Automata, Logic and Games

C.-H. L. Ong

March 6, 2013

Contents

Chapter 0

Automata, Logic and Games

0.1 Aims and Prerequisites

To introduce the mathematical theory underpinning the computer-aided verification of computing (more generally *reactive*) systems.

- *Automata* (on infinite words and trees) as a model of computation of state-based systems.
- *Logical systems* (such as temporal and modal logics) for specifying operational / correctness properties.
- *Two-person games* as a conceptual basis for understanding and representing *algorithmic* interactions between a system and its environment.

Prerequisites MSc (CS + MFoCS): *Foundations of Computer Science*. Main topics:

- (Finite) Automata Theory: 1st/2nd year *Models of Computation*
- Logic: 1st/2nd year *Logic and Proofs*, or *B1 Logic*
- Computability and Complexity: *Computational Complexity*

For information on the above courses, see http://www.cs.ox.ac.uk/teaching/courses/

Connexions with other DCS courses This course can be viewed as a follow-up of *Computer-Aided Formal Verification*, emphasising the logical and algorithmic foundations. In addition there are several points of contact with *Software Verification*, and *Theory of Data and Knowledge Bases*.

Bibliography Many papers (and some book chapters) in the list can be viewed on the Web.

- (Bradfield and Stirling, 2007) [Must-read for modal mu-calculus.]
- (Khoussainov and Nerode, 2001) [Useful general reference for Büchi automata and S1S.]
- (Grädel et al., 2002) [Encyclopaedic, but uneven quality.]
- (Stirling, 2001) [Good for modal mu-calculus and parity games.]
- (Thomas, 1990) [Quite standard reference, but a little dated.]
- (Thomas, 1997) [Excellent reference for the relevant parts of the course.]

- (Vardi, 1996) [Easy to read; covers Büchi automata and LTL, but takes a different approach.]

Course webpage Lecture slides, exercises for the problem classes, resources, and administrative details of the course will be posted at the course webpage

http://www.cs.ox.ac.uk/teaching/materials12-13/automatalogicgames/

0.2 Motivation

Reactive systems are computing systems that interact indefinitely with their environment. Typical examples are air traffic control systems, programs controlling mechanical devices such as trains and planes, ongoing processes such as nuclear reactors, operating systems and web servers.

Modelling Reactive Systems as Games There are different ways to model reactive systems. *Abstractly* we can model a reactive system by a two-player game:

- Player 0 (or Éloïse) representing the *System*
- Player 1 (or Abelard) representing the *Environment*

Desirable correctness properties of the System are coded as *winning conditions* for Éloïse. A strategy is *winning* for a player if it results in a win for the player, no matter what strategy is played by the other player. Winning strategies for Éloïse correspond to methods of constructing the System. Strategies are algorithms in abstract (and "neutral") form.

0.3 Example: Modelling a Lift Control

Assume a building of 8 levels.

Game perspective A 2-player game.

- Player 0 (Éloïse): Lift controller
- Player 1 (Abelard): Users

System state described by:

- A set of level numbers that have been requested, represented by a bit vector $(b_1, \cdots, b_8) \in \mathbb{B}^8$ whereby $b_i = 1$ iff level i has been requested.
- A level number $i \in \{1, \cdots, 8\}$ for the current position of the lift.
- A number (0 or 1) indicating whose turn it is.

State space of the system is $\mathbb{B}^8 \times \{1, \cdots, 8\} \times \{0, 1\}$.

State-transition graph A directed (bipartite) graph: vertices are states, and edges are transitions.

Transitions Two kinds: arrows from 0-states (Player 0's turn to play) to 1-states (Player 1's turn to play), and vice versa.

- $(b_1, \cdots, b_8, i, 0) \longrightarrow (b'_1, \cdots, b'_8, i', 1)$ such that $i \neq i', b'_{i'} = 0$ and $b'_j = b_j$ for $j \neq i'$.
  The actions involved are: door is closed, movement of lift, door is opened, and movement of people.
- $(b_1, \cdots, b_8, i, 1) \longrightarrow (b'_1, \cdots, b'_8, i, 0)$ such that $b_j \leq b'_j$ for all $j \in \{1, \cdots, 8\}$.
  The actions are: Users push buttons.

Winning conditions (= correctness properties)

Example properties

1. Every requested level will be served eventually.
2. The lift will return to level 1 infinitely often.
3. When the top level (where the boss lives!) is requested, the lift serves it immediately and expeditiously (i.e. does not stop on the way there).
4. While moving in one direction, the lift will stop at every requested level, unless the top level is requested.

Some key questions

1. Is there a lift-control (0-strategy) that can meet all the requirements (winning conditions)?
   Does a winning strategy exist?
   Correctness conditions are typically encoded as logical formulas.
2. How much memory does the control need? Is finite memory enough?
   Is there a finite-state[1] *winning strategy (i.e. one that uses only a finite amount of memory)?*
3. Is there a method that can automatically derive a lift control from a given state-transition graph and a given set of winning conditions?
   Is the winning strategy effectively constructible?
   Such an algorithm is often shown to be constructible by proving that it is a solution to an appropriate decision problem of some class of automata.

[1] A model of computation is *finite state* if it has only finitely many possible configurations. Thus finite-state automata are finite state, but pushdown automata and Turing machines are infinite state.

Chapter 1

Büchi Automata

Synopsis

Definition and examples. Büchi automata are not determinisable. Other acceptance conditions: Muller, Rabin, Streett and Parity. Closure properties of Büchi-recognisable languages. Büchi's proof of complementation via Ramsay's Theorem. McNaughton's Theorem. Transforming non-deterministic Büchi to deterministic Rabin automata: Safra trees and Schewe's history trees. Büchi's characterisation and ω-regular expressions. Decision problems and their complexity: non-emptiness is NL-complete, and universality is PSPACE-complete.

Notations Let U be a set.

- We write U^* to mean the set of finite sequences (or *words* or *strings*) of elements of U. The empty word is denoted ϵ. I.e. U^* is the free monoid over U: the associative binary operation is string concatenation $(u, v) \mapsto u \cdot v$ (or simply $u\, v$, eliding $\cdot$) and the identity is ϵ.
- We write U^ω to mean the set of infinite sequences (or ω-*words* or ω-*strings*) of elements of U. An ω-word is represented as a function from ω to U, ranged over by α, β, ρ, etc. Thus the map α represents the infinite word $\alpha(0)\ \alpha(1)\ \alpha(2) \cdots$.

Let $u \in U^*$ and $w \in (U^* \cup U^\omega)$. We say that u is a *prefix* of w, written $u \leq w$, just if $w = u\, v$ for some $v \in (U^* \cup U^\omega)$. We write $u < w$ just if $u \leq w$ and $u \neq w$.

Henceforth we assume a finite *alphabet* Σ i.e. a finite set of symbols (or letters). Subsets of Σ^* are called $*$-*languages*; subsets of Σ^ω are called ω-*languages*.

1.1 Definition and Examples

We use automata to define ω-languages.

Definition 1.1. A (non-deterministic) *Büchi automaton* is a quintuple $A = (Q, \Sigma, q_0, \Delta, F)$ where

- Q is a finite set of states
- Σ is a finite alphabet
- $q_0 \in Q$ is the initial state
- $\Delta \subseteq Q \times \Sigma \times Q$ is a transition relation

- $F \subseteq Q$ is the set of *final* (or *accepting*) states.

In case Δ is a function $Q \times \Sigma \longrightarrow Q$, we say that A is *deterministic*, and we write δ for the function.

It is helpful to think of a Büchi automaton as a finite, labelled directed graph: each edge is labelled with an element of Σ; and the vertex labels are "initial" (labelling a unique vertex) and "final" (labelling a subset of vertices).

Language Recognised by a Büchi Automaton

A *run* of A on an ω-word $\alpha \in \Sigma^\omega$ is an infinite sequence of states $\rho = \rho(0)\ \rho(1)\ \rho(2) \cdots$ such that $\rho(0) = q_0$, and for all $i \geq 0$, we have

$$(\rho(i), \alpha(i), \rho(i+1)) \ \in \ \Delta.$$

In words, a *run* on α is an infinite path in the directed graph A, starting from the initial vertex, whose labels on the edges trace out the ω-word α.

Note that if A is deterministic then every word has a *unique* run.

A run ρ on α is *accepting* just if there is a final state that occurs infinitely often in ρ; equivalently (because F is finite) $\mathsf{inf}(\rho) \cap F \neq \varnothing$, writing $\mathsf{inf}(\rho)$ for the set of states that occur infinitely often in ρ.

An ω-word α is *accepted* by an antomaton A just if there is an accepting run of A on α. The *language recognised* by A, written $L(A)$, is the set of ω-words accepted by A. An ω-language is *Büchi recognisable* just if it is recognised by some Büchi automaton.

Convention When drawing automata as graphs, we circle the final states, and indicate the initial state by an arrow.

Example 1.1. Set $\Sigma = \{ a, b, c \}$.

(i) $L_1 \subseteq \Sigma^\omega$ consists of ω-words in which after every occurrence of a there is some occurrence of b.

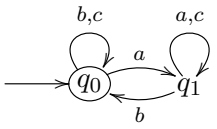

(ii) L_2 consists of ω-words in which between every two occurrences of a, there is an even number of b.

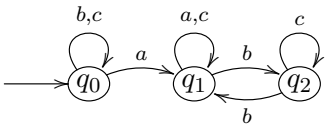

When the automaton reaches state q_1 (respectively q_2), it has read an even (respectively odd) number of b since the last a.

Is L_1 recognised by a deterministic automaton? What about L_2?

Example 1.2 (A Non-Determinisable Büchi Automaton)**.** The Büchi-recognisable language L_3 consisting of ω-words over $\{0,1\}$ that have only finitely many occurrences of 1 is not recognised by any deterministic Büchi automaton.

Suppose, for a contradiction, L_3 is recognised by a deterministic automaton

$$A = (Q, \{0,1\}, q_0, \delta, F).$$

It follows that δ extends to a function $Q \times \{0,1\}^* \longrightarrow Q$. Since A has an accepting run on 0^ω, we have $\delta(q_0, 0^{n_1}) \in F$ for some n_1. Let $u_1 \in Q^*$ be the "run" for 0^{n_1}.

Similarly, since A has an accepting run on $0^{n_1}10^\omega$, we have $\delta(q_0, 0^{n_1}10^{n_2}) \in F$ for some n_2. Let $u_2 \in Q^*$ be the "run" for $0^{n_1}10^{n_2}$. Note that $u_1 \leq u_2$.

In this fashion, we obtain an infinite sequence of numbers $n_1, n_2, \cdots$, and an infinite ascending chain $u_1 \leq u_2 \leq u_3 \cdots$ whose limit is an accepting run of A on the ω-word $0^{n_1}10^{n_2}10^{n_3}10^{n_4}1\cdots$, which is a contradiction. □

Where does the argument break down if A is not deterministic?

Exercise 1.1. Construct a Büchi automaton that recognises (i) L_3 (ii) $\overline{L_3} = \{0,1\}^\omega \setminus L_3$.

Example 1.3. Construct a Büchi automaton for the language L_4 consisting of ω-words α over $\{a,b,c\}$ such that α contains the segments $a\,b$ and $a\,c$ infinitely often, but $b\,c$ only finitely often.

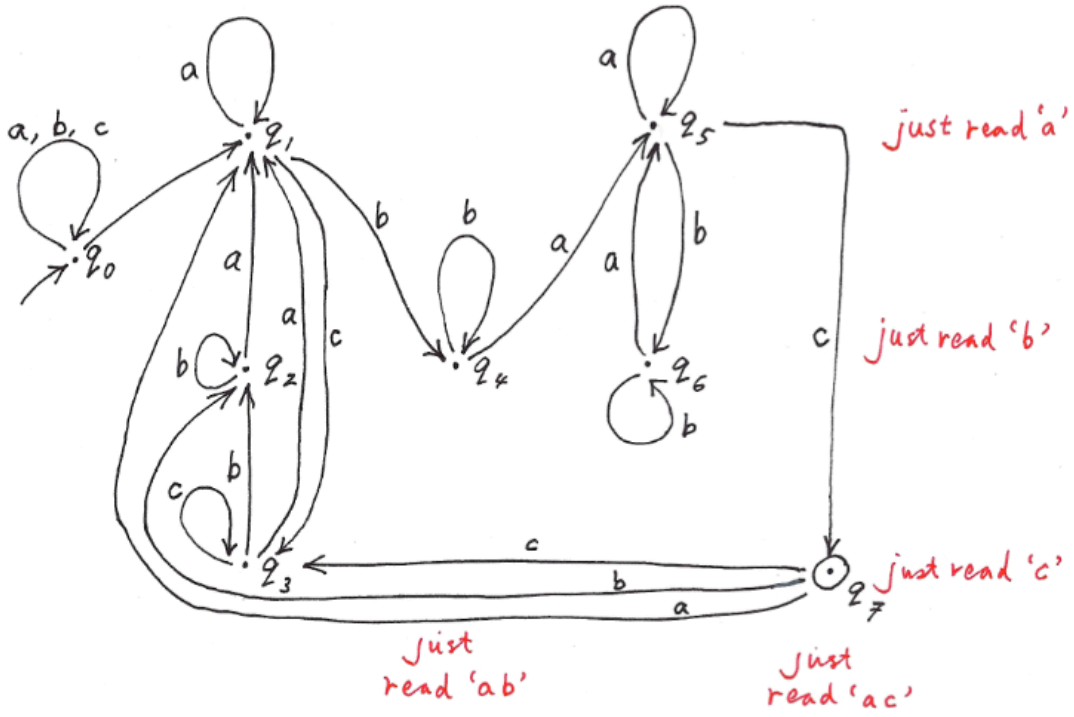

Figure 1.1: An automaton accepting infinitely many $a\,b$ and $a\,c$ but only finitely many $b\,c$.

We argue that the automaton in Figure 1.1 recognises L_4 as follows. By construction:

- when the automaton reaches the states q_1 and q_5, it has just read a

- when it reaches the states q_2, q_4 and q_6, it has just read b
- when it reaches the states q_3 and q_7, it has just read c.

Consequently, after leaving q_0, the automaton is unable to read $b\,c$. Further

- when the automaton reaches state q_4, it has just read $a\,b$
- when it reaches state q_7, it has just read $a\,c$.

It then remains to observe that q_7 is the (only) final state, and after leaving it, the automaton must visit q_4 before it is able to revisit q_7.

Other Acceptance Conditions

An *ω-automaton* is a quintuple $A = \langle\, Q,\ \Sigma,\ q_0 \in Q,\ \Delta \subseteq Q \times \Sigma \times Q,\ Acc \,\rangle$; the component Acc is its acceptance condition.

An ω-automaton is called

- *Büchi*: if Acc is of the form $F \subseteq Q$, and a run ρ is accepting just if $\mathsf{inf}(\rho) \cap F \neq \varnothing$.
- *Muller*: if Acc is of the form $\mathcal{F} = \{\, F_1, \cdots, F_k \,\}$ with each $F_i \subseteq Q$, and a run ρ is accepting just if $\mathsf{inf}(\rho) \in \mathcal{F}$.
- *Rabin*: if Acc is of the form $\{\, (E_1, F_1), \cdots, (E_k, F_k) \,\}$ with $E_i, F_i \subseteq Q$, and a run ρ is accepting just if

$$\exists i \in \{\, 1, \cdots, k \,\} \,.\, \mathsf{inf}(\rho) \cap E_i = \varnothing \ \wedge\ \mathsf{inf}(\rho) \cap F_i \neq \varnothing$$

  I.e. for some i, every state in E_i is visited only finitely often in ρ, but some state in F_i is visited infinitely often in ρ.
- *Streett*: if Acc is of the form $\{\, (E_1, F_1), \cdots, (E_k, F_k) \,\}$ with $E_i, F_i \subseteq Q$, and a run ρ is accepting just

$$\forall i \in \{\, 1, \cdots, k \,\} \,.\, \mathsf{inf}(\rho) \cap E_i = \varnothing \ \vee\ \mathsf{inf}(\rho) \cap F_i \neq \varnothing$$

- *Parity*: Acc is specified by a *priority function* $\Omega : Q \to \omega$, and a run ρ is accepting just if $\min\, \Omega(\mathsf{inf}(\rho))$ is even i.e. the least priority that occurs infinitely often is even.

Rabin condition is sometimes called *pairs condition*; Streett condition is sometimes called *complement pairs condition*; parity condition is sometimes called *Moskowski condition.*

It is straightforward to see that parity condition is closed under negation. Given a priority function $\Omega : Q \to \omega$, let ρ be a run that is not parity-accepting. I.e. $\min\, \Omega(\mathsf{inf}(\rho))$ is odd. Then ρ is parity-accepting w.r.t. the parity function $\Omega' : q \mapsto \Omega(q) + 1$. Note that the Muller condition is also closed under negation. The negation of a Rabin condition is a Streett condition, and vice versa. Given a set of pairs $\{\, (E_1, F_1), \cdots, (E_k, F_k) \,\}$, let ρ be a run that is not Rabin-accepting. I.e. we have $\neg(\exists i.\mathsf{inf}(\rho) \cap E_i = \varnothing \wedge \mathsf{inf}(\rho) \cap F_i \neq \varnothing)$, which is equivalent to $\forall i.(\mathsf{inf}(\rho) \cap E_i \neq \varnothing \ \vee\ \mathsf{inf}(\rho) \cap F_i = \varnothing)$. Thus ρ is Streett-accepting w.r.t. the set of pairs $\{\, (F_1, E_1), \cdots, (F_k, E_k) \,\}$.

Example 1.4 (Muller and Rabin Conditions)**.** Let $L_5 \subseteq \{\, a, b, c \,\}^\omega$ be the language consisting of all words α satisfying: if a occurs infinitely often in α, so does b.

Take the complete state-transition graph with state-set $Q = \{\, q_a, q_b, q_c \,\}$. The idea is that when the automaton reaches state q_a, it has just read a; similarly for q_b and q_c. The Muller and Rabin conditions for a 3-state automaton that recognises L_4 are as follows.

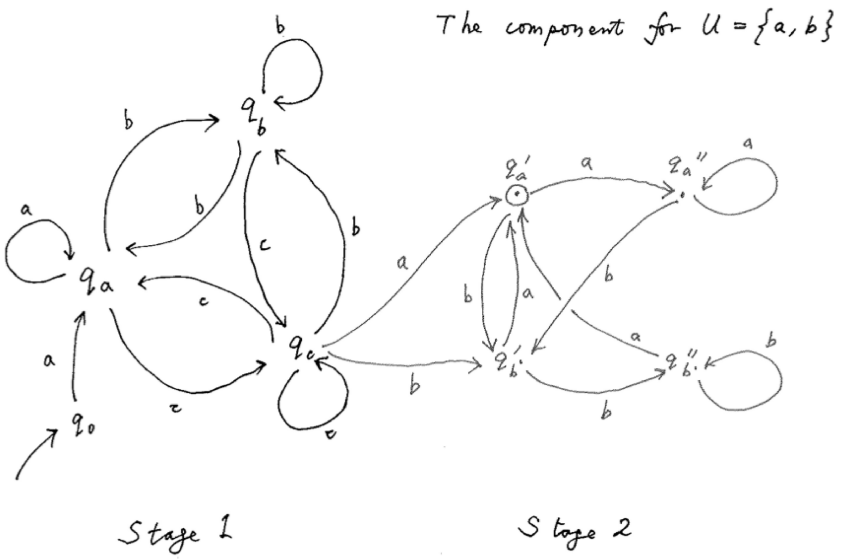

Figure 1.2: Transforming Muller to Büchi

- Muller: Take $\mathcal{F} = \{\, U \subseteq Q \mid q_a \in U \Rightarrow q_b \in U \,\}$.
- Rabin: Take $\Omega = \{\, (\{\, q_a \,\}, \{\, q_b, q_c \,\}), (\varnothing, \{\, q_b \,\}) \,\}$. Observe that $\alpha \in L_5$ iff a occurs only finitely often in α or b occurs infinitely often in α.

