Chapter I

Automata and infinite words

1 Introduction

This first chapter constitutes an introduction to the theory of automata on infinite words. It includes some basic definitions, such as ω-rational sets of infinite words. These sets are defined starting from the elements of the alphabet by making use of the four operations of *union*, *concatenation*, *finite iteration* and *infinite iteration*. These operations are denoted $\cup$, ., $*$, ω as indicated in Table 1.1.

Operation	**Symbol**
union	$\cup$
concatenation	$\cdot$
finite iteration	$*$
infinite iteration	ω

Table 1.1: The operations used in ω-rational expressions.

The expressions obtained are often called ω-rational expressions. The name of rational expressions is reserved for those expressions that do not use the symbol ω and thus define sets of finite words. The classical theorem of Kleene establishes the equivalence between rational expressions and finite automata. Its extension to infinite words is one of the results contained in this chapter. We shall see that a set is ω-rational if and only if it can be recognized by a finite automaton (Theorem 5.4). Recognizing infinite words with finite automata requires a convention. The simplest one, introduced by Büchi, consists in considering an infinite path in the automaton as successful if it starts in an initial state and passes infinitely often through a terminal state.

Another basic result of the theory of finite automata is given by the *determinizationdeterminization* algorithm which allows one to replace any finite automaton by an equivalent deterministic one. A consequence of this is the closure of the class of rational sets under complementation. The extension of these results is fraught with difficulties. The solution requires first the introduction of a more powerful acceptance mode than Büchi's one, since one has to specify the set of states met infinitely often on an infinite path. The acceptance mode, called Muller's mode, declares an infinite path p to be successful if the set $\mathrm{Inf}(p)$ of states met infinitely often on p belongs to a prescribed set $\mathcal{T}$ of sets of states. This constitutes a more constrained acceptance mode than Büchi's one, for which it suffices to check whether the set $\mathrm{Inf}(p)$ meets the set F of final states (see Table 1.2).

Acceptance modes	**Definition**
Büchi	$\mathrm{Inf}(p) \cap F \neq \emptyset$
Muller	$\mathrm{Inf}(p) \in \mathcal{T}$

Table 1.2: Büchi and Muller's acceptance modes.

The basic result of this theory, due to R. McNaughton, states that any finite automaton is equivalent on infinite words to a deterministic Muller automaton (Theorem 7.1). This implies in particular that the class of ω-rational sets is closed under complementation, a result proved for the first time by Büchi in a direct way (see the notes at the end of the chapter).

This chapter is organized as follows. Section 2 introduces our notation. Section 3 introduces ω-rational sets of infinite words. It contains the result that characterizes them as the finite unions of sets of the form XY^ω where X and Y are rational sets of finite words. Sections 4, 5, 6, 7.1 and 8 introduce the definitions of the various acceptance modes: Büchi's mode, Muller's mode, Rabin's mode and a transition mode. Section 9 contains a proof of McNaughton's theorem. This proof, discovered in 1989 by S. Safra has the advantage, compared to other possible ones, of being direct and of providing a better algorithm: starting from a nondeterministic n-state Büchi automaton, one obtains a deterministic Muller automaton with $O(n^n)$ states (the other constructions lead to a double exponential). We shall however see other proofs in the following chapters, especially one using ω-semigroups which makes various generalizations possible. The last section (Section 10) deals with computational issues concerning the transformations between various possible representations of ω-rational sets and the operations on them.

2 Words and trees

In this book, we are going to consider possibly infinite sequences of elements of a set called an *alphabet*. The elements of this set are called *letters* or also *symbols*. Most often, in the examples, the alphabet will be finite or even reduced to two elements. We shall however also consider countable alphabets. Recall that a set A is *countable* if there exists an injective map from A to $\mathbb{N}$.

A finite sequence of elements of A is called a *finite word* on A, or just a *word*. We denote by mere juxtaposition

$$a_0 a_1 \cdots a_n$$

the sequence $(a_0, a_1, \ldots, a_n)$. The set of words is endowed with the operation of *concatenation product* also called *product*, which associates with two words $x = a_0 a_1 \cdots a_p$ and $y = b_0 b_1 \cdots b_q$ the word $xy = a_0 a_1 \cdots a_p b_0 b_1 \cdots b_q$. This operation is associative. It has a neutral element, the *empty word*, denoted by 1 or ε and which is the empty sequence.

We denote by A^* the set of words on A and by A^+ the set of nonempty words. The set A^* (resp. A^+), equipped with the concatenation product is thus a monoid with neutral element 1 (resp. a semigroup). The set A^* is called the *free monoid* on A and A^+ the *free semigroup* on A. This terminology will be justified later.

If u is a word and a a letter, we denote by $|u|_a$ the number of occurrences of a in u. Thus, if $A = \{a, b\}$ and $u = abaab$, we have $|u|_a = 3$ and $|u|_b = 2$. The sum

$$|u| = \sum_{a \in A} |u|_a$$

is the *length* of the word u. Thus $|abaab| = 5$.

An *infinite word* on the alphabet A is an infinite sequence of elements of A, which we also denote by juxtaposition

$$u = a_0 a_1 \cdots a_n \cdots$$

This notation represents the mapping from $\mathbb{N}$ into A defined, for all $n \in \mathbb{N}$ by $u(n) = a_n$.

We also denote by $u[r, s]$ the word $u(r)u(r+1) \cdots u(s) = a_r \cdots a_s$. This notation is also used for finite words.

We denote indifferently by $A^{\mathbb{N}}$ or by A^ω the set of infinite words over the alphabet A and we let

$$A^\infty = A^* \cup A^\omega$$

which is thus the set of finite or infinite words on the alphabet A. The product of a finite word $u = a_0 a_1 \cdots a_n$ from A^* with an infinite word $v = b_0 b_1 \cdots$ of A^ω is the infinite word

$$uv = a_0 a_1 \cdots a_n b_0 b_1 \cdots$$

Let $u = a_0 a_1 \cdots a_n$ be a word in A^*. A word $x \in A^*$ is a *factor* of u if there exist integers r and s such that $0 \leqslant r \leqslant s \leqslant n$ and $x = u[r, s]$. This is equivalent to saying that there exist words v and w in A^* such that $u = vxw$. In the same way, we say that

x is a *left factor* or a *prefix* of u if there exists a word w in A^* such that $u = xw$,

x is a *right factor* or a *suffix* of u if there exists a word v in A^* such that $u = vx$,

x is a *proper factor* of u if there exist words v and w not both empty such that $u = vxw$,

x is *strict factor* of u if there exist words v and w in A^+ such that $u = vxw$.

For example, if $u = abaabab$, aba is a prefix, ab is a suffix, $abaab$ is a proper factor and $baaba$ is a strict factor u.

The relation "to be a left factor of" is a partial order on words called "prefix order" and sometimes denoted $\leqslant$. Thus,

$$1 \leqslant a \leqslant ab \leqslant abb \leqslant abba$$

If one starts with an ordered alphabet, a total order relation can be defined on words, called *lexicographic order* and denoted $\leqslant_{lex}$. It is the usual order in a dictionary. In formal terms, one has $u \leqslant_{lex} v$ if $u \leqslant v$ or if $u = xau'$, $v = xbv'$ with $x, u', v' \in A^*$, $a, b \in A$ and $a <_{lex} b$.

The notions of a factor and of a left factor extend without difficulty to infinite words: a word $x \in A^*$ is a *factor* of an infinite word u if there exist integers r and s such that $0 \leqslant r \leqslant s$ and $x = u[r, s]$, or in other terms, if there exist a finite word v and an infinite word w such that $u = vxw$. We say that x is a *left factor* or a *prefix* of u if there exists an infinite word w such that $u = xw$.

If x is a prefix of u, we denote by $x^{-1}u$ the unique word v such that $xv = u$. The prefix order can be extended to a partial order on A^∞ by setting for u and v in A^∞, $u \leqslant v$ if $u = v$ or if u is a finite prefix of v.

An integer $p > 0$ is a *period* of a finite word $u = a_0a_1 \cdots a_n$ if, for all k such that $k + p \leqslant n$, we have $a_k = a_{k+p}$. The smallest period of u is called *the period* of u. For instance, the period of the word $abcaabcaab$ is 4, since $abcaabcaab = (abca)^2ab$.

The notion of a period extends without difficulty to infinite words. We say that an infinite word is *periodic* if the set of its periods is not empty.

An integer $p > 0$ is an *ultimate period* of an infinite word $u = a_0a_1 \cdots$ if there is a $k_0 \geqslant 0$ such that for all $k \geqslant k_0$ we have $a_k = a_{k+p}$. If p and q are two ultimate periods of u, their gcd is still a ultimate period of u. The smallest ultimate period of an infinite word u is called *the ultimate period* of u. An infinite word which admits a ultimate period is called *ultimately periodic*. We know, for example, that rational integers are those real numbers that admit, in a given base, an ultimately periodic expansion.

The set A^ω is equipped with an internal operation called the *shift*, which is defined as the mapping $\sigma : A^\omega \to A^\omega$, associating to each infinite word $u \in A^\omega$ the infinite word $\sigma(u)$ defined, for all $n \in \mathbb{N}$, by

$$\sigma(u)(n) = u(n + 1).$$

Thus, for $u = a_0a_1 \cdots$, one has $\sigma(u) = a_1a_2 \cdots$, which means that the action of σ consists in shifting one place to the left all symbols of u. It is therefore a surjective application, which is not injective as soon as $\mathrm{Card}(A) \geqslant 2$.

If X is a set of infinite words, we let

$$\sigma(X) = \{\sigma(u) \mid u \in X\}$$

We say that X is *stable* (resp. *shift invariant*) if $\sigma(X) \subset X$ (resp. $\sigma(X) = X$).

It is easy to verify that an integer p is a period of an infinite word u if and only if $\sigma^p(u) = u$ and that u is ultimately of period p if and only if there is an integer q such that $\sigma^{p+q}(u) = \sigma^q(u)$.

A subset of A^* is said to be *prefix-free* or simply *prefix* if its elements are incomparable for the prefix order. For instance, if $A = \{a, b\}$, the set $\{a^n b \mid n \geqslant 0\}$ is prefix-free. As a dual definition, a subset of A^* is called *prefix-closed* if it contains the prefixes of all of its elements. In particular, if X is a subset of A^+ or of A^ω, the set $\mathrm{Pref}(X)$ of all prefixes of the elements of X is prefix-closed. The complement of a prefix-closed set is a *right ideal*: a subset R of A^* is indeed a right ideal if for all $u \in R$ and for all $v \in A^*$, one has $uv \in R$.

There is consequently a bijection between finite prefix-free sets and finite prefix-closed sets defined as follows: we associate with each prefix-free set P, the prefix-closed set formed of the prefixes of P; conversely, we associate with every prefix-closed set T, the prefix-free set formed by the elements of T which are maximal for the prefix ordering.

These notions are closely related with the notion of a *tree* that we introduce now. A tree is a tuple (N, r, p), where N is a nonempty set whose elements are called the *nodes* of the tree, r is a distinguished element of N, called the *root* of the tree and

$$p : N \setminus \{r\} \to N$$

is a mapping associating with each node distinct from the root a unique node called its *parent* and such that for each node n in N there is an integer $k \geqslant 0$ such that

$$p^k(n) = r$$

The terms employed for indicating kinship are commonly used for trees. Thus, we say that n is a *child* of $p(n)$. The notions of an *ancestor* and of a *descendant* are clear also. Formally, n is an ancestor of n' if there is an integer $k \geqslant 0$ such that $n = p^k(n')$. The node n' is then a descendant of n.

If n is a node in a tree (N, r, p), the mapping p induces a mapping

$$p_n : N_n \setminus \{n\} \to N_n,$$

where N_n is the set of descendants of n. The triple (N_n, n, p_n) is by definition the subtree of N rooted at n.

One associates in a natural way a tree with each nonempty prefix-closed set T using the elements of T as nodes, the empty word as a root and by defining for each word u

and each letter a such that $ua \in T$, $p(ua) = u$. Conversely, one may associate with each tree (N, r, p) a prefix-closed set T on the alphabet N by

$$T = \{n_1 \cdots n_k \mid k \geqslant 0, p(n_1) = r \text{ and for every } i \in \{2, \ldots, k\}, p(n_i) = n_{i-1}\}$$

A node u is thus an ancestor of v if u is a prefix of v. And the subtree rooted at a node u is the set $u^{-1}T = \{v \in A^* \mid uv \in T\}$. This allows a convenient graphical representation of prefix-closed sets. For instance, if $A = \{a, b\}$ and if $T = \{1, a, aa, ab, aba, abb, b, bb\}$, the set T is represented in Figure 2.1.

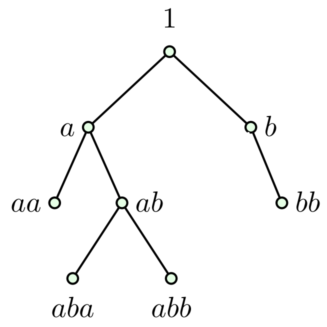

Figure 2.1: A prefix-closed set.

The number of children of a node is the *arity* of this node. The *arity of a tree* is the maximal arity of its nodes. In particular, the tree associated with a prefix-closed set on an alphabet with k elements is a tree of arity at most k. The free monoid A^* itself, which is a prefix-closed set, is associated with a tree represented in Figure 2.2. We shall have

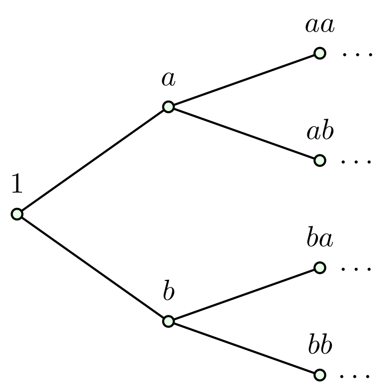

Figure 2.2: The tree of the free monoid $\{a, b\}^*$.

occasions to use three variants of trees. First of all, an *oriented tree* is a tree in which an order relation denoted $\leqslant$ is defined on the set of children of each node. An oriented tree can therefore be defined by a function

$$f : N \to N^*$$

associating to each node the ordered list of its children.

We say that two oriented trees $T = (N, r, f)$ and $T' = (N', r', f')$ are *equivalent* if there is a bijection σ from N onto N' (extending to a bijection from N^* onto N'^*) such that

$$\sigma(r) = r', \text{ and for every } n \in N, \sigma(f(n)) = f'(\sigma(n))$$

which means that σ preserves the order on the nodes (i.e. if $n_1 < n_2$ in T, then $\sigma(n_1) < \sigma(n_2)$). For example, the trees represented in Figure 2.3 are equivalent.

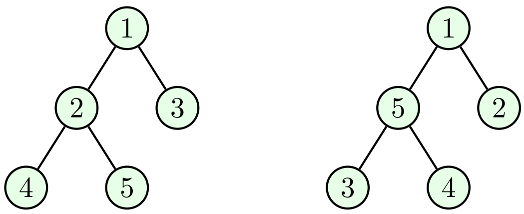

Figure 2.3: Two equivalent trees.

A *planar tree* is an equivalence class of this relation. Planar trees can be represented by a figure without mention of the set N. For example, the planar tree which is the equivalence class of the trees of Figure 2.3 is represented in Figure 2.4.

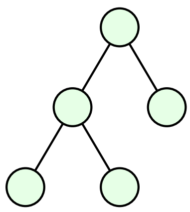

Figure 2.4: A planar tree.

It is relatively easy to compute the number of planar trees with n nodes.

Proposition 2.1. *For each $n > 0$, the number of planar trees with n nodes is the Catalan number $C_n = \dfrac{(2n-2)!}{n!(n-1)!}$.*

Proof. Let P_n be the number of planar trees with n nodes and let, by convention $P_0 = 0$. Let (N, r, f) be an oriented tree with at least two nodes and let z be the leftmost child of r. We obtain a partition of N by considering on one hand the set N_z of descendants of z and, on the other hand, the set $N' = N \setminus N_z$. These two sets define oriented trees (see Figure 2.5). Conversely given an oriented tree (N_z, z, f_z) with k nodes and an oriented tree (N', r, f') with $n-k$ nodes, we can, supposing N_z and N' disjoint, build an oriented tree with n nodes (N, r, f), where $N = N_z \cup N'$ and where

$$f(x) = \begin{cases} f_z(x) & \text{if } x \in N_f, \\ f'(x) & \text{if } x \in N_r \setminus \{r\}, \\ zf'(r) & \text{if } x = r. \end{cases}$$

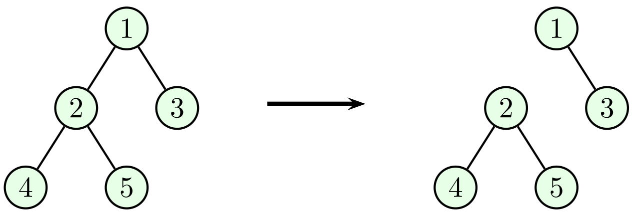

Figure 2.5: Decomposition of a tree.

Since this operation preserves the equivalence of trees we have for all $n \geqslant 2$:

$$P_n = \sum_{1 \leqslant k \leqslant n-1} P_k P_{n-k}$$

Thus, the generating series

$$P(X) = \sum_{n \geqslant 0} P_n X^n$$

satisfies the equation

$$P^2 - P + X = 0$$

The desired formula is obtained by expanding $\frac{1}{2}\big(1 - (1-4X)^{\frac{1}{2}}\big)$. □

Finally a *labeled tree* is a pair (T, e) of a tree T and a function

$$e : N \to V$$

from N into a set V of labels associating with each node $n \in N$ its label $e(n)$.

In particular, we may label the tree associated with a prefix-closed set T on an alphabet A by setting, for all $u \in T$, $e(u) = u$. In this case, a node u is at the left of a node v if $u \leqslant_{lex} v$ and if u is not an ancestor of v.

In a tree, a *path* from a node n_0 to a node n_k is a sequence $(n_0, \ldots, n_k)$ of nodes such that, for $0 \leqslant i \leqslant k-1$, n_i is the parent of n_{i+1}. In a similar way, an *infinite path* starting at n_0 is a sequence $(n_0, n_1, \ldots)$ of nodes such that for all $i \geqslant 0$, n_i is the parent of n_{i+1}.

The following result, known as *König's lemma*, is actually a compactness property (see Exercise III.6).

Proposition 2.2. *An infinite tree in which every node has finite arity contains an infinite path.*

Proof. Let (N, r, f) be an infinite tree in which every node is of finite arity. We build an infinite path starting at the root

$$(n_0, n_1, \ldots)$$

and such that for each k, the subtree rooted at n_k is infinite. Let first $n_0 = r$. Suppose then that n_k has already been chosen for $k \geqslant 0$. Since n_k is of finite arity, the set of its children is finite. There has to be at least one child such that the corresponding subtree is infinite. We choose this child for the node n_{k+1}. □

3 Rational sets of infinite words

In this section, we define the notion of an ω-rational set. It extends the corresponding notion of rational sets of finite words. Let us first recall that one defines on the subsets of A^* the rational operations as

(1) the set union $X \cup Y$,

(2) the set product

$$XY = \{xy \mid x \in X \text{ and } y \in Y\}$$

(3) the star operation defined by

$$X^* = \{x_1 \cdots x_n \mid n \geqslant 0 \text{ and } x_1, \ldots, x_n \in X\}$$

Thus X^* is the submonoid of A^* generated by X.

We also define the "plus" operation by

$$L^+ = \{u_1 \cdots u_n \mid n > 0 \text{ and } u_1, \ldots, u_n \in L\}$$

Thus L^+ is the subsemigroup generated by L. Obviously, $L^* = L^+ \cup \{1\}$ and $L^+ = LL^*$.

The class of *rational subsets* of A^* is the smallest class $\mathcal{R}$ of subsets of A^* such that

(a) $\mathcal{R}$ contains the empty set and for each $a \in A$, the singleton set $\{a\}$,

(b) $\mathcal{R}$ is closed under finite union, finite product and the operation $L \to L^*$.

It is an immediate consequence of the definitions that all finite subsets of A^* are rational. Indeed, since rational sets are closed under finite union, it suffices to verify that every singleton $\{u\}$ is rational. If $u = 1$, we observe that $\{1\} = \emptyset^*$, and if $u = a_1 \cdots a_n$, $\{u\} = \{a_1\} \cdots \{a_n\}$.

Rational sets are described by *rational expressions*, which are expressions using symbols from the alphabet A and the symbols $\cup$, $\cdot$, $^+$, *. To simplify the notation, we often denote a singleton set u instead of $\{u\}$ and the union $X \cup Y$ is also denoted $X + Y$. For instance, the expression $(ab + 1)^+$ is a shorthand for $(\{a\} \cdot \{b\} \cup \{1\})^+$.

Example 3.1. The set of words on the alphabet $A = \{a, b\}$ in which every a is always followed by a b can be described by the rational expression $b^*(ab^+)^*$. It is thus a rational subset of A^*. It is also described by the rational expression $(b + ab)^*$. This example shows that a rational set can be described by several rational expressions.

The definition of the product of two sets $X, Y \subset A^*$ can be extended to a more general case by setting, for $X \subset A^*$ and $Y \subset A^\infty$,

$$XY = \{xy \mid x \in X \text{ and } y \in Y\}$$

Then a restricted form of associativity holds: for each $X, Y \subset A^*$ and $Z \subset A^\infty$,

$$(XY)Z = X(YZ)$$

We now introduce an additional operation expressing infinite iteration. For each $X \subset A^*$, let

$$X^\omega = \{x_0 x_1 \cdots \mid \text{ for all } i \geqslant 0, x_i \in X \setminus \{1\}\}$$

and

$$X^\infty = X^* \cup X^\omega.$$

Thus, X^ω is the set of infinite words obtained by concatenating an infinite sequence of nonempty words of X. In particular, if $u = a_0 a_1 \cdots a_n$ and if $X = \{u\}$, we have $X^\omega = \{u^\omega\}$, where u^ω is the infinite word

$$a_0 a_1 \cdots a_n a_0 a_1 \cdots a_n a_0 a_1 \cdots a_n a_0 a_1 \cdots$$

obtained by repeating u an infinity of times.

The following proposition puts together some useful identities, whose formal proof is left to the reader.

Proposition 3.1. *For all $X, Y \subset A^*$, we have*

(1) $(X + Y)^\omega = (X^*Y)^\omega + (X + Y)^* X^\omega$,
(2) $(XY)^\omega = X(YX)^\omega$,
(3) *for all* $n > 0$, $(X^n)^\omega = (X^+)^\omega = X^\omega$,

(4) $XX^\omega = X^+X^\omega = X^\omega$.

