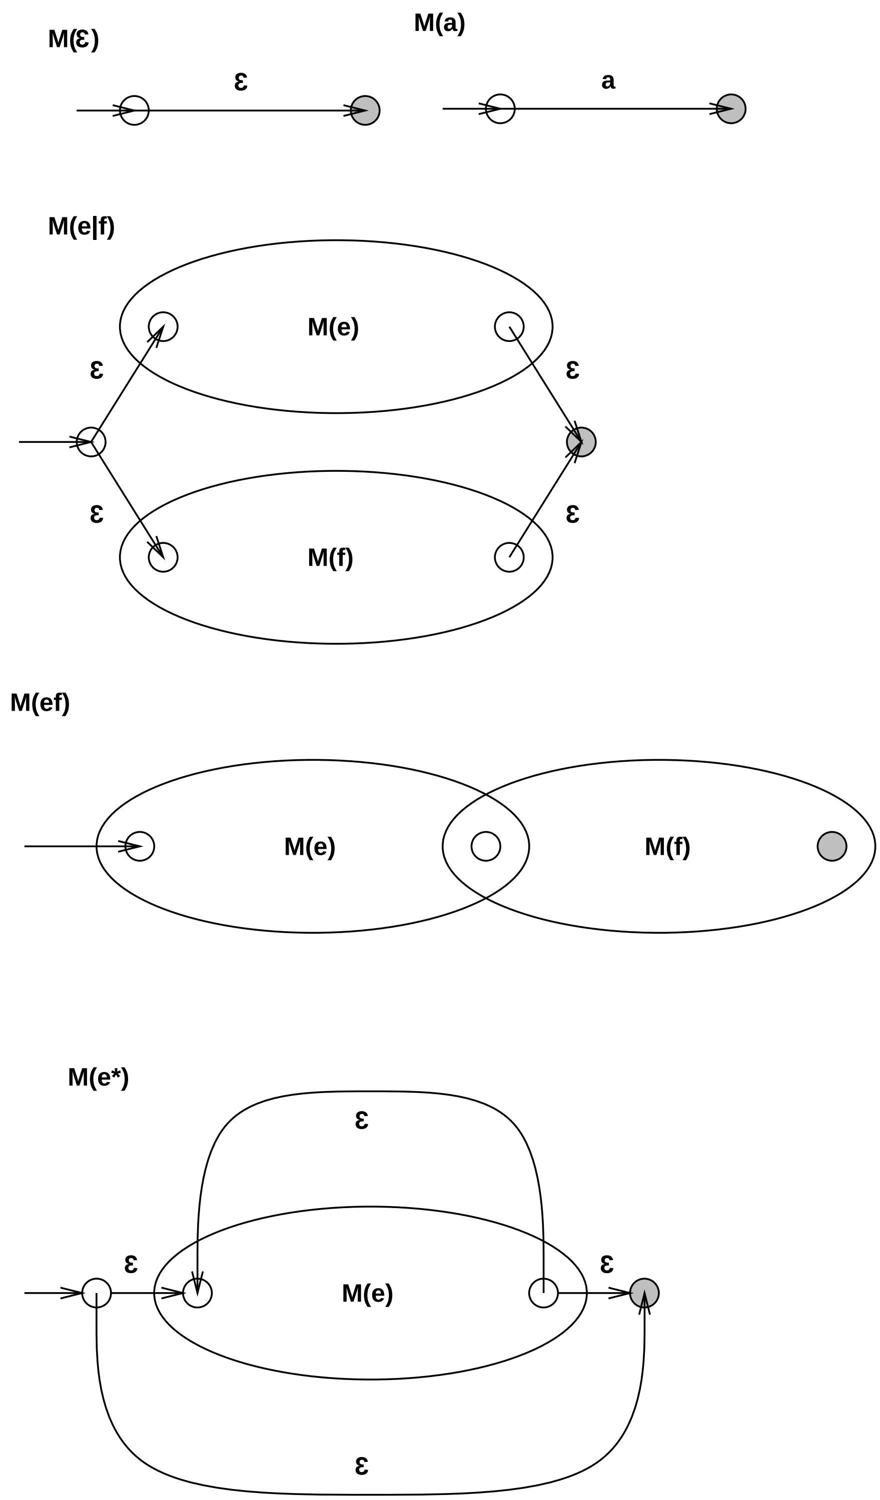

Figure 3: Building NFAs for regular expressions

`e*` Paths through `M(e*)` are of two sorts; the first is simply an ε, others begin with a path through `M(e)`, and continue with a path through `M(e*)`. In other words, paths through `M(e*)` go through `M(e)` zero or more times.