Example 1.5 (From Muller to Büchi)**.** Let $A_M = (\{\, q_a, q_b, q_c \,\}, \Sigma, q_0, \mathcal{F})$ be a Muller automaton recognising L_5. We can construct an equivalent Büchi automaton as follows.

Stage 1. Simulate a run ρ of A_M. Guess $\mathsf{inf}(\rho) = U$, for some $U \in \mathcal{F}$. At some point, guess that all states not in U have just been seen.

Stage 2. Check that henceforth:

(i) Every state reached is in U.

(ii) Every state in U is read infinitely often.

See Figure 1.2 for an illustration.

Let $U \subseteq \Sigma^*$.

$$\begin{aligned} U^* &:= \{\, w \in \Sigma^* \mid w = u_1\, u_2 \cdots u_n \text{ for some } n \geq 0, \text{ each } u_i \in U \,\} \\ U^+ &:= \{\, w \in \Sigma^* \mid w = u_1\, u_2 \cdots u_n \text{ for some } n \geq 1, \text{ each } u_i \in U \,\} \\ U^\omega &:= \{\, w \in \Sigma^\omega \mid w = u_1\, u_2 \cdots \text{ where each } u_i \in U \,\} \\ \lim U &:= \{\, w \in \Sigma^\omega \mid w(0) \cdots w(j) \in U \text{ for infinitely many } j \in \omega \,\} \end{aligned}$$

In words, $w \in \lim U$ just if U contains infinitely many prefixes of w.

Example 1.6. (i) Let $U_1 = 0110^* + (00)^+$. Then $\lim U_1$ consists of only two ω-words, namely, $0110000\cdots$ and $0000\cdots$.

(ii) Let $U_2 = 0^*1$. Then $\lim U_2 = \varnothing$.

1.2 Closure Properties

Büchi recognisable languages are closed under all boolean operations i.e. union, intersection and complementation.

Proposition 1. *(i) If $U \subseteq \Sigma^*$ is regular then U^ω is Büchi recognisable.*

(ii) If $U \subseteq \Sigma^*$ is regular and $L \subseteq \Sigma^\omega$ is Büchi recognisable then

$$U \cdot L := \{ u \cdot \alpha \mid u \in U, \alpha \in L \}$$

is Büchi recognisable.

(iii) If L_1 and L_2 are Büchi-recognisable ω-languages, so are $L_1 \cup L_2$ and $L_1 \cap L_2$.

Exercise 1.2. Prove (i) and (ii) of the proposition.

Closure under Union

As a first attempt, use the standard "union construction" in automata for finite words, but note that we can't use ϵ-label edges. Thus for $i = 1, 2$, suppose L_i is recognised by $A_i = (Q_i, \Sigma, q_0^i, \Delta_i, F_i)$. Assume Q_1 and Q_2 are disjoint. Then $L_1 \cup L_2$ is recognised by the Büchi automaton

$$(Q_1 \cup Q_2 \cup \{ q_0 \}, \Sigma, q_0, \Delta, F_1 \cup F_2)$$

where q_0 is a fresh state, and

$$\begin{aligned} \Delta \;&:=\; \Delta_1 \cup \Delta_2 \\ &\cup\; \{ (q_0, a, q) \mid a \in \Sigma, (q_0^1, a, q) \in \Delta_1 \} \\ &\cup\; \{ (q_0, a, q) \mid a \in \Sigma, (q_0^2, a, q) \in \Delta_2 \} \end{aligned}$$

I.e. for each a-transition from q_0^i to q, we add a fresh a-transition from q_0 to q (for $i = 1, 2$).

Closure under Intersection

Suppose L_i is accepted by $A_i = (Q_i, \Sigma, \Delta_i, q_0^i, F_i)$, for $i = 1, 2$. As a first attempt, run the two automata synchronously i.e. in lockstep. Following finite automata for finite words, construct the product automaton

$$(Q_1 \times Q_2, \Sigma, \Delta, (q_0^1, q_0^2), F_1 \times F_2)$$

where $((p, q), a, (p', q')) \in \Delta$ iff $(p, a, p') \in \Delta_1$ and $(q, a, q') \in \Delta_2$. This does not work because we cannot guarantee that the final states of A_1 and A_2 are visited infinitely often *simultaneously*.

The point is that we need to ensure infinite alternation of a F_1-state and a F_2-state. Thus we construct a product automaton and cycle through the following:

1. Wait for an F_1-state in first component.
2. When an F_1-state is encountered in first component, wait for an F_2-state in second component.
3. When an F_2-state is encountered in second component, go to 1.

The Modified Intersection Automaton Work with state-set $Q_1 \times Q_2 \times \{1,2\}$. Form modified product automaton:

$$A' := (Q_1 \times Q_2 \times \{1,2\},\ \Sigma,\ \Delta',\ (q_0^1, q_0^2, 1),\ Q_1 \times F_2 \times \{2\})$$

where: for every $(p,a,p') \in \Delta_1$ and every $(q,a,q') \in \Delta_2$, we have

- $((p,q,1),a,(p',q',1)) \in \Delta'$ if $p \notin F_1$
- $((p,q,1),a,(p',q',2)) \in \Delta'$ if $p \in F_1$
- $((p,q,2),a,(p',q',2)) \in \Delta'$ if $q \notin F_2$
- $((p,q,2),a,(p',q',1)) \in \Delta'$ if $q \in F_2$

It follows that a run ρ of A' on α simulates runs $\rho_1 = \pi_1^*(\rho)$ of A_1 on α and $\rho_2 = \pi_2^*(\rho)$ of A_2 on α—where $\pi_1 : (p,q,j) \mapsto p$ and π_1^* is the point-wise extension—such that ρ visits a state in $Q_1 \times F_2 \times \{2\}$ infinitely often iff ρ_1 visits F_1 infinitely often and ρ_2 visits F_2 infinitely often. □

Closure under Complementation

Theorem 1.1 (Büchi 1960)**.** *If $L \subseteq \Sigma^\omega$ is Büchi recognisable (by A say), so is $\Sigma^\omega \setminus L$. Further the automaton recognising $\Sigma^\omega \setminus L$ can be effectively constructed from A.*

As a first attempt, consider the standard method to complement a finite-state automaton for finite words: first determinise A, then "invert" the final states. Unfortunately, *Büchi automata are not determinisable.* As we have seen, there are non-deterministic Büchi automata (for example, any automaton that recognises L_3 of Example 1.2) that are not equivalent to any deterministic automata.

Büchi's proof We follow the account in (Thomas, 1990) of the proof by Büchi (1960b). Let L be recognisable by a Büchi automaton A. We aim to show that both L and $\Sigma^\omega \setminus L$ are representable as finite unions of sets of the form $L_1 \cdot L_2^*$ where L_1 and L_2 are regular $*$-languages. (Note that it is relatively straightforward to prove the result for L alone: *cf.* Proposition 2.)

We shall construct L_1 and L_2 as congruence classes. We say that a relation $\sim \subseteq \Sigma^* \times \Sigma^*$ is a *congruence* just if $\sim$ is an equivalence relation such that whenever $u \sim u'$ and $v \sim v'$ then $u \cdot v \sim u' \cdot v'$. Set

$$W_{q,q'} := \{\, w \in \Sigma^* \mid q \xrightarrow{w} q' \,\}$$
$$W^F_{q,q'} := \{\, w \in \Sigma^* \mid q \xrightarrow{w}_F q' \,\}$$

where $q \xrightarrow{a_1 \cdots a_n}_X q'$ with $X \subseteq Q$ means that exist $q_0, \cdots, q_n \in Q$ such that $q = q_0 \xrightarrow{a_1} q_1 \xrightarrow{a_2} \cdots \xrightarrow{a_n} q_n = q'$, and $\{\, q_0, \cdots, q_n \,\} \cap X \neq \varnothing$; and in case $X = Q$, we omit the subscript X from $q \xrightarrow{w}_X q'$. Define

$$w \sim_A w' := \forall q, q' \in Q.\, \big((w \in W_{q,q'} \leftrightarrow w' \in W_{q,q'}) \wedge (w \in W^F_{q,q'} \leftrightarrow w' \in W^F_{q,q'})\big)$$

It is straightforward to see that $\sim_A$ is an equivalence relation over Σ^* which has a finite index (because Q is a finite set). It is an easy exercise to show that $\sim_A$ is a congruence.

The equivalence classes can be described as follows: for $w \in \Sigma^*$

$$\begin{aligned}[w]_{\sim_A} \quad = \quad & \textstyle\bigcap_{q,q' \in Q, w \in W_{q,q'}} W_{q,q'} \ \cap \ \bigcap_{q,q' \in Q, w \notin W_{q,q'}} (\Sigma^* \setminus W_{q,q'}) \\ \cap \quad & \textstyle\bigcap_{q,q' \in Q, w \in W^F_{q,q'}} W^F_{q,q'} \ \cap \ \bigcap_{q,q' \in Q, w \notin W^F_{q,q'}} (\Sigma^* \setminus W^F_{q,q'})\end{aligned}$$

Since each $W^F_{q,q'}$ is regular, so is each equivalence class $[w]_{\sim_A}$.

Let $X \subseteq \Sigma^\omega$. We say that a congruence relation $\simeq \subseteq \Sigma^* \times \Sigma^*$ *saturates* X just if for all $u, v \in \Sigma^*$, if $[u]_\simeq \cdot [v]^\omega_\simeq \cap X \neq \varnothing$ then $[u]_\simeq \cdot [v]^\omega_\simeq \subseteq X$.

Lemma 1.1. *(i)* $\sim_A$ *saturates* L

(ii) $\sim_A$ *saturates* $\Sigma^\omega \setminus L$.

Exercise 1.3. Prove (i) and (ii) of the lemma.

Finally it suffices to prove the following.

Claim. Let $X \subseteq \Sigma^\omega$. If a congruence $\simeq$ saturates X and has finite index, then

$$X = \bigcup \{\, [u]_\simeq \cdot [v]^\omega_\simeq \mid [u]_\simeq \cdot [v]^\omega_\simeq \cap X \neq \varnothing \,\}.$$

Assume $\simeq$ saturates X. Then "$\supseteq$" follows from the definition of saturation.

To prove "$\subseteq$", let $w \in X$. Define an equivalence relation $\approx_w \subseteq D \times D$ where $D := \{\, (i,j) \in \omega^2 \mid i < j \,\}$ by

$$(i,j) \approx_w (i',j') \;:=\; w[i,j] \simeq w[i',j']$$

where $w[i,j] = a_i \cdots a_{j-1}$ with $w = a_0\, a_1\, \cdots$. The index of $\approx_w$ is finite because the index of $\simeq$ is finite by assumption.

Now it follows from Ramsey's Theorem[1] that there is an infinite set $H = \{\, i_0, i_1, i_2, \cdots \}$ with $i_0 < i_1 < i_2 < \cdots$ which is *homogeneous* for the map:

$$\{\, i, j \,\} \mapsto [(i,j)]_{\approx_w}$$

assuming $i < j$. I.e. there is a pair (i, i') such that whenever $k < l$ then $(i, i') \approx_w (i_k, i_l)$. In particular all pairs (i_k, i_{k+1}) are in $[(i,i')]_{\approx_w}$. Thus

$$w = w[0, i_0] \cdot w[i_0, i_1] \cdot w[i_1, i_2] \cdot \cdots \in [w[0,i_0]]_\simeq \cdot ([w[i,i']]_\simeq)^\omega$$

as required. This completes the proof of the Claim, and hence the proof of Theorem 1.1.

Given a Büchi automaton A with n states, there are n^2 different pairs (q, q') and hence $O(2^{2n^2})$ different $\sim_A$-classes.

Exercise 1.4. Show that Büchi's complement automaton has a size bound of $O(2^{4n^2})$ states. *Cf.* (Pécuchet, 1986; Sistla et al., 1987).

[1] Let A be a set. We write $(A)_n := \{\, B \subseteq A : |B| = n \,\}$.

Theorem 1.2 (Frank P. Ramsey 1930). *Suppose* $f : (\omega)_n \to \{\, 0, 1, \cdots, k-1 \,\}$. *Then there is an infinite set* $A \subseteq \omega$ *which is homogeneous for* f *i.e.* f *is constant on* $(A)_n$.

If we think of f as a k-colouring of the n-elements subsets of ω, then when A is homogeneous for f, in the sense that all n-element subsets of A have the same colour.

McNaughton's Theorem

Alternatively, following McNaughton (1966), given a Büchi-recognisable language $L(B)$, one can complement it by first transforming the Büchi automaton B to an equivalent deterministic Muller automaton M, and then complement M to get $\overline{M}$. Thus $\Sigma^\omega \setminus L(B) = \Sigma^\omega \setminus L(M) = L(\overline{M})$. Note that Muller acceptance condition is *closed under complementation*: A *deterministic* Muller automaton $(Q, \Sigma, \Delta, q_0, \mathcal{F})$ recognises L if and only if the Muller automaton $(Q, \Sigma, \Delta, q_0, \mathcal{P}(Q) \setminus \mathcal{F})$ recognises $\Sigma^\omega \setminus L$.

Theorem 1.3 (McNaughton 1966)**.** *The following are equivalent:*

(i) non-deterministic Büchi automata (NB)

(ii) deterministic Rabin automata (DR)

(iii) non-deterministic Rabin automata (NR)

(iv) deterministic Muller automata (DM)

(v) non-deterministic Muller automata (NM)

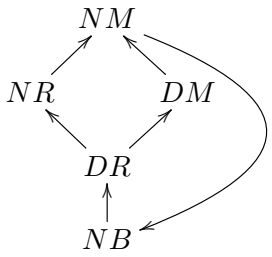

Proof. We write $\Rightarrow$ to mean "can be simulated by".

- "DR $\Rightarrow$ NR" and "DM $\Rightarrow$ NM" are immediate.
- "NB $\Rightarrow$ NR" and "DB $\Rightarrow$ DR": Büchi conditions are instances of Rabin: Given $F \subseteq Q$, the corresponding Rabin condition is $\{ (\varnothing, F) \}$.
- "DR $\Rightarrow$ DM" and "NR $\Rightarrow$ NM": Given Rabin $\{ (E_1, F_1), \cdots, (E_n, F_n) \}$. Set Muller

$$\mathcal{F} \;:=\; \{ U \subseteq Q \mid \bigvee_{i=1}^{n} (U \cap E_i = \varnothing) \wedge (U \cap F_i \neq \varnothing) \}$$

- "NM $\Rightarrow$ NB": (informal)

  Stage 1 Simulate a run ρ of the given Muller automaton. Guess $\mathsf{inf}(\rho) = U$, some $U \in \mathcal{F}$. At some point, guess that all states of ρ that are not in U have just been seen.

  Stage 2 Check that henceforth:

  (i) Every state reached is in U.

  (ii) Every state in U is read infinitely often.

- "NB $\Rightarrow$ DR" is the most difficult: it was first shown by McNaughton (1966), using a double exponential construction.

□

NB $\Rightarrow$ DR: Safra's Construction

Given a NB with n states, Safra's method (Safra, 1988) constructs an equivalent DR with at most $(12)^n\, n^{2n}$ states and $2n$ pairs in the acceptance condition. The key ideas:

(i) Reachable states are organised into trees, called *Safra trees*, carrying information on the history of possible runs reaching each state.

(ii) Nodes of Safra trees are named, allowing a dynamic reuse of the Rabin pairs.

See Lectures 26 and 27 in Kozen's book (Kozen, 2006) for an informal account of Safra's construction, and (Schewe, 2009) for a state-of-the-art NB-to-DR transformation based on a variant of Safra trees called *history trees*.

Recent Progress

(i) Piterman (2007) modified Safra trees to construct equivalent deterministic Parity automata of size at most $2n\,n^n\,n!$

(ii) By a finer analysis of Piterman, Liu and Wang (2009) obtained an upper bound of $2n\,(n!)^2$.

(iii) Schewe (2009) gave a NB-to-DR construction with state complexity $o(2.66\,n)^n$, but requiring 2^{n-1} Rabin pairs.

(iv) Colcombet and Zdanowski (2009) proved that Schewe's construction is optimal for state complexity.

1.3 ω-Regular Expressions

Proposition 2 (Büchi 1960)**.** *A language $L \subseteq \Sigma^\omega$ is Büchi recognisable if and only if L is a finite union of sets of the form $J \cdot K^\omega$, where $J, K \subseteq \Sigma^*$ are regular and $\varnothing \neq K \subseteq \Sigma^+$ (and we may assume $K \cdot K \subseteq K$).*

Proof. Suppose L is recognised by $A = (Q, \Sigma, \Delta, q_0, F)$. For $p, q \in Q$, let $A_{p\,q}$ be the finite automaton $(Q, \Sigma, \Delta, p, \{\,q\,\})$. Write $L^*(A_{q\,q'})$ for the *finite-word language* recognised by $A_{q\,q'}$. Then

$$\begin{aligned} & \alpha \in \Sigma^\omega \text{ is accepted by } A \\ \text{iff} \quad & \text{there exists a run } \rho \text{ with } \mathsf{inf}(\rho) \cap F \neq \varnothing \\ \text{iff} \quad & \text{there exists } q \in F \text{ and } \alpha = u_0 u_1 u_2 \cdots \text{ where } u_0 \text{ is accepted by } A_{q_0\,q} \\ & \text{and for each } i \geq 1,\ u_i \text{ is non-empty and accepted by } A_{q\,q}. \end{aligned}$$

Hence $L = \bigcup_{q \in F} L^*(A_{q_0\,q}) \cdot (L^*(A_{q\,q}))^\omega$. □

Regular Expressions: A Revision

Fix a finite alphabet Σ and let a range over Σ. Regular expressions e are defined by the grammar:

$$e \;::=\; \varnothing \mid \epsilon \mid a \mid e + e \mid e \cdot e \mid e^*$$

For simplicity $e \cdot f$ is often written as $e\,f$. We define the *denotation* of a regular expression $[\![\,e\,]\!] \subseteq \Sigma^*$ as follows.

$$\begin{array}{rclcrcl} [\![\,\epsilon\,]\!] & = & \{\,\epsilon\,\} & \qquad & [\![\,e \cdot f\,]\!] & = & [\![\,e\,]\!] \cdot [\![\,f\,]\!] \\ [\![\,\varnothing\,]\!] & = & \varnothing & & [\![\,e^*\,]\!] & = & [\![\,e\,]\!]^* \\ [\![\,a\,]\!] & = & \{\,a\,\} & & [\![\,e + f\,]\!] & = & [\![\,e\,]\!] \cup [\![\,f\,]\!] \end{array}$$

Let $w \in \Sigma^*$. We say that w *matches* e just if $w \in [\![\,e\,]\!]$.

Theorem 1.4 (Kleene)**.** *A set of finite words is recognisable by a finite-state automaton if, and only if, it is the denotation of a regular expression.* □

An *ω-regular expression* has the form

$$e_1 \cdot f_1^\omega + \cdots + e_n \cdot f_n^\omega$$

where $n \geq 0$, and $e_1, f_1, \cdots, e_n, f_n$ are regular expressions. The *denotation* of a ω-regular expression $[\![e]\!] \subseteq \Sigma^\omega$ is defined by the same clauses as regular expressions, and $[\![e^\omega]\!] := [\![e]\!]^\omega$. We say that an ω-language is *ω-regular* just if it is the denotation of a ω-regular expression.