Identity (1) is to be compared with the identity between subsets of A^*:

$$(X+Y)^* = (X^*Y)^*X^*$$

Identity (2) relates the operator $^\omega$ to the product. It is the counterpart of the identity

$$(XY)^* = 1 + X(YX)^*Y$$

which may also be written

$$(XY)^+ = X(YX)^*Y$$

Identity (3) expresses the fact that infinite iteration rules out multiplicities, in contrast with the analogous identity

$$X^* = (1 + X + X^2 + \cdots + X^{n-1})(X^n)^*$$

We can now give the definition of ω-rational subsets of A^∞. The class of *ω-rational* subsets of A^∞ is the smallest set $\mathcal{R}$ of subsets of A^∞ such that

(a) $\emptyset \in \mathcal{R}$ and for all $a \in A$, $\{a\} \in \mathcal{R}$,

(b) $\mathcal{R}$ is closed under finite union,

(c) for each subset X of A^* and for each subset Y of A^∞, $X \in \mathcal{R}$ and $Y \in \mathcal{R}$ imply $XY \in \mathcal{R}$,

(d) for every subset X of A^*, $X \in \mathcal{R}$ implies $X^* \in \mathcal{R}$ and $X^\omega \in \mathcal{R}$.

As a summary, the class of ω-rational subsets of A^∞ is the smallest class of subsets of A^∞ containing the finite subsets of A^* and closed under finite union, finite product and the operations $X \to X^*$ and $X \to X^\omega$.

In the sequel, we shall be especially interested in the ω-rational subsets of A^∞ which are contained in A^ω and which will be called ω-rational subsets of A^ω. There is a simple characterization of these subsets, which can also be used as a definition.

Theorem 3.2. *A subset of A^ω is ω-rational if and only if it is a finite union of sets of the form XY^ω where X and Y are rational subsets of A^*.*

Proof. We denote by $\mathcal{R}at(A^\omega)$ the class of subsets defined in the statement. It is clear that every element of $\mathcal{R}at(A^\omega)$ is an ω-rational subset of A^ω. To prove the converse, we establish a slightly more precise statement: if X is a rational subset of A^∞, then

(1) $X \cap A^*$ is a rational subset of A^*.

(2) $X \cap A^\omega \in \mathcal{R}at(A^\omega)$.

This property reduces obviously to $X \in \mathcal{R}at(A^\omega)$ when $X \subset A^\omega$. Let $\mathcal{E}$ be the class of subsets of A^∞ satisfying (1) and (2). We have successively

(a) $\emptyset \in \mathcal{E}$ and $\{a\} \in \mathcal{E}$ for every $a \in A$,
(b) $\mathcal{E}$ is closed under finite union,
(c) $\mathcal{E}$ is closed under product. In fact, if $X \subset A^*$ and $Y \subset A^\infty$, then $(XY \cap A^*) = X(Y \cap A^*)$, which is rational since Y satisfies condition (1). Also, $(XY \cap A^\omega) = X(Y \cap A^\omega)$, which is in $\mathcal{R}at(A^\omega)$ since, by condition (2), $Y \cap A^\omega \in \mathcal{R}at(A^\omega)$.
(d) $\mathcal{E}$ is closed under the operation $X \to X^*$.
(e) $\mathcal{E}$ is closed under the operation $X \to X^\omega$.

As a result, $\mathcal{E}$ contains the class of rational subsets of A^∞, which proves the theorem. □

Example 3.2. The set X of infinite words on the alphabet $\{a, b\}$ with only a finite number of occurrences of the symbol b is given by the ω-rational expression $X = (a + b)^* a^\omega$. The complement of X in A^ω, which is the set of words with an infinite number of occurrences of b is given by the expression $(a^*b)^\omega$, and is therefore also ω-rational.

Example 3.2 actually presents a particular case of a general result: the set $\mathcal{R}at(A^\omega)$ is closed under all boolean operations. This result will be proved later. The delicate point is the complement since, as it can be seen on the previous example, given an ω-rational expression for a set, it is not easy to find an ω-rational expression for its complement. Computing the intersection of two ω-rational sets is easier and can be done directly (see Exercise 7).

Let A and B be two alphabets. Any application $\varphi : A \to B^+$ defines a unique semigroup morphism $\varphi : A^+ \to B^+$, obtained by setting for a word $u = a_0 a_1 \cdots a_k$, $\varphi(u) = \varphi(a_0)\varphi(a_1) \cdots \varphi(a_k)$. It can also be turned into a monoid morphism from A^* into B^* by setting $\varphi(1) = 1$. It also extends to a mapping $\varphi : A^\infty \to B^\infty$, also called morphism and defined for an infinite word $u = a_0 a_1 \cdots$, by $\varphi(u) = \varphi(a_0)\varphi(a_1) \cdots$.

A semigroup morphism $\varphi : A^+ \to B^+$ will be called *alphabetic* if, for every $a \in A$, $\varphi(a) \in B$. If X is a subset of A^∞, we set

$$\varphi(X) = \{\varphi(u) \mid u \in X\}$$

Then the following formulas hold, where X is a subset of A^+ and where X_1 and X_2 are subsets of A^∞:

$$\begin{aligned} \varphi(X_1 \cup X_2) &= \varphi(X_1) \cup \varphi(X_2), \\ \text{if } X_1 \subset A^+, \quad \varphi(X_1 X_2) &= \varphi(X_1)\varphi(X_2), \\ \varphi(X^+) &= \varphi(X)^+, \\ \varphi(X^\omega) &= \varphi(X)^\omega. \end{aligned}$$

It follows immediately that ω-rational sets are stable under morphism.

Proposition 3.3. *Let A and B be two alphabets and let $\varphi : A^+ \to B^+$ be a morphism. If X is an ω-rational subset of A^∞ (resp. of A^*, A^ω), then $\varphi(X)$ is an ω-rational subset of B^∞ (resp. of B^*, B^ω).*

4 Automata

An *automaton* on the alphabet A is given by a set Q, called the set of *states*, and a subset E of $Q \times A \times Q$, called the set of *edges* or *transitions*. Some additional components may be added, and in particular a subset $I \subset Q$ of *initial* states and a subset $F \subset Q$ of *final* or *terminal* states. The automaton is often denoted as a tuple $= (Q, A, E)$ or $\mathcal{A} = (Q, A, E, I, F)$ if I and F are specified. Part of the components may always be omitted. In particular, we sometimes denote the automaton merely (E, I, F) when Q and A are unambiguously defined. Moreover, if $I = \{i\}$, or if $F = \{f\}$, the automaton is denoted (E, i, F) or (E, I, f).

An automaton is said to be *finite* (resp. *countable*) if both its alphabet and the sets of its states are finite (resp. countable).

Two transitions (p, a, q) and (p', a', q') are called *consecutive* if $q = p'$. A *path* in the automaton $\mathcal{A}$ is a finite sequence of consecutive transitions

$$e_0 = (q_0, a_0, q_1),\ \ e_1 = (q_1, a_1, q_2),\ \ \dots\ \ ,\ \ e_{n-1} = (q_{n-1}, a_{n-1}, q_n)$$

also denoted

$$q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \ \cdots\ q_{n-1} \xrightarrow{a_{n-1}} q_n \qquad \text{or} \qquad q_0 \xrightarrow{a_0 \cdots a_{n-1}} q_n.$$

The state q_0 is the *origin* of the path and the state q_n its *end*. One says that the path *passes through* (or *visits*) the states $q_0, q_1, \dots, q_n$. The word $x = a_0 a_1 \cdots a_n$ is the *label* of the path and the integer $n + 1$ its *length*. The set $\{q_0, q_1, \dots, q_n\}$ is the *content* of the path.

It is convenient to introduce, for each state $q \in Q$, an empty path with origin and end equal to q. Its label is the empty word and its length is 0.

An *infinite path* in the automaton $\mathcal{A}$ is an infinite sequence p of consecutive transitions

$$e_0 = (q_0, a_0, q_1),\ \ e_1 = (q_1, a_1, q_2),\ \ \dots$$

also denoted

$$q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \ \cdots$$

The state q_0 is the *origin* of the infinite path and the infinite word $a_0 a_1 \cdots$ is its *label*. We say that the path p *passes infinitely often* through a state q (or that p *visits* q infinitely often, or yet that q is *infinitely repeated* in p) if there are infinitely many integers n such that $q_n = q$. The set of *infinitely repeated* states in p is denoted by $\mathrm{Inf}(p)$.

We usually specify for each automaton a set of *successful* finite or infinite paths. This will be done for infinite paths in various ways in the next sections. For finite paths, there is just one notion.

A finite path in $\mathcal{A}$ is *initial* if its origin is in I and *final* if its end is in F. A path is *successful* if it is both initial and final. The set of words *recognized* by the automaton

$\mathcal{A}$ is the set, denoted by $L^*(\mathcal{A})$, of all labels of successful paths in $\mathcal{A}$. We also set $L^+(\mathcal{A}) = L^*(\mathcal{A}) \setminus \{1\}$.

A set of finite words X is *recognizable* if there exists a finite automaton $\mathcal{A}$ such that $X = L^*(\mathcal{A})$.

Example 4.1. Let $\mathcal{A} = (Q, A, E, I, F)$ where $Q = \{1, 2\}$, $A = \{a, b\}$ and

$$E = \{(1, a, 1), (2, b, 1), (1, a, 2), (2, b, 2)\},\ \ I = \{1\},\ \ F = \{2\}$$

This automaton is represented in Figure 4.1. According to a convention used in all the sequel, the initial states are indicated by an incoming arrow and the final states by an outgoing one.

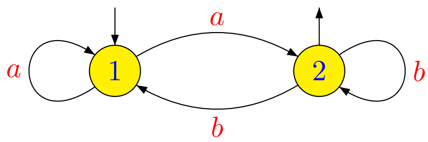

Figure 4.1: The automaton $\mathcal{A}$.

We have $L^+(\mathcal{A}) = a\{a, b\}^*$, which is the set of finite words beginning with an a. In this automaton, every finite word is the label of exactly two paths. Indeed, every letter determines the state it comes from and conversely, the transitions going out of a given state all have the same label. This automaton can thus be interpreted as a mechanism predicting the next symbol. Every infinite word is the label of exactly one path.

An automaton $\mathcal{A} = (Q, A, E, I, F)$ is said to have *deterministic transitions*, if, for every state $q \in Q$ and every letter $a \in A$, there is at most one state q' such that (q, a, q') is a transition. It is *deterministic* if it has deterministic transitions and if I is a singleton. It is *complete* if, for every state $q \in Q$ and every letter $a \in A$, there is at least one state q' such that (q, a, q') is a transition. If q_0 is the unique initial state, we adopt the notation (Q, A, E, q_0, F) instead of $(Q, A, E, \{q_0\}, F)$.

For instance, the automaton of Example 4.1 is neither complete nor deterministic. In contrast, the automaton represented in Figure 4.2 is complete and deterministic.

Example 4.2. The automaton given by Figure 4.2 is deterministic.

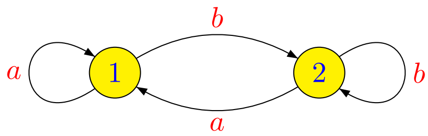

Figure 4.2: A complete deterministic automaton.

Each word is the label of exactly two paths, one going out of state 1 and the other from state 2. The automaton can be considered as "remembering the last symbol" since each state is accessible only by one symbol.

If $\mathcal{A} = (Q, A, E)$ is deterministic, we define a partial function from $Q \times A$ into Q by associating with each pair (q, a) in $Q \times A$ the unique state $q\cdot a$ (when it exists), such that $(q, a, q\cdot a) \in E$. If there is no q' such that $(q, a, q') \in E$, the image of (q, a) is not defined. The partial function $(q, a) \to q\cdot a$ thus defined is the *transition function* of the automaton. It is clear that a deterministic automaton is defined by its transition function.

The transition function can be extended to a partial function from $Q \times A^*$ into Q by setting $q\cdot 1 = q$ and, for every word $u \in A^+$ and for every symbol $a \in A$, $q\cdot(ua) = (q\cdot u)\cdot a$ if $(q\cdot u)$ and $(q\cdot u)\cdot a$ are defined. For example, if $\mathcal{A}$ is the automaton of Example 4.2, we have $1\cdot bbaba = 2\cdot bbaba = 1$.

One may also define the *transition function* of a nondeterministic automaton $\mathcal{A} = (Q, A, E)$. It is the function from $\mathcal{P}(Q) \times A^*$ into $\mathcal{P}(Q)$, traditionally denoted δ, defined by the following formula, where $S \subset Q$ and $u \in A^*$,

$$\delta(S, u) = \{q \in Q \mid q \text{ is the end of a path in } \mathcal{A} \text{ with label } u \text{ going out from some state of } S\}$$

It is clear that knowing E or δ is the same. We shall use one of either notation in the sequel. This transition function allows one to define the *deterministic version* of an automaton (Q, A, E, I, F) as the complete deterministic automaton

$$(\mathcal{P}(Q), A, \delta, I, \mathcal{F}) \text{ with } \mathcal{F} = \{P \subset Q \mid P \cap F \neq \emptyset\}$$

This construction is motivated by the following result, which shows that, for finite words, deterministic automata have the same expressive power as non deterministic ones. We shall see in the next sections that this elementary result does not extend easily to infinite words.

Proposition 4.1. *An automaton and its deterministic version recognize the same set of finite words.*

Proof. Let $\mathcal{A} = (Q, A, E, I, F)$ and let $\mathcal{B}$ be its deterministic version. If $u = a_0 \cdots a_{n-1}$ is recognized by $\mathcal{A}$, there is a successful path

$$c : q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \;\cdots\; q_{n-1} \xrightarrow{a_{n-1}} q_n$$

Define a sequence of subsets of Q by setting $P_0 = I$, $P_1 = P_0\cdot a_0$, $\ldots$, $P_n = P_{n-1}\cdot a_{n-1}$. Since c is a successful path, $q_0 \in I = P_0$ and $q_n \in F$. Suppose, by induction, that

$q_{i-1} \in P_{i-1}$. Then since $q_{i-1} \xrightarrow{a_i} q_i$ is a transition, $q_i \in P_i$. In particular, $q_n \in P_n \cap F$. Therefore $P_n \cap F \neq \emptyset$ and $P_n \in \mathcal{F}$. Consequently, u is accepted by $\mathcal{B}$.

Conversely, let $u = a_0 \cdots a_{n-1}$ be a word accepted by $\mathcal{B}$. Set, as above,

$$P_0 = I,\ P_1 = P_0 \cdot a_0,\ \ldots, P_n = P_{n-1} \cdot a_{n-1}.$$

Since $P_n \in \mathcal{F}$, one may choose an element q_n in $P_n \cap F$, and, for $0 \leqslant i \leqslant n$, an element $q_i \in P_i$ such that $q_i \xrightarrow{a_i} q_{i+1}$ is a transition in $\mathcal{A}$. Since $q_0 \in I$ and $q_n \in F$, the path $q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \ \cdots\ q_{n-1} \xrightarrow{a_{n-1}} q_n$ is successful, and u is accepted by $\mathcal{A}$. □

An important consequence of Proposition 4.1 is that recognizable sets are closed under complementation. Actually, they are closed under any boolean operation, but we need an auxiliary definition before proving this result. The *product* of two automata $\mathcal{A} = (Q, A, E)$ and $\mathcal{A}' = (Q', A, E')$ is the automaton

$$\mathcal{A} \times \mathcal{A}' = (Q \times Q', A, P)$$

with

$$P = \big\{\big((p, p'), a, (q, q')\big) \mid (p, a, q) \in E \text{ and } (p', a, q') \in E'\big\}.$$

Proposition 4.2. *Recognizable subsets of A^* are closed under finite union, finite intersection and set difference.*

Proof. Let X and X' be two recognizable subsets of A^*. By Proposition 4.1, we may assume that $X = L^*(\mathcal{A})$ and $X' = L^*(\mathcal{A}')$ for some deterministic complete automata $\mathcal{A} = (Q, A, \cdot, i, F)$ and $\mathcal{A}' = (Q', A, \cdot, i', F')$. Let $\mathcal{B}$ be the product of $\mathcal{A}$ and $\mathcal{A}'$. Equipped with the initial state (i, i'), $\mathcal{B}$ is a deterministic automaton. Let R be a set of final states for $\mathcal{B}$. If $R = (F \times Q') \cup (Q \times F')$, we have $L^*(\mathcal{B}) = X \cup X'$. If $R = F \times F'$, we have $L^*(\mathcal{B}) = X \cap X'$. Finally, if $R = F \times (Q \setminus F')$, we have $L^*(\mathcal{B}) = X \setminus X'$. □

The *minimal deterministic automaton* of a set of finite words is the deterministic automaton $\mathcal{A} = (Q, A, E, i, T)$ where Q is the set of nonempty classes of the *Nerode equivalence* $\sim$, defined for $u, v \in A^*$ by $u \sim v$ if, for all $w \in A^*$, one has

$$uw \in X \implies vw \in X.$$

The initial state i is the class of the empty word and a state q is final if all its elements are in X. Finally, for every $u \in A^*$ and $a \in A$, there is an edge labeled by a from the class of u to the class of ua.

Let $\mathcal{A} = (Q, A, E)$ be an automaton. The *reversed* automaton of $\mathcal{A}$ is the automaton $\mathcal{A}^r = (Q, A, E^r)$, where

$$E^r = \{(q, a, p) \mid (p, a, q) \in E\},$$

obtained by reverting the arrows of $\mathcal{A}$. An automaton $\mathcal{A}$ has *co-deterministic* transitions if the automaton $\mathcal{A}^r$ has deterministic transitions. Thus, the automaton of Figure 4.1 has co-deterministic transitions.

An automaton $\mathcal{A} = (Q, A, E)$ is *unambiguous* if, for each pair of states $(p, q) \in Q \times Q$ and for each word x, there is at most one path from p to q with label x. It is easy to verify that any automaton with deterministic or co-deterministic transitions is unambiguous, but the converse is not true as shown by the following example. Let $\mathcal{A} = (Q, A, E)$ be an automaton. We define its *transition matrix* as follows. It is the $Q \times Q$-matrix T with coefficients in $\mathcal{P}(A^*)$ defined by

$$T_{p,q} = \{a \in A \mid (p, a, q) \in E\}.$$

The set $\mathcal{P}(A^*)$ of subsets of A^* comes with two operations: the union, which we shall denote additively and the product

$$XY = \{xy \mid x \in X \text{ and } y \in Y\}$$

The set $\mathcal{P}(A^*)^{Q\times Q}$ of $Q \times Q$-matrices with coefficients in $\mathcal{P}(A^*)$ comes in turn with an addition and a product by setting, for any $M, N \in \mathcal{P}(A)^{Q\times Q}$ and $p, q \in Q$

$$(M + N)_{p,q} = M_{p,q} + N_{p,q}$$

and

$$(MN)_{p,q} = \sum_{r\in Q} M_{p,r} N_{r,q}$$

The fact that the above union may have an infinite number of terms is not an obstacle, since infinite unions are well-defined. We note that the identity matrix defined as

$$1_{p,q} = \begin{cases} \{1\} & \text{if } p = q, \\ \emptyset & \text{otherwise} \end{cases}$$

is the neutral element of the product.

Given a transition matrix T, we define a new matrix T^* by

$$T^* = \sum_{n\geqslant 0} T^n = 1 + T + T^2 + \cdots$$

The reader can verify that

$$T^*_{p,q} = \{u \in A^* \mid \text{ there exists a path from } p \text{ to } q \text{ with label } u\}.$$

It can be shown that if a matrix T has all its coefficients in $\mathcal{P}(A^+)$, as it is the case here, then $X = T^*$ is the unique solution of the equation

$$X = 1 + XT. \tag{4.1}$$

We recall the following statement, which is classical in automata theory.

Proposition 4.3. *If $\mathcal{A}$ is a finite automaton, the coefficients of T^* are rational subsets of A^*.*

Proof. We use an induction on the number n of states of $\mathcal{A}$. The case $n = 0$ is trivial. If $n = 1$, we have $T^*_{1,1} = B^*$, with

$$B = \{a \mid (1, a, 1) \in E\}.$$

Since E is finite, B is finite and $T_{1,1}$ is a rational subset of A^*. If $n > 1$, let us consider a partition of the matrix T into blocks

$$\begin{pmatrix} U & V \\ W & Z \end{pmatrix}$$

where U and Z are square matrices. Let, for the same partition,

$$T^* = \begin{pmatrix} U' & V' \\ W' & Z' \end{pmatrix}.$$

Then

$$\begin{aligned} U' &= (U + VZ^*W)^* & V' &= U'VZ^* \\ W' &= Z'WU^* & Z' &= (Z + WU^*V)^* \end{aligned}$$

since these formulas allow to verify that T^* is a solution of Equation 4.1. As a result, the coefficients of the matrices U', V', W', Z' are rational, and thus so are those of T^*. □

Example 4.3. The matrices T and T^* corresponding to Example 4.2 are

$$T = \begin{pmatrix} a & b \\ a & b \end{pmatrix} \quad \text{and}$$

$$T^* = \begin{pmatrix} (a + bb^*a)^* & (a + bb^*a)^*bb^* \\ (b + aa^*b)^*aa^* & (b + aa^*b)^* \end{pmatrix} = \begin{pmatrix} (b^*a)^* & (b^*a)^*bb^* \\ (a^*b)^*aa^* & (a^*b)^* \end{pmatrix}$$

We now recall the statement of Kleene and give a sketch of the proof.

Theorem 4.4. (Kleene) *A set of finite words is recognizable if and only if it is rational.*

Proof. By Proposition 4.3 we already know that every recognizable set is rational. To prove the converse, we prove a lemma which will also be used for the case of infinite words.

An automaton is said to be *normalized* if it has only one initial state i and only one final state f and if no transition ends in i or starts from f. The following lemma shows that, for subsets of A^+, one may always replace an automaton by a normalized one.

Lemma 4.5. *For any finite automaton $\mathcal{A}$, there is a finite normalized automaton $\mathcal{A}'$ such that $L^+(\mathcal{A}) = L^+(\mathcal{A}')$.*

Proof. Let $\mathcal{A} = (Q, A, E, I, F)$ and $\mathcal{A}' = (Q \cup \{i', f'\}, A, E', \{i'\}, \{f'\})$, where i' and f' are two new states and where $E' = E \cup E_0 \cup E_1 \cup E_2$ with

$$\begin{aligned} E_0 &= \{(i', a, f') \mid \text{ there exists } i \in I \text{ and } f \in F \text{ such that } (i, a, f) \in E\} \\ E_1 &= \{(i', a, q) \mid \text{ there exists } i \in I \text{ such that } (i, a, q) \in E\} \\ E_2 &= \{(q, a, f') \mid \text{ there exists } f \in F \text{ such that } (q, a, f) \in E\} \end{aligned}$$

Then $\mathcal{A}'$ is normalized by construction. Moreover, if a nonempty word $u = a_1 \cdots a_n$ is recognized by $\mathcal{A}$, there exists a nonempty path

$$q_0 \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \ \cdots \ q_{n-1} \xrightarrow{a_n} q_n$$

starting in I and ending in F. Thus,

$$i' \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \ \cdots \ q_{n-1} \xrightarrow{a_n} f'$$

is a path in $\mathcal{A}'$, which shows that u is recognized by $\mathcal{A}'$. Conversely, if u is recognized by $\mathcal{A}'$, there exists a path in $\mathcal{A}'$ of the form

$$i' \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \ \cdots \ q_{n-1} \xrightarrow{a_n} f'$$

Thus there exist $i \in I$ and $f \in F$ such that

$$i \xrightarrow{a_1} q_1 \xrightarrow{a_2} q_2 \ \cdots \ q_{n-1} \xrightarrow{a_n} f$$

is a path in $\mathcal{A}$ and u is recognized by $\mathcal{A}$. □

Consider now two subsets X and X' of A^*, recognized respectively by the normalized finite automata $\mathcal{A} = (E, i, f)$ and $\mathcal{A}' = (E', i', f')$. We may suppose that E and E' are disjoint. Let then $\mathcal{B}$ be the automaton $\mathcal{B} = (E \cup E', \{i, i'\}, \{f, f'\})$ represented in Figure 4.3. This automaton recognizes $X \cup X'$.