The machine for the pattern `(ab|ba)*` is given by

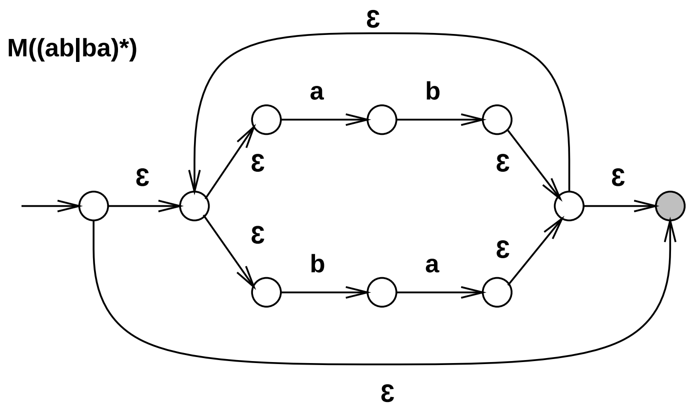

The Miranda description of the construction is given in `build_nfa.m`. At the top level the function

```
build :: reg -> nfa num
```

does the recursion. For the base case,

```
build (Literal c)
  = NFA
    (makeset [0..1])
    (sing (Move 0 c 1))
    0
    (sing 1)
```

The definition of `build Epsilon` is similar. In the other cases we define

```
build (Or r1 r2)    = m_or (build r1) (build r2)
build (Then r1 r2)  = m_then (build r1) (build r2)
build (Star r)      = m_star (build r)
```

in which the functions `m_or` and so on build the machines from their components as illustrated.

We make certain assumptions about the NFAs we build. We take it that the states are numbered from `0`, with the final state having the highest number. Putting the machines together will involve adding various new states and transitions, and renumbering the states and moves in the constituent machines. An example program is

```
m_or :: nfa num -> nfa num -> nfa num

m_or (NFA states1 moves1 start1 finish1)
     (NFA states2 moves2 start2 finish2)
   = NFA
     (states1' $union states2' $union newstates)
     (moves1' $union moves2' $union newmoves)
     0
     (sing (m1+m2+1))
     where
     m1 = card states1
     m2 = card states2
     states1' = mapset (renumber 1) states1
     states2' = mapset (renumber (m1+1)) states2
     newstates = makeset [0,(m1+m2+1)]
     moves1'  = mapset (renumber_move 1) moves1
     moves2'  = mapset (renumber_move (m1+1)) moves2
     newmoves = makeset [ Emove 0 1 ,
                          Emove 0 (m1+1) ,
                          Emove m1 (m1+m2+1) ,
                          Emove (m1+m2) (m1+m2+1) ]
```

The function `renumber` renumbers states and `renumber_move` moves.

9. Deterministic machines

A deterministic finite automaton is an NFA which

- contains no ε-moves, and
- has at most one arrow labelled with a particular symbol leaving any given state.

The effect of this is to make operation of the machine deterministic – at any stage there is at most one possible move to make, and so after reading a sequence of characters, the machine can be in one state at most.

Implementing a machine of this sort is much simpler than for an general NFA: we only have to keep track of a single position. Is there a general mechanism for finding a DFA corresponding to a regular expression? In fact, there is a general technique for transforming an arbitrary NFA into a DFA, and this we examine now.

The conversion of an NFA into a DFA is based on the implementation given in Section 6. The main idea there is to keep track of a *set* of states, representing all the possible positions after reading a certain amount of input. This set itself can be thought of as a state of another machine, which will be *deterministic*: the moves from one *set* to another are completely deterministic.

We show how the conversion works with the machine `P`. The start state of the machine will be the closure of the set $\{$`0`$\}$, that is

`A = `$\{$`0,1,2,4`$\}$

Now, the construction proceeds by finding the sets accessible from `A` by moves on `a` and on `b` – all the characters in the *alphabet* of the machine `P`. These sets are states of the new machine; we then repeat the construction with these new states, until no more states are produced by the construction.