Corollary 1.1. *A set of ω-words is Büchi recognisable if and only if it is ω-regular.*

Proof. Immediate consequence of Proposition 2. □

Example 1.7. (i) A regular expression for L_3 (i.e. the set of binary words containing only finitely many 1s) is $(0+1)^*0^\omega$.

(ii) A regular expression for $\overline{L_3}$ is $(0^*1)^\omega$.

For ω-regular expressions e and f, we say $e \equiv f$ just if $[\![e]\!] = [\![f]\!]$.

Lemma 1.2. *For $X, Y \subset \Sigma^*$*

(i) $(X+Y)^\omega \equiv (X^*Y)^\omega + (X+Y)^*X^\omega$

(ii) $(XY)^\omega \equiv X(YX)^\omega$

(iii) For all $n > 0$, $(X^n)^\omega \equiv (X^+)^\omega \equiv X^\omega$.

(iv) $X^\omega \equiv X^+X^\omega$.

Proof. Exercise □

Research Problem Interlude: The Star Height Problem

The *star height* of a regular expression is the maximum nesting depth of the stars that occur in the expression. Formally the *star height* of E, $h(E)$, is defined inductively as follows.

$$\begin{aligned} &h(\epsilon) := 0 \qquad h(\varnothing) := 0 \qquad h(a) := 0 \quad \text{for } a \in \Sigma \\ &h(e+f) := \max(h(e), h(f)) \qquad h(e \cdot f) := \max(h(e), h(f)) \\ &h(e^*) := h(e) + 1 \end{aligned}$$

The *star height*, $h(L)$, of a regular language L is the minimum star height among all regular expressions representing L. For example, the star height of $e = (b + a\,a^*\,b)^*a\,a^*$ is 2, but the star height of the language $[\![e]\!]$ is 1 because e is equivalent to $(a+b)^*a$.

Question Eggan (1963): *Is there an algorithm that computes the star height of a regular language?*

Two break-through results:

(i) Hashiguchi (1988) published a non-elementary algorithm. The method is impractical. For example, to determine the star height of a certain 4-state automaton A would require the following number of languages to be tested for equivalence against A

$$(10^{10^{10}})^{\left(10^{10^{10}}\right)^{\left(10^{10^{10}}\right)}}$$

(and equivalence of finite-state automata is PSPACE-complete). N.B. $10^{10^{10}}$ in decimal has more than 10 billion zeros.

(ii) Kirsten (2005) introduced a double-exponential space algorithm, and takes an NFA as input. His method is based on a new type of automata called *distance desert automata*.

1.4 Decision Problems and their Complexity

Decision problems for Büchi automata are worth studying because algorithms for solving these problems are basic building blocks for the construction of algorithmic solutions to the complex problems that arise in the verification of computing systems.

Non-Emptiness Problem Given a Büchi automaton A, is $L(A) \neq \varnothing$?

Proposition 3. *The non-emptiness problem for Büchi automata $A = (Q, \Sigma, \Delta, q_0, F)$ is decidable in time $O(|Q| + |\Delta|)$.*

Proof. We have:

$L(A) \neq \varnothing$
iff there is a path from q_0 to some $q \in F$, and there is a path from q back to itself
iff automaton A (*qua* digraph) has a *non-trivial* SCC which is reachable from q_0 and contains a final state q

Recall that a *strongly connected component* (SCC) of a directed graph is a maximal subgraph such that for every pair of vertices in the subgraph, there is a directed path from one vertex to the other.

There is a simple algorithm to decide the Non-Emptiness Problem. The idea is to find a "lasso" in the graph underlying the automaton: the base of the lasso is the initial state, and the loop must include a final state.

Algorithm: *Finding a Lasso*
Input: Büchi automaton $A = (Q, \Sigma, \Delta, q_0, F)$.
Output: YES, if $L(A) \neq \varnothing$; NO otherwise.

1. Determine the set Q_0 of states reachable from q_0 using (say) depth-first search.
2. Generate all non-trivial SCCs over Q_0; at the same time, check for containment of a final state.
3. If there is a non-trivial SCC that contains a final state, return YES; otherwise return NO.

Stages 1 and 2 require time $O(|Q| + |\Delta|)$ (Tarjan, 1972). □

In fact the non-emptiness problem is *complete for NL* (Non-deterministic Logspace).

An Interlude: NL and NL-Completeness

Recall that a non-deterministic Turing machine *accepts* an input word just if there is a computation path from the initial configuration to an accepting configuration.

Definition 1.2. A decision problem is *NL-complete* just if it is

(i) solvable in **NL** (i.e. decidable by a non-deterministic Turing machine using $O(\log n)$ space on a work tape, where n is the size of the input), and

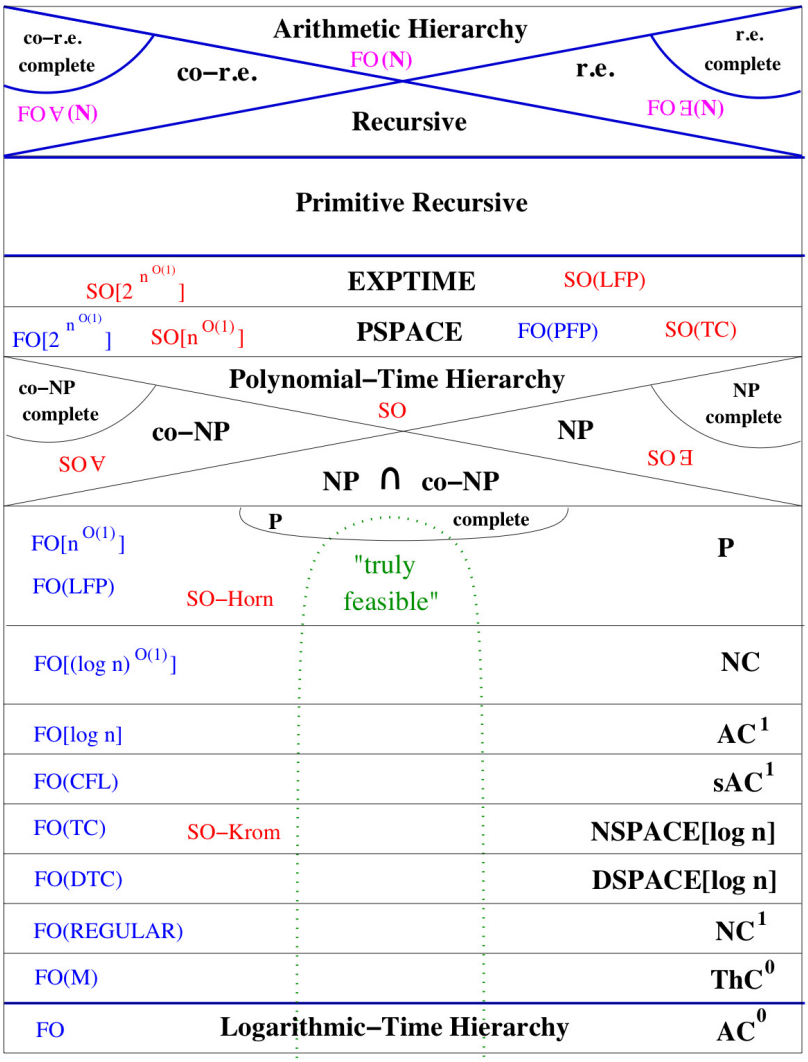

Figure 1.3: Immerman's World of Complexity Classes (Immerman, 1999)

(ii) NL-hard (i.e. every NL-solvable problem is logspace-reducible to it).

Intuitively **L** (Logspace) is the collection of problems that are solvable using (i) a constant number of pointers into the input (because each number in $\{0, \ldots, n-1\}$ can be represented in binary in at most $\log n$ bits) (ii) and a logarithmic number of boolean flags. See Michael Sipser's book (Sipser, 2005) and Immerman's book (Immerman, 1999) for a systematic treatment.

Proposition 4. *The Non-Emptiness Problem for Büchi automata is NL-complete.*

Proof. We give an algorithm in NL that checks if there is a final state which is reachable from the initial state q_0, and reachable from itself. To do this, we first guess a final state f (say), and then a path from q_0 to f, and from f to f.

To guess a path from states x to y:

1. Make x the current state.

2. Guess a transition from the current state, and make the target of the transition the new current state.
3. If the current state is y, STOP; otherwise repeat from step 2.

The algorithm is in **NL**: at each stage, only 3 states are remembered (by 3 pointers into the input string).

NL-hardness is proved by reduction from the Graph Reachability Problem (Given nodes x and y in a finite directed graph, is y reachable from x?), which is NL-complete. □

Universality Problem Given a Büchi automaton A over Σ, is $L(A) = \Sigma^\omega$?

Proposition 5. *The Universality Problem is PSPACE-complete.*

To decide non-universality, given a Büchi automaton A, we could construct the complement automaton $\overline{A}$ (i.e. $L(\overline{A}) = \Sigma \setminus L(A)$), then use the non-emptiness algorithm on $\overline{A}$. Unfortunately this would give an algorithm that is exponential in space and time!

To show PSPACE decidability, we determinise the automaton (which may be of exponential size) but calculate the states only *on demand*, and look for a word that is not recognised.

PSPACE Hardness Proof We present a proof by Sistla et al. (1987), which is by reduction from the Universality Problem for Finite-State Automata (Given a FSA A, is $L^*(A) = \Sigma^+$?). The latter problem is PSPACE-complete (Meyer and Stockmeyer, 1972).

The idea is to define a transformation $L \subseteq \Sigma^* \mapsto L' \subseteq \Sigma^\omega$ such that whenever A is a FSA then $L(A)'$ is Büchi-recognisable; further $L(A) = \Sigma^*$ if and only if $L(A)' = \Sigma^\omega$.

Proof. Fix $\Sigma = \{ a^1, \cdots, a^n \}$. Given FSA $A = \langle \Sigma, Q, \Delta, q_0, F \rangle$, define two alphabets $\Sigma_i = \{ a_i^1, \cdots, a_i^n \}$ $(i = 1, 2)$. Consider automata $A_i = \langle \Sigma_i, Q, \Delta_i, q_0, F \rangle$ such that for $i = 1, 2$:

$$\forall q, q', j : \ (q, a_i^j, q') \in \Delta_i \quad \leftrightarrow \quad (q, a^j, q') \in \Delta$$

Thus A_1 and A_2 recognise the image of $L^*(A)$ over Σ_1 and Σ_2 respectively. Now define $L' \subseteq (\Sigma_1 \cup \Sigma_2)^\omega$ by

$$\begin{aligned} L' \ &:= \ (L_1\ L_2)^\omega \ \cup \ (L_1\ L_2)^* L_1^\omega \ \cup \ (L_1\ L_2)^* L_2^\omega \\ &\cup \ (L_2\ L_1)^\omega \ \cup \ (L_2\ L_1)^* L_2^\omega \ \cup \ (L_2\ L_1)^* L_1^\omega \end{aligned}$$

where $L_i := L^*(A_i)$.

Exercise 1.5. *Construct a Büchi automaton A' that recognises L', with size linear in that of A.*

Claim. *The FSA A is universal (i.e. $L^*(A) = \Sigma^+$) if and only if the Büchi automaton A' is universal.*

$\Rightarrow$: Assume A is universal. Then L_i contains every non-empty word over Σ_i (for $i = 1, 2$). Now every ω-word over $\Sigma_1 \cup \Sigma_2$ is either

(i) entirely over Σ_1 or entirely over Σ_2, or

(ii) alternates between Σ_1 and Σ_2 and then entirely over one of the two

(iii) alternates between Σ_1 and Σ_2 infinitely.

These cases are covered by L', by definition.

$\Leftarrow$: Assume A' is universal. Take $w \in \Sigma^+$. Let w_i be the image of w in Σ_i. Now, by definition of L', $(w_1\ w_2)^\omega \in (L_1 L_2)^\omega$, because $(w_1\ w_2)^\omega$ cannot belong to the other five components of L'. Further, by construction of A', this implies that $w_i \in L(A_i)$ (for $i = 1, 2$). Hence $w \in L(A)$ as required. □

In the preceding proof, note that taking L' to be L^ω where $L = L^*(A)$ does not work, because L^ω is universal does not imply that L is universal: just take $L = \{ a, b \}$ over alphabet $\{ a, b \}$.

Problems

1.1 Let Σ be a finite alphabet. Prove that every $w \in \Sigma^\omega$ can be factorised as $w = u\,v$ where $u \in \Sigma^*$ and $v \in \Sigma^\omega$ and each letter in v occurs infinitely often in w.

1.2 Construct Büchi automata that recognise the following ω-languages over $\Sigma = \{\, a, b, c \,\}$:

(a) The set of words in which after each a, there is a b.

(b) The set of words in which a appears only at odd, or only at even positions.

1.3 Construct Büchi automata that recognise the following ω-languages over $\Sigma = \{\, a, b, c \,\}$:

(a) The set of ω-words in which *abc* appears as a segment at least once.

(b) The set of ω-words in which *abc* appears as a segment infinitely often.

(c) The set of ω-words in which *abc* appears as a segment only finitely often.

1.4 Prove that every nonempty Büchi-recognisable language contains an *ultimately periodic* word (i.e. an infinite word of the form $u\,v^\omega$ for finite words u and v).

1.5 Prove or disprove the following: for $U, V \subseteq \Sigma^+$

(a) $(U \cup V)^\omega = U^\omega \cup V^\omega$

(b) $\lim(U \cup V) = \lim U \cup \lim V$

(c) $U^\omega = \lim U^+$

(d) $\lim(U \cdot V) = U \cdot V^\omega$.

(e) $(U + V)^\omega \equiv (U^*V)^\omega + (U + V)^*U^\omega$

(f) $(UV)^\omega \equiv U(VU)^\omega$

(g) For all $n > 0$, $(U^n)^\omega \equiv (U^+)^\omega \equiv U^\omega$

(h) $U^\omega \equiv U^+U^\omega$.

1.6 Prove that the ω-language $L = \{\, u^\omega : u \in \{\, 0, 1 \,\}^+ \,\}$ is not recognised by any Büchi automaton.

[*Hint*. Consider the word $(01^n)^\omega$ where n is a number greater than the number of states of A.]

1.7 Prove the following (from first principles):

(a) If $U \subseteq \Sigma^*$ is regular then U^ω is Büchi-recognisable.

(b) If $U \subseteq \Sigma^*$ is regular and $L \subseteq \Sigma^\omega$ is Büchi-recognisable then $U \cdot L$ is Büchi-recognisable.

1.8 Prove Lemma 1.1.

1.9 Prove that an ω-language is deterministic Büchi-recognisable iff it is of the form $\lim U$ for some regular U.

1.10 (***Hard***) A quasi order (i.e. reflexive and transitive binary relation) $\lesssim$ over a set X is called a *well quasi ordering* (w.q.o.) if every infinite sequence $a_1, a_2, \cdots$ from X is saturated, meaning that there exist $i < j$ such that $a_i \lesssim a_j$.

Let Σ be a finite alphabet. The *subword ordering* $\lesssim \subseteq \Sigma^* \times \Sigma^*$ is defined as: $u_1 \cdots u_m \lesssim v_1 \cdots v_n$ just if there exist $1 \leq i_1 < i_2 < \cdots < i_m \leq n$ such that for each $1 \leq j \leq m$, $u_j = v_{i_j}$. Prove that $(\Sigma^*, \lesssim)$ is a w.q.o.

[*Hint.* Suppose, for a contradiction, there is an infinite sequence of words $w_1, w_2, \cdots$ that is unsaturated. For an appropriate notion of "minimal", choose a minimal such sequence. Then consider the derived sequence $v_1, v_2, \cdots$ whereby $w_i = a_i v_i$ and $a_i \in \Sigma$, for each i.]

1.11 Consider the ω-language

$$L \;:=\; \{\, \alpha \in \{\, 0, 1 \,\}^\omega \mid \alpha \text{ contains } 00 \text{ infinitely often, but } 11 \text{ only finitely often} \,\}.$$

(a) Construct a Büchi automaton that recognises L. Explain why it works.

(b) Show that L is not recognisable by a deterministic Büchi automaton.

(c) We say that a ω-automaton *co-Büchi recognises* an ω-word α if there is a run ρ of the automaton on α such that from some point onwards, only final states will be visited i.e. there is an $n \geq 0$ such that for every $i > n$, $\rho(i)$ is a final state.

Is L recognisable by a deterministic co-Büchi automaton? Justify your answer.

1.12

(a) Let L be an ω-language over the alphabet Σ. Define *right-congruence* $\sim_L \subseteq \Sigma^* \times \Sigma^*$ by

$$u \sim_L v \;:=\; \forall \alpha \in \Sigma^\omega . u\,\alpha \in L \leftrightarrow v\,\alpha \in L.$$

Prove that every deterministic Muller automaton that recognises L needs at least as many states as there are $\sim_L$-equivalence classes.

Show that there is a ω-language L, which is not ω-regular, such that $\sim_L$ has only finitely many equivalence classes.

Hence, or otherwise, state (without proof) a result about regular $*$-languages (i.e. sets of finite words) that does not generalise to ω-regular ω-languages.

(b) Is it true that an ω-language is ω-regular if and only if it is expressible as a Boolean combination of languages of the form $\lim U$ where U is a regular $*$-language? Justify your answer.

Chapter 2

Linear-time Temporal Logic

Synopsis[1]

Kripke structures. Examples of correctness properties of reactive systems. LTL: syntax and semantics. Transformation of LTL formulas to generalised Büchi automata. LTL model checking is PSPACE-complete: Savitch's algorithm; encoding polynomial-space Turing machines in LTL. Expressivity of LTL: Kamp's theorem.

2.1 Motivating Example: Mutual Exclusion Protocol

The model checking problem: Given a system *Sys* and a specification *Spec* on the runs of the system, does *Sys* satisfy *Spec*?

Example 2.1 (Mutual exclusion protocol)**.** A MUX protocol is modelled by a transition system over state-space $\mathbb{B}^5$:

```
Process 0:  Repeat
00:  <non-critical region 0>
01:  wait unless turn = 0
10:  <critical region 0>
11:  turn := 1

Process 1:  Repeat
00:  <non-critical region 1>
01:  wait unless turn = 1
10:  <critical region 1>
11:  turn := 0
```

A *state* is a bit-vector "$a_1\, a_2\, b_1\, b_2\, t$" where $a_1\, a_2$ are $b_1\, b_2$ are line no. of processes 0 and 1 respectively, and t is the value of shared variable `turn`; the initial state is 00000. Some examples of correctness properties *Spec*:

(i) *Safety*: The state $1010t$ is never reached.

[1] The contributions of past guest lecturers, Matthew Hague and Anthony Lin, are gratefully acknowledged.

(ii) *Liveness*: It is always the case that whenever $01b_1b_2t$ is reached, $10b'_1b'_2t'$ is eventually reached (similarly for a_1a_201t and $a'_1a'_210t'$).

Temporal Logic in Computer Science

Amir Pnueli (1941–2009) won the *ACM Turing Award 1996*

> "For seminal work introducing temporal logic into computing science and for outstanding contributions to program and system verification."

A landmark publication is (Pnueli, 1977).

2.2 Kripke Structures

Fix a set $\{ p_1, \cdots, p_n \}$ of atomic propositions. We use Kripke structures $\mathcal{K}$ to model reactive systems.