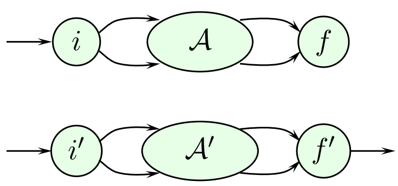

Figure 4.3: An automaton recognizing the union.

To build an automaton $\mathcal{C}$ recognizing the product XX' we merge f and i', as shown in Figure 4.4.

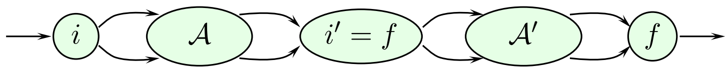

Figure 4.4: An automaton recognizing the product.

We finally build an automaton recognizing X^* by merging i and f in the automaton $\mathcal{A}$, as represented in Figure 4.5.

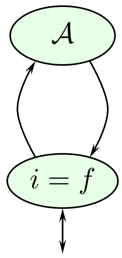

Figure 4.5: An automaton recognizing X^*.

Lemma 4.5 shows that every recognizable set of A^+ is recognized by a normalized automaton. Furthermore, we have seen that the union, the product and the star of such sets is still recognizable.

Working with normalized automata simplified the previous constructions. However, the price to pay is a special treatment for the empty word. We first observe that the set $\{1\}$ is recognized by the automaton with a single state, both initial and final, and no transition, represented in figure 4.6.

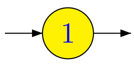

Figure 4.6: An automaton recognizing $\{1\}$.

It follows, by Proposition 4.2, that a set X of finite words is recognizable if and only if $X \setminus \{1\}$ is recognizable. We can now conclude that the class of recognizable sets of A^* is closed under finite union, product and star. For instance, for the product it follows from the following elementary formula, in which $\varepsilon(X) = X \cap \{1\}$:

$$XX' = (X \cap A^+)(X' \cap A^+) \cup \varepsilon(X)(X' \cap A^+) \cup (X \cap A^+)\varepsilon(X') \cup \varepsilon(X)\varepsilon(X')$$

Finally, the empty set is recognized by the empty automaton and if a is a letter, the set $\{a\}$ is recognized by the automaton (E, I, F) with $E = \{(1, a, 2)\}$, $I = \{1\}$ and $F = \{2\}$, represented in Figure 4.7.

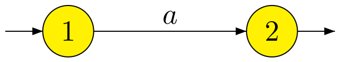

Figure 4.7: An automaton recognizing $\{a\}$.

Thus any rational subset of A^* is recognizable and this concludes the proof of Kleene's theorem. □

5 Büchi automata

We now introduce Büchi automata, which correspond to the simplest recognizing mode for infinite words. A *Büchi automaton* is a 5-tuple $\mathcal{A} = (Q, A, E, I, F)$ where

(1) (Q, A, E) is an automaton,

(2) I and F are subsets of Q, called resp. set of *initial states* and set of *final states*.

Let $\mathcal{A} = (Q, A, E, I, F)$ be a Büchi automaton. We say that an infinite path in $\mathcal{A}$ is *initial* if its origin is in I and *final* if it visits F infinitely often. It is *successful* if it is initial and final. The set of infinite words *recognized* by $\mathcal{A}$ is the set, denoted by $L^\omega(\mathcal{A})$, of labels of infinite successful paths in $\mathcal{A}$. In the case where F is *finite* and in particular if $\mathcal{A}$ is a finite automaton, $L^\omega(\mathcal{A})$ is also the set of labels of infinite initial paths p in $\mathcal{A}$ and such that $\mathrm{Inf}(p) \cap F \neq \emptyset$.

A set X of infinite words is *recognizable* if there exists a finite Büchi automaton $\mathcal{A}$ such that $X = L^\omega(\mathcal{A})$.

A Büchi automaton $\mathcal{A} = (Q, A, E, I, F)$ is said to be *deterministic* if it has deterministic transitions and if I is a singleton, i.e. if $\mathcal{A}$ contains exactly one initial state i. In this case, every word in A^+ (resp. A^ω) is the label of at most one initial path. In particular, every word in $L^\omega(\mathcal{A})$ is the label of exactly one initial path.

A Büchi automaton $\mathcal{A} = (Q, A, E, I, F)$ is called *co-deterministic* if it has co-deterministic transitions and if any word in A^ω is the label of at most one final path. It is *co-complete* if any word in A^ω is the label of at least one final path.

More generally, a Büchi automaton $\mathcal{A} = (Q, A, E, I, F)$ is said to be ω*-unambiguous* if every word in A^ω is the label of at most one infinite successful path. In particular, every word in $L^\omega(\mathcal{A})$ defines a unique successful infinite path of which it is the label. It is clear that any deterministic or co-deterministic Büchi automaton is ω-unambiguous, but the converse is not true. The various terms are summarized in Table 5.3.

Det. transitions	**Co-det. transitions**	**Unambiguous**
Forbidden configuration: 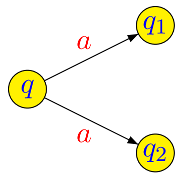 where a is a letter.	Forbidden configuration: 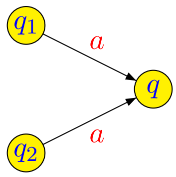 where a is a letter.	Forbidden configuration: 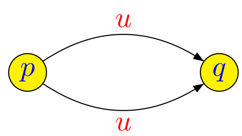 where u is a word.
Deterministic	**Co-deterministic**	**ω-unambiguous**
Deterministic transitions and exactly one initial state.	Co-det. transitions and two final paths with the same label are equal.	Every infinite word is the label of at most one successful path.
Complete	**Co-complete**	
Every word is the label of some initial path.	Every word is the label of some final path.	

Table 5.3: Summary of the definitions.

Example 5.1. Let $\mathcal{A}$ be the Büchi automaton obtained from the automaton of Example 4.1.

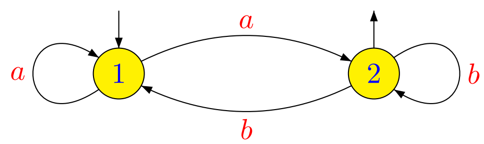

Figure 5.1: A co-deterministic Büchi automaton.

We have $L^\omega(\mathcal{A}) = a(a^*b)^\omega$, which is the set of infinite words starting with a and containing an infinite number of occurrences of b.

Example 5.2. Let $\mathcal{A}$ be the Büchi automaton represented in Figure 5.2.

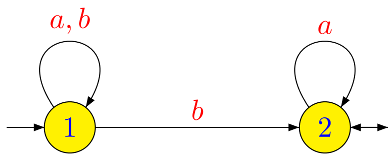

Figure 5.2: Another co-deterministic Büchi automaton.

We have now $L^+(\mathcal{A}) = \{a, b\}^+$ and $L^\omega(\mathcal{A}) = \{a, b\}^*a^\omega$. This automaton is also co-deterministic.

Example 5.3. The Büchi automaton represented in Figure 5.3 is ω-unambiguous but it is neither deterministic nor co-deterministic.

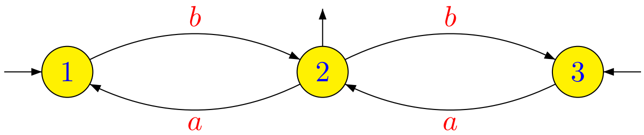

Figure 5.3: An ω-unambiguous automaton.

It is not deterministic since it has two initial states and it is not co-deterministic since the infinite word $(ab)^\omega$ is the label of two final paths, one starting at state 2, the other one at state 3.

We shall see later on that a recognizable set of infinite words is completely determined by its ultimately periodic words. The following lemma proves the easy part of this result.

Lemma 5.1. *Any nonempty recognizable subset of A^ω contains an ultimately periodic word.*

Proof. Let X be a nonempty recognizable subset of A^ω recognized by a Büchi automaton $\mathcal{A} = (E, I, F)$. Since X is nonempty, there exists a path of the form $p = p_0p_1p_2\cdots$, where p_0 starts in I, ends in a state $q \in F$, and where $p_1, p_2, \cdots$ are paths from q to q. The path $p_0p_1p_1p_1\cdots$ is thus also a successful path and its label is an ultimately periodic word. □

Let $\mathcal{A} = (E, I, F)$ be a Büchi automaton. A state q is called *accessible* if there is a (possibly empty) finite initial path in $\mathcal{A}$ ending in q. A state q is called *coaccessible* if there exists an infinite final path starting at q. Finally, $\mathcal{A}$ is *trim* if all its states are both accessible and coaccessible.

Suppressing all "useless" states of a Büchi automaton always gives a trim automaton. More formally, one has the following result.

Proposition 5.2. *With any Büchi automaton $\mathcal{A}$, one may associate a trim Büchi automaton $\mathcal{A}'$ such that*

(1) *The automata $\mathcal{A}$ and $\mathcal{A}'$ recognize the same subset of A^ω,*
(2) *if $\mathcal{A}$ is deterministic, so is $\mathcal{A}'$,*
(3) *if $\mathcal{A}$ is finite, so is $\mathcal{A}'$.*

Proof. (1) Let $\mathcal{A} = (Q, A, E, I, F)$ be a Büchi automaton and let P be the set of states of $\mathcal{A}$ that are both accessible and coaccessible. Let $\mathcal{A}' = (P, A, E', I \cap P, F \cap P)$ where $E' = E \cap (P \times A \times P)$. It is clear that $L^\omega(\mathcal{A}') \subset L^\omega(\mathcal{A})$. Conversely, let $u = a_0 a_1 \cdots \in L^\omega(\mathcal{A})$. There is a final path

$$p = q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \ \cdots$$

with label u such that $q_0 \in I$. The states $q_0, q_1, \ldots$ are then both accessible and coaccessible. Thus, p is a path in $\mathcal{A}'$ and $u \in L^\omega(\mathcal{A}')$. We conclude that $L^\omega(\mathcal{A}) = L^\omega(\mathcal{A}')$.

(2) If $\mathcal{A}$ is deterministic, the automaton (P, A, E') is deterministic. Moreover, if the set recognized by $\mathcal{A}$ is nonempty, there exists a final path starting at the unique initial state. Thus, the initial state is coaccessible and $I \cap P$ is a singleton. Therefore $\mathcal{A}'$ is deterministic.

(3) Finally it is clear that these constructions preserve the finiteness of the automaton. □

A Büchi automaton can, in the same way, be made complete.

Proposition 5.3. *A subset of A^ω recognized by a Büchi automaton $\mathcal{A}$ can be recognized by a complete Büchi automaton $\mathcal{A}'$ such that if $\mathcal{A}$ is finite (resp. deterministic), then $\mathcal{A}'$ is also finite (resp. deterministic).*

Proof. Let $\mathcal{A} = (Q, A, E, I, F)$ be a Büchi automaton recognizing a subset X of A^ω. If $\mathcal{A}$ is not complete, we add a new state p and we create a transition (q, a, p) if there is no transition of the form (q, a, q') in $\mathcal{A}$. More formally, let $\mathcal{A}' = (Q \cup \{p\}, A, E', I, F)$, where p is a new state and $E' = E \cup E_1 \cup E_2$ with

$$E_1 = \{(p, a, p) \mid a \in A\}$$
$$E_2 = \{(q, a, p) \mid q \in Q, a \in A \text{ and } (\{q\} \times \{a\} \times Q) \cap E = \emptyset\}$$

The automaton $\mathcal{A}'$ still recognizes X, is complete and it is deterministic (resp. finite) if $\mathcal{A}$ is deterministic (resp. finite). □

Example 5.4. The Büchi automaton of Example 5.2 can be completed as follows

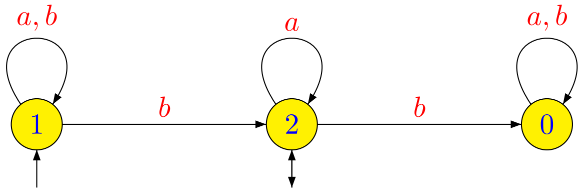

Figure 5.4: Completion of the automaton of Figure 5.2.

We now prove the analogue of Kleene's theorem for infinite words.

Theorem 5.4. *A subset of A^ω is recognizable if and only if it is ω-rational.*

Proof. First consider a recognizable subset X of A^ω and let $\mathcal{A} = (E, I, F)$ be a finite Büchi automaton recognizing X. Then

$$X = L^\omega(E, I, F) = \bigcup_{i \in I} \bigcup_{f \in F} L^*(E, i, f)\big(L^+(E, f, f)\big)^\omega$$

which shows that X is ω-rational, since, by Kleene's theorem, the sets $L^*(E, i, f)$ and $L^+(E, f, f)$ are rational subsets of A^*.

Conversely, let us first consider an ω-rational set Y of the form $X(X')^\omega$ with X, X' rational subsets of A^*. Let $\mathcal{A} = (Q, A, E, i, f)$ and $\mathcal{A}' = (Q', A, E', i', f')$ be two normalized automata such that $X = L^+(\mathcal{A})$ and $X' = L^+(\mathcal{A}')$. We build an automaton $\mathcal{B}$ by merging the states f, i' and f' as indicated in Figure 5.5.

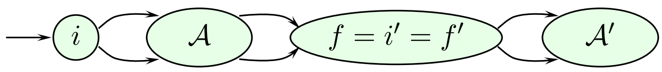

Figure 5.5: An automaton recognizing $X(X')^\omega$.

Formally, we have $\mathcal{B} = ((Q \cup Q') \setminus \{i', f'\}, T, i, f)$, where $T = E \cup E_0 \cup E_1 \cup E_2$ with

$$E_0 = \{(f, a, f) \mid (i', a, f') \in E'\}$$
$$E_1 = \{(f, a, q) \mid q \in Q' \setminus \{i', f'\} \text{ and } (i', a, q) \in E'\}$$
$$E_2 = \{(q, a, f) \mid q \in Q' \setminus \{i', f'\} \text{ and } (q, a, f') \in E'\}$$

This shows that $X(X')^\omega$ is recognizable.

To complete the proof of the theorem, we only have to prove that the class of recognizable subsets of A^ω is stable under finite union.

Let then Y and Y' be two recognizable subsets of A^ω, recognized resp. by the finite Büchi automata $\mathcal{A} = (Q, A, E, I, F)$ and $\mathcal{A}' = (Q', A, E', I', F')$. We may suppose that

Q and Q' are disjoint and thus we may identify E and E' with subsets of $(Q \cup Q') \times A \times (Q \cup Q')$. With this convention, we have the formula

$$Y \cup Y' = L^\omega(E, I, F) \cup L^\omega(E', I', F') = L^\omega(E \cup E', I \cup I', F \cup F')$$

and thus $Y \cup Y'$ is recognized by the finite automaton $(Q \cup Q', A, E \cup E', I \cup I', F \cup F')$. $\square$

We conclude this section with an additional closure property of the class of recognizable sets of A^ω. We have already seen (Proposition 3.3) that the class of ω-rational sets is closed under morphism. We consider now inverse morphisms. Let $\varphi : A \to B^+$ be a function and let $\varphi : A^\omega \to B^\omega$ be the morphism induced by φ.

Proposition 5.5. *If X is a recognizable subset of B^ω, then $\varphi^{-1}(X)$ is a recognizable subset of A^ω.*

Proof. Let $\mathcal{B} = (Q, B, E, I, F)$ be a Büchi automaton recognizing X. Let

$$\mathcal{A} = (Q \times \{0, 1\}, A, E', I \times \{0\}, Q \times \{1\}),$$

where $E' = E_1 \cup E_2$, with

$$E_1 = \Big\{ \big((q, i), a, (q', 0)\big) \mid i \in \{0, 1\} \text{ and there exists a path of } \mathcal{B} \text{ from } q \text{ to } q', \\ \text{labeled } \varphi(a) \text{, visiting no state of } F \Big\}$$

and

$$E_2 = \Big\{ \big((q, i), a, (q', 1)\big) \mid i \in \{0, 1\} \text{ and there exists a path of } \mathcal{B} \text{ from } q \text{ to } q', \\ \text{labeled } \varphi(a) \text{, visiting a state of } F \Big\}.$$

Let $u = a_0 a_1 \cdots$ be a word in A^ω and let for all $n \geqslant 0$, $v_n = \varphi(a_n)$. Then u is accepted by $\mathcal{A}$ if and only if there exists a successful path in $Q \times \{0, 1\}$ and thus a path

$$(q_0, 0) \xrightarrow{a_0} (q_1, \varepsilon_1) \xrightarrow{a_1} (q_2, \varepsilon_2) \cdots$$

labeled u passing infinitely often in $Q \times \{1\}$. By the definition of E', this means that there exists a sequence of consecutive paths in $\mathcal{B}$:

$$q_0 \xrightarrow{v_0} q_1 \xrightarrow{v_1} q_2 \cdots$$

such that an infinite number of the paths $q_n \xrightarrow{v_n} q_{n+1}$ visits a state of F. This is equivalent to the fact that the infinite path $q_0 \xrightarrow{v_0} q_1 \xrightarrow{v_1} q_2 \cdots$ is successful in $\mathcal{B}$, i.e. that $\varphi(u)$ is in X. $\square$

Example 5.5. Let $\varphi : \{a, b, c\}^+ \to \{a, b\}^+$ be the semigroup morphism defined by $\varphi(a) = a$, $\varphi(b) = ab$, $\varphi(c) = babaab$ and let $X = (ababa)^\omega$. The set X is recognized by the Büchi automaton represented in Figure 5.6:

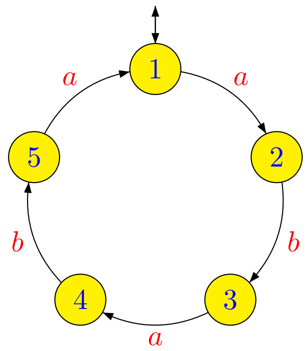

Figure 5.6: An automaton recognizing $(ababa)^\omega$.

Applying the construction described above, we obtain after deletion of some useless states, the following automaton recognizing $\varphi^{-1}(X) = \{acba, bba\}^\omega$.

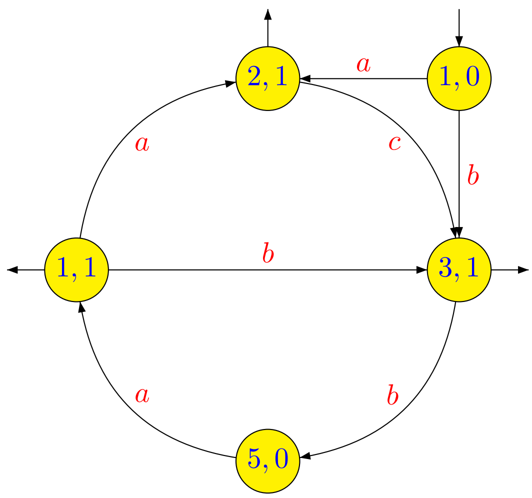

Figure 5.7: An automaton recognizing $\varphi^{-1}((ababa)^\omega)$.

6 Deterministic Büchi automata

It is well known that a set of finite words recognized by a finite automaton can also be recognized by a deterministic one. The situation is quite different for infinite words and we will see that (as soon as A has at least two symbols) there are recognizable subsets of A^ω which cannot be recognized by a finite deterministic Büchi automaton. Actually, this difference between deterministic and nondeterministic automata even holds for countable automata.

The description of the subsets of A^ω recognized by deterministic Büchi automata comes with the introduction of a new operator. For a subset L of A^*, let

$$\overrightarrow{L} = \{u \in A^\omega \mid u \text{ has infinitely many prefixes in } L\}.$$

The operator $L \to \overrightarrow{L}$ plays a role similar to that of the operator $L \to L^\omega$, since it allows one to define infinite words from finite ones. Comparing the properties of these two operators takes an important part in what follows.

The following example gives the value of $\overrightarrow{L}$ for simple sets L and it can help the reader to get more familiar with this operator.

Example 6.1.

(a) If $L = a^*b$, then $\overrightarrow{L} = \emptyset$.
(b) If $L = (ab)^+$, then $\overrightarrow{L} = (ab)^\omega$.
(c) If $L = (a^*b)^+ = (a+b)^*b$, that is if L is the set of words ending with b, then $\overrightarrow{L} = (a^*b)^\omega$, which is the set of infinite words containing an infinity of occurrences of b.

We now give a simple example showing that not every set of words can be written in the form $\overrightarrow{L}$.

Example 6.2. The set $X = (a+b)^*a^\omega$ of words with a *finite* number of occurrences of b is not of the form $\overrightarrow{L}$. Otherwise, the word ba^ω would have a prefix $u_1 = ba^{n_1}$ in L, the word $ba^{n_1}ba^\omega$ would have a prefix $u_2 = ba^{n_1}ba^{n_2}$ in L, etc. and the infinite word $u = ba^{n_1}ba^{n_2}ba^{n_3}\cdots$ would have an infinity of prefixes in L. This word would therefore be in $\overrightarrow{L}$, which is impossible, since u contains infinitely many b's.

The following statement shows that the operator $L \to \overrightarrow{L}$, just as the operator $L \to L^\omega$, preserves the class of recognizable sets.

Proposition 6.1. *Let $\mathcal{A}$ be a deterministic Büchi automaton. Then $L^\omega(\mathcal{A}) = \overrightarrow{L^+(\mathcal{A})}$.*

Proof. Let $\mathcal{A} = (Q, A, E, i, F)$. If $u \in L^\omega(\mathcal{A})$, then u is the label of a path

$$p = (q_0, a_0, q_1)(q_1, a_1, q_2)\cdots$$

such that $q_0 = i$ and such that there exists a subsequence $n_0 < n_1 < \cdots$ satisfying $q_{n_0}, q_{n_1}, \ldots \in F$. By construction, the words $u_k = a_0a_1\cdots a_{n_k-1}$ are in $L^+(\mathcal{A})$ and are prefixes of u. Thus $L^\omega(\mathcal{A}) \subset \overrightarrow{L^+(\mathcal{A})}$.

Conversely, if $u \in \overrightarrow{L^+(\mathcal{A})}$, u has infinitely many prefixes in $L^+(\mathcal{A})$. And since $\mathcal{A}$ is deterministic, we deduce that u is the label of an initial path passing infinitely often in F. Thus $u \in L^\omega(\mathcal{A})$. □

We now study the subsets of A^ω which can be recognized by a deterministic Büchi automaton, which will be called *deterministic*. Informally, a subset X of A^ω is deterministic if testing whether a word belongs to X can be made in one left-to-right pass. A first characterization of deterministic sets is given below. We shall give others later, in particular in connection with topology.