From `A` on the symbol `a` we can move to `3` from `2`. Closing under ε-moves we have the set $\{$`1,2,3,4,6,7`$\}$, which we call `B`

`B = `$\{$`1,2,3,4,6,7`$\}$

$\mathtt{A} \xrightarrow{\mathtt{a}} \mathtt{B}$

In a similar way, from `A` on `b` we have

`C = `$\{$`1,2,4,5,6,7`$\}$

$\mathtt{A} \xrightarrow{\mathtt{b}} \mathtt{C}$

Our new machine so far looks like

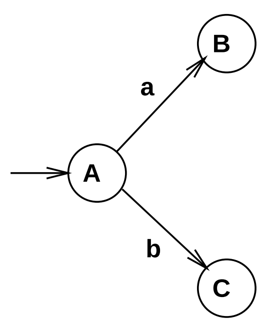

We now have to see what is accessible from `B` and `C`. First `B`.

`D = `$\{$`1,2,3,4,6,7,8`$\}$

$\mathtt{B} \xrightarrow{\mathtt{a}} \mathtt{D}$

which is another new state. The process of generating new states must stop, as there is only a finite number of sets of states to choose from $\{0,1,2,3,4,5,6,7,8\}$. What happens with a `b` move from `B`?

$\texttt{B} \xrightarrow{\texttt{b}} \texttt{C}$

This gives the partial machine

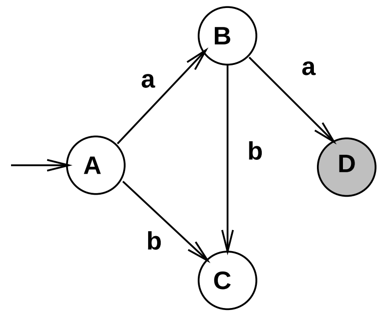

Similarly,

$\texttt{C} \xrightarrow{\texttt{a}} \texttt{D}$
$\texttt{C} \xrightarrow{\texttt{b}} \texttt{C}$
$\texttt{D} \xrightarrow{\texttt{a}} \texttt{D}$
$\texttt{D} \xrightarrow{\texttt{b}} \texttt{C}$

which completes the construction of the DFA

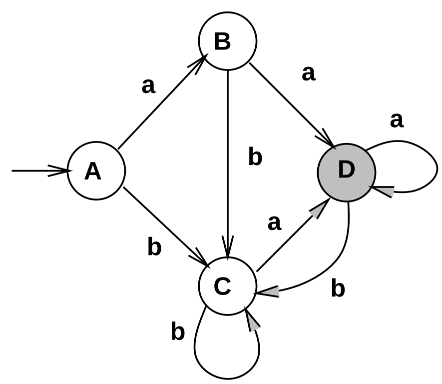

Which of the new states is final? One of these sets represents an accepting state exactly when it contains a final state of the original machine. For `P` this is `8`, which is contained in the set `D` only. In general there can be more than one accepting state for a machine. (This need not be true for NFAs, since we can always add a new final state to which each of the originals is linked by an ε-move.)

10. Transforming NFAs to DFAs

The Miranda code to covert an NFA to a DFA is found in the file `nfa_to_dfa.m`, and the main function is

```
make_deterministic :: nfa num -> nfa num

make_deterministic = number . make_deter
```

A deterministic version of an NFA with numeric states is defined in two stages, using

```
make_deter :: nfa num -> nfa (set num)

number :: nfa (set num) -> nfa num
```

`make_deter` does the conversion to the deterministic automaton with sets of numbers as states, `number` replaces sets of numbers by numbers (rather than capital letters, as was done above). States are replaced by their position in a list of states – see the file for more details.