Definition 2.1. A *Kripke structure* over a fixed set of atomic propositions $\{ p_1, \cdots, p_n \}$ is a quadruple (S, R, λ, s_0) with

- a finite state-set S, and $s_0 \in S$ is the initial state
- a transition relation $R \subseteq S \times S$, and
- a labelling function $\lambda : S \to \mathcal{P}(\{ p_1, \cdots, p_n \})$, associating with each $s \in S$ the set of those p_i that are satisfied at s.

A Kripke structure is just a directed graph whose nodes are labelled by elements of the power set, $\mathcal{P}(\{ p_1, \cdots, p_n \})$, as given by λ.

Notation We often write $\lambda(s)$ as a bit vector $\begin{pmatrix} b_1 \\ \vdots \\ b_n \end{pmatrix} \in \mathbb{B}^n$ such that $b_i = 1$ iff $p_i \in \lambda(s)$.

A *path* through a Kripke structure (S, R, λ, s_0) is an infinite sequence of states, $s_0 s_1 s_2 \cdots$, where for each $i \geq 0$, $(s_i, s_{i+1}) \in R$. The corresponding *label sequence* is the ω-word over the alphabet $\mathbb{B}^n$: $\lambda(s_0)\,\lambda(s_1)\,\lambda(s_2) \cdots$.

Example 2.2. Fix atomic propositions p_1 and p_2.

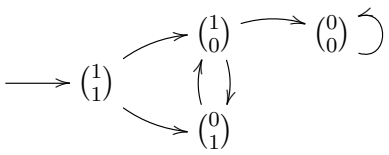

Example label sequences:

(i) $\binom{1}{1}\binom{1}{0}\binom{0}{1}\binom{1}{0}\binom{0}{0}\binom{0}{0}\cdots$

(ii) $\binom{1}{1}\binom{0}{1}\binom{1}{0}\binom{0}{1}\binom{1}{0}\binom{0}{0}\binom{0}{0}\cdots$

What are the differences between Büchi automata and Kripke structures?

Correctness Properties of Reactive Systems: Examples

When a reactive system is modelled as a Kripke structure, runs of the system correspond to label sequences of $\mathcal{K}$, which are ω-words over $(\mathbb{B}^n)^\omega$. Correctness properties of the reactive system are thus naturally expressed as properties of ω-words. In other words, they are path properties.

The model checking problem asks: given a correctness property φ expressed as a property of ω-words, does every label sequence of $\mathcal{K}$ satisfy φ?

Example 2.3 (MUX protocol revisited)**.** For $i = 0, 1$, let

- p_{i+1} stand for "Process i is waiting (to enter the critical region)"
- p_{i+3} stand for "Process i is in critical region"

Consider the following φ:

(i) "It is always the case that when p_1 holds then sometime later p_3 holds" which means: for any label sequence, when letter $(1, b_2, b_3, b_4)$ occurs, subsequently a letter $(b'_1, b'_2, 1, b'_4)$ occurs.

(ii) "p_3 and p_4 never hold simultaneously" which means: no label sequence contains the letter $(b_1, b_2, 1, 1)$.

Example 2.4 (Sequence properties)**.** Fix state properties p_1 and p_2. Label sequences are ω-words over $\mathbb{B}^2 = \{ \binom{0}{0}, \binom{0}{1}, \binom{1}{0}, \binom{1}{1} \}$.

(i) *Recurrence*: "p_1 holds again and again (i.e. infinitely often)."

(ii) *Periodicity*: "p_1 is true initially and precisely at every third moment."

(iii) *Request-response*: "It is *always* the case that whenever p_1 holds, p_2 will hold sometime later."

(iv) *Obligation*: "p_1 eventually holds but p_2 never does."

(v) *Until condition*: "It is always the case that when p_1 holds, sometime later p_1 will be true again, and in the meantime p_2 is always true."

(vi) *Fairness*: "If p_1 is true again and again (infinitely often), then the same is true of p_2".

2.3 Syntax and Semantics

In our present modelling framework, correctness properties are path properties. We present Linear-time Temporal Logic, a logical system for expressing properties of ω-words.

LTL-formulas, over atomic propositions $p_1, \cdots, p_n$, are defined by the grammar:

$$\begin{array}{rcll} \varphi & ::= & p_i & \text{atomic proposition} \\ & | & \neg\varphi & \text{negation} \\ & | & \varphi \wedge \psi & \text{conjunction} \\ & | & \varphi \vee \psi & \text{disjunction} \\ & | & \boldsymbol{X}\varphi & \textit{next} \\ & | & \varphi\ \boldsymbol{U}\ \psi & \textit{until} \end{array}$$

Intuitively

$\boldsymbol{X}\varphi$ "φ is true at the *next* time-step"
$\varphi\ \boldsymbol{U}\ \psi$ "φ is true *until* ψ is true (and ψ holds eventually)"

(Picture of time-line)

Two additional constructs

$\boldsymbol{F}\varphi$ "φ is *eventually* true"
i.e. φ is true at *some* point in the future (starting from the present)
$\boldsymbol{G}\varphi$ "φ is *always* true"
i.e. φ is true at *every* point in the future (including the present)

They are expressible in LTL by

$$\begin{array}{rcl} \boldsymbol{F}\varphi & := & \mathsf{true}\ \boldsymbol{U}\ \varphi \\ \boldsymbol{G}\varphi & := & \neg(\boldsymbol{F}\neg\varphi) \end{array}$$

(Henceforth we regard the above as definitions.)

LTL-formulas over atomic propositions $p_1, \cdots, p_n$ are interpreted as sets of ω-words α over the alphabet $\mathbb{B}^n$.

Notation Let $\alpha = \alpha(0)\ \alpha(1)\ \alpha(2) \cdots \in (\mathbb{B}^n)^\omega$:

- α^i stands for $\alpha(i)\ \alpha(i+1)\ \alpha(i+2) \cdots$, so $\alpha = \alpha^0$.
- $(\alpha(i))_j$ is the j-th component of the vector $\alpha(i)$.

Definition 2.2 (Satisfaction)**.** Let $i \geq 0$. Define $\alpha^i \vDash \varphi$ by recursion over the syntax of φ:

$$\begin{array}{rcl} \alpha^i \vDash p_j & := & (\alpha(i))_j = 1 \\ \alpha^i \vDash \neg\varphi & := & \neg(\alpha^i \vDash \varphi) \\ \alpha^i \vDash \varphi \vee \psi & := & \alpha^i \vDash \varphi\ \vee\ \alpha^i \vDash \psi \\ \alpha^i \vDash \varphi \wedge \psi & := & \alpha^i \vDash \varphi\ \wedge\ \alpha^i \vDash \psi \\ \alpha^i \vDash \boldsymbol{X}\varphi & := & \alpha^{i+1} \vDash \varphi \\ \alpha^i \vDash \varphi\ \boldsymbol{U}\ \psi & := & \exists j \geq i : \big(\alpha^j \vDash \psi\ \wedge\ \forall i \leq k \leq j-1 : \alpha^k \vDash \varphi\big) \end{array}$$

We say that $\alpha \vDash \varphi$, read α *satisfies* φ, just if $\alpha^0 \vDash \varphi$.

Examples of LTL-definable Correctness Properties

Example 2.5 (Sequence properties revisited)**.** (i) *Recurrence*: p_1 holds again and again (i.e. infinitely often).

$$\boldsymbol{G}(\boldsymbol{F}p_1)$$

(ii) *Periodicity*: p_1 is true initially and precisely at every third moment.

$$p_1 \ \wedge \ \boldsymbol{X}\neg p_1 \ \wedge \ \boldsymbol{X}\boldsymbol{X}\neg p_1 \ \wedge \ \boldsymbol{G}(p_1 \leftrightarrow \boldsymbol{X}\boldsymbol{X}\boldsymbol{X}p_1)$$

(iii) *Request-response*: It is *always* the case that whenever p_1 holds, p_2 will hold sometime later.

$$\boldsymbol{G}(p_1 \rightarrow \boldsymbol{X}\boldsymbol{F}p_2)$$

(iv) *Obligation*: Eventually p_1 holds but p_2 never does.

$$\boldsymbol{F}p_1 \wedge \neg\boldsymbol{F}p_2$$

(v) *Until condition*: It is always the case that when p_1 holds, sometime later p_1 will hold again, and in the meantime p_2 is always true.

$$\boldsymbol{G}(p_1 \rightarrow \boldsymbol{X}(p_2 \ \boldsymbol{U}p_1))$$

(vi) *Fairness*: If p_1 is true again and again, then the same is true of p_2.

$$\boldsymbol{G}\boldsymbol{F}p_1 \rightarrow \boldsymbol{G}\boldsymbol{F}p_2$$

Exercise 2.1. Verify the following:

(i) $\alpha^i \vDash \boldsymbol{F}\varphi \ \leftrightarrow \ \exists j \geq i : \alpha^j \vDash \varphi$

(ii) $\alpha^i \vDash \boldsymbol{G}\varphi \ \leftrightarrow \ \forall j \geq i : \alpha^j \vDash \varphi$

Definition 2.3. (i) An ω-language $L \subseteq (\mathbb{B}^n)^\omega$ is *LTL-definable* just if there is an LTL-formula φ over $p_1, \cdots, p_n$ such that $L = \{\, \alpha \in (\mathbb{B}^n)^\omega \mid \alpha \vDash \varphi \,\}$. We say that L is definable by φ.

(ii) We say that two LTL-formulas φ and ψ are *equivalent*, written $\varphi \equiv \psi$, if they define the same ω-language.

(iii) A Kripke structure $\mathcal{K} = (S, R, \lambda, s_0)$ *satisfies* an LTL-formula φ, written $\mathcal{K} \vDash \varphi$, just if every label sequence of $\mathcal{K}$ satisfies φ.

Translating LTL formulas into Büchi automata We consider the translation of LTL formulas into equivalent Büchi automata by examples.

Example 2.6.

$\alpha \vDash \boldsymbol{F}(p_1 \wedge \boldsymbol{X}(\neg p_2 \ \boldsymbol{U}p_1))$

iff for some $j \geq 0$: $\alpha^j \vDash p_1$ and $\alpha^{j+1} \vDash \neg p_2 \ \boldsymbol{U}p_1$

iff for some $j \geq 0$: $\alpha^j \vDash p_1$ and for some $j' \geq j+1$: $\alpha^{j'} \vDash p_1$ and for all $j+1 \leq k \leq j'-1$: $\alpha^k \vDash \neg p_2$

iff for some j and some $j' > j$: $\alpha(j)$ and $\alpha(j')$ have 1 in the 1st component, and for all $j < k < j'$, $\alpha(k)$ has 0 in 2nd component

iff α has two occurrences of $\binom{1}{*}$ between which only letters of the form $\binom{*}{0}$ occur.

Exercise Draw a Büchi automaton that recognises the same ω-language.

Example 2.7. (i) $\boldsymbol{G}(\boldsymbol{F}p_1)$

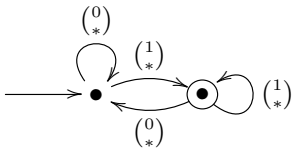

(ii) $p_1 \wedge \boldsymbol{X}\neg p_1 \wedge \boldsymbol{X}\boldsymbol{X}\neg p_1 \wedge \boldsymbol{G}(p_1 \leftrightarrow \boldsymbol{X}\boldsymbol{X}\boldsymbol{X}p_1)$

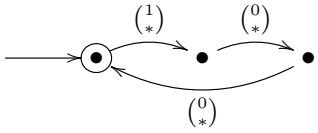

(iii) $\boldsymbol{G}(p_1 \rightarrow \boldsymbol{X}\boldsymbol{F}p_2)$

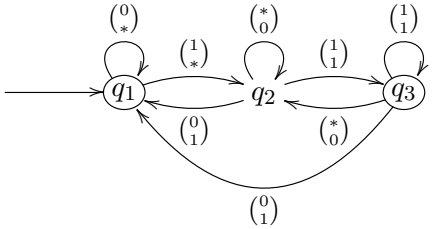

At state q_1, the automaton has no obligation to read a $\binom{*}{1}$; q_2 means that it is obliged to, but has not yet, read a $\binom{*}{1}$ since the last $\binom{1}{*}$; q_3 is reached after reading $\binom{1}{1}$.

Exercise 2.2. Translate the following to Büchi automata:

(i) $(\boldsymbol{F}p_1) \wedge \neg\boldsymbol{F}p_2$

(ii) $\boldsymbol{G}(p_1 \rightarrow \boldsymbol{X}(p_2 \,\boldsymbol{U}\, p_1))$

2.4 Translating LTL to Generalised Büchi Automata

We can systematically translate a given LTL formula to an equivalent (generalised) Büchi automaton. In fact, we shall utilise such an automaton construction to design a decision procedure for the LTL model checking problem.

Definition 2.4. A *generalised Büchi automaton* (GBA) is a 5-tuple

$$(Q, \Sigma, \Delta, q_0, \{\, F_0, \cdots, F_{l-1} \,\})$$

with final state-sets $F_0, \cdots, F_{l-1} \subseteq Q$. A run ρ is *accepting* just if for each i, there is some state in F_i which occurs infinitely often in ρ i.e. $\bigwedge_i(\mathsf{inf}(\rho) \cap F_i \neq \varnothing)$.

Proposition 6. *Given a generalised Büchi automaton* $A = (Q, \Sigma, \Delta, q_0, \{\, F_0, \cdots, F_{l-1} \,\})$, *define Büchi automaton*

$$A' = (Q \times \{\, 0, 1, \cdots, l-1 \,\}, \Sigma, \Delta', (q_0, 0), F_0 \times \{\, 0 \,\})$$

with Δ' *consisting of*

- $((p,i),a,(q,i))$ *if* $p \notin F_i$
- $((p,i),a,(q,(i+1) \bmod l))$ *if* $p \in F_i$

assuming that $(p,a,q) \in \Delta$. *Prove that* A *and* A' *recognise the same* ω*-language over* Σ.

Exercise 2.3. *Prove the proposition.*

The rest of the section is concerned with the proof of the following theorem.

Theorem 2.1 (Translating LTL to GBA)**.** *Let* φ *be an LTL formula over* $p_1, \cdots, p_n$. *Suppose* m *is the number of distinct non-atomic subformulas of* φ. *There is a generalised Büchi automaton* A_φ *with state-set* $\{ q_0 \} \cup \mathbb{B}^{n+m}$ *that is equivalent to* φ *i.e. the language definable by* φ *coincides with the language recognised by* A_φ. *Further the translation* $\varphi \mapsto A_\varphi$ *is effective.*

Evaluating LTL-formula φ **over** $\alpha \in (\mathbb{B}^n)^\omega$ Given ω-word α over $(\mathbb{B}^n)^\omega$, and LTL-formula φ over $p_1, \cdots, p_n$. Define formulas $\varphi_1, \varphi_2, \cdots, \varphi_{n+m}$ where

- $\varphi_1 = p_1, \cdots, \varphi_n = p_n$, and
- $\varphi_{n+1}, \cdots, \varphi_{n+m} = \varphi$ are all the distinct *non-atomic* subformulas of φ, listed in non-decreasing order of size.

We construct a two-dimensional semi-infinite array of truth values, $\beta \in (\mathbb{B}^{n+m})^\omega$, defined by: $(\beta(i))_j = 1$ (i.e. j-th row, i-th column is 1) if and only if $\alpha^i \vDash \varphi_j$. In particular $\alpha \vDash \varphi \leftrightarrow (\beta(0))_{m+n} = 1$. We call $\beta \in (\mathbb{B}^{n+m})^\omega$ the φ*-expansion of* α.

Example 2.8. Take $\varphi = \boldsymbol{F}(\neg p_1 \wedge \boldsymbol{X}(\neg p_2 \, \boldsymbol{U} \, p_1))$ and $\alpha = \binom{1}{0} \binom{0}{1} \binom{1}{1} \binom{0}{0} \binom{1}{0} \binom{0}{1} \cdots$. We construct $\beta \in (\mathbb{B}^{2+6})^\omega$, the φ-expansion of α:

$\varphi_1 = p_1$	1	0	1	0	1	0	$\cdots$
$\varphi_2 = p_2$	0	1	1	0	0	1	
$\varphi_3 = \neg p_1$	0	1	0	1	0	1	$\cdots$
$\varphi_4 = \neg p_2$	1	0	0	1	1	0	$\cdots$
$\varphi_5 = \neg p_2 \, \boldsymbol{U} \, p_1$	1	0	1	1	1	0	$\cdots$
$\varphi_6 = \boldsymbol{X}(\neg p_2 \, \boldsymbol{U} \, p_1)$	0	1	1	1	0	$\cdot$	$\cdots$
$\varphi_7 = \neg p_1 \wedge \boldsymbol{X}(\neg p_2 \, \boldsymbol{U} \, p_1)$	0	1	0	1	0	$\cdot$	$\cdots$
$\varphi_8 = \underbrace{\boldsymbol{F}(\neg p_1 \wedge \boldsymbol{X}(\neg p_2 \, \boldsymbol{U} \, p_1))}_{\varphi}$	1	1	1	1	$\cdot$	$\cdot$	$\cdots$

Note that the 3rd (resp. 4th) row is the negation of the 1st (resp. 2nd) row.

Finite characterisation of the φ**-expansion of an** ω**-word** α The semantics of an LTL formula φ over an ω-word α is captured by the φ*-expansion of* α, which is an infinite object. We first characterise φ-expansions by a finite set of compatibility conditions, and then construct a generalised Büchi automaton that *guesses* the φ-expansion of α, as α is read. We divide these rules into local and global as follows.

Local compatibility conditions: *Local* in the sense that the conditions relate contiguous letters (*qua* column vectors) of $\beta \in (\mathbb{B}^{m+n})^\omega$.

Cases	*Local conditions*
$\varphi_j = \neg(\varphi_k)$	$(\beta(i))_j = 1 \leftrightarrow (\beta(i))_k = 0$
$\varphi_j = \varphi_k \wedge \varphi_l$	$(\beta(i))_j = 1 \leftrightarrow [(\beta(i))_k = 1$ and $(\beta(i))_l = 1]$
$\varphi_j = \varphi_k \vee \varphi_l$	$(\beta(i))_j = 1 \leftrightarrow [(\beta(i))_k = 1$ or $(\beta(i))_l = 1]$
$\varphi_j = \boldsymbol{X}\varphi_k$	$(\beta(i))_j = 1 \leftrightarrow (\beta(i+1))_k = 1$
$\varphi_j = \varphi_k \ \boldsymbol{U}\varphi_l$	$(\beta(i))_j = 1 \leftrightarrow (\beta(i))_l = 1$ or $[(\beta(i))_k = 1$ and $(\beta(i+1))_j = 1]$

The last clause can be explained by the equivalence: $\varphi \ \boldsymbol{U}\psi \equiv \psi \vee (\varphi \wedge \boldsymbol{X}(\varphi \ \boldsymbol{U}\psi))$.

Global compatibility condition $\varphi_j = \varphi_k \ \boldsymbol{U}\varphi_l$: *There is no m such that for all $n \geq m$, we have $(\beta(n))_j = 1$ and $(\beta(n))_l = 0$.*

If $\alpha \in (\mathbb{B}^n)^\omega$ and $\gamma \in (\mathbb{B}^m)^\omega$, let $\alpha /\!\!/ \gamma$ denote the ω-word over $(\mathbb{B}^{n+m})^\omega$ obtained by stacking α *on top of* γ.

Lemma 2.1. *$\beta := \alpha /\!\!/ \gamma \in (\mathbb{B}^{n+m})^\omega$ satisfies the compatibility conditions if and only if it is the φ-expansion of α.*

Proof. "$\Leftarrow$" direction is obvious. "$\Rightarrow$": Let B_j be the statement $\forall i \geq 0 : (\beta(i))_j = 1 \leftrightarrow \alpha^i \models \varphi_j$. We prove $\forall j \geq 1 : B_j$ by induction on j.