Theorem 6.2. *Let X be a subset of A^ω. The following conditions are equivalent:*

(1) *X can be recognized by a deterministic Büchi automaton,*

(2) *there exists a set L of A^+ such that $X = \overrightarrow{L}$.*

If, moreover, the alphabet A is countable, these conditions are equivalent to

(3) *X can be recognized by a countable deterministic Büchi automaton.*

Proof. If $\mathcal{A}$ is a deterministic Büchi automaton recognizing X, then $X = L^\omega(\mathcal{A}) = \overrightarrow{L^+(\mathcal{A})}$ by Proposition 6.1. Thus (1) implies (2).

Let now L be a subset of A^+ such that $X = \overrightarrow{L}$. Then L is recognized by the deterministic automaton $\mathcal{A} = (A^*, A, \cdot, 1, L)$, where the transition function is defined, for all $u \in A^*$ and for every $a \in A$, by $u.a = ua$. By Proposition 6.1, we have $L^\omega(\mathcal{A}) = \overrightarrow{L^+(\mathcal{A})} = \overrightarrow{L} = X$. Thus (2) implies (1). If, moreover, the alphabet A is countable, $\mathcal{A}$ is a countable automaton. Thus, in this case, (2) implies (3), which establishes the equivalence of the three conditions since (3) obviously implies (1). □

The analogous statement for recognizable sets is given below.

Corollary 6.3. *Let X be a subset of A^ω. The following conditions are equivalent:*

(1) *X is recognized by a finite deterministic Büchi automaton.*

(2) *There exists a recognizable subset L of A^+ such that $X = \overrightarrow{L}$.*

Proof. The proof of Theorem 6.2 can be adapted for finite automata. In the proof that (1) implies (2), it suffices to observe that if $\mathcal{A}$ is a finite automaton, then $L^+(\mathcal{A})$ is recognizable. For (2) implies (1), we choose for $\mathcal{A}$ a finite deterministic automaton recognizing L. □

The class of deterministic sets is closed under finite union and under finite intersection. We shall see that it is not closed under complementation.

Proposition 6.4. *Any finite union (resp. intersection) of deterministic sets is deterministic.*

Proof. For finite union, the result follows from the formula

$$\overrightarrow{\bigcup_{1 \leqslant i \leqslant n} L_i} = \bigcup_{1 \leqslant i \leqslant n} \overrightarrow{L_i}$$

and one may apply Theorem 6.2.

For intersection, we need a special construction. Consider two deterministic Büchi automata $\mathcal{A}_1 = (Q_1, A, E_1, i_1, F_1)$ and $\mathcal{A}_2 = (Q_2, A, E_2, i_2, F_2)$. Let

$$Q = Q_1 \times Q_2 \times \{0, 1\}$$

and let $\mathcal{A} = (Q, A, E, (i_1, i_2, 1), F)$ be the deterministic automaton defined by

$$E = \Big\{ \big((p_1, p_2, s), a, (q_1, q_2, t)\big) \mid (p_1, a, q_1) \in E_1, (p_2, a, q_2) \in E_2 \text{ and}$$
$$t = 0 \text{ if and only if } (s = 1, p_2 \in F_2 \text{ and } q_1 \notin F_1) \text{ or } (s = 0 \text{ and } q_1 \notin F_1) \Big\}$$

and

$$F = Q_1 \times F_2 \times \{1\}.$$

Intuitively, the component s of a state (q_1, q_2, s) memorizes the fact that the current path has visited a state of F_1 since its last visit to F_2. In fact, if $(p_1, p_2, s) \xrightarrow{a} (q_1, q_2, t)$ is a transition, and if $p_2 \in F_2$, we have

$$t = \begin{cases} 1 & \text{if } q_1 \in F_1 \\ 0 & \text{if } q_1 \notin F_1 \end{cases}$$

But, if $p_2 \notin F_2$, we have

$$t = \begin{cases} 1 & \text{if } s = 1 \text{ (“a state in } F_1 \text{ has already been visited”) or} \\ & \text{if } q_1 \in F_1 \text{ (“first visit of a state in } F_1\text{.”)} \\ 0 & \text{if } s = 0 \text{ and if } q_1 \notin F_1 \text{ (“no state in } F_1 \text{ has been visited”)} \end{cases}$$

We shall show that $L^\omega(\mathcal{A}) = L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$. Let

$$c = (q_{0,1}, q_{0,2}, s_0) \xrightarrow{a_0} (q_{1,1}, q_{1,2}, s_1) \xrightarrow{a_1} (q_{2,1}, q_{2,2}, s_2) \cdots$$

be a path in $\mathcal{A}$ labeled $u = a_0 a_1 a_2 \cdots$. By projection, c defines a path c_1 of $\mathcal{A}_1$ and a path c_2 of $\mathcal{A}_2$:

$$c_1 = q_{0,1} \xrightarrow{a_0} q_{1,1} \xrightarrow{a_1} q_{2,1} \cdots$$
$$c_2 = q_{0,2} \xrightarrow{a_0} q_{1,2} \xrightarrow{a_1} q_{2,2} \cdots$$

Suppose that c is a successful path of $\mathcal{A}$. Then $q_{0,1} = i_1$, $q_{0,2} = i_2$ and $s_0 = 1$. Moreover c passes infinitely often by a state of F. Thus, there exists an infinite sequence $(n_k)_{k \geqslant 0}$ such that $q_{n_k,2} \in F_2$ and $s_{n_k} = 1$. As an immediate consequence, c_2 is a successful path of $\mathcal{A}_2$. The path c_1 is also successful since condition $s_{n_k} = 1$ indicates that, since the last visit to a state of F_2, at least one state of F_1 has been visited. Thus $u \in L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$.

Conversely, if $u \in L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$, u is the label of a successful path c_1 of $\mathcal{A}_1$ and c_2 of $\mathcal{A}_2$, say

$$c_1 = q_{0,1} \xrightarrow{a_0} q_{1,1} \xrightarrow{a_1} q_{2,1} \cdots$$
$$c_2 = q_{0,2} \xrightarrow{a_0} q_{1,2} \xrightarrow{a_1} q_{2,2} \cdots$$

The word u defines thus an initial path c in $\mathcal{A}$. Now, if $q_{n,2} \in F_2$, we have $q_{n',1} \in F_1$ for some $n' \geqslant n$. If n'' is the smallest integer strictly larger than n' such that $q_{n'',2} \in F_2$, we have $(i_1, i_2, 1) \cdot a_0 a_1 \cdots a_{n''} = (q_{n'',1}, q_{n'',2}, 1) \in F$, showing that c is a successful path and that $u \in L^\omega(\mathcal{A})$. □

Example 6.2 shows that the set $X = (a + b)^* a^\omega$ is not of the form $\overrightarrow{L}$. But X is recognized by the Büchi automaton of Example 5.2. Thus, there exist recognizable subsets of A^ω which cannot be recognized by a deterministic Büchi automaton. This contrasts with the case of finite words.

Corollary 6.3 leaves open the problem of whether there exists a recognizable set which could be recognized by a deterministic Büchi automaton but not by a finite one (or, equivalently, if there exist recognizable sets of the form $\overrightarrow{L}$ for which L cannot be chosen recognizable). We shall see later that this cannot happen (Theorem 9.9).

7 Muller and Rabin automata

7.1 Muller automata.

Contrary to what happens for finite words, deterministic Büchi automata are not able to recognize all recognizable subsets of A^ω. This is the motivation for introducing Muller automata which are, by definition, deterministic, but with a more powerful acceptance mode. The result that all recognizable sets can be recognized by Muller automata (McNaughton's theorem) will be established in the next section.

A *Muller automaton* is a 5-tuple $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ where (Q, A, E) is a *finite deterministic* automaton, i is the initial state and $\mathcal{T}$ is a set of subsets of Q, called the *table* of the automaton. We often denote a Muller automaton by $\mathcal{A} = (E, i, \mathcal{T})$.

A path p in $\mathcal{A}$ is *initial* if it starts in the initial state and *final* if $\mathrm{Inf}(p) \in \mathcal{T}$, that is, if the set of infinitely visited states is an element of the table. In this way, one may both specify the infinitely visited states and those that are not. A path is *successful* if it is both initial and final. The set of infinite words recognized by $\mathcal{A}$ is the set, denoted by $L^\omega(\mathcal{A})$, of labels of infinite successful paths in $\mathcal{A}$.

Example 7.1. The automaton represented in Figure 7.1,

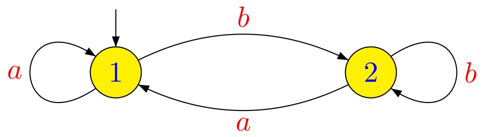

Figure 7.1: A Muller automaton.

with the initial state $i = 1$ and the table $\mathcal{T} = \{\{2\}\}$, recognizes the set $(a+b)^* b^\omega$ of infinite words with a finite number of a's.

We first show that Muller automata are not more powerful than finite Büchi automata. The proof does not make use of the fact that Muller automata are, by definition, supposed to be deterministic (see Exercise 15).

Theorem 7.1. *The set of infinite words recognized by a Muller automaton is ω-rational.*

Proof. Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton. Since

$$L^\omega(\mathcal{A}) = \bigcup_{T \in \mathcal{T}} L^\omega(E, i, \{T\})$$

and since the class of recognizable sets is closed under finite union, we may assume that $\mathcal{T}$ is a singleton. Let thus $\mathcal{T} = \{T\}$ where $T = \{t_0, t_1, \cdots, t_k\}$. Let $D = E \cap (T \times A \times T)$, $X = L^+(E, i, t_0)$, $Y_0 = L^+(D, \{t_0\}, \{t_1\})$, $Y_1 = L^+(D, \{t_1\}, \{t_2\})$, $\ldots$, $Y_k = L^+(D, \{t_k\}, \{t_0\})$. Thus, X is the set of labels of the paths going from the initial state to t_0. Next, Y_i, for $0 \leqslant i < k$ (resp. $i = k$) is the set of labels of paths going from t_i to t_{i+1} (resp. from t_k to t_0) and visiting only states in T. We claim that $L^\omega(\mathcal{A}) = X(Y_0 Y_1 \cdots Y_k)^\omega$, which proves that $L^\omega(\mathcal{A})$ is ω-rational.

First, if $u \in X(Y_0 Y_1 \cdots Y_k)^\omega$, then u is the label of an initial path p such that $\mathrm{Inf}(p) = T$. Consequently $u \in L^\omega(\mathcal{A})$.

Conversely, let $u \in L^\omega(\mathcal{A})$. Then u is the label of an initial path

$$p = (q_0, a_0, q_1)(q_1, a_1, q_2), \cdots$$

such that $\mathrm{Inf}(p) = T$. In particular, there is an integer n_0 such that $q_{n_0} = t_0$ and such that, for all $n \geqslant n_0$, one has $q_n \in T$. One may therefore find an infinite sequence of integers $n_0 < n_1 < n_2 \cdots$ such that the sequence $q_{n_0}, q_{n_1}, \ldots$ is equal to the periodic sequence $t_0, t_1, \ldots, t_k, t_0, t_1, \ldots, t_k, t_0, \ldots$. We have then $a_0 a_1 \cdots a_{n_0 - 1} \in X$ and for all $r \geqslant 0$,

$$a_{n_r} a_{n_r+1} \cdots a_{n_{r+1}-1} \in Y_{\bar{r}},$$

where $\bar{r}$ is the rest of the division of r by $k+1$. This shows that $u \in X(Y_0 \ldots Y_k)^\omega$, proving the claim. □

Example 7.2. Consider again the automaton of Figure 7.1, but this time with $\mathcal{T} = \{\{1,2\}\}$. The set X recognized by this automaton is formed of all infinite words having an infinite number of occurrences of a and of b. The method used in the proof of Theorem 7.1 leads to the rational expression

$$X = (a+b)^*a\big((a+b)^*b(a+b)^*a\big)^\omega.$$

We now consider the various reductions and modifications that one may operate on a Muller automaton. Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton. A state q is called *accessible* if there exists a finite initial path (possibly empty) ending in q. A state q is called *coaccessible* if there exists an infinite final path starting in q. A subset T of Q is called *admissible* if there exists an infinite initial path p such that $\mathrm{Inf}(p) = T$. Finally, $\mathcal{A}$ is *trim* if all its states are accessible and coaccessible and if all the elements of $\mathcal{T}$ are admissible. In practice, it is easy to verify if a subset of Q is admissible. One computes first the graph G with vertex set Q and with edges

$$R = \{(q, q') \mid \text{ there exists } a \in A \text{ such that } (q, a, q') \in E\}.$$

Then a subset T of Q is admissible if the restriction of G to T is strongly connected and accessible from i, or equivalently if there exists a path between two arbitrary elements of T, and a path from i to some vertex in T.

We now establish two equivalence results for Muller automata. The first statement shows that one may always suppose a Muller automaton to be trim.

Proposition 7.2. *Any nonempty subset of A^ω recognized by a Muller automaton can be recognized by a trim Muller automaton.*

Proof. Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton recognizing a nonempty subset X of A^ω. Let P be the set of states of $\mathcal{A}$ that are both accessible and coaccessible. Since X is nonempty, the initial state i is in P. Let $\mathcal{B} = (P, A, E', i, \mathcal{T}')$ where $\mathcal{T}'$ is the set of admissible T in $\mathcal{T}$ and let $E' = E \cap (P \times A \times P)$. It is clear that $L^\omega(\mathcal{B}) \subset L^\omega(\mathcal{A})$. Conversely, let $u = a_0a_1 \cdots \in L^\omega(\mathcal{A})$. There exists then an initial path

$$p = q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \ \cdots$$

labeled u such that $\mathrm{Inf}(p) \in \mathcal{T}$. The states $q_0, q_1, \ldots$ are both accessible and coaccessible and $\mathrm{Inf}(p)$ is admissible. Thus p is a path in $\mathcal{B}$ and $u \in L^\omega(\mathcal{B})$. Finally $L^\omega(\mathcal{A}) = L^\omega(\mathcal{B})$. □

The second statement shows that one may always suppose a Muller automaton to be complete.

Proposition 7.3. *Any subset of A^ω recognized by a Muller automaton can be recognized by a complete Muller automaton.*

Proof. Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton recognizing a subset X of A^ω. If $\mathcal{A}$ is not complete, we add a new state p and we "complete" the automaton by creating a transition (q, a, p) if there is no transition of the form(q, a, q') in $\mathcal{A}$. More formally, let $\mathcal{A}' = (Q \cup \{p\}, A, E', i, \mathcal{T})$, where p is a new state and $E' = E \cup E_1 \cup E_2$ with

$$E_1 = \{(p, a, p) \mid a \in A\} \tag{7.1}$$
$$E_2 = \{(q, a, p) \mid q \in Q, a \in A \text{ and } (\{q\} \times \{a\} \times Q) \cap E = \emptyset\}. \tag{7.2}$$

The Muller automaton $\mathcal{A}'$ is complete and recognizes X. □

The following two results, combined with McNaughton's theorem to be proved in Section 9, constitute basic properties of recognizable subsets of A^ω.

Proposition 7.4. *Let $\mathcal{A} = (E, i, \mathcal{T})$ be a complete Muller automaton. Then the automaton $\mathcal{B} = (E, i, \mathcal{P}(Q) \setminus \mathcal{T})$, obtained by changing $\mathcal{T}$ into its complement, recognizes $A^\omega \setminus L^\omega(\mathcal{A})$.*

Proof. Since $\mathcal{A}$ is complete, every infinite word u is the label of one and only one initial path in $\mathcal{A}$. If $\mathrm{Inf}(p) \in \mathcal{T}$, then $u \in L^\omega(\mathcal{A})$, and otherwise $u \in A^\omega \setminus L^\omega(\mathcal{A})$. □

Proposition 7.5. *Let $\mathcal{A}_1 = (Q_1, A, E_1, i_1, \mathcal{T}_1)$ and $\mathcal{A}_2 = (Q_2, A, E_2, i_2, \mathcal{T}_2)$ be two Muller automata. Let π_1 and π_2 be the projections from $Q_1 \times Q_2$ onto Q_1 (resp. Q_2). The automaton $\mathcal{A} = (Q_1 \times Q_2, A, E, (i_1, i_2), \mathcal{T})$, with*

$$\mathcal{T} = \{R \subset Q_1 \times Q_2 \mid \pi_1(R) \in \mathcal{T}_1 \text{ and } \pi_2(R) \in \mathcal{T}_2\}$$
$$E = \Big\{\big((q_1, q_2), a, (q_1', q_2')\big) \mid (q_1, a, q_1') \in E_1 \text{ and } (q_2, a, q_2') \in E_2\Big\}$$

recognizes $L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$.

Proof. The projections $\pi_1 : Q_1 \times Q_2 \to Q_1$ and $\pi_2 : Q_1 \times Q_2 \to Q_2$ induce a function, also denoted π_1 (resp. π_2) from the set of paths in $\mathcal{A}$ to the set of paths in $\mathcal{A}_1$ (resp. $\mathcal{A}_2$). Let u be the label of an infinite path p in $\mathcal{A}$. Let us show that $\pi_1(\mathrm{Inf}(p)) = \mathrm{Inf}(\pi_1(p))$ (resp. $\pi_2(\mathrm{Inf}(p)) = \mathrm{Inf}(\pi_2(p))$). In fact, if $q_1 \in \pi_1(\mathrm{Inf}(p))$, there exists a state $q_2 \in Q_2$ such that $(q_1, q_2) \in \mathrm{Inf}(p)$ and thus $q_1 \in \mathrm{Inf}(\pi_1(p))$. Conversely, if $q_1 \in \mathrm{Inf}(\pi_1(p))$, there exists an infinite sequence of states $(q_{2,n})_{n \geqslant 0}$ of Q_2 such that for every $n \geqslant 0$,the path p passes by $(q_1, q_{2,n})$. Since Q_2 is finite, there exists $q_2' \in Q_2$ such that $(q_1, q_2') \in \mathrm{Inf}(p)$ and thus $q_1 \in \pi_1(\mathrm{Inf}(p))$. This proves the equality $\pi_1(\mathrm{Inf}(p)) = \mathrm{Inf}(\pi_1(p))$ and the equality corresponding to π_2 can be proved in the same way.

Now, if p is a successful path, it starts in (i_1, i_2) and there exists $R \in \mathcal{T}$ such that $\mathrm{Inf}(p) = R$. We have then, by definition of $\mathcal{T}$, $\pi_1(R) \in \mathcal{T}_1$ and $\pi_2(R) \in \mathcal{T}_2$. Thus $\mathrm{Inf}(\pi_1(p)) = \pi_1(\mathrm{Inf}(p)) = \pi_1(R) \in \mathcal{T}_1$ and $\mathrm{Inf}(\pi_2(p)) \in \pi_2(R)$ whence $\pi_1(p)$ (resp. $\pi_2(p)$) is a successful path in $\mathcal{A}_1$ (resp. $\mathcal{A}_2$). Finally $u \in L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$.

Conversely, let $u \in L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$. The word u is then the label of a successful path p_1 (resp. p_2) of $\mathcal{A}_1$ (resp. $\mathcal{A}_2$). Let p be the unique path in $\mathcal{A}$ such that $\pi_1(p) = p_1$ and $\pi_2(p) = p_2$. This path starts in (i_1, i_2) and is labeled u, whence, by the above discussion $\pi_1(\mathrm{Inf}(p)) = \mathrm{Inf}(\pi_1(p)) \in \mathcal{T}_1$ and $\pi_2(\mathrm{Inf}(p)) = \mathrm{Inf}(\pi_2(p)) \in \mathcal{T}_2$. This implies $\mathrm{Inf}(p) \in \mathcal{T}$, showing that p is a successful path and that $u \in L^\omega(\mathcal{A})$. □

Proposition 7.6. *For any subset X of A^ω, the following conditions are equivalent.*

(1) *X is recognizable by a Muller automaton.*

(2) *X is of the form*

$$X = \bigcup_{1 \leqslant i \leqslant n} (U_i \setminus V_i) \tag{7.3}$$

where the U_i and the V_i are subsets of A^ω recognizable by finite deterministic Büchi automata.

(3) *X is a finite boolean combination of subsets of A^ω recognizable by finite Büchi deterministic automata.*

Proof. (1) $\Longrightarrow$ (2). Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton. Since

$$L^\omega(\mathcal{A}) = \bigcup_{T \in \mathcal{T}} L^\omega(E, i, \{T\})$$

we may suppose that $\mathcal{T}$ contains just one element T. Now, the connection between Muller's and Büchi's acceptance conditions is summarized in the formula

$$L^\omega(E, i, \{T\}) = \bigcap_{t \in T} L^\omega(E, i, t) \setminus \bigcup_{t \notin T} L^\omega(E, i, t)$$

Since, by Proposition 6.4, the class of deterministic sets is closed under finite union and under finite intersection, $L^\omega(\mathcal{A})$ is of the form 7.3.

(2) $\Longrightarrow$ (3) is obvious.

(3) $\Longrightarrow$ (1). If X is recognized by a finite deterministic Büchi automaton $\mathcal{A} = (Q, A, E, i, F)$, it is also recognized by the Muller automaton $(Q, A, E, i, \mathcal{T})$ with

$$\mathcal{T} = \{T \subset Q \mid T \cap F \neq \emptyset\}.$$

Furthermore, Propositions 7.4 and 7.5 show that the class of sets recognized by a Muller automaton is a boolean algebra. Thus any finite boolean combination of sets recognized by a Muller automaton is recognized by a Muller automaton. □

7.2 Rabin automata.

The expression of a recognizable set as a union of differences as in Condition (2) of Proposition 7.6 motivates the introduction of an additional class of automata. A *Rabin* *automaton* is a tuple $\mathcal{A} = (Q, A, E, i, \mathcal{R})$ where (Q, A, E) is deterministic automaton, i is the initial state and $\mathcal{R} = \{(L_j, U_j) \mid j \in J\}$ is a family of pairs of sets of states. An infinite path p is successful if there exists an index $j \in J$ such that p visits U_j infinitely often and visits L_j only finitely often. When the automaton is finite, this is equivalent to

$$\mathrm{Inf}(p) \cap L_j = \emptyset \quad \text{and} \quad \mathrm{Inf}(p) \cap U_j \neq \emptyset.$$

An infinite word is recognized by $\mathcal{A}$ if it is the label of a successful path in $\mathcal{A}$.

Any deterministic Büchi automaton (Q, i, F) can be considered as a particular Rabin automaton with $\mathcal{R} = \{(\emptyset, F)\}$. Muller automata are equivalent to finite Rabin automata, as shown in the following proposition, in which the term "recognizable deterministic set" means a subset of A^ω recognized by a finite deterministic Büchi automaton.