`make_deter` is a special case of the function

```
deterministic :: nfa num -> [char] -> nfa (set num)

make_deter mach = deterministic mach (alphabet mach)
```

The process of adding state sets is repeated until no more sets are added. This is a version of taking a limit, given by the `nfa_limit` function, which acts as the usual limit function, except that it checks for equality of NFAs as collections of sets.

```
deterministic mach alpha
  = nfa_limit (addstep mach alpha) startmach
    where
    startmach = NFA
                (sing starter)
                empty
                starter
                finish
```

```
      starter = closure mach (sing start)
      finish  = empty
                  , if eqset (term $inter starter) empty
              = sing starter      , otherwise
      (NFA sts mvs start term) = mach
```

The start machine, `startmach`, consists of a single state, the ε-closure of the start state of the original machine. `addstep mach alpha` takes a partially built DFA and adds the state sets of `mach` accessible by a single move on any of the characters in `alpha`, the alphabet of `mach`.

```
addstep
  :: nfa num -> [char] -> nfa (set num) -> nfa (set num)
addstep mach alpha dfa
  = add_aux mach alpha dfa (flatten states)
    where
    (NFA states m s f) = dfa
    add_aux mach alpha dfa [] = dfa
    add_aux mach alpha dfa (st:rest)
        = add_aux mach alpha (addmoves mach st alpha dfa) rest
```

This involves iterating over the state sets in the partially built DFA, which is done using `addmoves`. `addmoves mach x alpha dfa` will add to `dfa` all the moves from state set `x` over the alphabet `alpha`.

```
addmoves :: nfa num -> set num -> [char]
               -> nfa (set num) -> nfa (set num)

addmoves mach x [] dfa    = dfa

addmoves mach x (c:r) dfa
  = addmoves mach x r (addmove mach x c dfa)
```

In turn, `addmoves` iterates along the alphabet, using `addmove`. `addmove mach x c dfa` will add to `dfa` the moves from state set `x` on character `c`.

```
addmove :: nfa num -> set num -> char
```

```
                 -> nfa (set num) -> nfa (set num)

addmove mach x c (NFA states moves start finish)
  = NFA states' moves' start finish'
    where
    states' = states $union (sing new)
    moves'  = moves  $union (sing (Move x c new))
    finish' = finish $union (sing new)
                 , if ~ (eqset empty (term $inter new))
            = finish  , otherwise
    new = onetrans mach c x
    (NFA s m q term) = mach
```

The new state set added by `addmove` is defined using the `onetrans` function first defined in the simulation of the NFA.

11. Minimising a DFA

In building a DFA, we have produced a machine which cam be implemented more efficiently. We might, however, have more states in the DFA than necessary. This section shows how we can *optimise* a DFA so that it contains the minimum number of states to perform its function of recognising the strings matching a particular regular expression.

Two states `m` and `n` in a DFA are *distinguishable* if we can find a string `st` which reaches an accepting state from `n` but not from `n` (or vice versa). Otherwise, they can be treated as the same, because no string makes them behave differently — putting it a different way, no *experiment* makes the two different.

How can we tell when two states are different? We start by dividing the states into two *partitions*: one contains the accepting states, and the other the remainder, or non-accepting states. For our example, we get the partition

```
I:  D
II: A,B,C
```

Now, for each set in the partition, we check whether the elements in the set can be further divided. We look at how each of the states in the set behaves *relative to the previous partition*. In pictures,

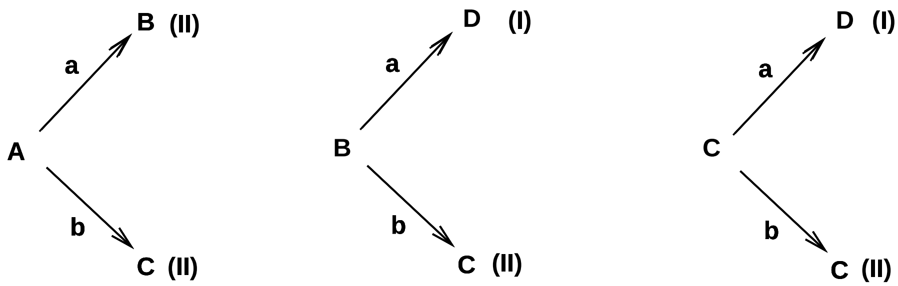

This means that we can re-partition thus:

```
I:   D
II:  A
III: B,C
```

We now repeat the process, and examine the only set which might be further subdivided, giving