Base case: $\varphi_j = p_j$. Vacuously true.

Inductive case: The only less obvious case is when $\varphi_j = \varphi_k \ \boldsymbol{U}\varphi_l$. If $(\beta(i_0))_l = 1$ then for all $i \leq i_0$, the entry $(\beta(i))_j$ is "correct", because of $\varphi_k \ \boldsymbol{U}\varphi_l = \varphi_l \vee (\varphi_k \wedge \boldsymbol{X}(\varphi_k \ \boldsymbol{U}\varphi_l))$ and the induction hypothesis as $k, l < j$. It follows that if $(\beta(i))_l = 1$ for infinitely many i, then the entry $(\beta(i))_j$ is "correct" for all $i \geq 0$. Now suppose for some i_0, we have $(\beta(i))_l = 0$ for all $i \geq i_0$. We claim that for all $i \geq i_0$, $(\beta(i))_j = 0$, and hence the entry is correct; for otherwise we have $(\beta(i))_j = 1$ for *all* $i \geq i_1$, for some $i_1 \geq i_0$ (because of local compatibility for until formulas), and so, violating global compatibility. □

Proof of Theorem 2.1

Lemma 2.2. *The generalised Büchi automaton*

$$A_\varphi \ = \ (\{\, q_0 \,\} \cup \mathbb{B}^{n+m},\ \mathbb{B}^n,\ \Delta,\ q_0,\ \{\, F_1, \cdots, F_p \,\}),$$

defined as follows, accepts $\alpha \in (\mathbb{B}^n)^\omega$ if and only if $\alpha \vDash \varphi$.

Proof. Write non-initial state $\overline{xy} = (x_1, \cdots, x_n, y_1, \cdots, y_m)$. The transition relation Δ is defined as follows: for $\overline{xy}$ and $\overline{x'y'}$ ranging over $\mathbb{B}^{n+m}$

- $q_0 \xrightarrow{\overline{x}} \overline{xy}$ provided $\overline{xy}$ satisfies the local compatibility conditions and $y_m = 1$
- $\overline{xy} \xrightarrow{\overline{x'}} \overline{x'y'}$ provided $\overline{xy}$ and $\overline{x'y'}$ satisfy the local compatibility conditions (i.e. $\overline{xy}$ corresponds to the i-column and $\overline{x'y'}$ to the $(i+1)$-column in the table of local compatibility conditions).

For each until subformula $\varphi_j = \varphi_k \ \boldsymbol{U}\varphi_l$, a final state-set F containing all states with j-component $= 0$ or l-component $= 1$. Let $F_1, \cdots, F_p$ be all such sets, one for each until subformula of φ. Thus we have

A_φ accepts $\alpha \in (\mathbb{B}^n)^\omega$
iff { Definition of acceptance }
for some A_φ-run $\rho \in (\mathbb{B}^{n+m})^\omega$ on α, each F_j is visited infinitely often
iff { Lemma 2.1 }
ρ is the φ-expansion of α, and $(\rho(0))_{m+n} = 1$
iff { Definition of φ-expansion of α }
$\alpha = \alpha^0 \vDash \varphi_{m+n} = \varphi$

as desired. □

2.5 The LTL Model Checking Problem and its Complexity

Definition 2.5. A Kripke structure $\mathcal{K} = (S, R, \lambda, s_0)$ over $AP = \{p_1, \ldots, p_n\}$ *satisfies* an LTL-formula φ over AP, written $\mathcal{K} \vDash \varphi$, if every label sequence of $\mathcal{K}$ satisfies φ.

LTL Model-Checking Problem Given a Kripke structure $\mathcal{K} = (S, R, \lambda, s_0)$ over the atomic propositions $p_1, \cdots, p_n$, and an LTL-formula φ, does $\mathcal{K}$ *satisfy* φ?

The approach is to verify the negation: Is there a label sequence through K that does *not* satisfy φ? Note that $P \subseteq Q$ iff $P \cap \overline{Q} = \varnothing$.

Label sequences of a given Kripke structure are Büchi recognisable Given a Kripke structure $\mathcal{K} = (S, R, \lambda, s_0)$ over $p_1, \cdots, p_n$. Construct a Büchi automaton $A_\mathcal{K} = (S, \mathbb{B}^n, s_0, \Delta, S)$ whereby

$$(s, (b_1, \cdots, b_n), s') \in \Delta \quad \leftrightarrow \quad (s, s') \in R \text{ and } \lambda(s) = (b_1, \cdots, b_n)$$

Thus each transition of $A_\mathcal{K}$ has the label of the source state, and every state is final. Then $A_\mathcal{K}$ recognises the language of label sequences of $\mathcal{K}$.

Example 2.9. The Büchi automaton on the right recognises the set of label sequences of the Kripke structure on the left.

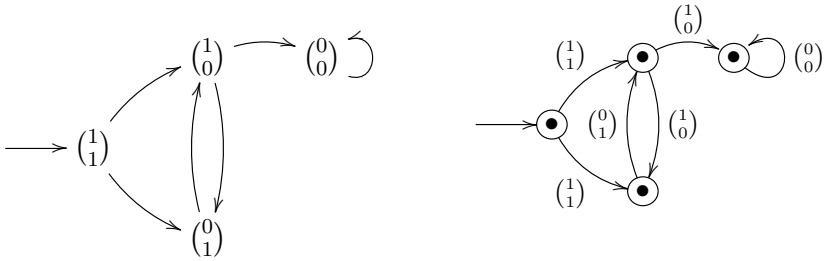

Proposition 7. *The LTL Model Checking Problem is solvable in time polynomial in the size of the Kripke structure $\mathcal{K}$ and exponential in the size of the formula φ.*

Proof. We give the model checking algorithm as follows.

LTL Model Checking: Given a Kripke structure $\mathcal{K}$ and an LTL formula φ, does $\mathcal{K} \vDash \varphi$?

Algorithm:

1. Construct a Büchi automaton $A_{\mathcal{K}}$ that recognises the ω-language of all label sequences through $\mathcal{K}$.
2. Construct a generalised Büchi automaton $A_{\neg\varphi}$ that recognises the ω-language of all label sequences that do not satisfy φ.
3. Construct the *intersection automaton* $A_{\mathcal{K}} \times A_{\neg\varphi}$ i.e. the Büchi automaton that recognises $L(A_{\mathcal{K}}) \cap L(A_{\neg\varphi})$.
4. Check for *non-emptiness* of $A_{\mathcal{K}} \times A_{\neg\varphi}$.

Stages 1, 3 and 4 are all polytime. Stage 2 require exponential time in the size of φ.

□

Theorem 2.2 (Sistla and Clarke 1985)**.** *The LTL Model Checking Problem is PSPACE-complete in the size of the formula.*

We present a proof of a result due to Sistla and Clarke (1985). To prove that LTL model checking is solvable in PSPACE, we improve the EXPTIME algorithm of Theorem 2.1. For PSPACE-hardness, we encode polynomial space Turing machines.

An Interlude: Savitch's Algorithm We review the famous result of Savitch (1970); see (Sipser, 2005, Ch. 8 Space Complexity) or (Papadimitriou, 1994).

In time complexity, non-determinism is exponentially more expensive than determinism. But in space complexity, thanks to Savitch, non-determinism is only quadratically more expensive than determinism. Savitch proved that if a nondeterministic Turing machine can solve a problem using $f(n)$ space, then a deterministic Turing machine can solve the same problem in the square of that space bound.

Savitch's insight lies in a method to decide graph reachability which, though wasteful in time, is highly efficient in space. The well-known depth-first and breadth-first graph search algorithms are linear in the size of the graph. Savitch's algorithm could be viewed as "middle-first search" based on the fact that every path of length 2^i has a mid-way point which is reachable from the start, and from which the end is reachable, in no more than 2^{i-1} steps.

Savitch's Algorithm
Input: A finite digraph $G = (V, E)$, vertices $u, v \in V$, $i \in \mathbb{N}$
Output: YES iff there is a path in G from u to v of length at most 2^i

```
Path(G, u, v, i) =
  if i = 0
    if u = v or (u, v) ∈ E
      return YES
      else return NO
  for all vertices w ∈ V
    if Path(G, u, w, i − 1) and Path(G, w, v, i − 1)
      return YES
  return NO
```

Theorem 2.3 (Savitch 1970)**.** *Reachability (given a graph $G = (V, E)$ and vertices $u, v \in V$, is there a path from u to v?) can be solved by calling $Path(G, u, v, \log |V|)$, which is computable in space $O(\log^2 |V|)$.*

To obtain the $O(\log^2 |V|)$ space bound, we use a Turing machine to implement the recursive program $Path(G, u, v, \log |V|)$, with its work tape acting like the stack of activation records. At any time, the work tape contains $\log |V|$ or fewer triples of the form (x, y, j) where $x, y \in V$ and $j \leq \log |V|$, where each triple has length at most $3 \log |V|$. For a proof, see for example (Papadimitriou, 1994, p. 149-150).

Corollary 2.1 (Savitch 1970)**.** *For every function $f(n) \geq \log(n)$*

$$NSPACE(f(n)) \subseteq DSPACE(f(n)^2).$$

It follows that **PSPACE** = **NPSPACE**.

Proof. Let P be a problem in $NSPACE(f(n))$. Let M be a nondeterministic Turing machine with space usage bounded by $f(n)$ and accepting P. To determine whether $x \in P$, check whether the configuration graph of M has a path of length at most $2^{O(f(|x|))}$ from the initial to an accepting configuration. This can be done in $DSPACE(O(f(|x|)^2))$. □

LTL Model Checking is in PSPACE

The idea is to use the algorithm of Proposition 7 without building the intersection automaton $A_{\mathcal{K}} \times A_{\neg\varphi}$ in full; rather we compute the states of the automaton *on demand*. From Savitch's algorithm, we know that if the space required to store a state and decide a given transition $q \to q'$ is polynomial in the size of the transition system, so is the space required to decide reachability $q \to^* q'$.

States of the intersection automaton $A_{\mathcal{K}} \times A_{\neg\varphi}$, which are elements of the shape

$$(s, \overline{x}\,\overline{y}, i) \;\in\; Q \times \mathbb{B}^{n+m} \times \{1, \cdots, l\},$$

can be stored in space polynomial in the size of the formula $|\varphi|$. Note that $m, l = O(|\varphi|)$. To decide $(s, \overline{x}\,\overline{y}, i) \to (s', \overline{x}'\,\overline{y}', j)$, we need to verify:

- $\overline{x}\ \overline{y}$ and $\overline{x}'\ \overline{y}'$ satisfy the local compatibility conditions, such as $[(\beta(i))_k = 1$ or $(\beta(i))_l = 1]$; since there are only linearly many conditions, they are easy to check.
- i, j are as determined by the global compatibility condition
- Finally $s \to s'$ is a transition in $A_{\mathcal{K}}$.

Thus transitions can be decided in space polynomial in $|\varphi|$. It follows from Savitch that we can decide $(s, \overline{x}\ \overline{y}, i) \to^* (s', \overline{x}'\ \overline{y}', j)$ in polynomial space.

To decide non-emptiness of $L(A_{\mathcal{K}} \times A_{\neg\varphi})$, we seek a "lasso" on a final state $(s, \overline{x}\ \overline{y}, i)$ in the intersection automaton i.e.

$$(s_0, q_0, 0) \to^* (s, \overline{x}\ \overline{y}, i) \to^+ (s, \overline{x}\ \overline{y}, i)$$

by the following algorithm:

```
for all (s, x̄ ȳ, i)
  if (s, x̄ ȳ, i) is final and (s₀, q₀, 0) →* (s, x̄ ȳ, i)
    for all (s, x̄ ȳ, i) → (s', x̄' ȳ', j)
      if (s', x̄' ȳ', j) →* (s, x̄ ȳ, i)
        return YES
return NO
```

Thus we conclude that LTL model checking is in PSPACE.

LTL Model Checking is PSPACE-hard

Let T be a Turing machine with space usage bounded by a polynomial function $s(n)$. WLOG, assume that T loops at each accepting configuration. We shall build a Kripke structure $\mathcal{K}$ and an LTL formula φ such that $\mathcal{K} \not\models \varphi$ iff T can reach an accepting state.

(1) Runs of T can be represented as ω-words.

(2) $\mathcal{K}$ tries to construct all runs of T.

(3) The LTL formula φ asserts that the run in question is *non*-accepting (or malformed).

Runs as ω-words An accepting run of T can be written as a sequence of configurations c_i, separated by a marker $[]$ as follows.

$$[]\ c_0\ []\ c_1\ []\ \cdots\ []\ c_k\ []\ c_k\ []\ c_k\ []\ c_k\ \cdots$$

Each c_i, which has the shape $a_0\ a_1 \cdots (q, a_j) \cdots a_{s(n)}$ where $0 \le j \le s(n)$ and a_i ranging over input symbols, represents the configuration comprising the tape

$$a_0\ a_1\ \cdots\ a_{s(n)-1}\ a_{s(n)}\ \square\ \square\ \square\ \square\ \cdots$$

with control state q and tape head position j. The initial configuration, c_0, is the sequence $\underbrace{(q_0, \square)\ \square \cdots \square}_{s(n)}$. Further, for all i, c_{i+1} follows from c_i, and the control state in c_k is accepting.
Thus an accepting run *is an ω-word* of a certain shape.

$\mathcal{K}$ **constructs every run** by generating *every* possible ω-word.

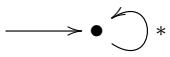

Recognising a bad run We want φ to characterise exactly the non-accepting or malformed ω-words. Such a word is

(i) *either* not a sequence of configurations. For example $a \,[\!]\, (q,b) \,[\!]\, [\!]\, a\ a \,[\!]\, (q,a)\ (q,b)\ \cdots$

(ii) *or* it does not start with the initial configuration

(iii) *or* it does not reach a final configuration

(iv) *or* for some i, c_{i+1} cannot follow from from c_i.

Then φ is the disjunction of the formulas describing the respective cases above. We consider them in turn.

Not a sequence of configurations The ω-word is not of the form $[\!]\underbrace{\cdots}_{s(n)}[\!]\underbrace{\cdots}_{s(n)}[\!]\cdots$

$$\neg\left[[\!] \ \wedge\ \boldsymbol{G}\left([\!] \Rightarrow \left(\boldsymbol{X}^{s(n)+1}[\!] \wedge \bigwedge_{1\le i\le s(n)} \boldsymbol{X}^i \neg[\!]\right)\right)\right]$$

or some configuration does not contain exactly one head.

$$\neg\boldsymbol{G}\,[[\!] \Rightarrow \boldsymbol{X}\,(\mathsf{Cell}\ \boldsymbol{U}\,(\mathsf{Head} \wedge \boldsymbol{X}\,(\mathsf{Cell}\ \boldsymbol{U}\,[\!])))]$$

where

$$\begin{aligned} \mathsf{Head} &:= \bigvee_{q,a}(q,a) \\ \mathsf{Cell} &:= \bigvee_{a} a \end{aligned}$$

Initial and final conditions The ω-word:

- does not start with the initial configuration $(q_0, \square)\ \square\ \cdots\ \square$.

$$\neg\left(\boldsymbol{X}\,(q_0,\square) \wedge \bigwedge_{2\le i\le s(n)} \boldsymbol{X}^i\,\square\right)$$

(Why don't we test for $[\!]$ characters?)

- or does not reach a final configuration.

$$\neg\boldsymbol{F}\left(\bigvee_{q\ \text{final},a}(q,a)\right)$$

Some c_{i+1} does not follow from c_i Let r range over the transitions of T.

$$\neg \boldsymbol{G}\left([\!] \Rightarrow \bigvee_r \mathsf{Follows}_r \right)$$

Let (q, a) be the head of r, q' the next state, b the character to write, $d \in \{-1, 0, 1\}$ is the direction of the head movement. Then

$$\mathsf{Follows}_r := \bigvee_{1 \le i \le s(n)} \left[\begin{pmatrix} \boldsymbol{X}^i (q, a) \wedge \\ \boldsymbol{X}^{s(n)+1+i} \mathsf{Char}(b) \wedge \\ \boldsymbol{X}^{s(n)+1+i+d} \mathsf{State}(q') \end{pmatrix} \wedge \bigwedge_{j \neq i} \bigvee_a \begin{pmatrix} \boldsymbol{X}^j \mathsf{Char}(a) \wedge \\ \boldsymbol{X}^{s(n)+1+j} \mathsf{Char}(a) \end{pmatrix} \right]$$

where

$$\begin{aligned} \mathsf{Char}(b) &:= b \vee \bigvee_{q''} (q'', b) \\ \mathsf{State}(q') &:= \bigvee_c (q', c) \end{aligned}$$

Putting all of the above together, we have $\mathcal{K} \nvDash \varphi$ iff T has an accepting run.

Further considerations:

- What if the formula is fixed?
- What if the model is fixed?

2.6 Expressive Power of LTL

We say that $L \subseteq \Sigma^\omega$ *non-counting* just if there is an $n_0 \geq 0$ such that for every $n \geq n_0$ and for every $u, v \in \Sigma^*$ and $\beta \in \Sigma^\omega$, we have $u\, v^n\, \beta \in L \;\leftrightarrow\; u\, v^{n+1}\, \beta \in L$. I.e. if L contains an infinite word that embeds a finite word repeated sufficiently often, then for every n larger than the threshold, L contains such an embedded word in which the finite word is repeated n-times.

For example, $(0\,0)^*\, 1^\omega$ is counting.

Proposition 8. *Every LTL-definable ω-language is non-counting. It follows that there are Büchi recognisable ω-languages that are not LTL-definable.* □

Some FAQs

(1) What does "linear time" in LTL mean?

Linear-time specifications set same conditions on every infinite path through system modelled by Kripke structure. *Branching-time specifications* are conditions on the structure of tree formed by all paths through a Kripke structure. Well-known logics for describing branching-time properties are computational tree logic (CTL) and CTL*. See (Vardi, 2001) for a readable study on linear-time versus branching-time logics.

(2) Is LTL a "robust" logic? Are there nice characterisations of the LTL-definable languages?

A *star-free regular expression* over Σ is an expression built up using ϵ, symbols $a \in \Sigma$, concatenation, union and complementation with respect to Σ^*. A *star-free regular language* is a language that matches a star-free regular expression.

Theorem 2.4 (Characterisations of LTL Definability)**.** *Let $L \subseteq \Sigma^\omega$. The following are equivalent.*

(i) *L is definable in LTL*

(ii) *L is star-free ω-regular i.e. a finite union of ω-languages of the form $L_1 \cdot L_2^\omega$ where $L_1, L_2 \subseteq \Sigma^*$ are star-free regular*

(iii) *L is a finite union of ω-languages of the form $\lim L_1 \cap (\Sigma^\omega \setminus \lim L_2)$ where $L_1, L_2 \subseteq \Sigma^*$ are star-free regular.* □

(3) What is Kamp's Theorem?

In his UCLA PhD thesis, Kamp (1968) proved that an ω-language is LTL-definable if and only if it is definable in $FO(<, (P_a)_{a\in\Sigma})$ i.e. first-order logic with a binary predicate symbol $<$ and a unary predicate P_a for each $a \in \Sigma$. We write $\mathcal{M}_\alpha = (\omega, <, (P_a)_{a\in\Sigma})$ for the obvious structure determined by $\alpha \in \Sigma^\omega$ where for each $a \in \Sigma$, $n \in P_a \leftrightarrow \alpha(n) = a$.