Proposition 7.7. *Any finite Rabin automaton is equivalent to a Muller automaton. Conversely, any Muller automaton is equivalent to a finite Rabin automaton. Moreover, a subset of A^ω is recognized by a Rabin automaton with n pairs if and only if it is a union of n differences of recognizable deterministic subsets of A^ω.*

Proof. Let $\mathcal{A} = (Q, A, E, i, \mathcal{R})$ be a finite Rabin automaton and let

$$\mathcal{T} = \{T \subset Q \mid T \cap L = \emptyset \text{ and } T \cap U \neq \emptyset \text{ for some } (L, U) \in \mathcal{R}\}. \tag{7.4}$$

Then the Muller automaton $(Q, A, E, i, \mathcal{T})$ recognizes the same set of infinite words as $\mathcal{A}$. Thus every finite Rabin automaton is equivalent to a Muller automaton.

Let $\mathcal{E}_n$ be the class of subsets of A^ω of the form $\bigcup_{1 \leqslant i \leqslant n}(U_i \setminus V_i)$ where the U_i and the V_i are recognizable deterministic subsets of A^ω. We have to show that $\mathcal{E}_n$ is also the class of sets recognized by a Rabin automaton with n pairs. We shall obtain as a consequence that any Muller automaton is equivalent to a finite Rabin automaton, since Proposition 7.6 shows that any subset recognized by a Muller automaton belongs to some $\mathcal{E}_n$.

Let $\mathcal{A} = (Q, A, E, i, \mathcal{R})$ be a Rabin automaton with n pairs, say $\mathcal{R} = \{(L_j, U_j) \mid 1 \leqslant j \leqslant n\}$. On the one hand

$$L^\omega(\mathcal{A}) = \bigcup_{1 \leqslant j \leqslant n} L^\omega(\mathcal{A}_j) \tag{7.5}$$

where $\mathcal{A}_j = (Q, A, E, i, \mathcal{R}_j)$ is the Rabin automaton defined by $\mathcal{R}_j = \{(L_j, U_j)\}$ and on the other hand

$$L^\omega(\mathcal{A}_j) = L^\omega(Q, A, E, i, U_j) \setminus L^\omega(Q, A, E, i, L_j)$$

showing that $L^\omega(\mathcal{A}) \in \mathcal{E}_n$.

Let now X_1 and X_2 be two recognizable deterministic subsets recognized respectively by the finite deterministic Büchi automata $\mathcal{A}_1 = (Q_1, A, E_1, i_1, F_1)$ and $\mathcal{A}_2 = (Q_2, A, E_2, i_2, F_2)$. By Proposition 5.3, we may suppose that $\mathcal{A}_1$ and $\mathcal{A}_2$ are complete. Then $X_1 \setminus X_2$ is recognized by the one pair Rabin automaton $\mathcal{A} = (Q_1 \times$

$Q_2, A, E, (i_1, i_2), \mathcal{R})$, with

$$E = \Big\{ \big((q_1, q_2), a, (q_1', q_2')\big) \mid (q_1, a, q_1') \in E_1 \text{ and } (q_2, a, q_2') \in E_2 \Big\}.$$

$$\mathcal{R} = \Big\{ \big((Q_1 \times F_2), (F_1 \times Q_2)\big) \Big\}.$$

This shows that $\mathcal{E}_1$ is equal to the class of sets recognized by a one pair Rabin automaton.

Let us now consider a set X in $\mathcal{E}_n$. By definition, X is the union of n sets X_j from $\mathcal{E}_1$ and by the above, every X_j is recognized by a one pair Rabin automaton $\mathcal{A}_j = (Q_j, A, E_j, i_j, \mathcal{R}_j)$, with $\mathcal{R}_j = \{(L_j, U_j)\}$. Then the n pair Rabin automaton

$$\mathcal{A} = (Q_1 \times Q_2 \times \cdots \times Q_n, A, E, (i_1, i_2, \cdots, i_n), \mathcal{R})$$

defined by

$$\begin{aligned} E = &\big\{ \big((q_1, q_2, \ldots, q_n), a, (q_1', q_2', \ldots, q_n')\big) \\ &\qquad \mid (q_1, a, q_1') \in E_1,\ (q_2, a, q_2') \in E_2, \ldots, (q_n, a, q_n') \in E_n \big\} \\ \mathcal{R} = &\big\{ (Q_1 \times \cdots \times Q_{j-1} \times L_j \times Q_{j+1} \times \cdots \times Q_n, \\ &\qquad Q_1 \times \cdots \times Q_{j-1} \times U_j \times Q_{j+1} \times \cdots \times Q_n) \mid 1 \leqslant j \leqslant n, (L_j, U_j) \in \mathcal{R}_j \big\} \end{aligned}$$

recognizes X. Let in fact

$$p = (q_{0,1}, \ldots, q_{0,n}) \xrightarrow{a_0} (q_{1,1}, \ldots q_{1,n}) \xrightarrow{a_1} (q_{2,1}, \ldots q_{2,n}) \cdots$$

be a path in $\mathcal{A}$ labeled $u = a_0 a_1 a_2 \cdots$. By projection, p defines, for $1 \leqslant j \leqslant n$, a path p_j in $\mathcal{A}_j$

$$p_j = q_{0,j} \xrightarrow{a_0} q_{1,j} \xrightarrow{a_1} q_{2,j} \cdots$$

Let $(L, U) = (Q_1 \times \cdots \times Q_{j-1} \times L_j \times Q_{j+1} \times \cdots \times Q_n, Q_1 \times \cdots \times Q_{j-1} \times U_j \times Q_{j+1} \times \cdots \times Q_n)$ be a pair of $\mathcal{R}$. We have then $(q_{i,1}, \ldots, q_{i,n}) \notin L$ if and only if $q_{i,j} \notin L_j$ and $(q_{i,1}, \ldots, q_{i,n}) \in U$ if and only if $q_{i,j} \in U_j$. Thus p is a successful path in $\mathcal{A}$ if and only if one among the p_j is a successful path in $\mathcal{A}_j$. □

A direct construction of a Rabin automaton equivalent to a given Muller automaton is also possible (see Exercise 17).

The notion of Streett automaton is the dual to that of Rabin automaton. A *Streett automaton* is a tuple

$$\mathcal{A} = (Q, A, E, i, \mathcal{S})$$

where (Q, A, E) is a deterministic automaton, $i \in Q$ is the initial state and $\mathcal{S} = \{(L_j, U_j) \mid j \in J\}$ is a family of pairs of sets of states. An infinite path p is successful if it starts in the initial state and if, for each $j \in J$, either it visits L_j a finite number of times or it visits U_j an infinite number of times.

Negating the acceptance condition shows that the complement of a set recognized by a Streett automaton can be recognized by a Rabin automaton with the same number of pairs.

7.3 Muller automata with a full table

We end this section by an important result characterizing those Muller automata which are equivalent to a finite deterministic Büchi automaton. Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton. We say that the table $\mathcal{T}$ is *full* if for every admissible set $T \in \mathcal{T}$ and for every admissible set T' containing T, we have $T' \in \mathcal{T}$.

Theorem 7.8. *A Muller automaton is equivalent to a finite deterministic Büchi automaton if and only if its table is full.*

The 'if' part is a consequence of the following statement.

Proposition 7.9. *Any Muller automaton with a full table is equivalent to a finite deterministic Büchi automaton.*

Proof. Let $\mathcal{A}' = (Q, A, E, i, \mathcal{T})$ be a Muller automaton with a full table. We define a finite deterministic Büchi automaton

$$\mathcal{A} = (\mathcal{P}(Q) \times Q, (\emptyset, i), F)$$

where $F = \{(\emptyset, q) \mid q \in Q\}$ and where, for every $(S, q) \in \mathcal{P}(Q) \times Q$ and for every $a \in A$, the transition function is given by the following formula.

$$(S, q) \cdot a = \begin{cases} (\emptyset, q \cdot a) & \text{if } S \cup \{q \cdot a\} \text{ contains an element of } \mathcal{T} \\ (S \cup \{q \cdot a\}, q \cdot a) & \text{otherwise.} \end{cases}$$

Additionally, $(S, q) \cdot a$ is undefined if $q \cdot a$ is undefined.

Let u be a word accepted by $\mathcal{A}$. Then u defines in $\mathcal{A}$ an initial path p. Let $T \in \mathcal{T}$ be such that all its states are repeated infinitely often in p. Thus $T \subset \mathrm{Inf}(p)$ and since the table of $\mathcal{A}'$ is full, we have also $\mathrm{Inf}(p) \in \mathcal{T}$. Thus u is accepted by $\mathcal{A}'$.

Conversely, if u is accepted by $\mathcal{A}'$, u defines in $\mathcal{A}'$ a (unique) initial path p' such that $\mathrm{Inf}(p') \in \mathcal{T}$. Let p be the initial path defined by u in $\mathcal{A}$. This path is final. Indeed, on the contrary, there would exist a prefix v of u, such that, if vw is a prefix of u, the state $(\emptyset, i) \cdot vw$ is not a final state of $\mathcal{A}$. Let $(\emptyset, i) \cdot v = (S, q)$ and let us choose w in such a way that the path of $\mathcal{A}'$ starting at q and defined by w visits all the states of $\mathrm{Inf}(p)$. We have then $(\emptyset, i) \cdot vw = (T, q \cdot w)$ where T is a set containing $\mathrm{Inf}(p)$. But $\mathrm{Inf}(p) \in \mathcal{T}$, whence $T \in \mathcal{T}$, which is impossible. □

The construction described in the proof above can be paraphrased intuitively as follows. To simulate a Muller automaton with a full table by a finite deterministic Büchi automaton, one builds a new automaton which sends signals. It accumulates in a box the states met since the last signal. As soon as the box contains an element of the table, it sends a signal and empties the box. We shall use a similar idea for the proof of McNaughton's theorem in Section 9.

Example 7.3. The automaton of Example 7.2 has a full table. Applying the above construction, one obtains the automaton of Figure 7.2.

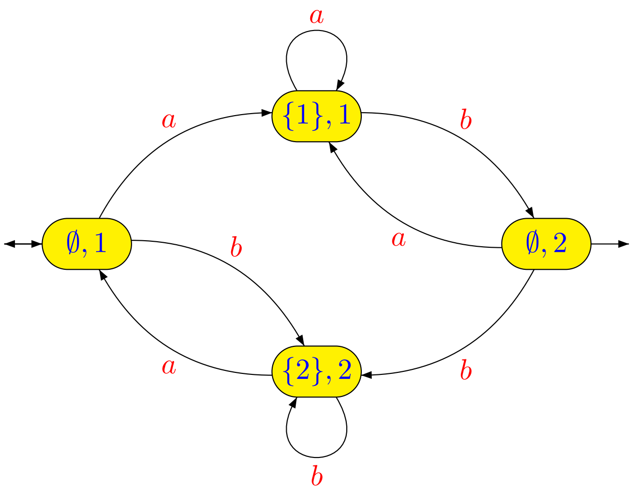

Figure 7.2: A Büchi automaton.

The necessity of the condition in Theorem 7.8 follows from the next result. Note that the finiteness of the automaton is not used in the proof. We shall return to this point in a while (Theorem 9.9).

Proposition 7.10. *A Muller automaton equivalent to a deterministic Büchi automaton has a full table.*

Proof. Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton. Suppose that the set $X = L^\omega(\mathcal{A})$ can be recognized by a deterministic Büchi automaton $\mathcal{B}$. Then there exist a subset L of A^+ such that $X = \overrightarrow{L}$. Let T and T' be two admissible subsets of Q such that $T \in \mathcal{T}$ and $T \subset T'$. We have to show that $T' \in \mathcal{T}$. Since T is admissible, there exists an infinite path p starting in i such that $\mathrm{Inf}(p) = T$. We fix a state $t \in T$. There exists a finite word u such that $i \cdot u = t$ and an infinite word v such that the infinite path starting at t defined by v does not leave T but visits each state of T infinitely often. We build then by induction two sequences of words $u_0, u_1, \cdots$ and $v_0, v_1, \cdots$ satisfying for every $k \geqslant 0$, the following properties:

(a) $u_0 v_0 u_1 v_1 \cdots u_k v_k \in L$,

(b) if $q_k = i.u_0 v_0 u_1 v_1 \cdots u_{k-1} v_{k-1}$, then $q_k u_k = t$ and the path starting at q_k defined by u_k passes at least once by each state of T' and stays in T'.

For this, let $u_0 = u$ and knowing the sequences up to rank $k-1$, we choose u_k in such a way that (b) is satisfied - which is possible, since T' is admissible. The word

$u_0v_0u_1v_1 \cdots u_{k-1}v_{k-1}u_kv$ is accepted by $\mathcal{A}$ and there exists a prefix v_k of v such that condition (a) is satisfied.

Finally, let $w = u_0v_0u_1v_1 \cdots$. We have $w \in \overrightarrow{L}$ by construction, but simultaneously, the initial path s defined by w satisfies $\mathrm{Inf}(s) = T'$. We thus obtain that $T' \in \mathcal{T}$. $\square$

The notion of a Muller automaton with a full table appears as the first level of a hierarchy of Muller automata using the notion of a chain. This leads to the notion of a *Rabin chain automaton* (see Chapter V).

8 Transition automata

It is sometimes convenient to use a variant of automata in which a set of final transitions is specified, instead of the usual set of final states.

Formally, a *transition Büchi automaton* is a 5-tuple $\mathcal{A} = (Q, A, E, I, F)$ where (Q, A, E) is a finite automaton, $I \subseteq Q$ is the set of initial states and $F \subseteq E$ is the set of final transitions. A path is *initial* if it starts in some initial state, *final* if it goes through F infinitely often and *successful* if it is initial and final. The notions of complete (resp. co-complete) automata are defined in the usual way.

Example 8.1. The transition automaton represented in Figure 8.1 recognizes the set of infinite words containing infinitely many a's. The final transitions are circled.

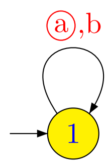

Figure 8.1: A transition automaton recognizing $(\{a, b\}^*a)^\omega$.

Similarly, a *transition Muller automaton* is a 5-tuple $\mathcal{A} = (Q, A, E, I, \mathcal{T})$ where (Q, A, E) is a finite deterministic automaton, i is the initial state and $\mathcal{T}$ is a set of subsets of E, called the *table* of the automaton. If p is an infinite path, the set of transitions which occur infinitely often in p is denoted by $\mathrm{Inf}_T(p)$. A path p is *successful* if it is an initial path and if $\mathrm{Inf}_T(p) \in \mathcal{T}$, that is, if the set of transitions occurring infinitely often in p is an element of the table.

We now establish the equivalence of these notions with the standard ones. Recall that a Büchi automaton is co-deterministic (resp. co-complete) if any infinite word labels at most (resp. at least) one final path.

Proposition 8.1. *Every Büchi automaton is equivalent to a transition Büchi automaton. Conversely, every transition Büchi automaton is equivalent to a Büchi automaton.*

Furthermore, if one of these automata is co-deterministic (resp. co-complete) one can choose the other one to have the same property.

Proof. Let $\mathcal{A} = (Q, A, E, I, F)$ be a Büchi automaton. Then $\mathcal{A}$ is equivalent to the transition Büchi automaton $\mathcal{A}' = (Q, A, E, I, F')$ where $F' = \{(p, a, q) \in E \mid p \in F\}$. Clearly, if $\mathcal{A}$ is deterministic (resp. co-deterministic, complete, co-complete), then so is $\mathcal{A}'$.

Conversely, let $\mathcal{A} = (Q, A, E, I, F)$ be a transition Büchi automaton. Then $\mathcal{A}$ is equivalent to the Büchi automaton $\mathcal{A}' = (Q', A, E', I', F')$ defined by $Q' = Q \times \{0, 1\}$, $I' = I \times \{0, 1\}$ and

$$\begin{aligned}
F' &= \{(p, 1) \in Q' \mid (p, a, q) \in F \text{ for some } q \in Q \text{ and } a \in A\} \\
E' &= \Big\{\big((p, 1), a, (q, \varepsilon)\big) \mid (p, a, q) \in F,\ \varepsilon \in \{0, 1\}\Big\} \\
&\quad \cup \Big\{\big((p, 0), a, (q, \varepsilon)\big) \mid (p, a, q) \in E \setminus F,\ \varepsilon \in \{0, 1\}\Big\}
\end{aligned}$$

If $\mathcal{A}$ is co-deterministic, (resp. co-complete), then so is $\mathcal{A}'$. □

Proposition 8.2. *Every Muller automaton is equivalent to a transition Muller automaton. Conversely, every transition Muller automaton is equivalent to a Muller automaton.*

Proof. Let $\mathcal{A} = (Q, A, E, i, \mathcal{T})$ be a Muller automaton. We claim that $\mathcal{A}$ is equivalent to the transition Muller automaton $\mathcal{A}' = (Q, A, E, j, \mathcal{T}')$, where j is a new state $Q' = E \times \{j\}$, $\mathcal{T}' = \mathcal{T}$ and

$$\begin{aligned}
E' &= \Big\{\big(j, a, (i, a, q)\big) \mid (i, a, q) \in E\Big\} \\
&\quad \cup \Big\{\big(q, a, q'), b, (q', b, q'')\big) \mid (q, a, q') \in E \text{ and } (q', a, q'') \in E\Big\}
\end{aligned}$$

Indeed, to any infinite path of $\mathcal{A}$ starting at i

$$p = (i, a_0, q_1)(q_1, a_1, q_2) \cdots$$

corresponds an infinite path of $\mathcal{A}'$ with the same label starting at j

$$p' = \big(j, a_0, (i, a_0, q_1)\big)\big((i, a_0, q_1), a_1, (q_1, a_1, q_2)\big) \cdots$$

and conversely, every infinite path of $\mathcal{A}'$ starting at j arises this way. Furthermore, the transition (q, a, q') occurs in p if and only if p' visits the state (q, a, q'). Therefore $\mathrm{Inf}_T(p) = \mathrm{Inf}(p')$ and $\mathcal{A}$ and $\mathcal{A}'$ are equivalent. □

9 McNaughton's theorem

The aim of this section is to prove the following result, due to R. McNaughton.

Theorem 9.1. *Any recognizable subset of A^ω can be recognized by a Rabin automaton.*

The proof that we are going to present relies on a determinization algorithm due to S. Safra. It computes a Rabin automaton equivalent to a given Büchi automaton. Using Proposition 7.7 allows one to obtain the conclusion. The states of the Rabin automaton are labeled trees.

We first give an informal description of the construction. Consider a finite Büchi automaton

$$\mathcal{A} = (Q, A, E, I, F)$$

and a successful path p labeled $u \in A^\omega$. There exist states $i \in I$ and $f \in F$ and a factorization $p = p_0 p_1 p_2 \cdots$ where p_0 is a path from i to f labeled u_0 and for every $n > 0$, p_n is a path from f to f labeled u_n.

$$i \xrightarrow{u_0} f \xrightarrow{u_1} f \xrightarrow{u_2} f \ \cdots$$

The usual determinization algorithm does not work with infinite words. One obtains indeed a deterministic automaton $\mathcal{B} = (\mathcal{P}(Q), A, \cdot, \{I\}, \mathcal{F})$, where $\mathcal{F} = \{P \mid P \cap F \neq \emptyset\}$ with transitions given by

$$S.a = \delta(S, a)$$

for $S \subset Q$ and $a \in A$. Processing u in $\mathcal{B}$ gives a path p'

$$I \xrightarrow{u_0} S_0 \xrightarrow{u_1} S_1 \xrightarrow{u_2} S_2 \ \cdots$$

Each S_i contains f, but this is not enough to make sure that such a path is successful, since nothing says that the state f appearing in S_i comes from the state f appearing in S_{i-1}. Thus, one cannot define as a table $\mathcal{T} = \{P \subset Q \mid P \cap F \neq \emptyset\}$. For example, if $\mathcal{A}$ is the automaton of Example 5.2, recognizing the set of infinite words containing a finite number of b's,

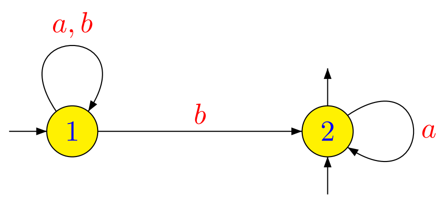

the path $\{1\} \xrightarrow{b} \{1,2\} \xrightarrow{b} \{1,2\} \xrightarrow{b} \{1,2\} \ \cdots$ would be successful in $\mathcal{B}$, although b^ω is not recognized by $\mathcal{A}$. Actually, the automaton obtained by this algorithm, once made trim, contains only one state

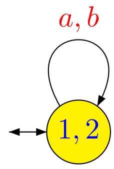

Figure 9.1: The automaton obtained by determinization.

and thus recognizes A^ω whatever be the acceptance mode.

The idea is to look for a path $I \xrightarrow{u_0} S_0 \xrightarrow{u_1} S_1 \xrightarrow{u_2} S_2 \ \cdots$ such that the following two conditions are satisfied:

(1) $S_0 \subset \delta(I, u_0)$, and, for every $n \geqslant 0$, $S_{n+1} \subset \delta(S_n, u_{n+1})$

(2) for every $n \geqslant 0$ and every $q \in S_{n+1}$, there is a state $p \in S_n$ and a path $p \xrightarrow{u_{n+1}} q$ in $\mathcal{A}$ *passing through a final state*.

To find such a path, we are going to build an automaton memorizing the occurrences of final states. The states of this automaton are oriented trees whose nodes are labeled by the sets S_i mentioned above. We then apply the usual determinization algorithm, taking care of adding the new final states that appear as the label of a new child of the vertex.

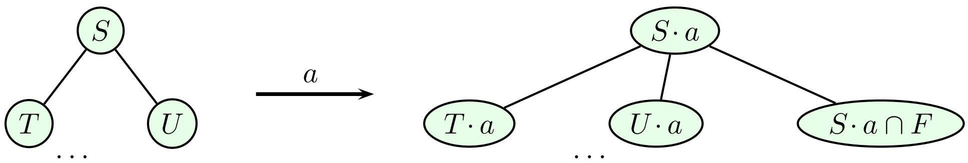

Figure 9.2: The action of letter a.

When all the states in the label S of a vertex have already visited a final state, that is when they all appear in the children of the node, this node is marked and all its descendants disappear.

We now proceed to the formal description of the construction. Let $\mathcal{A} = (Q, A, E, I, F)$ be a finite Büchi automaton with

$$Q = \{1, 2, \cdots, n\} \quad \text{and} \quad V = \{1, 2, \cdots, 2n\}.$$

We build a deterministic Rabin automaton $\mathcal{D}$ as follows. Its *states* are labeled oriented trees with marks on some nodes. Formally the states are tuples (T, f, e, M) where

(1) the set of nodes T is a subset of V,

(2) $f : T \to T^*$ is a function mapping each node on the ordered sequence of its children.

(3) e is a function from T into the set of nonempty subsets of Q, mapping each node to its *label*.

(4) $M \subset T$ is the set of *marked nodes*.

These trees should also satisfy the following conditions:

(5) The root of the tree is 1.

(6) The marked nodes have to be leaves in the tree.

(7) For every node v, the union of the labels of its children is a strict subset of $e(v)$.