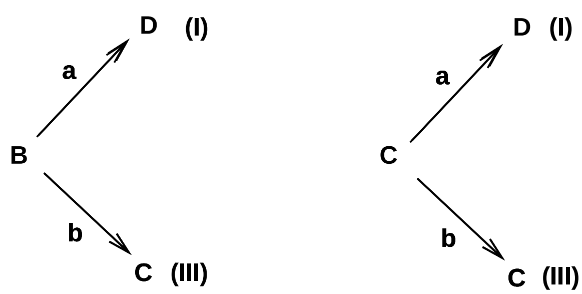

This shows that we don't have to re-partition any further, and so that we can stop now, and collapse the two states `B` and `C` into one, thus:

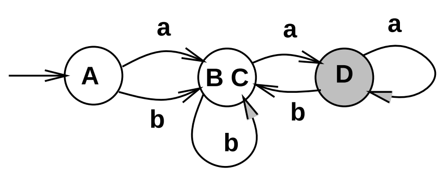

The Miranda implementation of this process is in the file `minimise_dfa.m`.

Exercises

13. For the regular expression `b(ab|ba)*a`, find the corresponding NFA.
14. For the NFA of question 1, find the corresponding (non-optimised) DFA.
15. For the DFA of question 2, find the optimised DFA.

12. Regular definitions

A *regular definition* consists of a number of *named* regular expressions. We are allowed to use the defined names on the right-hand sides of definitions *after* the definition of the name. For example,

```
alpha    -> [a-zA-Z]
digit    -> [0-9]
alphanum -> alpha | digit
ident    -> alpha | alphanum*
digits   -> digit+
fract    -> (.digits)?
num      -> digits fract
```

Because of the stipulation that a definition precedes the *use* of a name, we can expand each right-hand side to a regular expression involving no names.

We can build machines to recognise strings from a number of regular expressions. Suppose we have the patterns

```
p1:  a
p2:  abb
p3: a*b*
```

We can build the three NFAs thus:

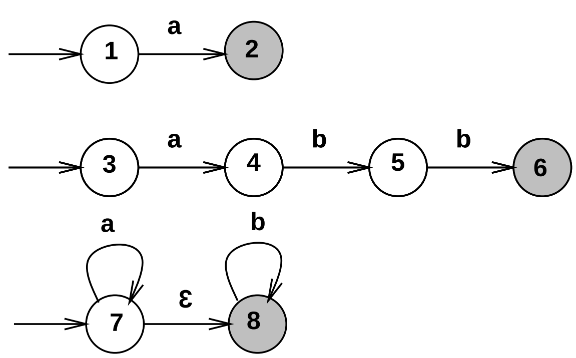

and then they can be joined into a single machine, thus

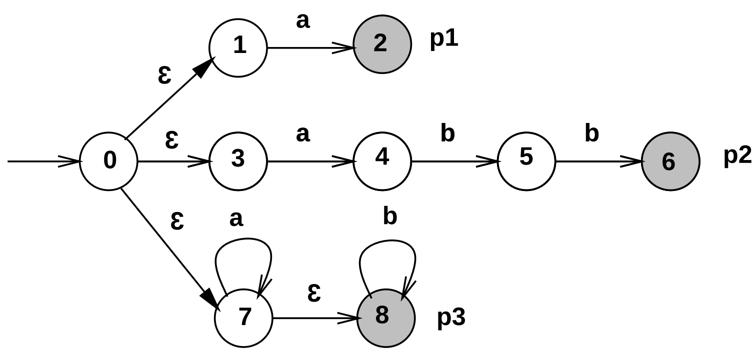

In using the machine we look for the *longest* match against any of the patterns:

```
      0,1,3,7,8(p3)
a     2(p1),4,7,8(p3)
a     7,8(p3)
b     8(p3)
a     -
```

In the example, the segment of `aab` matches the pattern `p3`.

Exercises

16. Fully expand the names `digits` and `num` given above.

17. Build a Miranda program to recognise strings according to a set of regular definitions, as outlined in this section.

Bibliography

[Aho *et. al.*] Aho, A.V., Sethi, R. and Ullman, J.D., *Compilers: Principles, Techniques and Tools,* Addison-Wesley, Reading, MA, USA, 1986.

[Thompson] Thompson, S., *Miranda The Craft of Functional Programming*, Addison-Wesley, Wokingham , Berks, UK, 1995.