Theorem 2.5 (Kamp 1968)**.** *Let α be an ω-word over Σ and $n \in \omega$.*

(i) *For each LTL formula φ there exists a formula $\chi_\varphi(x)$ in $FO(<, (P_a)_{a\in\Sigma})$ with a free variable x such that*

$$\alpha^n \vDash \varphi \;\leftrightarrow\; \mathcal{M}_\alpha \vDash \chi_\varphi(\underline{n}).$$

(ii) *For each formula $\chi(x)$ in $FO(<, (P_a)_{a\in\Sigma})$, there exists an LTL formula φ_χ such that*

$$\mathcal{M}_\alpha \vDash \chi(\underline{n}) \;\leftrightarrow\; \alpha^n \vDash \varphi_\chi.$$

It follows that for each FO sentence χ there exists an LTL formula φ_χ such that $\mathcal{M}_\alpha \vDash \chi \leftrightarrow \alpha^0 \vDash \varphi_\chi$. □

See (Rabinovich, 2012) for a new proof of Kamp's theorem.

(4) Can we extend LTL to make it *equi-expressive* with Büchi automata (for ω-languages)?

Yes. μLTL: LTL augmented by (a least μ, and hence also greatest ν) fixpoint operators.

Theorem 2.6 (Characterisations of ω-Regularity)**.** *Let $L \subseteq \Sigma^\omega$. The following are equivalent.*

(i) *L is ω-regular*

(ii) *L is definable in μLTL*

(iii) *L is definable in S1S (see the following chapter)*

(iv) *L is definable in Weak S1S.* □

(5) Is there a characterisation of the subclass of Büchi automata that is equivalent to LTL?

The automata that are equi-expressive with LTL formulas are called *linear weak alternating automata*. For details see (Löding and Thomas, 2000; Gastin and Oddoux, 2001; Hammer et al., 2005).

Problems

2.1 Consider the following properties for the lift system introduced in the introductory chapter:

(A1) Every requested level will be served eventually.

(A2) The lift will return to level 1 again and again.

(A3) Whenever the top level is requested, the lift serves it immediately and does not stop on the way there.

(A4) It is always the case that while moving in one direction, the lift will stop at every requested level, unless the top level is requested.

Assume that the lift serves only four levels. By introducing appropriate atomic propositions (ten should suffice), describe the above properties as LTL-formulas. You should begin by constructing the state-transition graph.

2.2 Let φ, ψ and χ be LTL-formulas. We say that two formulas are *equivalent* if they define the same language. For each of the following, prove or disprove each of the two implications:

(a) $\boldsymbol{F}\,\boldsymbol{G}\,\varphi \equiv \boldsymbol{G}\,\boldsymbol{F}\,\varphi$

(b) $\boldsymbol{X}\,(\varphi \wedge \psi) \equiv \boldsymbol{X}\,\varphi \wedge \boldsymbol{X}\,\psi$

(c) $(\varphi \vee \psi)\ \boldsymbol{U}\,\chi \equiv (\varphi\ \boldsymbol{U}\,\chi) \vee (\psi\ \boldsymbol{U}\,\chi)$

(d) $(\varphi\ \boldsymbol{U}\,\psi)\ \boldsymbol{U}\,\chi \equiv \varphi\ \boldsymbol{U}\,(\psi\ \boldsymbol{U}\,\chi)$

2.3 Let φ and ψ be LTL-formulas. Consider the following temporal operators:

(a) "at next" $\varphi\,\boldsymbol{A}\boldsymbol{X}\,\psi$: When ψ next[2] holds (if it does), so does φ.
(Note that ψ may never hold.)

(b) "while" $\varphi\ \boldsymbol{W}\psi$: φ holds for at least as long as ψ does.
(Note that ψ is not assumed to hold at the beginning.)

(c) "before" $\varphi\,\boldsymbol{B}\,\psi$: When ψ next holds (if it does), φ does so before.
(Note that ψ may never hold.)

For each construct, find an equivalent LTL-formula.

[2] We do *not* mean "when φ hold at the next time step", but rather "when φ holds at some point in the future".

2.4 Translate the following LTL formulas to Büchi automata:

(a) $\boldsymbol{F} p_1 \wedge \neg \boldsymbol{F} p_2$

(b) $\boldsymbol{G}(p_1 \rightarrow \boldsymbol{X}(p_2 \ \boldsymbol{U} p_1))$

(c) $\boldsymbol{G}\,\boldsymbol{F} p_1 \rightarrow \boldsymbol{F}\,\boldsymbol{G} p_2$.

2.5 We define a sublogic $\mathcal{T}(\boldsymbol{U})$ of LTL consisting of formulas that are built up from the atomic propositions, using conjunction, negation and the until-operator $\varphi \ \boldsymbol{U} \psi$. (Thus we may write LTL as $\mathcal{T}(\boldsymbol{X}, \boldsymbol{U})$.)

Suppose there is only one atomic proposition, p_1. Consider the label sequence

$$\alpha \ = \ (1)(1)(0)(0)\cdots$$

(a) Prove, by structural induction on formulas, that for all $\varphi \in \mathcal{T}(\boldsymbol{U})$, we have $\alpha^0 \vDash \varphi$ iff $\alpha^1 \vDash \varphi$.

(b) Find an LTL-formula ψ satisfying $\alpha^0 \vDash \psi$ and $\alpha^1 \nvDash \psi$.

(c) Hence prove that $\mathcal{T}(\boldsymbol{U})$ is strictly less expressive than LTL.

2.6 Prove that for each generalised Büchi automaton

$$A = (Q, \Sigma, \Delta, q_0, \{ F_0, \cdots, F_{l-1} \})$$

there is an equivalent Büchi automaton A' i.e. they recognise the same ω-language.

2.7 Consider the LTL-formula $\varphi = p_1 \ \boldsymbol{U}(\boldsymbol{X} p_2)$.

(a) Let $\alpha \in (\{ 0, 1 \}^2)^\omega$. Formulate the compatibility conditions for the φ-expansion of α.

(b) Construct a generalised Büchi automaton A equivalent to φ. What are the final states of A?

(c) Construct directly a Büchi automaton recognising $L = \{ \alpha \in (\{ 0, 1 \}^2)^\omega : \alpha \vDash \varphi \}$.

2.8

(a) Show that the ω-language $L_1 = \{ (01)^\omega \}$ is non-counting.

(b) Show that the ω-language is $L_2 = \{ 01(0101)^* 0^\omega \}$ counting.

2.9 An ω-language L over an alphabet Σ is said to be *stuttering* if for each letter $a \in \Sigma$, we have

$$u\,a\,\beta \in L \ \leftrightarrow \ u\,a\,a\,\beta \in L$$

(where u ranges over Σ^* and β over Σ^ω). Which of the following are true? Justify your answers.

(a) If φ is an LTL formula then $L(\varphi)$ is stuttering.

(b) If φ is an LTL formula without $\boldsymbol{X}$(-) then $L(\varphi)$ is stuttering.

2.10 Let p be the only atomic proposition. For $n \geq 0$, let γ_n be the infinite word $1^n 01^\omega$. Thus if $i \neq n$ then $\gamma_n^i \vDash p$, and $\gamma_n^n \nvDash p$.

Let φ be an LTL-formula. Prove, by structural induction, that if φ has no more than n occurrences of $\boldsymbol{X}(-)$, then for all $i, j > n$, we have $\gamma_j \vDash \varphi$ iff $\gamma_i \vDash \varphi$.

[*Hint.* For the inductive case of $\varphi = \varphi_1 \, \boldsymbol{U} \, \varphi_2$: $\gamma_i \vDash \varphi$ means that for some $l \geq 0$, we have $\gamma_i^l \vDash \varphi_2$, and $\gamma_i^k \vDash \varphi_1$ for all $0 \leq k < l$. Consider the cases of $l \leq i$ and $l > i$ in turn. In the former case, analyse the cases of $i - l > n$ and $i - l \leq n$. Note that $\gamma_i^l = \gamma_{i-l}$ if $l \leq i$. What is γ_i^l if $l > i$?]

2.11 † We define a sublogic $\mathcal{T}(\boldsymbol{X}, \boldsymbol{G})$ of LTL consisting of formulas that are built up from the atomic propositions, using conjunction, negation, next-time operator $\boldsymbol{X}\varphi$, and the always-modality $\boldsymbol{G}\varphi$. We shall prove:

Theorem 2.7. *$\mathcal{T}(\boldsymbol{X}, \boldsymbol{G})$ is strictly less expressive than LTL.*

Consider a Kripke structure (over atomic propositions p_1, p_2) that has states $s_0, \cdots, s_{4m-1}$ and the transition relation is specified by:

$$s_i \to s_j \quad \Longleftrightarrow \quad j = i + 1 \text{ or } (i = 4m - 1 \text{ and } j = 2m).$$

The proposition p_1 is assumed to hold for all states except s_{3m}, and p_2 is assumed to hold for just the states s_{2m-1} and s_{4m-1}. Let α be the uniquely determined label sequence starting in state s_0. I.e.

$$\alpha \;=\; \underbrace{\binom{1}{0}\cdots\binom{1}{0}\binom{1}{1}}_{2m}\left(\underbrace{\underbrace{\binom{1}{0}\cdots\binom{1}{0}}_{m}\binom{0}{0}\binom{1}{0}\cdots\binom{1}{0}\binom{1}{1}}_{2m}\right)^{\omega}$$

(a) List the elements of the set $\{\, 0 \leq l \leq 4m - 1 : \alpha^l \vDash p_1 \, \boldsymbol{U} \, p_2 \,\}$.

(b) Prove that for all $\varphi \in \mathcal{T}(\boldsymbol{X}, \boldsymbol{G})$ containing fewer than $m - 1$ occurrences of $\boldsymbol{X}$, we have

$$\alpha^0 \vDash \varphi \quad \Longleftrightarrow \quad \alpha^{2m} \vDash \varphi.$$

(c) Observe that $\alpha^0 \vDash p_1 \, \boldsymbol{U} p_2$ and $\alpha^{2m} \nvDash p_1 \, \boldsymbol{U} p_2$. Hence or otherwise prove that $\mathcal{T}(\boldsymbol{X}, \boldsymbol{G})$ is strictly less expressive than LTL.

2.12 The aim of the problem is to establish the co-NP-hardness of the LTL Model Checking Problem by reducing (the co-NP-complete problem) 3-UNSAT to it.

Precisely, prove that a propositional formula ψ in conjunctive normal form (with three literals per clause, where a literal is a variable or the negation of a variable) can be transformed in polynomial time into a Kripke structure $\mathcal{K}_\psi$ and an LTL-formula F_ψ such that ψ

is satisfiable if, and only if, the pair $(\mathcal{K}_\psi, F_\psi)$ is a no-instance of the LTL Model Checking Problem.

Take $\psi = (x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3)$ which can be satisfied with the assignment $x_1 \mapsto 1, x_2 \mapsto 0, x_3 \mapsto 0$.

Consider the following Kripke structure $\mathcal{K}_\psi$

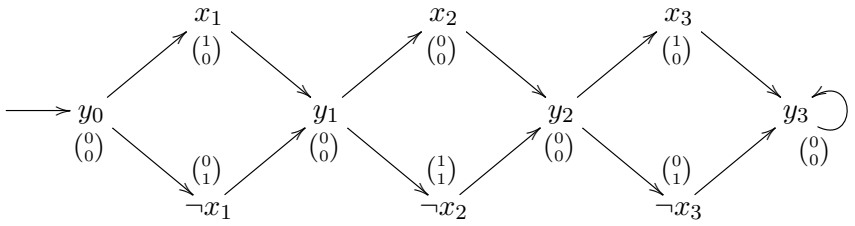

The labels are defined as follows: the j-component of the label of the literal ℓ is 1 just if ℓ occurs in the j-clause of ψ. The idea is that an assignment determines a path through the Kripke structure.

(a) What is F_ψ for the above example?

(b) By giving the construction $\psi \mapsto \mathcal{K}_\psi$ and the corresponding formula F_ψ, prove that there is a polynomial reduction of the desired kind.

Thus conclude that the LTL Model Checking Problem is co-NP-hard.

Chapter 3

S1S

Synopsis

Examples. Syntax and semantics. Büchi-recognisable languages are S1S-definable. S1S-definable languages are Büchi-recognisable. Church's Synthesis Problem (Overview).

References (Büchi, 1960a; Elgot, 1961; Büchi, 1962; Muller, 1963)

3.1 Introduction

How would you decide the following kinds of sentences?

(i) $\forall x \,.\, \exists y, z \,.\, (y + 1 \leq x \;\; \vee z + 1 \leq y \;\; \wedge \; 2x + 1 \leq z)$

(ii) "For every (non-zero) natural number n, there are powers of two, x and y, such that $x \leq n < y$."

$$\forall n \,.\, \exists x, y \,.\, \mathit{PowsOf2}(x) \wedge \mathit{PowsOf2}(y) \wedge x \leq n < y$$

(iii) A given subset X of natural numbers contains only even numbers.

The logical system S1S The above can be described as sentences of monadic second-order logic of one-successor, or S1S. *Second-order* means that we allow quantification over *relations*; *monadic* means that quantification is restricted to *monadic* (or *unary*) relations, namely, sets. The *theory S1S* is the set of true statements about $\omega = \{\, 0, 1, 2, \cdots \}$ expressible in a language that has:

- a unary function symbol **s** for successor
- a binary set-membership predicate $\in$
- first-order quantification (over elements of ω) and second-order quantification (over subsets of ω).

The theory S1S is *decidable, but not in elementary time* i.e. there is no (fixed) $h \geq 0$ such that the theory is decidable in time bounded by the tower-of-exponentials function of height h, $\mathit{exp}_h(n)$ where

$$\mathit{exp}_0(n) := n \qquad \mathit{exp}_{h+1}(n) := 2^{\mathit{exp}_h(n)}.$$

Why study S1S?

- *Historical importance*: the first time automata and logic connection is established and exploited.
- *Powerful decidable theory*: many decisions problems can be shown decidable by reduction to S1S.
- MSO is considered a gold standard of logics for describing correctness properties of reactive systems.

3.2 The logical system S1S

The *vocabulary* consists of a unary function symbol $\mathbf{s}$ and a binary predicate symbol $\in$. The corresponding *logical structure* is $(\omega, \mathbf{s}, \in)$ where[1] $\mathbf{s}$ is the successor function $x \mapsto x + 1$, and $\in \subseteq \omega \times 2^\omega$ is the standard membership relation between elements and sets.

The logical system S1S is defined as follows.

- *Variables.* 1st-order variables (x, y, z, etc.) range over natural numbers (regarded as positions in ω-words). 2nd-order variables (X, Y, Z, etc.) range over sets of natural numbers.
- *Terms.* 1st-order variables are terms. If t is a term, so is $\mathbf{s}\,t$.
- *Formulas. Atomic formulas* are of the shape $t \in X$ where t is a term and X is a 2nd-order variable.
  S1S formulas are built up from atomic formulas using standard Boolean connectives, with $\forall$- and $\exists$-quantifications over 1st and 2nd-order variables.

Constructs definable in S1S Note that 0 and the atomic formulas $s = t$ and $s < t$ are definable in terms of set-membership and successor.

- "$x = y$" $:=\ \forall X.x \in X \leftrightarrow y \in X$
- "$X \subseteq Y$" $:=\ \forall x.x \in X \rightarrow x \in Y$
- "$X = Y$" $:=\ X \subseteq Y \wedge Y \subseteq X$
- "$x = 0$" $:=\ \forall y.\neg(x = \mathbf{s}\,y)$ "x has no predecessor"
- "$x = 1$" $:=\ x = \mathbf{s}\,0$
- "$x \leq y$" $:=\ \forall X.(x \in X \wedge (\forall z.z \in X \rightarrow \mathbf{s}\,z \in X)) \rightarrow y \in X$
  "Every set X that contains x and is closed under successor (in particular, the *smallest* such X) also contains y."
- "X is finite" $:=\ \exists x.\forall y.(y \in X \rightarrow y \leq x)$

Subsystems of S1S

- *First-order fragment: $S1S_1$.* Formulas are built up from atomic formulas using boolean connectives and first-order quantifiers.

[1] We use the same LaTeX-symbol for the *function symbol* $\mathbf{s}$ (in the vocabulary) and its *interpretation* (in the logical structure) – it should be clear from the context which is intended; similarly for $\in$.

- *Existential S1S.* Formulas are $\mathrm{S1S}_1$-formulas preceded by a block $\exists Y_1 \cdots Y_m$ of existential second-order quantifiers.

3.3 Semantics of S1S

Write $\varphi(x_1, \cdots, x_m, X_1, \cdots, X_n)$ to mean: φ has free 1st-order variables from $x_1, \cdots, x_m$ and free 2nd-order variables from $X_1, \cdots, X_n$. Let $a_i \in \omega$ and $P_j \subseteq \omega$. For $\overline{a} = a_1, \cdots, a_m$ and $\overline{P} = P_1, \cdots, P_n$, we write

$$(\omega, \mathbf{s}, \in); \overline{a}; \overline{P} \ \vDash \ \varphi(x_1, \cdots, x_m, X_1, \cdots, X_n)$$

or simply $\overline{a}, \overline{P} \vDash \varphi$, to mean " the structure $(\omega, \mathbf{s}, \in)$ with the assignment $\overline{x} \mapsto \overline{a}; \overline{X} \mapsto \overline{P}$ satisfies φ".

Definition 3.1. We define the *satisfaction relation*

$$\overline{a}, \overline{P} \vDash \varphi(\overline{x}, \overline{X})$$

by recursion over the syntax of φ:

$$\begin{array}{rcl}
\overline{a}; \overline{P} \vDash \underbrace{\mathbf{s}(\cdots(\mathbf{s}}_{k}\ x_i)\cdots) \in X_j & := & a_i + k \in P_j \\
\overline{a}; \overline{P} \vDash \neg\varphi(\overline{x}, \overline{X}) & := & \overline{a}; \overline{P} \nvDash \varphi(\overline{x}, \overline{X}) \\
\overline{a}; \overline{P} \vDash \varphi_1(\overline{x}, \overline{X}) \vee \varphi_2(\overline{x}, \overline{X}) & := & \overline{a}; \overline{P} \vDash \varphi_1(\overline{x}, \overline{X}) \quad \text{or} \quad \overline{a}; \overline{P} \vDash \varphi_2(\overline{x}, \overline{X}) \\
\overline{a}; \overline{P} \vDash \exists y.\varphi(\overline{x}, y, \overline{X}) & := & \overline{a}, b; \overline{P} \vDash \varphi(\overline{x}, y, \overline{X}) \text{ for some } b \in \omega \\
\overline{a}; \overline{P} \vDash \exists Z.\varphi(\overline{x}, \overline{X}, Z) & := & \overline{a}; \overline{P}, Q \vDash .\varphi(\overline{x}, \overline{X}, Z) \text{ for some } Q \subseteq \omega
\end{array}$$

Standardly $\varphi_1 \wedge \varphi_2$ is equivalent to $\neg\varphi_1 \vee \neg\varphi_2$, and $\forall X.\varphi$ and $\forall x.\varphi$ are equivalent to $\neg(\exists X.\neg\varphi)$ and $\neg(\exists x.\neg\varphi)$ respectively.

Representing a set of natural numbers as an infinite word We represent any $P \subseteq \omega$ by its *characteristic word*, written $\ulcorner P \urcorner \in \mathbb{B}^\omega$, defined by

$$\ulcorner P \urcorner(i) = 1 \quad \leftrightarrow \quad i \in P.$$

E.g. the characteristic words of the set of multiples of 3 and the set of prime numbers are respectively:

$$1001001001001001001001\cdots$$
$$0011010100010100010100\cdots$$

We represent $a \in \omega$ by the characteristic word of the singleton set $\{\, a \,\}$.