(8) If v is not an ancestor of w and if w is not an ancestor of v, then $e(v) \cap e(w) = \emptyset$.

The set $\mathcal{T}_n$ of all trees defined in this way is finite. More precisely, the following result holds:

Proposition 9.2. *A tree in $\mathcal{T}_n$ has at most n nodes.*

Proof. We associate with each node $v \in T$, the set

$$r(v) = e(v) \setminus \bigcup_{w \text{ child of } v} e(w)$$

By condition (7), $r(v)$ is not empty and, if v_1 and v_2 are distinct, we have $r(v_1) \cap r(v_2) = \emptyset$. This follows from condition (7) if v_1 is an ancestor of v_2 and from condition (8) in the other cases, since $r(v) \subset e(v)$. The sets $r(v)$ are thus pairwise distinct and we obtain

$$\text{Card}(T) = \sum_{v \in T} 1 \leqslant \sum_{v \in T} \text{Card}(r(v)) \leqslant \text{Card}(Q) = n$$

establishing the proposition. □

In an oriented tree, the children of a given node are ordered. These local orders can be extended to a partial order on the set of nodes as follows. Given two nodes m and n which are not ancestor of one another, let p be their least common ancestor and let m' (resp. n') be the child of p which is an ancestor of m (resp. n). We say that m is *on the left* of n if $m' < n'$, as illustrated in Figure 9.3.

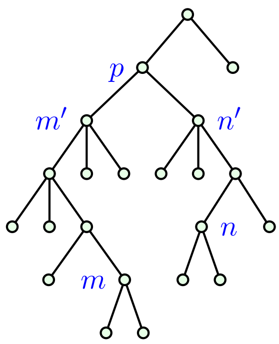

Figure 9.3: The node m is on the left of n.

We return to the construction of the automaton $\mathcal{D}$. The set of its states is thus $\mathcal{T}_n$ and its transition function Δ is defined as follows. Let $R = (T, f, e, M)$ be a tree in $\mathcal{T}_n$ and let a be a letter from A. The state $\Delta(R, a)$ is obtained by the following steps.

(1) We perform the transition by a on the labels of each node and we erase the marks. For this, we build the tree (T, f, e_1, M_1), with $M_1 = \emptyset$, and, for each $v \in T$,

$$e_1(v) = \delta(e(v), a)$$

(2) We add to each node v a new child placed at the right of all children of v and labeled $e(v) \cap F$. This new node is marked and taken arbitrarily among the available nodes (in practice, we take the smallest available node). Formally, we choose an injection from T into $V \setminus T$ associating with each node $v \in T$ a node denoted $\overline{v}$. This is possible since T has at most n elements. Let $\overline{T} = \{\overline{v} \mid v \in T\}$ and consider the tree (T_2, f_2, e_2, M_2) where

$$T_2 = T \cup \overline{T}, \quad M_2 = \overline{T}$$

and, for every $v \in T$,

$$f_2(v) = f(v)\overline{v}, \qquad f_2(\overline{v}) = \varepsilon$$
$$e_2(v) = e_1(v), \qquad e_2(\overline{v}) = e_1(v) \cap F.$$

(3) In the label of each node v, we suppress the states appearing in the label of a node at the left of v. For this, we build e_3 defined for each node $v \in T_2$ by,

$$e_3(v) = e_2(v) \setminus \bigcup_{w \text{ to the left of } v} e_2(w)$$

(4) We suppress the nodes with an empty label and we update the function f and the marks accordingly. This operation is represented in Figure 9.4.

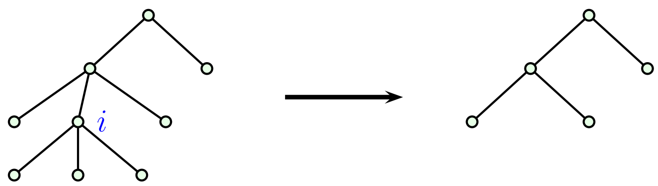

Figure 9.4: Suppressing a node with an empty label.

Formally, we change to the tree (T_4, f_4, e_4, M_4) where

$$T_4 = \{v \in T_2 \mid e_3(v) \neq \emptyset\},$$

$M_4 = M_2 \cap T_4$, $e_4(v)$ is the restriction of $e_3(v)$ to T_4 and, for each node $v \in T_4$, the word $f_4(v)$ is obtained by erasing the symbols of $T_2 \setminus T_4$.

(5) We mark all nodes with a label equal to the union of the labels of their children, i.e. such that

$$e(v) = \bigcup_{w \text{child of } v} e(w).$$

and we suppress all their descendants.

We finally obtain a state $(T_5, f_5, e_5, M_5) = \Delta(R, a)$ which is an element of $\mathcal{T}_n$.

The initial state of $\mathcal{D}$ is the tree reduced to an unmarked node labeled I if $I \cap F = \emptyset$, to a marked node labeled I if $I \subset F$ and to a node labeled I with a marked child labeled $I \cap F$ in all other cases.

There remains to specify the set $\mathcal{R}$ defining the acceptance condition. Let

$$\mathcal{R} = \{(L_v, U_v) \mid v \in V\}$$

where

$$L_v = \{R \in \mathcal{T}_n \mid v \text{ is not a node of } R\}$$
$$U_v = \{R \in \mathcal{T}_n \mid v \text{ is a marked node of } R\}.$$

Thus, a path in $\mathcal{D}$ is successful if there exists an element $v \in V$ such that, ultimately, the path uses states in which v is a node and infinitely often states in which v is marked.

Before proving that this Rabin automaton recognizes the same set of infinite words as the automaton we started from, we are going to illustrate the construction by some examples. In these examples, the states are represented by labeled oriented trees and marked nodes are indicated by a double circle. An arrow of the form $\xrightarrow{(i)}$ indicates that step i of the algorithm has been performed.

Example 9.1. Consider the automaton represented in Figure 9.5, which recognizes the set of words having a finite nonzero number of b's.

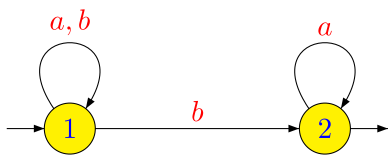

Figure 9.5: A Büchi automaton.

We detail the steps of Safra's algorithm. The initial state is the tree with a single node of Figure 9.6.

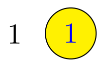

Figure 9.6: The initial state.

The action of the letters a and b on the initial state are represented in Figure 9.7 and Figure 9.8, respectively.

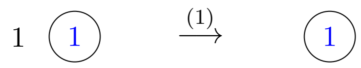

Figure 9.7: The action of a on the initial state.

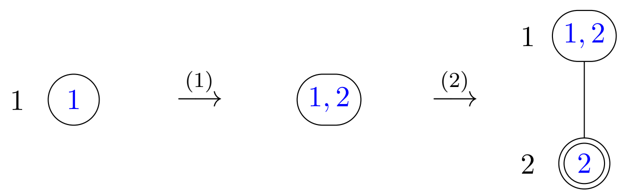

Figure 9.8: The action of b on the initial state.

A new state has now been created. The actions of the letters a and b on this new state are represented in Figure 9.9 and Figure 9.10, respectively. Thus another new state has been created. The action of the letters a and b on this new state are easily derived from the ones represented in Figures 9.9 and 9.10 by exchanging the names 2 and 3 in every place. After renaming the states, we obtain the automaton of Figure 9.11.

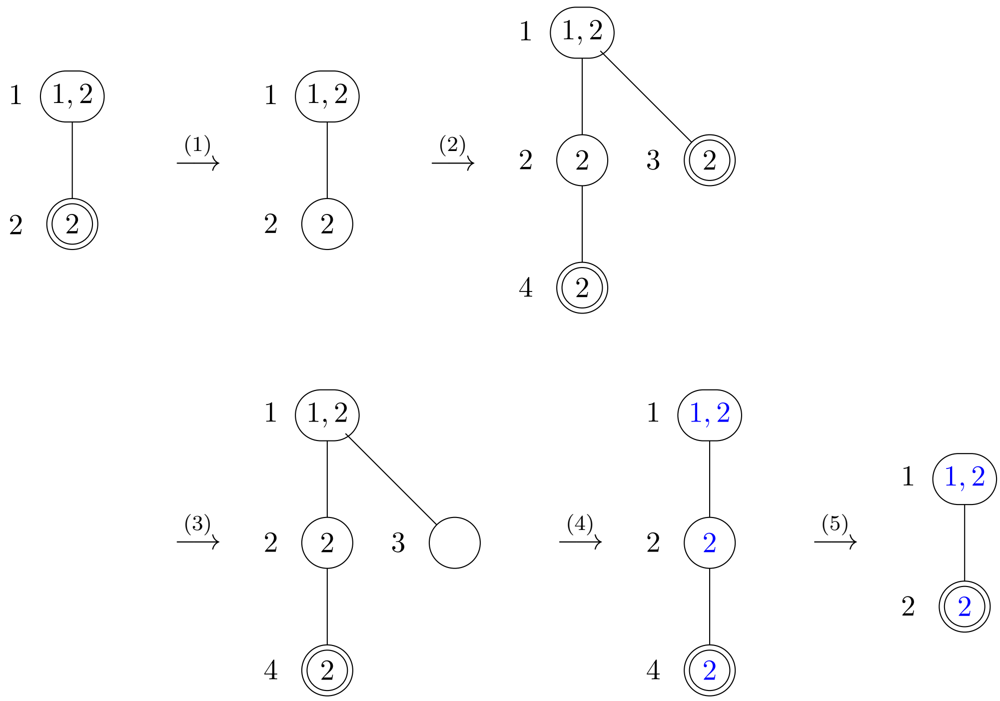

Figure 9.9: The action of a on the new state.

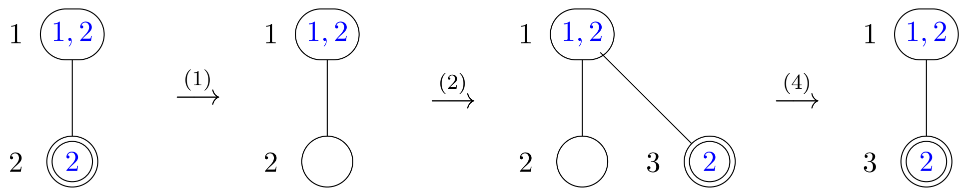

Figure 9.10: The action of b on the new state.

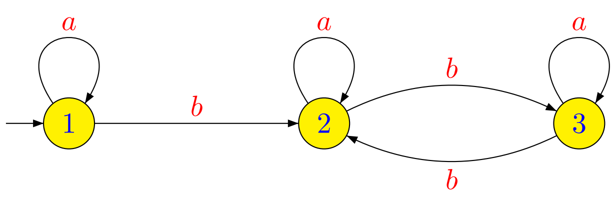

Figure 9.11: The Rabin automaton obtained by Safra's algorithm.

We have $L_1 = \emptyset$, $L_2 = \{1,3\}$, $L_3 = \{1,2\}$, $U_1 = \emptyset$, $U_2 = \{2\}$, $U_3 = \{3\}$. Thus the accepting pairs are $(\{1,3\},\{2\})$ and $(\{1,2\},\{3\})$. Note that, by Formula (7.4), these Rabin pairs are equivalent to the table $\mathcal{T} = \big\{\{2\},\{3\}\big\}$.

Example 9.2. The Büchi automaton represented in Figure 9.12 recognizes the set of words with a finite number of b.

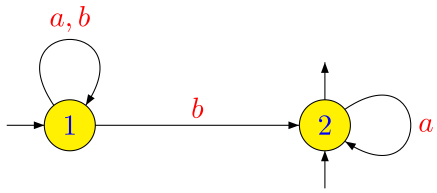

Figure 9.12: A Büchi automaton for the set A^*a^ω.

The deterministic automaton obtained by Safra's algorithm is represented in Figure 9.13.

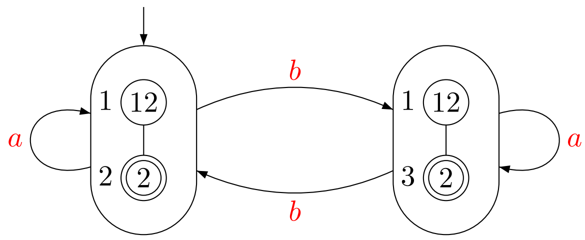

Figure 9.13: The deterministic automaton obtained by Safra's algorithm.

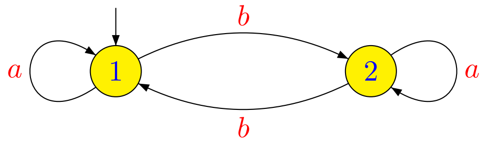

Figure 9.14: The same automaton after renaming the states.

We recognize the automaton computing the parity of the number of b. The acceptance conditions are, in Rabin's form, the pairs $\Big\{\big(\{1\},\{2\}\big),\big(\{2\},\{1\}\big)\Big\}$, which gives the table $\big\{\{1\},\{2\}\big\}$.

Example 9.3. Consider the set $X = (\{b,c\}^*a \cup b)^\omega$. A Büchi automaton recognizing X is given in Figure 9.15:

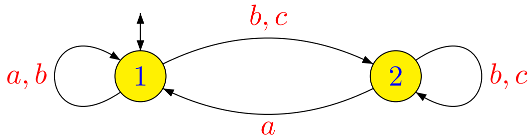

Figure 9.15: A Büchi automaton for $(\{b,c\}^*a \cup b)^\omega$.

The application of Safra's algorithm gives the deterministic automaton of Figure 9.16, in which the set of states is $\{I, II, III, IV, V\}$:

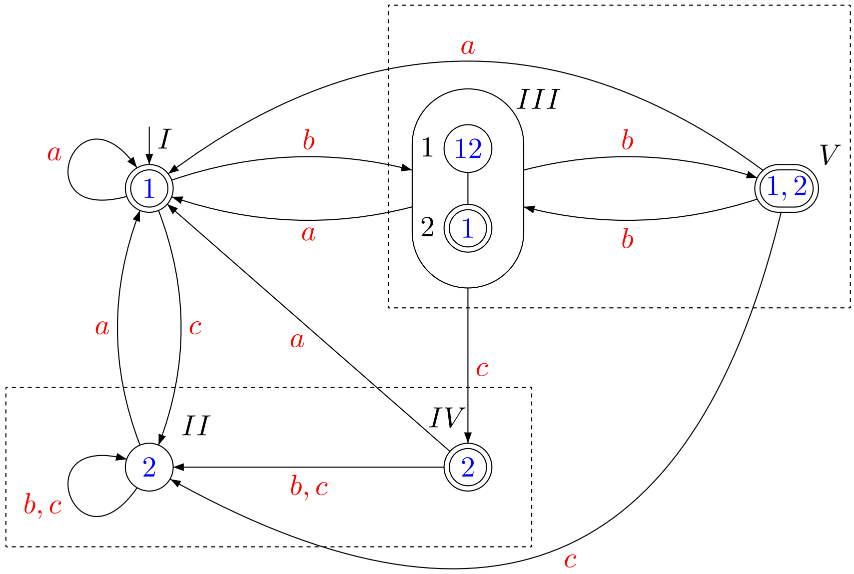

Figure 9.16: A deterministic Muller automaton for $(\{b,c\}^*a \cup b)^\omega$.

The Rabin pairs are $(\emptyset, \{I, IV, V\})$ and $(\{I, II, IV, V\}, \{III\})$. Therefore the table of the corresponding Muller automaton is

$$\mathcal{T} = \{T \subset Q \mid T \text{ contains either } I, IV \text{ or } V\} \cup \{\{III\}\}$$

This table is full. Indeed, if a set T contains I, IV or V, any superset of T has the same property. If $T = \{III\}$, any superset of T contains I, IV or V, or is equal to the set $\{II, III\}$. But this latter set is not admissible.

One can in fact obtain a smaller automaton by merging the states grouped inside each dashed rectangle. After renaming the states, the resulting automaton is represented in Figure 9.17:

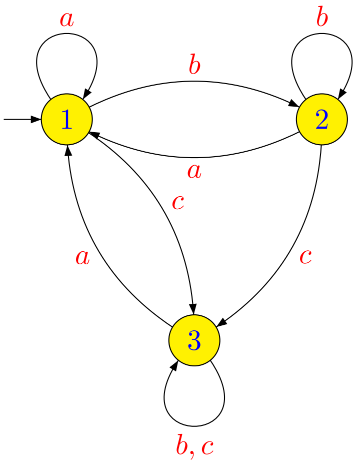

Figure 9.17: Applying Safra's construction.

The table is $\mathcal{T} = \{\{1\}, \{2\}, \{1,2\}, \{1,3\}, \{1,2,3\}\}$, which is the complement of the table $\{\{3\}\}$. Since the table is full, X is deterministic. We have actually $X = \overrightarrow{\{a,b,c\}^*ab^* \cup b^+}$.

We are now going to prove that the deterministic automaton $\mathcal{D}$ is equivalent to the automaton $\mathcal{A}$ we started from. We shall need a lemma which makes more precise the behavior of $\mathcal{D}$. Let $u = a_1 \cdots a_n$ be a finite word and let R_0 be a state of $\mathcal{D}$ containing a marked node v labeled S_0. We suppose that, for $1 \leqslant i \leqslant n$, the states $R_i = \Delta(R_0, a_1 \cdots a_i)$ also contain the node v, with a label S_i, but that this node is marked only for $i = n$. The hypotheses are represented in Figure 9.18.

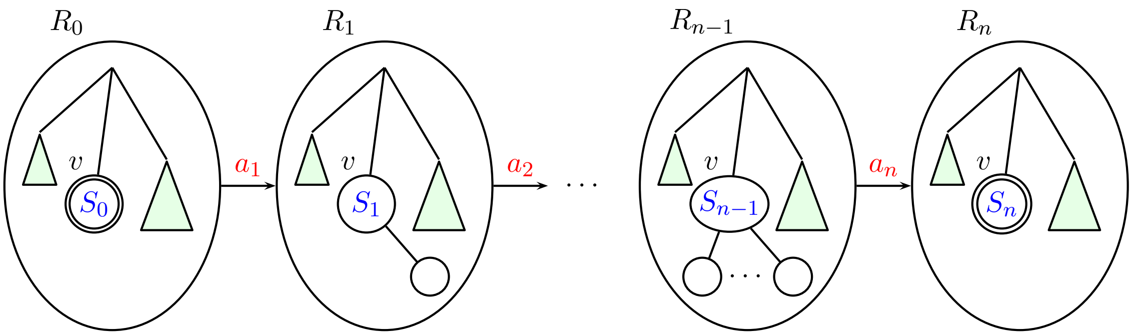

Figure 9.18: The states R_i.

Lemma 9.3. *For $0 \leqslant i \leqslant n-1$, S_{i+1} is contained in $\delta(S_i, a_{i+1})$. Moreover, for every $q \in S_n$, there is a path in $\mathcal{A}$ starting in S_0, ending at q, labeled u and visiting at least one final state after its origin.*

Proof. We follow the construction step by step. We first compute at step 1 the set $S_{i+1} = \delta(S_i, a_{i+1})$, then we suppress some states of S_{i+1} during step 3. The first part of the lemma follows.

Let us show by induction on i that, for $0 \leqslant i \leqslant n-1$ and for every state q_i appearing in the label of a descendant of v in R_i, there exists a path in $\mathcal{A}$ starting in S_0, ending with q_i, labeled $a_1 \cdots a_i$ and passing at least once more by a final state. The result holds for $i = 0$, since v, which is marked has to be a leaf of R_0 and has no descendants. On the other hand, if q_{i+1} appears in the label of a descendant of v in R_{i+1}, either $q_{i+1} \in \delta(q_i, a_{i+1})$ for some q_i appearing in the label of a strict descendant of v in R_i, and we conclude by induction, or q_{i+1} appears in a label created at step 2 and thus $q_{i+1} \in F$, which also allows one to conclude.

Finally, since v is marked in R_n, it received its mark at step 5. Thus if $q \in S_n$, either $q \in \delta(S_{n-1}, a_n) \cap F$, or q belongs to the union of the $\delta(q_{n-1}, a_n)$ where q_{n-1} appears in the label of a descendant of v in R_{n-1}. In the first case, there exists a path labeled u, starting in S_0 and ending with q, which is a final state. In the second case, we use the conclusion of the above induction: there exists a path in $\mathcal{A}$ starting in S_0, ending with q_{n-1}, labeled $a_1 \cdots a_{n-1}$ and passing at least once more by a final state. The lemma follows immediately. □

Consider now a successful path c in $\mathcal{D}$ and let $u \in A^\omega$ be the label of c. There exists a $v \in V$ such that, ultimately, the path visits only states in which v is a node and infinitely often states in which v is a marked node. Setting $S_0 = I$, there exists by Lemma 9.3 a factorization $u = u_0u_1u_2 \cdots$ and subsets S_n of Q, such that

(a) For every $n \geqslant 0$, $S_{n+1} \subset \delta(S_n, u_n)$

(b) For every $n > 0$ and for every $q \in S_{n+1}$, there exists a path in $\mathcal{A}$ starting in S_n, ending with q, labeled u_n and visiting at least one final state after its origin.

In order to apply König's lemma, we build a tree (N, r, p) as follows. The set of nodes is

$$N = \{r\} \cup \{(q, n) \mid q \in S_n, n \in \mathbb{N}\}$$

The parent of each node of the form $(q, 0)$ is r and, for $n > 0$, the parent of each node of the form $(q, n+1)$ is chosen among the states (q', n) such that there is a path in $\mathcal{A}$ starting in q', ending in q, labeled u_n and visiting at least one final state after its origin. Conditions (a) and (b) guarantee the possibility of such a construction. Since the tree thus obtained is infinite and since each child has only a finite number of children, it contains an infinite path by König's lemma. This implies the existence of an infinite path in $\mathcal{A}$, labeled u, starting in I and passing infinitely often through a final state. Thus u is accepted by $\mathcal{A}$.

Conversely, let us consider a successful path c of the automaton $\mathcal{A}$

$$c : q_0 \xrightarrow{a_0} q_1 \xrightarrow{a_1} q_2 \ \cdots$$

There is a unique initial path in $\mathcal{D}$ with the same label u

$$d : I = R_0 \xrightarrow{a_0} R_1 \xrightarrow{a_1} R_2 \cdots$$

Each of the states q_i belongs to the label of the root of the trees R_i. This root is never suppressed and it is thus a fixed element v_0 of V. If v_0 is marked infinitely often in the R_i's, the path is successful in $\mathcal{D}$ and the word u is accepted. Otherwise, there is a largest integer n such that v_0 is marked in R_n. Let n_0 be this integer and let us consider the smallest integer $m > n_0$ such that q_m is an infinitely repeated final state. Since q_m is final, it appears in a child of the root, and from some time $n_1 \geqslant m$ on, each q_n with $n \geqslant n_1$ appears in a fixed child v_1 of the root of R_n. Indeed, if q_n occurs in the label of given node v, then q_{n+1} occurs again in the label of v at the next step, unless it occurs on the left of v (step 3). But such a left shift can occur only a finite number of times. If v_1 is marked infinitely often, the path is successful in $\mathcal{D}$. Otherwise, we repeat the same process, replacing v_0 by v_1. Since the tree has a finite height, we always find some node which is marked infinitely often.