More generally the *characteristic word* of a tuple

$$(a_1, \cdots, a_m, P_1, \cdots, P_n) \ \in \ \omega^m \times (2^\omega)^n$$

written $\ulcorner a_1, \cdots, a_m, P_1, \cdots, P_n \urcorner$, is an infinite word over the alphabet $\mathbb{B}^{m+n}$ such that each of the $m + n$ tracks (or rows) is the characteristic word of the corresponding component of the tuple $(\overline{a}, \overline{P})$.

Defining ω-languages by S1S formulas We say $L \subseteq \mathbb{B}^\omega$ is *S1S-definable* by $\varphi(X)$ just if $L = \{ \ulcorner P \urcorner \in \mathbb{B}^\omega : P \vDash \varphi(X) \}$. I.e. Each P that satisfies $\varphi(X)$ contains exactly the numbers denoting the positions of '1' in an ω-word in $L \subseteq \mathbb{B}^\omega$.

Example 3.1. (i) The set $L_1 = \{ \alpha \in \mathbb{B}^\omega : \alpha \text{ has infinitely many 1s} \}$ is first-order definable by $\varphi_1(X) = \forall x.\exists y.x < y \wedge y \in X$.

(ii) $(00)^*1^\omega$ is definable by

$$\varphi_2(X) \;=\; \exists Y.\exists x.\left(\begin{array}{ll} & 0 \in Y \\ \wedge & \forall y.y \in Y \leftrightarrow \mathbf{s}\, y \notin Y \\ \wedge & x \in Y \\ \wedge & \forall z.z < x \rightarrow z \notin X \\ \wedge & \forall z.x \leq z \rightarrow z \in X \end{array}\right)$$

(What is Y?) Recall that $(00)^*1^\omega$ is a "counting" language, hence not LTL-definable.

Translating LTL to S1S Fix atomic formulas $p_1, \cdots, p_n$. We say S1S-formula $\varphi(X_1, \cdots, X_n)$ is *equivalent* to an LTL-formula ψ just if $[\![\psi]\!] = \{ \ulcorner \overline{P} \urcorner \in (B^n)^\omega : \overline{P} \vDash \varphi(\overline{X}) \}$.

Example 3.2. (i) $\boldsymbol{X}\,\boldsymbol{X}(p_2 \rightarrow \boldsymbol{F}\,p_1)$

$$\varphi_3(X_1, X_2) \;=\; \mathbf{s}\,\mathbf{s}\,0 \in X_2 \rightarrow \exists x.(\mathbf{s}\,\mathbf{s}\,0 \leq x \wedge x \in X_1)$$

(ii) $\boldsymbol{F}(p_1 \wedge \boldsymbol{X}(\neg p_2\ \boldsymbol{U}\, p_1))$

$$\begin{array}{ll} & \varphi_4(X_1, X_2) \\ = & \exists x.\left(\begin{array}{ll} & x \in X_1 \\ \wedge & \exists y.\left(\begin{array}{ll} & \mathbf{s}\,x \leq y \\ \wedge & y \in X_1 \\ \wedge & \forall z.(\mathbf{s}\,x \leq z \wedge z < y) \rightarrow \neg(z \in X_2) \end{array}\right) \end{array}\right) \end{array}$$

3.4 Büchi-Recognisable Languages are S1S-Definable

Definition 3.2. An ω-language $L \subseteq (\mathbb{B}^n)^\omega$ is *S1S definable* just if there is an S1S-formula $\varphi(X_1, \cdots, X_n)$ such that $L = \{ \ulcorner P_1, \cdots, P_n \urcorner \in (\mathbb{B}^n)^\omega : \overline{P} \vDash \varphi(\overline{X}) \}$.

Theorem 3.1 (Büchi)**.** *For every Büchi automaton A over the alphabet $\mathbb{B}^n$, there is an S1S formula $\varphi_A(X_1, \cdots, X_n)$ such that for every $(P_1, \cdots, P_n) \in (2^\omega)^n$, we have $\overline{P} \vDash \varphi_A(\overline{X})$ if and only if A accepts $\ulcorner P_1, \cdots, P_n \urcorner$.*

The proof idea is simple: take a Büchi automaton $A = (Q, \Sigma, q_1, \Delta, F)$ where $\Sigma = \mathbb{B}^n$, and construct an S1S-formula $\varphi_A(X_1, \cdots, X_n)$ that asserts "there is an accepting run of A on input given by the characteristic word of $(X_1, \cdots, X_n)$".

Proof. The aim is to code a run. Suppose $Q = \{ q_1, \cdots, q_m \}$. A run $\rho(0)\rho(1)\cdots \in Q^\omega$ is coded by m subsets of ω, namely $Y_1, \cdots, Y_m$, such that

$$i \in Y_k \quad \leftrightarrow \quad \rho(i) = q_k$$

Clearly $Y_1, \cdots, Y_m$ form a *partition* of ω. Each tuple "$Y_1, \cdots, Y_m$" describes an infinite word over Q, α, whereby each

$$Y_j = \{\, \text{positions in } \alpha \text{ with symbol } q_j \,\}.$$

Define $\mathit{partition}(Y_1, \cdots, Y_m)$ to be

$$\forall x. \left(\bigvee_{i=1}^{m} x \in Y_i \right) \quad \wedge \quad \neg \left(\exists y. \bigvee_{i \neq j} (y \in Y_i \wedge y \in Y_j) \right)$$

Coding letters of the alphabet $\mathbb{B}^n$ For $a = \begin{pmatrix} b_1 \\ \vdots \\ b_n \end{pmatrix} \in \mathbb{B}^n$, introduce shorthand

$$x \in X_a \;:=\; [b_1](x \in X_1) \wedge [b_2](x \in X_2) \wedge \;\cdots\; \wedge [b_n](x \in X_n)$$

where

$$[b_i](x \in X_i) \;:=\; \begin{cases} x \in X_i & \text{if } b_i = 1 \\ \neg(x \in X_i) & \text{otherwise} \end{cases}$$

Given a Büchi automaton $A = (\{\, 1, \cdots, m \,\}, \mathbb{B}^n, 1, \Delta, F)$, define $\varphi_A(X_1, \cdots, X_n)$ to be

$$\exists Y_1 \cdots Y_m \;. \begin{pmatrix} & \mathit{partition}(Y_1, \cdots, Y_m) \\ \wedge & 0 \in Y_1 \\ \wedge & \forall x. \bigvee_{(i,a,j) \in \Delta} (x \in Y_i \;\wedge\; x \in X_a \;\wedge\; \mathbf{s}\, x \in Y_j) \\ \wedge & \forall x. \exists y. (x < y \;\wedge\; \bigvee_{i \in F} y \in Y_i) \end{pmatrix}$$

Thus for every $(P_1, \cdots, P_n) \in (2^\omega)^n$, A accepts $\ulcorner P_1, \cdots, P_n \urcorner$ iff $\overline{P} \vDash \varphi_A(X_1, \cdots, X_n)$. $\square$

Observer that $\varphi_A(X_1, \cdots, X_m)$ is an existential S1S-formula. It follows that Büchi-recognisable ω-languages are *existential S1S-definable*.

3.5 S1S-Definable Languages are Büchi-Recognisable

Theorem 3.2 (Büchi)**.** *For every S1S formula $\varphi(x_1, \cdots, x_m, X_1, \cdots, X_n)$, there is an equivalent non-determinstic Büchi automaton A_φ over alphabet $\mathbb{B}^{m+n}$, in the sense that*

$$L(A_\varphi) \;=\; \{\, \ulcorner a_1, \cdots, a_m, P_1, \cdots, P_n \urcorner \in (\mathbb{B}^{m+n})^\omega \mid \overline{a}, \overline{P} \vDash \varphi \,\}$$

Proof. The proof is by induction on the size of φ. An *atomic formula* has the form $\underbrace{\mathbf{s}\,(\mathbf{s} \cdots (\mathbf{s}}_{k}\, x_i) \cdots) \in X_j$. We build a Büchi automaton to read the tracks i and $m + j$ only (corresponding to x_i and X_j respectively), performing the following check: if the unique 1 of the x_i-track is at position l (say), then the X_j-track has a 1 in position $l + k$.

Disjunction: Consider $\varphi_1(\overline{x}, \overline{X}) \vee \varphi_2(\overline{x}, \overline{X})$. By the induction hypothesis, suppose automata A_{φ_1} and A_{φ_2} are equivalent to φ_1 and φ_2 respectively. Set $A_{\varphi_1 \vee \varphi_2}$ to be the Büchi automaton that accepts the union of (the respectively ω-languages of) A_{φ_1} and A_{φ_2}.

Negation: Consider $\neg\varphi(\overline{x}, \overline{X})$. By the induction hypothesis, suppose A_φ is equivalent to φ. Set $A_{\neg\varphi}$ to be the automaton that recognises the complement of $L(A_\varphi)$.

Second-order existential quantification: Consider $\exists Y.\varphi(\overline{x}, \overline{X}, Y)$. By the induction hypothesis, suppose A_φ is the Büchi automaton equivalent to $\varphi(\overline{x}, \overline{X}, Y)$. I.e. for all $(\overline{a}, \overline{P}, Q) \in \omega^m \times (2^\omega)^{n+1}$,

$$\overline{a}, \overline{P}, Q \vDash \varphi \quad \leftrightarrow \quad A_\varphi \text{ accepts } \ulcorner \overline{a}, \overline{P}, Q \urcorner$$

We construct $A_{\exists Y.\varphi}$ by replacing each transition label $\begin{pmatrix} b_1 \\ \vdots \\ b_{m+n} \\ b \end{pmatrix}$ in A_φ by $\begin{pmatrix} b_1 \\ \vdots \\ b_{m+n} \end{pmatrix}$ thus introducing (further) non-determinacy.

Consequently a transition via $\begin{pmatrix} b_1 \\ \vdots \\ b_{m+n} \end{pmatrix}$ in $A_{\exists Y.\varphi}$ corresponds to a transition via $\begin{pmatrix} b_1 \\ \vdots \\ b_{m+n} \\ 0 \end{pmatrix}$ or $\begin{pmatrix} b_1 \\ \vdots \\ b_{m+n} \\ 1 \end{pmatrix}$ in A_φ. Hence

$$\begin{array}{ll} & A_{\exists Y.\varphi} \text{ accepts } \ulcorner \overline{a}, \overline{P} \urcorner \in (\mathbb{B}^{m+n})^\omega \\ \text{iff} & \text{for some } c_0 c_1 \cdots \in \mathbb{B}^\omega, \text{ we have } A_\varphi \text{ accepts } \binom{\alpha(0)}{c_0} \binom{\alpha(1)}{c_1} \cdots \\ & (\text{Suppose } c_0 c_1 \cdots = \ulcorner Q \urcorner \text{ and } \alpha = \ulcorner \overline{a}, \overline{P} \urcorner.) \\ \text{iff} & \text{for some } Q \subseteq \omega \text{ we have } \overline{a}, \overline{P}, Q \vDash \varphi(\overline{x}, \overline{X}, Y) \\ \text{iff} & \overline{a}, \overline{P} \vDash \exists Y.\varphi(\overline{x}, \overline{Y}) \end{array}$$

First-order existential quantification: Consider $\exists y.\varphi(\overline{x}, y, \overline{X})$. Exactly the same as above. □

3.6 The Synthesis Problem

Church (1903-1995)

Identify (infinite) binary sequences α, β with subsets of $\mathbb{N}$. For example we identify $101011\cdots$ with $\{0, 2, 4, 5, \cdots\}$. Specifications $\varphi(X, Y)$ are given in monadic second-order logic (or S1S), where variables X and Y range over subsets of $\mathbb{N}$.

The Synthesis Problem (Alonzo Church, 1962) Construct a procedure that transforms a given logical specification $\varphi(X, Y)$ to a *finite-state automaton with output*, which, for every input sequence α, produces an output sequence β such that $P, Q \vDash \varphi(X, Y)$ holds, where $\lceil P \rceil = \alpha$ and $\lceil Q \rceil = \beta$.

Example 3.3. *Informal specification:* α = input stream; β = output stream

(i) "*1-Reflecting*": If current input symbol is 1, so is the next output symbol.

(ii) *Output stream β has no successive 0s.*

(iii) "*0-Fairness*": If α has infinitely many 0s, so has β.

MSO formula $\varphi(X, Y)$: Write $X(t)$ to mean "$t \in X$".

$$\begin{array}{rcl} \varphi(X,Y) & := & \forall z.(X(z) \to Y(z)) \\ & \wedge & \neg(\exists x.\neg Y(x) \wedge \neg Y(\mathbf{s}\, x)) \\ & \wedge & (\forall x.\exists y > x.\neg X(y) \ \to \ \forall x.\exists y > x.\neg Y(y)) \end{array}$$

A Mealy-automaton solution (= finite-state strategy for Éloïse):

- If current input is 1, then output 1
- If current input is 0, then output the opposite of the preceding output.

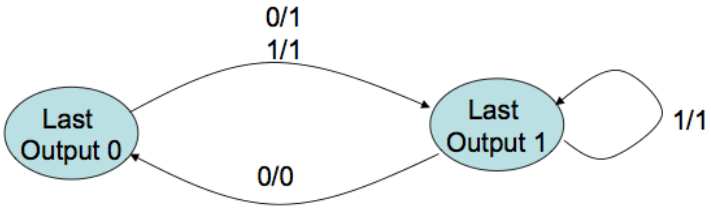

Church's Problem Recast as a Game Problem *A 2-Person Infinite 0/1-Game:*

- *Players:* Abelard and Éloïse.
- Abelard starts with $\alpha(0)$. Éloïse responds with $\beta(0)$.
- Thereafter the players alternate, producing the infinite sequence

$$\alpha(0),\ \beta(0),\ \alpha(1),\ \beta(1), \alpha(2),\ \beta(2),\ \cdots$$

- *Winning condition given by "specification" $\varphi(X, Y)$:* Éloïse wins just if $P, Q \vDash \varphi(X, Y)$ holds where $\lceil P \rceil = \alpha$ and $\lceil Q \rceil = \beta$.

  Fix a language $\mathcal{L}$ for describing φ.

The $\mathcal{L}$-Synthesis Problem: A Modern Game Version Construct a procedure that, given a 2-person infinite 0/1-game with winning condition $\varphi(X, Y)$ in $\mathcal{L}$, decides if there is a (finite-state) winning strategy for Éloïse, and if so, constructs it.

From MSO Specification to Muller Games An important first step to solving Church's Problem: translate the "external" specification φ to an "internal" winning condition of an appropriate game.

1. *From MSOL to Muller Automata*

   Theorem 3.3 (Büchi 1960, McNaughton 1966)**.** *Every MSO-formula $\varphi(X, Y)$ can be transformed to an equivalent* deterministic Muller automaton A_φ *i.e. for each $P_1, P_2 \subseteq \omega$, we have $\varphi(P_1, P_2)$ holds iff the infinite sequence of binary words, $\lceil P_1, P_2 \rceil$, is recognised by A_φ.*

2. *From Muller Automata to Muller Games*
   A deterministic Muller automaton can equivalently be re-presented (by "splitting the labelled transitions") as a *Muller game*.

A Solution to the MSOL-Synthesis Problem Büchi and Landweber solved the Synthesis Problem in 1969, using highly complex machineries.

A Modern Game Perpective:

Theorem 3.4 (Büchi-Landweber 1969)**.** *Given a Muller game with n states, one can effectively determine if Éloïse has a winning strategy from a given state, and if so, construct a finite-state winning strategy using $n! \cdot n$ control states.*

Proof. (Outline)

(i) Transform the MSO-formula $\varphi(X, Y)$ first to a deterministic Muller automaton.

(ii) Re-present the Muller automaton as a Muller game.

(iii) Transform the Muller game to a *parity game* that *simulates* it.

(iv) Solve the parity game; each of its (positional) winning strategy induces a finite-state winning strategy in the simulated Muller game.

□

For an exposition of Church's Synthesis Problem, see Thomas' tutorial Thomas (2008).

Problems

3.1 Consider the following Büchi automaton A:

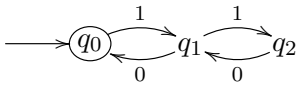

(a) Construct an Existential S1S-formula equivalent to A.

(b) Construct an LTL-formula equivalent to A.

(c) A *star-free ω-regular language* over an alphabet Σ is a finite union of ω-languages of the form $U \cdot V^\omega$, where U and V are (regular) languages constructed from a finite set of finite words over Σ using the Boolean operations, namely, complementation, union and concatenation.

Prove that A recognises a star-free regular ω-language.

3.2 Let A be the following Büchi automaton A:

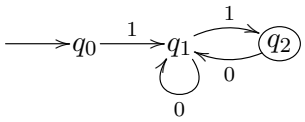

Construct an S1S-formula $\varphi(X)$ such that $\alpha \in \mathbb{B}^\omega$ satisfies φ iff A accepts α.

3.3

(a) An ω-language $L \subseteq \mathbb{B}^\omega$ is said to be *definable* by an S1S formula $\chi(X)$ just if $L = \{ \ulcorner P \urcorner \in \mathbb{B}^\omega \mid P \vDash \chi(X) \}$ where $\ulcorner P \urcorner$ is the characteristic word of $P \subseteq \omega$.

Let p be the only atomic proposition. Prove that for every LTL formula φ there is a formula $\widetilde{\varphi}(X)$ in $\mathrm{S1S}_1$ such that φ and $\widetilde{\varphi}(X)$ define the same ω-language. You should state the predicate symbols you assume in the vocabulary of S1S.

(b) Prove that *equi-cardinality* of sets, that is, the predicate

$$\mathit{EqCard}(A, B) \;:=\; A, B \subseteq \omega \text{ have the same cardinality}$$

cannot be expressed in S1S.

3.4 Give S1S-formulas $\varphi_1(X_1, X_2)$ and $\varphi_2(X_1, X_2)$ for the following ω-languages:

(a) $L_1 = \binom{1}{1}\binom{1}{0}^*\binom{1\ 0}{1\ 0}^\omega$

(b) $L_2 = \binom{1\ 1\ 1}{1\ 1\ 1}\binom{0}{1}^\omega$

Explain the purpose of the main subformulas of $\varphi_1(X_1, X_2)$ and $\varphi_2(X_1, X_2)$.

3.5 Show that (natural numbers) addition $x = y + z$ is not definable in S1S.

[*Hint.* Show that S1S-definability of addition would imply that the language $\{ a^n b^n c^\omega : n \geq 0 \}$ is Büchi recognizable.]

3.6 *Presburger arithmetic* is first-order logic over the structure $(\omega, +)$ where

$$+ \;=\; \{ (a, b, c) \in \omega^3 : a + b = c \}.$$

A number can be represented as a finite set of numbers corresponding to the positions of 1s in its binary representation.

(a) Show that there is an S1S formula $\varphi(X, Y, Z)$ asserting that the numbers a, b and c represented respectively by the finite sets X, Y and Z are related by the equation $a+b = c$.

[*Hint.* Recall the idea of full adder in Digital Hardware.]

(b) Deduce that formulas of Presburger arithmetic can be translated into S1S.

Hence prove that Presburger arithmetic is decidable.

Why does this not contradict the preceding question?

3.7 *Weak monadic second-order theory of one successor*, WS1S, is defined in the same way as S1S except that second-order variables range over only finite sets of natural numbers.

(a) Fix a deterministic Muller automaton A. Since it is not possible to say anything in WS1S about any complete run directly, we restrict ourselves to prefixes of runs. Note that every ω-word that is accepted by a deterministic Muller automaton has a *unique* accepting run.

Give a WS1S-formula that defines the ω-language recognized by A.