We shall see later other proofs of McNaughton's theorem which cast a different light upon it (see Section II.9). Among its numerous consequences, we begin with the most important one, known as Büchi's theorem.

Theorem 9.4. *The class of recognizable subsets of A^ω is closed under complement.*

Proof. By McNaughton's theorem, any recognizable set can be recognized by a Muller automaton. This automaton can be supposed to be complete by Proposition 7.3. Conversely, by Theorem 7.1 and Theorem 5.4 any set recognized by a Muller automaton is recognizable. The result follows from the fact that, by Proposition 7.6, the class of sets recognized by Muller automata is closed under all boolean operations. □

Büchi's theorem can also be proved directly using congruences (see Chapter II). But the size of the automaton for the complement given by Safra's algorithm is asymptotically optimal, as will be shown in Section 10.5 using the following result.

Theorem 9.5. *For each $n > 0$, there exists a set L_n of infinite words recognized by a Büchi automaton with $n + 2$ states, such that any Büchi automaton recognizing the complement of L_n has at least $n!$ states.*

Proof. Let $A_n = \{0, 1, \ldots, n\}$ and let $\mathcal{A}_n$ be the automaton on the alphabet A_n represented in Figure 9.19 and let $L_n = L^\omega(\mathcal{A}_n)$.

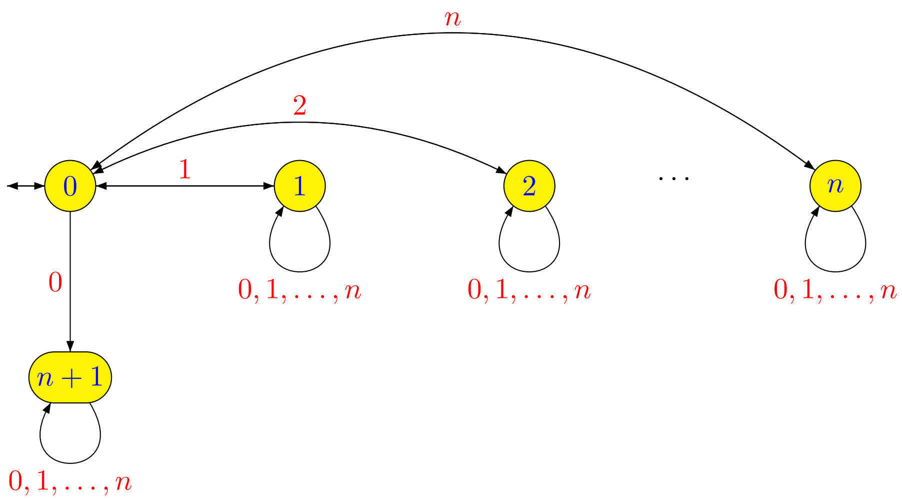

Figure 9.19: A Büchi automaton recognizing L_n.

One could of course describe precisely L_n, but two weaker lemmas will be sufficient for our purpose. We start with a sufficient condition for a word to be in L_n.

Lemma 9.6. *Let $\{i_1, i_2, \ldots, i_k\}$ be a subset of $\{1, 2, \ldots, n\}$. If an infinite word u contains infinitely many occurrences of each of the factors i_1i_2, i_2i_3, ..., i_ki_1, and if in $\mathcal{A}_n$, there is a finite path from 1 to i_1 labeled by a prefix of u, then $u \in L_n$.*

Proof. It suffices to describe a successful path of label u in $\mathcal{A}_n$. By hypothesis, there is a path from 1 to i_1 labeled by a prefix of u. We then stay in state i_1 until the next occurrence of i_1i_2, that is used to produce the transitions $i_1 \xrightarrow{i_1} 0 \xrightarrow{i_2} i_2$. Then we stay in state i_2 until the next occurrence of i_2i_3, that is used to produce the transitions $i_2 \xrightarrow{i_2} 0 \xrightarrow{i_3} i_3$, etc. This process, repeated infinitely often on the cycle (i_1i_2, i_2i_3, ..., i_ki_1), produces the desired successful path. □

With each permutation σ of $\{1, 2, \ldots, n\}$, associate the infinite word $u_\sigma = (\sigma(1) \cdots \sigma(n)0)^\omega$.

Lemma 9.7. *For any permutation σ of $\{1, \ldots, n\}$, the infinite word u_σ is not in L_n.*

Proof. Clearly, $L_n \subset K^\omega$, where $K = \bigcup_{1 \leqslant i \leqslant n} iA_n^* i$. Therefore, if $u_\sigma \in L_n$, $u = u_1u_2\ldots$, where each u_i is in K. It follows that $\sigma(1)$ is the first and last letter of u_1, $\sigma(2)$ is the first and last letter of u_2, and $\sigma(n)$ is the first and last letter of u_n. Consequently, the first letter of u_{n+1} is 0, a contradiction, since $u_{n+1} \in K$. □

Let now $\mathcal{B}$ be a Büchi automaton accepting the complement of L_n. By Lemma 9.7, each word u_σ is accepted by $\mathcal{B}$. Therefore, there is in $\mathcal{B}$ a successful path p_σ of label u_σ. We claim that if $\sigma \neq \sigma'$, then $\mathrm{Inf}(p_\sigma) \cap \mathrm{Inf}(p_{\sigma'}) = \emptyset$. Assume by contradiction that some state q belongs to both $\mathrm{Inf}(p_\sigma)$ and $\mathrm{Inf}(p_{\sigma'})$. Using the two paths, we build a new path

p in $\mathcal{B}$ which, at the beginning, follows a prefix of p_σ of length at least $n(n+1)$ until it reaches q. Then p enters a loop which is repeated infinitely often. This loop consists of two parts that we also take of length at least $n+1$: in the first part, p follows a portion of p_σ to go from q to q after visiting at least once all states of $\mathrm{Inf}(p_\sigma)$ and in the second part, p follows a portion of $p_{\sigma'}$ to go from q to q after visiting at least once all states of $\mathrm{Inf}(p_{\sigma'})$ (see Figure 9.20). Then $\mathrm{Inf}(p)$ contains $\mathrm{Inf}(p_\sigma)$ (and $\mathrm{Inf}(p_{\sigma'})$) and in particular contains a final state, since p_σ is successful. It follows that p is successful and thus its label u is not in L_n.

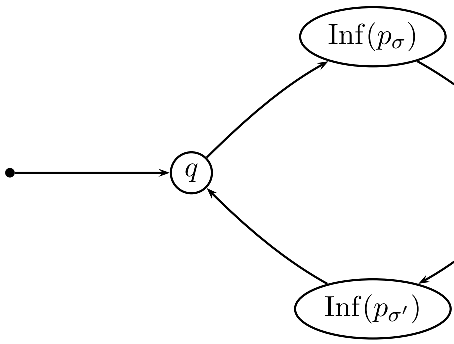

Figure 9.20: The path p.

We shall arrive to a contradiction by showing that u satisfies the conditions of Lemma 9.6, and therefore belongs to L_n.

We first verify the existence of a cycle of infinitely repeated factors of length two. Let k be the smallest integer such that $\sigma(k) \neq \sigma'(k)$. Then $\sigma'(k) = \sigma(l)$ for some $l > k$ and $\sigma(k) = \sigma'(m)$ for some $m > k$. Since u is a concatenation of factors of length at least $n+1$ of u_σ and $u_{\sigma'}$, each of the factors $\sigma(k)\sigma(k+1)$, $\sigma(k+1)\sigma(k+2)$, $\ldots$, $\sigma(\ell-1)\sigma(\ell)(=\sigma(\ell-1)\sigma'(k))$, $\sigma'(k)\sigma'(k+1)$, $\ldots$, $\sigma'(m-1)\sigma'(m)(=\sigma'(m-1)\sigma(k))$ occur infinitely often in u.

It suffices now to verify that the state $\sigma(k)$ is reachable in $\mathcal{A}_n$ by a path labelled by a prefix of u. By construction, the word $(\sigma(1)\cdots\sigma(n)0)^n$ is a prefix of u. Therefore, the path

$$0 \xrightarrow{\sigma(1)} \sigma(1) \xrightarrow{\sigma(2)\cdots\sigma(n)0} \sigma(1) \xrightarrow{\sigma(1)} 0 \xrightarrow{\sigma(2)} \sigma(2) \cdots \sigma(k-1) \xrightarrow{\sigma(k-1)} 0 \xrightarrow{\sigma(k)} \sigma(k)$$

is suitable for our purpose.

This proves the claim, and since there are $n!$ permutations on $\{1,\ldots,n\}$, there are at least $n!$ disjoint sets of the form $\mathrm{Inf}(p_\sigma)$, which clearly implies that $\mathcal{B}$ has at least $n!$ states. □

As announced above, recognizable sets are determined by the ultimately periodic words they contain.

Corollary 9.8. *Let X and Y be two recognizable subsets of A^ω. Let $U \subset A^\omega$ be the set of ultimately periodic words. If $X \cap U \subset Y$, then $X \subset Y$. In particular $X = Y$ if and only if X and Y contain the same ultimately periodic words, i.e. if $X \cap U = Y \cap U$.*

Proof. In fact, if X is not contained in Y, the set $X \setminus Y$ is, by Büchi's theorem, a nonempty recognizable subset of A^ω. By Lemma 5.1, there exists an ultimately periodic word which is in X but not in Y. □

We now turn to another consequence of McNaughton's theorem, which solves a subtle point raised in Section 6.

Theorem 9.9. *A subset of A^ω is recognizable by a finite deterministic Büchi automaton if and only if it is both deterministic and recognizable.*

Proof. Any set recognized by a finite deterministic Büchi automaton satisfies certainly these two conditions. Conversely, let X be a subset of A^ω satisfying the two conditions. By McNaughton's theorem, the set X is recognized by a Muller automaton $\mathcal{A} = (Q, A, E, i, \mathcal{T})$. But since X is deterministic, the table $\mathcal{T}$ is full by Proposition 7.10. Finally, Proposition 7.9 shows that X can be recognized by a finite deterministic Büchi automaton. □

Corollary 9.10. *It is decidable whether a given recognizable subset of A^ω is deterministic or not.*

Proof. Let X be a recognizable subset of A^ω. We may build, using the previously described algorithms, a Muller automaton recognizing X. Proposition 7.10 allows one to conclude. □

Example 9.4. The set $X = (\{b, c\}^* a \cup b)^\omega$ of Example 9.3 is deterministic. On the contrary, the set $Y = (a\{b, c\}^* \cup b)^\omega$ is not deterministic. In fact Y is recognized by the Büchi automaton represented in Figure 9.21.

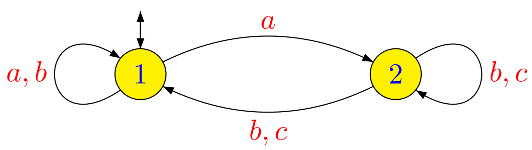

Figure 9.21: A co-deterministic but non deterministic automaton.

The deterministic automaton obtained by Safra's algorithm is represented in Figure 9.22

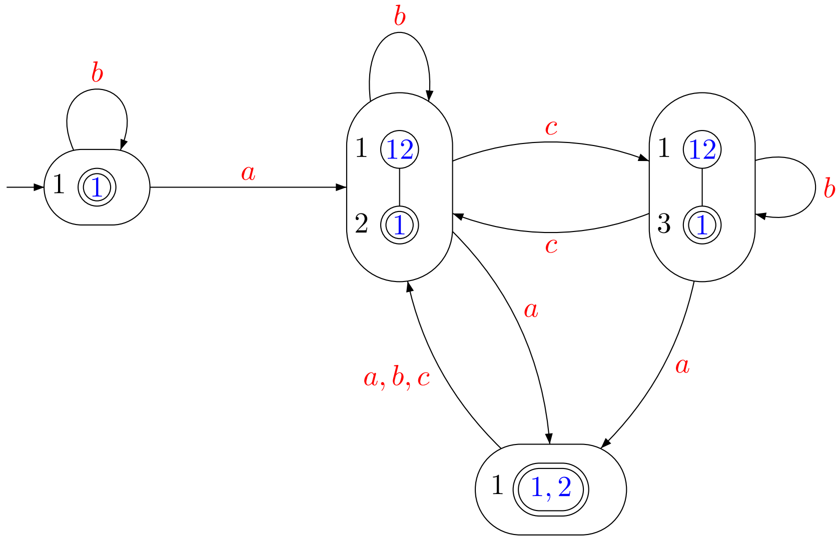

Figure 9.22: The resulting automaton.

One can check directly that Y is not deterministic, by imitating the construction used in Example 6.2. Let us indeed suppose that $Y = \overrightarrow{L}$. Since $acb^\omega \in Y$, there is an integer n_1 such that $acb^{n_1} \in L$. Again, since $acb^{n_1}cb^\omega \in Y$, there is an integer n_2 such that $acb^{n_1}cb^{n_2} \in L$, etc. and the infinite word $u = acb^{n_1}cb^{n_2}cb^{n_3}\cdots$ has an infinite number of prefixes in L. This implies that $u \in Y$, which is impossible since u contains infinitely many c's but a finite number of a's.

10 Computational complexity aspects

In this section, we address the problem of the computational complexity of the various transformations introduced in this chapter. The results are summarized in Figure 10.1. The nodes of this graph illustrate various representations of sets of infinite words, such as ω-rational expression, Büchi automaton, etc. An arrow between two nodes indicates an algorithm to convert one representation into another one. The label of the arrow indicates the complexity of the corresponding algorithm. The label P stands for a polynomial time algorithm and Exp for an exponential one.

The size of the various objects is defined according to the following conventions. As a general rule, we consider the cardinality of the alphabet as being a constant.

The *size* of an ω-rational expression is the number of symbols that it involves, without the parenthesis but taking the dot into account for the product. Thus $\text{size}(\varepsilon) = \text{size}(\{a\}) = 1$ and $\text{size}(X+Y) = \text{size}(XY) = \text{size}(X)+\text{size}(Y)+1$, $\text{size}(X^*) = \text{size}(X)+1$.

The size of a Büchi automaton $\mathcal{A} = (Q, E, I, T)$ is $\max(\text{Card}(Q), \text{Card}(E))$. It is thus at most equal to $\text{Card}(Q)^2 \times \text{Card}(A)$.

The size of a Muller automaton $\mathcal{A} = (Q, E, i, \mathcal{T})$ is

$$\max(\mathrm{Card}(Q), \mathrm{Card}(E), \mathrm{Card}(\mathcal{T}))$$

It is thus bounded by $\max(\mathrm{Card}(Q) \times \mathrm{Card}(A), 2^{\mathrm{Card}(Q)})$. Note that it may be exponential in $\mathrm{Card}(Q)$.

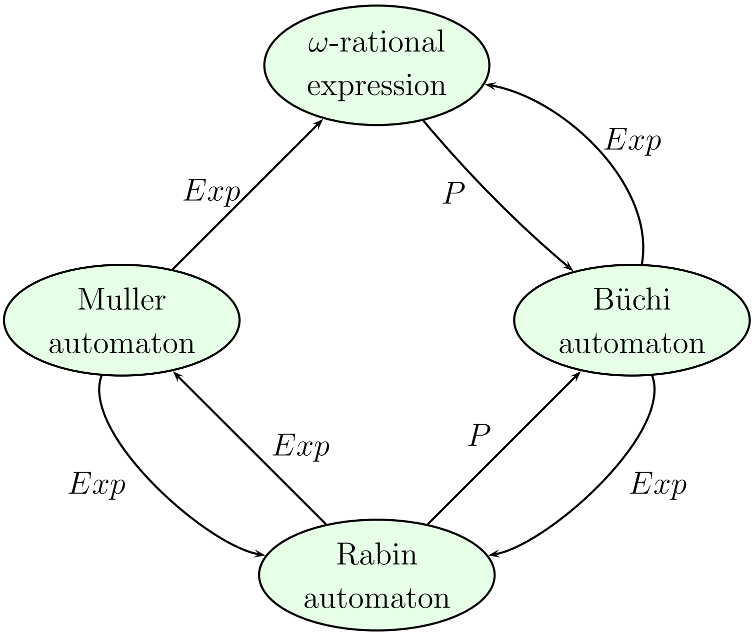

Figure 10.1: Summary of transformations.

We describe separately the transformation associated with each arrow in Figure 10.1. The algorithms associated with these transformations have the same complexity as the size of the resulting objects.

10.1 From ω-rational expressions to Büchi automata and back.

The following result shows that ω-rational expressions and Büchi automata are objects of essentially equivalent computational complexity.

Proposition 10.1. *For any ω-rational expression of size n describing a set $X \subset A^\omega$, there is a Büchi automaton of size $O(n)$ recognizing X.*

Proof. Given an ω-rational expression of size s, the size of a Büchi automaton built by the method described in Section 5 is bounded by $2s$.

Conversely, given a Büchi automaton of size s, one can compute a corresponding ω-rational expression of size bounded by some linear function of s. □

10.2 From Büchi automata to Rabin automata.

The following result implies that, given a Büchi automaton of size n, there is an equivalent Rabin automaton of size $2^{O(n \log n)}$.

Proposition 10.2. *For any Büchi automaton of size n, there exists an equivalent Rabin automaton with $2^{O(n \log n)}$ states and n pairs.*

Given a Büchi automaton with n states, the algorithm of Section 9 builds a Rabin automaton on a set of states $\mathcal{T}_n$ with $O(n)$ pairs. Thus the result follows from the following proposition.

Proposition 10.3. *The set $\mathcal{T}_n$ satisfies* $\ln(\mathrm{Card}(\mathcal{T}_n)) = O(n \ln n)$.

Proof. Let, as in the proof of Proposition 9.2, $r(v)$ be the set of states which appear in the label of v but in none of its children. The relation $s = r^{-1}$, is a function from Q onto T which completely determines e, since we have

$$e(v) = \bigcup_{w \text{ ancestor of } v} s^{-1}(w).$$

An element of $\mathcal{T}_n$ is described by the tuple (T, f, s, M) where $(T, 1, f)$ is an oriented tree with at most n nodes, s is a partial function from Q onto T and M is a subset of T. By Proposition 2.1, the number of planar trees with k nodes is

$$C_k = \frac{(2k-2)!}{k!(k-1)!}.$$

To obtain the number of oriented trees with k nodes chosen in V, we multiply by the number of injective functions from a k element set to V, which is

$$I_k = \frac{(2n)!}{(2n-k)!}.$$

We can then bound the number S_k of partial surjective functions from Q to a k-element set by the total number of partial functions from Q to a k-element set which is $(k+1)^n$. Finally, the number of subsets of a k-element set is 2^k. We obtain

$$\begin{aligned}
\mathrm{Card}(\mathcal{T}_n) &\leqslant \sum_{1 \leqslant k \leqslant n} C_k I_k S_k 2^k \leqslant n C_n I_n S_n 2^n \\
&\leqslant n \frac{(2n-2)!}{n!(n-1)!} \frac{(2n)!}{n!} (n+1)^n 2^n \\
&\leqslant \frac{(2n-2)!}{(n-1)!^2} \frac{(2n)!}{n!^2} (n-1)!(n+1)^{n+1} 2^n
\end{aligned}$$

and, observing that $\frac{(2n)!}{(n)!^2} \leqslant 4^n$, we conclude that

$$\ln(\mathrm{Card}(\mathcal{T}_n)) \leqslant (2n-1)\ln 4 + \ln((n-1)!) + (n+1)\ln(n+1) + n\ln 2$$

whence the result since $\ln(n!) = O(n \ln n)$. □

10.3 From Rabin automata to Muller automata.

The following result shows that the transformation from a Rabin automaton into a Muller automaton involves an exponential blow-up in the size of the automaton.

Proposition 10.4. *For any Rabin automaton with n states, there is an equivalent Muller automaton with n states and $O(2^n)$ accepting sets.*

Proof. A Rabin automaton with n states and m pairs can be transformed into a Muller automaton with n states and 2^n accepting sets of states, as shown in the proof of Proposition 7.7. □

The method indicated in Exercise 17 shows that, conversely, any Muller automaton with n states and m accepting sets can be converted into an equivalent Rabin automaton with $2^{O(n)}$ states and m pairs. Thus, the size of the Rabin automaton is exponential in the size of the original Muller automaton.

10.4 From Rabin automata to Büchi automata.

The following result shows that there is a polynomial algorithm transforming a Rabin automaton into an equivalent Büchi automaton.

Proposition 10.5. *For any Rabin automaton $\mathcal{A}$ with n states and m accepting pairs, there is a Büchi automaton $\mathcal{B}$ of size $O(nm)$ such that $L^\omega(\mathcal{A}) = L^\omega(\mathcal{B})$.*

Proof. Let $\mathcal{A} = (Q, i, \mathcal{R})$ be a Rabin automaton with $\mathcal{R} = \{(L_i, U_i) \mid 1 \leqslant i \leqslant m\}$. The set of states of the automaton $\mathcal{B}$ is the union of Q and of the set

$$\{(q, i) \in Q \times \{1, 2, \ldots, m\} \mid q \notin L_i\}$$

The initial state is the initial state of $\mathcal{A}$. Each transition $p \xrightarrow{a} q$ of $\mathcal{A}$ gives rise to the additional transitions $p \xrightarrow{a} (q, i)$ for each i such that $q \notin L_i$ and $(p, i) \xrightarrow{a} (q, i)$ for each i such that $p \notin L_i$ and $q \notin L_i$. The terminal states are the (p, i)'s such that $p \in U_i$. Thus a successful path in $\mathcal{B}$ begins with a path of $\mathcal{A}$ and chooses nondeterministically to avoid some L_i, while checking that it meets U_i infinitely often. It is clear that $\mathcal{A}$ and $\mathcal{B}$ are equivalent and that $\mathcal{B}$ has at most $n + nm = O(nm)$ states. □

10.5 From Streett automata to Büchi automata.

We describe a construction which will be used below to compute a Büchi automaton recognizing the complement.

Proposition 10.6. *Any Streett automaton with n states and m accepting pairs is equivalent to a Büchi automaton of size $n2^{O(m)}$.*

Proof. Let $\mathcal{A} = (Q, i, \mathcal{P})$ be a Streett automaton with a set $\mathcal{P} = \{(L_i, U_i) \mid 1 \leqslant i \leqslant m\}$ of accepting pairs. We construct an equivalent Büchi automaton $\mathcal{B} = (U, j, T)$ with U as set of states, j as initial state and T as set of terminal states as follows. The states have the form (q, S_1, S_2) where $q \in Q$ is a state of $\mathcal{A}$, and S_1, S_2 are finite sets of integers either reduced to $\{0\}$ or equal to a subset of $\{1, 2, \ldots, m\}$. The initial state is $j = (i, 0, 0)$. The transitions are of two kinds. There is a first set of initial transitions formed of all transitions $(p, 0, 0) \xrightarrow{a} (q, 0, 0)$ or $(q, \emptyset, \emptyset)$ where $p \xrightarrow{a} q$ is a transition of $\mathcal{A}$. The second set consists of transitions of the form $(p, S_1, S_2) \xrightarrow{a} (q, S'_1, S'_2)$ where $p \xrightarrow{a} q$ is a transition of $\mathcal{A}$ and

$$S'_1 = \begin{cases} S_1 \cup \{i\} & \text{if } q \in U_i \\ S_1 & \text{otherwise} \end{cases} \quad \text{and} \quad S'_2 = \begin{cases} \emptyset & \text{if } S_1 \subset S_2 \\ S_1 \cup \{i\} & \text{if } q \in L_i \text{ and } S_1 \not\subset S_2 \\ S_1 & \text{otherwise} \end{cases}$$

Thus the automaton $\mathcal{B}$ first behaves as $\mathcal{A}$ during some initial period. In the rest of the time, it uses the sets S_1, S_2 to remember which of the sets L_i, U_i were visited. The set T of terminal states is formed of the states (q, S_1, S_2) such that $S_2 = \emptyset$. A path β of the Büchi automaton $\mathcal{B}$ is of the form

$$\beta : (i, 0, 0) \xrightarrow{a_0} \ldots (p_n, S_n, T_n) \xrightarrow{a_n} (p_{n+1}, S_{n+1}, T_{n+1}) \xrightarrow{a_{n+1}} \ldots$$

where S_n is ultimately equal to some subset S of $\{1, 2, \ldots, m\}$ (unless, of course, $S_n = \{0\}$ for all n). It corresponds to a path

$$\gamma : i \xrightarrow{a_0} \ldots p_n \xrightarrow{a_n} p_{n+1} \xrightarrow{a_{n+1}} \ldots$$

of $\mathcal{A}$. The path β is successful if and only if $T_n = \emptyset$ for infinitely many n. This happens if and only if $S \subset T_n$ for infinitely many n. This last condition is equivalent to $p_n \in L_i$ for infinitely many n and for each $i \in S$. This is again equivalent to the fact that for each $i \in \{1, 2, \ldots, m\}$, either $i \notin S$, i.e. $\mathrm{Inf}(\gamma) \cap U_i = \emptyset$, or $i \in S$ and $\mathrm{Inf}(\gamma) \cap L_i \neq \emptyset$. This shows that β is successful if and only if γ is successful. Thus $\mathcal{A}$ and $\mathcal{B}$ are equivalent. □

In the case where $m = O(\log n)$, we obtain as a corollary.

Corollary 10.7. *Any Streett automaton with n states and $O(\log n)$ accepting pairs is equivalent to a Büchi automaton of size $O(n)$.*

Finally, we obtain the following result.

Theorem 10.8. *For any set $X \subset A^\omega$ recognized by an n-state Büchi automaton, there is an automaton of size $2^{O(n \log n)}$ recognizing $A^\omega \setminus X$.*

Proof. Indeed, let $\mathcal{A}$ be a Büchi automaton of size n recognizing X. By Proposition 10.2, we can first build a Rabin automaton $\mathcal{B}$ with $2^{O(n \log n)}$ states and n accepting pairs equivalent to $\mathcal{A}$. Then, considered as a Streett automaton, $\mathcal{A}$ recognizes $A^\omega \setminus X$. By Corollary 10.7, there is an equivalent Büchi automaton of size $2^{O(n \log n)}$. □

Let us formulate two comments about Theorem 10.8. First of all, the complexity of computing the complement of an ω-rational set is not considerably higher than for a rational set of finite words. Indeed, we obtain $2^{O(n \log n)}$ instead of 2^n. Second, this bound is essentially optimal. Indeed, Theorem 9.5 shows that the size of a Büchi automaton for the complement can be $n! = 2^{\Theta(n \log n)}$.

10.6 Complexity of algorithms on automata.

We consider now the algorithmic complexity of some algorithms on ω-rational sets. For a recognizable set $X \subset A^\omega$, let $\zeta(X)$ be the minimal size of a Büchi automaton recognizing X. The following result gives an evaluation of the complexity of some of the basic operations on ω-rational sets.

Proposition 10.9. *Let X_1 and X_2 be recognizable subsets of A^ω and let $m_1 = \zeta(X_1)$, $m_2 = \zeta(X_2)$. Then*

(1) $\zeta(X_1 \cup X_2) = O(m_1 + m_2)$.

(2) $\zeta(X_1 \cap X_2) = O(m_1 m_2)$.

(3) $\zeta(A^\omega \setminus X_1) = m_1^{O(m_1)}$.

Proof. The construction of an automaton recognizing the union of two recognizable sets has been described in the proof of Kleene's theorem (Theorem 5.4). The number m of edges of the new automaton satisfies $m \leqslant (m_1 + 2) + (m_2 + 2)$.

The construction of a Büchi automaton recognizing $X_1 \cap X_2$ is given in Exercise 7. The number m of edges satisfies $m \leqslant 4 m_1 m_2$.

The last statement is Theorem 10.8. □

In the following statement, we put together the consequences of several algorithms seen before in terms of decidability. The complexity of these algorithms will be considered afterwards.

Proposition 10.10. *One may effectively decide*

(1) *the emptiness of a recognizable subset of A^ω,*

(2) *the inclusion of two recognizable subsets of* A^ω,
(3) *the equality of two recognizable subsets of* A^ω.

Proof. (1) Let $\mathcal{A} = (Q, A, E, I, F)$ be a finite Büchi automaton. By Proposition 5.2, we may suppose that $\mathcal{A}$ is trim. In this case, $L^\omega(\mathcal{A})$ is empty if and only if Q is empty.

(2) The inclusion $X \subset Y$ is equivalent to $X \setminus Y = \emptyset$. We may thus apply (1) and Proposition 10.9.

(3) The equality of two subsets X and Y of A^ω is equivalent to the double inclusion $X \subset Y$ and $Y \subset X$ and is thus decidable. □

We now consider the algorithmic complexity of the basic problems on Büchi automata. We will use here some notation and results from the theory of algorithms (see the Notes Section for references). The *nonemptiness problem* for a Büchi automaton is to decide if $L^\omega(\mathcal{A}) \neq \emptyset$. A problem is in the class NL if it can be solved by a nondeterministic algorithm operating with a logarithmic amount of space.

Proposition 10.11. *The nonemptiness problem for Büchi automata is decidable in linear time. It is* NL*-complete.*

Proof. The set $L^\omega(\mathcal{A})$ is nonempty if and only if there is a cycle in the underlying graph of the automaton which is accessible from the initial state. A depth-first search in this graph can verify this in linear time.

To show that the nonemptiness problem is NL-complete, we first have to show that it is in the class NL and then that it is NL-hard, i.e. that any problem in NL can be reduced to it. We shall prove the first part only. See the Notes Section for a reference to a proof of the second part.

The nondeterministic algorithm to verify nonemptiness consists in choosing nondeterministically at each step a state of the automaton $\mathcal{A} = (Q, i, T)$ of size n. If, from state p, one chooses q, we check that there is an edge connecting p to q. If this is the case, we can forget p and continue. Otherwise, we choose another state q'. When we reach a terminal state $t \in T$, we remember it. We then continue until we find t again. We thus verify nondeterministically the existence of a successful path in $\mathcal{A}$. The space needed consists in memorizing two states (the current state on the path and the terminal state when reached). Since $\log n$ bits are enough to represent an element of Q, the proof is complete. □

The *nonuniversality problem* for a Büchi automaton $\mathcal{A}$ is to decide whether $L^\omega(\mathcal{A}) \neq A^\omega$. A problem is in the class $PSPACE$ if it can be solved by an algorithm using a polynomial amount of space.

Proposition 10.12. *The nonuniversality problem for automata is decidable in exponential time. It is* $PSPACE$*-complete.*

Proof. The nonuniversality problem for $\mathcal{A}$ is the same as the nonemptiness problem for the complement of $L^\omega(\mathcal{A})$. Since the size of an automaton $\bar{\mathcal{A}}$ for the complement is exponential, the first assertion follows easily. The second assertion requires more attention. First, by Proposition 10.11, we can verify the nonemptiness of $\bar{\mathcal{A}}$ with a nondeterministic polynomial space algorithm. Second, by a result known as Savitch theorem (see the Notes Section for references), any problem solvable by a nondeterministic polynomial space algorithm is in $PSPACE$. We again leave the proof of the $PSPACE$-hardness to be found in the references of the Notes Section. □

11 Exercises

11.1 Words and trees.

Exercise 1. Let A be a finite alphabet. Show that every word $x \in A^\omega$ can be factorized as $x = yz$, where $y \in A^*$, $z \in A^\omega$, and each letter of z occurs infinitely often in x.

11.2 Rational sets.

Exercise 2. Let X, X_1 and X_2 be subsets of A^∞. Verify the following equalities.

(1) $Pref(Pref(X)) = Pref(X)$,
(2) $Pref(X_1 \cup X_2) = Pref(X_1) \cup Pref(X_2)$,
(3) $Pref(X_1X_2) = Pref(X_1) \cup (X_1 \cap A^*)Pref(X_2)$,
(4) $Pref(X^\omega) = Pref(X^+) = (X \cap A^+)^*Pref(X)$.

Conclude that if X is an ω-rational subset of A^∞, $Pref(X)$ is a rational subset of A^*.

Exercise 3. Verify the following identities, where K and L are subsets of A^*:

$$(K + L)^* = (K^*L)^*K^*$$
$$(KL)^* = 1 + K(LK)^*L$$

Exercise 4. Prove that any subset of A^+ can be recognized by a (possibly infinite) deterministic automaton.

11.3 Büchi automata.

Exercise 5. Show that any subset of A^ω can be recognized by a (possibly infinite) Büchi automaton in which all the states are final.

Exercise 6. Prove that the following conditions are equivalent for a subset X of A^ω.

(1) X is stable by shift, i.e. $\sigma(X) \subset X$.

(2) If X is recognized by a Büchi automaton $\mathcal{A} = (Q, A, E, I, F)$, X is also recognized by the automaton (Q, A, E, P, F) where P is the set of accessible states of $\mathcal{A}$.
(3) X is recognized by a Büchi automaton in which all the states are initial.

Exercise 7. Let $\mathcal{A}_1 = (Q_1, A, E_1, i_1, F_1)$ and $\mathcal{A}_2 = (Q_2, A, E_2, i_2, F_2)$ be two Büchi automata with a single initial state. The aim of this problem is to build a Büchi automaton recognizing $L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$. Let

$$Q = Q_1 \times Q_2 \times \{0, 1\}$$

and consider the automaton $\mathcal{A} = (Q, A, E, (i_1, i_2, 1), F)$ where

$$E = \Big\{ \big((p_1, p_2, s), a, (q_1, q_2, t)\big) \mid (p_1, a, q_1) \in E_1, (p_2, a, q_2) \in E_2 \text{ and}$$
$$t = 0 \text{ if and only if } (s = 1, p_2 \in F_2 \text{ and } q_1 \notin F_1) \text{ or } (s = 0 \text{ and } q_1 \notin F_1 \Big\}$$

and

$$F = Q_1 \times F_2 \times \{1\}.$$

Show that $L^\omega(\mathcal{A}) = L^\omega(\mathcal{A}_1) \cap L^\omega(\mathcal{A}_2)$.

11.4 Deterministic Büchi automata.

Exercise 8. Let $\mathcal{A}$ be a deterministic Büchi automaton. The set of infinite words recognized *in the weak sense* by $\mathcal{A}$ is the set, denoted $L^\omega_w(\mathcal{A})$ of all words which are the label of an infinite path starting in the initial state and passing at least once by a final state.

Show that a set of infinite words can be recognized in the weak sense by a complete deterministic Büchi automaton if and only if it is of the form XA^ω for some $X \subset A^*$.

Exercise 9. Show that if a subset X of A^+ is a finite union of prefix-free sets (and thus in particular either if X is finite or if X is prefix), then $\overrightarrow{X} = \emptyset$.

Exercise 10. Let X be a rational subset of A^+. Show that $\overrightarrow{X} = \emptyset$ if and only if X is a finite union of prefix-free sets.

Show that the above result is not always true if X is not rational (consider the set $X = \{a^n b^m \mid 0 < m \leqslant n\}$).

Exercise 11. Let X be a subset of A^ω. Show that the following conditions are equivalent:
(1) the set X can be recognized by a deterministic Büchi automaton in which all states are final,

(2) there exists a subset P of A^* such that X is the set of infinite words having all its prefixes in P,
(3) there exists a prefix-closed subset P of A^* such that X is the set of infinite words having all its prefixes in P,
(4) there exists a subset R of A^* such that X is the set of infinite words having no prefix in R,
(5) there exists a prefix-closed subset P of A^* such that $X = \overrightarrow{P}$,
(6) the equality $X = \overrightarrow{\mathrm{Pref}(X)}$ holds.

These are the closed sets of the natural topology on A^ω (see chapter III).

Exercise 12. A subset of A^ω is called *open* if its complement satisfies the equivalent conditions of Exercise 11. The link with topology will, here again, be covered in Chapter III.

Show that any recognizable open subset of A^ω is deterministic.

It is also possible to give a characterization of open subsets in terms of automata (see Exercise 8).

Exercise 13. Let X be a subset of A^ω. Prove that the following conditions are equivalent.

(1) X is closed and shift invariant,
(2) X is recognized by a Büchi automaton with deterministic transitions with all its states initial and final.
(3) There exists a subset P of A^* such that X is the set of infinite words having all its factors in P.

11.5 Muller automata.

Exercise 14. Let $\mathcal{A}_1 = (Q_1, A, E_1, i_1, \mathcal{T}_1)$ and $\mathcal{A}_2 = (Q_2, A, E_2, i_2, \mathcal{T}_2)$ be two Muller automata. Build a Muller automaton recognizing $L^\omega(\mathcal{A}_1) \cup L^\omega(\mathcal{A}_2)$.

Exercise 15. A *non deterministic Muller automaton* is a tuple

$$\mathcal{A} = (Q, A, E, I, \mathcal{T})$$

where (Q, A, E) is a finite automaton, I is a subset of Q whose elements are the initial states and $\mathcal{T}$ is a family of subsets of Q, called the *table* of the automaton.

An infinite path p in $\mathcal{A}$ is *initial* if it starts in the initial state and *final* if $\mathrm{Inf}(p) \in \mathcal{T}$. It is *successful* if it is both initial and final. The set of infinite words recognized by $\mathcal{A}$ is the set of labels of infinite successful paths in $\mathcal{A}$.

Show that any subset of A^ω recognized by a non deterministic Muller automaton can be also recognized by a deterministic one. (Hint: the proof of Theorem 7.1 does not use the hypothesis that the automata are deterministic).

Exercise 16. Let $\mathcal{A} = (Q, A, E, i, F)$ and $\mathcal{A}' = (Q', A, E', i', F')$ be two deterministic Büchi automata. Our aim is to build directly a Rabin $\mathcal{B}$ automaton recognizing $L^*(\mathcal{A})L^\omega(\mathcal{A}')$. The idea is to use the automaton $\mathcal{A}$ starting a new computation of $\mathcal{A}'$ each time $\mathcal{A}$ visits a final state. To make the total size finite, we keep only the earliest copy of $\mathcal{A}'$ when two of them are in the same state. As an acceptance condition, we require a final state of a fixed copy of $\mathcal{A}'$ be visited infinitely often.

For this, we first define the set of *simple* words on Q' as the elements of Q'^* in which every state $q \in Q'$ appears at most once. For each $v \in Q'^*$, we denote by $[v]$ the simple word obtained by keeping only the leftmost occurrence of each state.

For a simple word $v = q_1 q_2 \cdots q_n \in Q'^*$ and a letter $a \in A$, we define the action of a on v by

$$v \cdot a = (q_1 \cdot a)(q_2 \cdot a) \cdots (q_n \cdot a)$$

The word $v \cdot a$ may perhaps not be simple. Let k be the largest integer at most equal to n such that the states $q_1 \cdot a, \ \ldots \ , \ q_k \cdot a$ are distinct. The integer k will be denoted $\ell_{v,a}$. If v is the empty word, we let $\ell_{v,a} = 0$.

We define now a Rabin automaton $\mathcal{B}$ by choosing the state set as being the set T of triples

$$(q, v, \ell)$$

with $q \in Q$, v a simple word on Q' and ℓ an integer such that $0 \leqslant \ell \leqslant |v|$. The states are thus formed of a state of $\mathcal{A}$ and a sequence of distinct states of $\mathcal{A}'$ marked at a position ℓ which is used to indicate the position of the last deleted state. The initial state is $(i, 1, 0)$. The transitions are defined by

$$(q, v, \ell) = \begin{cases} (q \cdot a, [v \cdot a], \ell_{v,a}) & \text{if } q \cdot a \notin F, \\ (q \cdot a, [(v \cdot a)i'], \ell_{v,a}) & \text{if } q \cdot a \in F. \end{cases}$$

The acceptance is defined by the family $\mathcal{R} = \big\{(L_k, U_k) \mid 1 \leqslant k \leqslant \text{Card}(Q)\big\}$, with

$$L_k = \{(q, v, \ell) \in T \mid \ell > 0 \text{ and the } k\text{-th symbol of } v \text{ is in } F'\},$$
$$U_k = \{(q, v, \ell) \in T \mid k \leqslant \ell\}.$$

Show that the Rabin automaton $\mathcal{B}$ built in this way recognizes $L^*(\mathcal{A})L^\omega(\mathcal{A}')$.

11.6 McNaughton's theorem.

Exercise 17. Let $\mathcal{A} = (Q, E, i, \mathcal{T})$ be a Muller automaton, with

$$\mathcal{T} = \{T_0, \ldots, T_{k-1}\}$$

Let

$$\mathcal{A}' = (Q', E', i', \mathcal{R}')$$

be the Rabin automaton defined by $Q' = \mathcal{P}(T_0) \times \ldots \times \mathcal{P}(T_{k-1}) \times Q$, $i' = (\emptyset, \ldots, \emptyset, i)$ with transitions defined by $(U_0, \ldots, U_{k-1}, q)\cdot a = (U'_0, \ldots, U'_{k-1}, q')$ where $q' = q\cdot a$,

$$U'_i = \begin{cases} \emptyset & \text{if } U_i = T_i \\ T_i \cap (U_i \cup \{q'\}) & \text{otherwise.} \end{cases}$$

and $\mathcal{R} = \{(L_i, R_i) \mid 0 \leqslant i \leqslant k-1\}$ with $L_i = \{(U_0, \ldots, U_{k-1}, q) \in Q' \mid q \notin T_i\}$ and $R_i = \{(U_0, \ldots, U_{k-1}, q) \in Q' \mid U_i = T_i\}$. Show that $L^\omega(\mathcal{A}) = L^\omega(\mathcal{A}')$.

Exercise 18. Verify Example 9.4 by developing in detail Safra's algorithm.

Exercise 19. Consider a Büchi automaton in which all states are final. Show that the algorithm described in Section 9 coincides with the usual determinization algorithm.

12 Notes.

The use of finite automata to recognize infinite words appears for the first time in the work of Büchi [45] and not much later in a paper of Muller [213]. Kleene's seminal paper on automata [162] already considers the case of infinite words, but on the left side. The motivations of Büchi came from the investigation of decidable logical theories. Those of Muller arose from the study of asynchronous circuits.

Looking for earlier origins of this kind of question leads one to the foundations of topology and measure theory. This dates back to the beginning of the twentieth century, with the work of Borel and Lebesgue and, after them, the Polish school, in particular Suslin, who developed the foundations of analysis (see Bourbaki [37]). We shall come back to this point in the chapter on topology (Chapter III). Even closer to our subject, infinite sequences have been studied as sequences of events in classical probability theory. And, for example, the lemma of Borel-Cantelli makes use of the favorite predicates used in this chapter: "there is an infinity of occurrences of some event" (cf. Feller [111]). This aspect of the study of infinite sequences leads to ergodic theory and symbolic dynamics.

Earlier introductory presentations of automata on infinite words include chapters in books on automata theory (e.g. chapter XIV of [98]) or in multi-author Handbooks (like the survey of Thomas in [336] or [338] and the survey of Staiger in [307]). The recent multi-author monograph [129] gives a consolidated overview of the recent research results achieved in the theory.

The notation used in this chapter is, in general, the one in use elsewhere in automata theory. In some cases, however, we follow the terminology introduced by Eilenberg [98]. Thus, we use systematically the terms "recognizable" and "rational" instead of "regular". The use of the term "rational" deserves a comment. One may, as first shown by Schützenberger, generalize the theory of Kleene to formal series in noncommuting variables. In this framework, rational series appear as the natural generalization of the classical notion of rational series in one variable.

Besides Kleene's theorem, for which we have followed Conway's presentation [80], the main results of this chapter are due to Büchi and to McNaughton. The chronology places first the work of Büchi. He introduced ω-rational sets and automata in connection with the logical monadic second-order theory of $(\mathbb{N}, <)$. We shall come to this theory in detail in Chapter VIII. His main result says that the class of recognizable sets is closed under complement. His original proof differs from the one given in this chapter. It uses semigroup congruences which are introduced in Chapter II. Example 9.4 is due to Gire (personal communication). See also Staiger [304, Example 2, p. 489]. Theorem 7.8 is due to Landweber [171]. It is the essential part of Theorem 9.9.

The notation $\overrightarrow{L}$ is from Eilenberg [98], whereas the classical notation X^ω is borrowed from the theory of ordinals.

McNaughton's theorem was conjectured by Muller [213], who first introduced the automata that bear his name. McNaughton's original proof [192] uses a construction which has been the object of several further studies. Variants of it were described by Rabin [256], Choueka [68] and independently by Eilenberg and Schützenberger [98] in two successive steps:

(1) any recognizable subset of A^ω can be written as a finite union of sets of the form $K\overrightarrow{L}$, where K and L are recognizable subsets of A^+,

(2) if K and L are recognizable subsets of A^+, it is possible to recognize $K\overrightarrow{L}$ by a deterministic Muller automaton.

Part (1) will be developed in Chapter II with the syntactic analysis of recognizable subsets. Part (2) is the subject of Exercise 16.

The proof presented here is due to Safra [272] (see also the tutorial of [268]. It represents a genuinely new proof both because of its reduced computational complexity and because it is a direct construction. Theorem 10.8 is an unpublished result of Michel, presented in the survey of Thomas [338] (see also [110]).

For a general reference concerning the complexity of algorithms, see [120] or [4] or the volume on algorithms of the Handbook of Theoretical Computer Science [345]. Proposition 10.6 is from [272], where it is credited to Vardi. Theorem 10.8 is also from [272]. Safra has obtained in [273] a construction which generalizes Proposition 10.2. It shows that for any nondeterministic Streett automaton with n states and m pairs, there is an equivalent (deterministic) Rabin automaton with $2^{nm\cdot\log(nm)}$ states and nm pairs (see the tutorial in [282].

Two earlier references had obtained a construction of an automaton of exponential size for the complement (Proposition 10.6). Both constructions, by Sistla, Vardi and Wolper [296] and by Pécuchet [229] make use of semigroups (see Proposition VI.4.5). Other constructions are possible, either working directly on Büchi automata, or using