(b) Hence deduce that an ω-language is S1S-definable iff it is WS1S-definable.

Chapter 4

Modal Mu-Calculus

Synopsis

Knaster-Tarski fixpoint theorem. Syntax and semantics. Syntactic approximants via infinitary syntax. Intuitions from examples. A branching-time temporal logic: computational tree logic (CTL).

References (Bradfield and Stirling, 2001, 2007; Stirling, 2001, 1997) For a primer on ordinals and Knaster-Tarski Theorem, see (Kozen, 2006, pp. 35-43).

Background Modal mu-calculus's defining feature – use of least and greatest fixpoint operators. The idea goes back a long way:

- Fixpoints in program logics: de Bakker, Park and Scott (late 60s).
- Fixpoints in modal logics of programs: Pratt (1980), Emerson and Clark (1980), Kozen (1983).

Formulas of modal mu-calculus are notoriously hard to read. A good intuitive appreciation is essential for understanding the theory.

4.1 Knaster-Tarski Fixpoint Theorem

Posets, Supremums and Infimums: A Revision A *partially-ordered set* is a pair $\langle L, \leq \rangle$ such that $\leq$ is a binary relation over L that is

(i) reflexive: for every $x \in L$, $x \leq x$

(ii) antisymmetric: for every $x, y \in L$, if $x \leq y$ and $y \leq x$ then $x = y$

(iii) transitive: for every $x, y, z \in L$, if $x \leq y$ and $y \leq z$ then $x \leq z$

Let $M \subseteq L$. An element $l \in L$ is the *least upper bound* (LUB, or supremum) of M, written $\bigvee M$, just if:

(i) for all $x \in M$, $x \leq l$

(ii) for all $y \in L$, if for all $x \in M$ we have $x \leq y$, then $l \leq y$.

Similarly for *greatest lower bound* (GLB, or infimum), written $\bigwedge M$.

Complete lattices and monotone functions A *complete lattice* is a partially-ordered set $\langle L, \leq \rangle$ in which every subset $M \subseteq L$ has a least upper bound $\bigvee M$ and a greatest lower bound $\bigwedge M$ in L. Every such L has a greatest $\bigvee L$ $(= \bigwedge \varnothing)$ and least element $\bigwedge L$ $(= \bigvee \varnothing)$.

Example 4.1. (i) Is $\langle \omega, \leq \rangle$ a complete lattice? No, because $\bigvee \omega = \omega \notin \omega$.

(ii) $\langle \mathcal{P}(S), \subseteq \rangle$ is a complete lattice. For every $\mathcal{U} \subseteq \mathcal{P}(S)$, we have

$$\bigvee \mathcal{U} \;=\; \bigcup \mathcal{U} \;:=\; \{\, s \in S : \exists U \in \mathcal{U} \,.\, s \in U \,\}.$$

The least and greatest elements are $\varnothing$ and S respectively.

A function from L to L is said to be *monotone* just if $f(x) \leq f(y)$ whenever $x \leq y$. An element $x \in L$ is a *fixpoint* of f just if $f(x) = x$. We say x is a *postfixed point* of f just if $x \leq f(x)$; x is a *prefixed point* of f if $f(x) \leq x$.

Example 4.2. (i) Reals: $\langle \mathbb{R}, \leq \rangle$

(ii) $\langle \mathbb{N} \cup \{\, \omega \,\}, \leq \rangle$

(iii) Divisibility: $\langle \mathbb{Z} \setminus \{\, 0 \,\}, \_ \mid \_ \rangle$ where $x \mid y := \exists u \,.\, x \times u = y$.

(iv) Words ordered by prefix: $\langle \Sigma^*, \leq_{\text{pref}} \rangle$. E.g. $a\,b <_{\text{pref}} a\,b\,b\,a$.

(v) Lexicographical: $\langle \Sigma^*, \leq_{\text{lexico}} \rangle$. E.g. $a\,b\,b\,a <_{\text{lexico}} a\,b\,c$ (assuming that Σ is linearly ordered).

(vi) Subword: $\langle \Sigma^*, \leq_{\text{subw}} \rangle$, with $a\,b\,a \leq_{\text{subw}} b\,a\,a\,b\,b\,a\,a$

	Poset	Complete Lattice
$\langle \mathbb{R}, \leq \rangle$	Y	N
$\langle \mathbb{N} \cup \{\, \omega \,\}, \leq \rangle$	Y	Y
$\langle \mathbb{Z} \setminus \{\, 0 \,\}, \_ \mid \_ \rangle$	N	N
$\langle \Sigma^*, \leq_{\text{pref}} \rangle$	Y	N
$\langle \Sigma^*, \leq_{\text{lexico}} \rangle$	Y	N
$\langle \Sigma^*, \leq_{\text{subw}} \rangle$	Y	N

Least prefixed and greatest postfixed points

Lemma 4.1. *Let L be a complete lattice and $f : L \to L$ be a monotone function.*

(i) The least prefixed point *of f, denoted $lpr(f)$, exists, and is $\bigwedge\{\, x \in L : f(x) \leqslant x \,\}$.*

(ii) The greatest postfixed point *of f exists and is $\bigvee\{\, x \in L : x \leqslant f(x) \,\}$.*

Proof. (i) Let $pref(f) := \{\, x \in L : f(x) \leqslant x \,\}$ be the set of prefixed points of f. It suffices to show that $\bigwedge pref(f)$ is a prefixed point. Let $x \in pref(f)$. Then $\bigwedge pref(f) \leq x$. Since f is monotone we have $f(\bigwedge pref(f)) \leq f(x)$, but $f(x) \leq x$ because $x \in pref(f)$. Hence $f(\bigwedge pref(f)) \leq x$ for every $x \in pref(f)$. Since $f(\bigwedge pref(f))$ is an lower bound, we have $f(\bigwedge pref(f)) \leq \bigwedge pref(f)$ as required. (ii) Exercise. □

Construction of fixpoints Let $\langle L, \leqslant \rangle$ be a complete lattice, and $f : L \longrightarrow L$ is monotone. Define, by transfinite induction, a family of elements of L, indexed by *ordinals*:

$$\begin{array}{rcl} f^{\alpha+1} & := & f(f^\alpha) \\ f^\lambda & := & \bigvee_{\alpha<\lambda} f^\alpha \quad \text{for } \lambda \text{ a limit ordinal} \end{array}$$

We then set $f^* := \bigvee_{\alpha \in \mathbf{Ord}} f^\alpha$.

(Base case — $f^0 = \bot$ — is included in the case for limit ordinal.)

Lemma 4.2. *If $\alpha \leq \beta$ then $f^\alpha \leq f^\beta$.*

Proof. We prove by transfinite induction on α. Two cases:

(i) $\alpha = \alpha_0 + 1$ is successor ordinal. Two cases of β:

- If $\beta = \beta_0 + 1$ then $f^\alpha := f(f^{\alpha_0}) \leq f(f^{\beta_0}) = f^\beta$, by monotonicity of f and IH.
- If $\beta = \bigvee_{\beta_0<\beta} \beta_0$, then $\alpha_0 \leq \beta_0$ for some $\beta_0 < \beta$, and so by monotonicity of f and IH

$$f^\alpha := f^{\alpha_0+1} \leq f^{\beta_0+1} \leq \bigvee_{\delta<\beta} f^\delta = f^\beta.$$

(ii) α is a limit ordinal. For each $\alpha_0 < \alpha < \beta$, by IH, $f^{\alpha_0} \leq f^\beta$. Because f^β is an upper bound, we have

$$f^\alpha := \bigvee_{\alpha_0<\alpha} f^{\alpha_0} \leq f^\beta.$$

□

Thanks to the Lemma, we have a chain:

$$\bot = f^0 \leq f^1 \leq f^2 \leq \cdots \leq f^n \leq \cdots \leq f^\omega \leq \cdots$$

Since **Ord** is a class (not a set), the map **Ord** $\longrightarrow L$ defined by $\alpha \mapsto f^\alpha$ cannot be injective, and so, the chain must "plateau out" at some point.

The *closure ordinal* of f is defined to be the smallest ordinal κ such that $f^\kappa = f^{\kappa+1}$.

Hence $f^* = f^\kappa$ where κ is the closure ordinal.

Theorem 4.1 (Knaster-Tarski)**.** *Let $\langle L, \leqslant \rangle$ be a complete lattice. If $f : L \longrightarrow L$ is a monotone function, the* least prefixed point *of f, written $lpr(f)$, is f^*.*

Proof. Recall $f^* := \bigvee_{\alpha \in \mathbf{Ord}} f^\alpha = f^\kappa$, where κ is the closure ordinal.

"$lpr(f) \leq f^*$": It suffices to prove that f^* is a prefixed point of f. We have $f(f^\kappa) = f^{\kappa+1} = f^\kappa$.

"$lpr(f) \geq f^*$": We shall prove by transfinite induction that $f^\alpha \leq lpr(f)$, for all ordinals α. This is sufficient, since $f^* := \bigvee_{\alpha \in \mathbf{Ord}} f^\alpha \leq lpr(f)$. For successor ordinals $\alpha + 1$:

$$\begin{array}{rcll} f^{\alpha+1} & = & f(f^\alpha) & \\ & \leq & f(lpr(f)) & \text{induction hypothesis} \\ & \leq & lpr(f) & \text{definition of } lpr(f) \end{array}$$

For limit ordinals λ, we have $f^\alpha \leq lpr(f)$ for all $\alpha < \lambda$ by the IH; therefore

$$f^\lambda := \bigvee_{\alpha<\lambda} f^\alpha \leq lpr(f)$$

□

Exercise 4.1. State and prove a corresponding version of the Theorem for greatest postfixed points.

Lemma 4.3. *$lpr(f)$ exists, and coincides with the* least fixpoint *of f, denoted $lfp(f)$. Similarly greatest fixpoint exists and coincides with greatest postfixed point.*

Proof. By definition $f(lpr(f)) \leq lpr(f)$. Since f monotone, $f(lpr(f))$ is also a prefixed point. Hence $lpr(f) \leq f(lpr(f))$. I.e. $lpr(f)$ is a fixpoint of f. That $lpr(f)$ is the least fixpoint follows from the fact that every fixpoint is also a prefixed point. □

Fixpoints as recursion Given a state transition system (graph) $\langle\, S, R \subseteq S \times S \,\rangle$. We give the semantics of a basic state-based modal logic by the mapping

$$\varphi \;\mapsto\; \|\varphi\| \;:=\; \{\, s \in S : s \vDash \varphi \,\}$$

I.e. denotation of a formula is an element of $\mathcal{P}(S)$. Hence $\varphi(Z)$, with a free second-order variable Z that ranges over $\mathcal{P}(S)$, can be viewed as a function $f_\varphi : \mathcal{P}(S) \longrightarrow \mathcal{P}(S)$. Recall:

(i) $\langle\, \mathcal{P}(S), \subseteq \,\rangle$ is a complete lattice.

(ii) If $f_\varphi : \mathcal{P}(S) \longrightarrow \mathcal{P}(S)$ is monotone, by Knaster-Tarski, f_φ has unique least and greatest fixpoints, denoted μf_φ and νf_φ respectively.

Thus we can extend modal logic by

- *least fixpoint operator*, $\mu Z.\varphi(Z)$, interpreted as μf_φ, and
- *greatest fixpoint operator*, $\nu Z.\varphi(Z)$, interpreted as νf_φ.

Fixpoints give a semantics of recursion, as in domain theory. Recursive modal logic formulas give *succinct* expressions of the usual operators of temporal logic.

4.2 Syntax of the Modal Mu-Calculus

Given

- Var, a set of 2nd-order variables, ranged over by X, Y, Z, etc.
- $Prop$, a set of atomic propositions, ranged over by P, Q, etc.
- $\mathcal{L}$, a set of labels, ranged over by a, b, etc.

modal mu-calculus formulas are defined by the grammar:

$$\varphi \quad ::= \quad P \;\mid\; Z \;\mid\; \varphi_1 \wedge \varphi_2 \;\mid\; [a]\varphi \;\mid\; \neg\varphi \;\mid\; \nu Z.\varphi$$

The *last case*, namely, formation of $\nu Z.\varphi$, is subject to the requirement that each free occurrence of Z in φ be *positive* i.e. in the scope of an even number of negations (so that $\varphi(Z)$ denotes a monotone function in Z).

Notation. If φ is written $\varphi(Z)$, subsequent writing of $\varphi(\psi)$ means "φ with ψ substituted for all free occurrences of Z".

Positive, and positive normal forms *Derived operators*:

$$\begin{array}{rcl} \varphi_1 \vee \varphi_2 & := & \neg(\neg\varphi_1 \wedge \neg\varphi_2) \\ \langle a\rangle\varphi & := & \neg([a]\neg\varphi) \\ \mu Z.\varphi & := & \neg\nu Z.\neg\varphi(\neg Z) \end{array}$$

A modal mu-calculus formula is in *positive form* if it is written, using derived operators where necessary, so that $\neg$ is always applied to atomic propositions. A formula φ is in *positive normal form* if, in addition, all bound variables are distinct. I.e. if $\sigma X.\psi$ and $\sigma' Y.\chi$ are distinct subterms of φ (where $\sigma X.-$ and $\sigma' Y.-$ are fixpoint operators), then $X \neq Y$. E.g. $\mu Z.\neg P \vee \nu Y.(Y \wedge \langle a\rangle Z)$

Operator precedence

$$\begin{matrix} [a]- \\ \langle a\rangle- \end{matrix} \quad > \quad \text{Booleans} \quad > \quad \begin{matrix} \mu Z.- \\ \nu Z.- \end{matrix}$$

Reading: If $[a]-$ and $\wedge$ contend for a formula, then $[a]-$ wins. For example $[a]\varphi \wedge \psi$ means $([a]\varphi) \wedge \psi$, and $\mu Z.Z \wedge \varphi$ means $\mu Z.(Z \wedge \varphi)$. Thus the scope of a fixpoint extends as far right as possible.

Note that every formula can be converted to positive normal form using *de Morgan laws*, and α*-conversion*: replacing a bound name by a fresh name.

4.3 Labelled Transition Systems

A modal mu-calculus structure over $(Prop, \mathcal{L})$ is a *labelled transition system* (LTS), namely, a triple $T = \langle\, S, \to, \rho \,\rangle$ with

- a set S of states
- a transition relation $\to \subseteq S \times \mathcal{L} \times S$ (as usual we write $s \xrightarrow{a} t$ to mean $(s, a, t) \in \to$)
- a function $\rho : Prop \longrightarrow \mathcal{P}(S)$ interpreting the atomic propositions.

Observe that an LTS is just a directed graph, whose edges are labelled by elements of $\mathcal{L}$, and vertices are labelled by elements of $\mathcal{P}(Prop)$ i.e. each $x \in S$ is labelled by $\{\, P \in Prop : x \in \rho(P) \,\}$.

Definition 4.1. Given an LTS T, and a *valuation* (or *assignment*) $V : Var \longrightarrow \mathcal{P}(S)$, we define:

$$\begin{array}{rcl} \|P\|^T_V & := & \rho(P) \\ \|Z\|^T_V & := & V(Z) \\ \|\neg\varphi\|^T_V & := & S \setminus \|\varphi\|^T_V \\ \|\varphi_1 \wedge \varphi_2\|^T_V & := & \|\varphi_1\|^T_V \cap \|\varphi_2\|^T_V \\ \|[a]\varphi\|^T_V & := & \{\, s \mid \forall t \in S \,.\, s \xrightarrow{a} t \implies t \in \|\varphi\|^T_V \,\} \\ \|\nu Z.\varphi\|^T_V & := & \bigcup\{\, U \subseteq S \mid U \subseteq \|\varphi\|^T_{V[Z\mapsto U]} \,\} \end{array}$$

where the valuation $V[Z \mapsto U]$ is defined by:

$$V[Z \mapsto U](X) \;:=\; \begin{cases} U & \text{if } X = Z \\ V(X) & \text{if } X \neq Z. \end{cases}$$

N.B. $\bigcup\{\, U_i \subseteq S : i \in I \,\} := \{\, x \in S : x \in U_i \text{ for some } i \in \mathcal{I} \,\}$

Remark 4.1. (i) Generally let $K \subseteq \mathcal{L}$, define

$$\|[K]\varphi\|_V^T := \{\, s \mid \forall a \in K\,.\,\forall t \in S\,.\, s \xrightarrow{a} t \implies t \in \|\varphi\|_V^T \,\}$$

Write $[-]\varphi$ to mean $[\mathcal{L}]\varphi$, similarly for $\langle - \rangle \varphi$.

(ii) Equivalently we can define

$$\|\nu Z.\varphi\|_V^T := gfp(f_{\varphi,Z,V})$$

where $f_{\varphi,Z,V} : \mathcal{P}(S) \longrightarrow \mathcal{P}(S)$ is the function $U \mapsto \|\varphi\|^T_{V[Z\mapsto U]}$.
Note that $f_{\varphi,Z,V}$ is monotone. By two lemmas (greatest fixpoint equals greatest postfixed point, which is the supremum of all postfixed points), we have

$$\begin{aligned} gfp(f_{\varphi,Z,V}) &= \bigvee\{\, U \in \mathcal{P}(S) : U \leq f_{\varphi,Z,V}(U) \,\} \\ &= \bigcup\{\, U \subseteq S : U \subseteq \|\varphi\|^T_{V[Z\mapsto U]} \,\}. \end{aligned}$$

Exercise 4.2. Prove the following.

$$\begin{aligned} \|\varphi_1 \vee \varphi_2\|_V^T &= \|\varphi_1\|_V^T \cup \|\varphi_2\|_V^T \\ \|\langle a \rangle \varphi\|_V^T &= \{\, s \mid \exists t \in S\,.\, s \xrightarrow{a} t \ \wedge\ t \in \|\varphi\|_V^T \,\} \\ \|\mu Z.\varphi\|_V^T &= \bigcap\{\, U \subseteq S \mid \|\varphi\|^T_{V[Z\mapsto U]} \subseteq U \,\} \\ &= \bigwedge\{\, U \in \mathcal{P}(S) : f_{\varphi,Z,V}(U) \leq U \,\} = lfp(f_{\varphi,Z,V}) \end{aligned}$$

Note that $\bigcap\{\, U_i \subseteq S : i \in I \,\} := \{\, x \in S : x \in U_i \text{ for every } i \in \mathcal{I} \,\}$

Since $\mu Z.\varphi = \neg\nu Z.\neg\varphi(\neg Z)$, we have

$$\begin{aligned} \|\mu X.\varphi\|_V^T &= \overline{\bigcup\{\, U : U \subseteq \|\neg\varphi(\neg Z)\|^T_{V[Z\mapsto U]} \,\}} \\ &= \bigcap\{\, \overline{U} : U \subseteq \|\neg\varphi(Z)\|^T_{V[Z\mapsto \overline{U}]} \,\} \\ &= \bigcap\{\, \overline{U} : \|\varphi(Z)\|^T_{V[Z\mapsto \overline{U}]} \subseteq \overline{U} \,\} \\ &= \bigcap\{\, U : \|\varphi(Z)\|^T_{V[Z\mapsto U]} \subseteq U \,\} \end{aligned}$$

Notations

(i) We sometimes write $s \vDash^T_V \varphi := s \in \|\varphi\|_V^T$.
In case T is understood, and φ is closed, we simply write $s \vDash \varphi$.

(ii) We often write $\mathsf{t} := \nu Z.Z$ and $\mathsf{f} := \neg\mathsf{t}$.
What are $\|\nu Z.Z\|_V^T$ and $\mathsf{f} \equiv \mu Z.Z$?

(iii) For closed formulas φ, and ψ, we write $\varphi \equiv \psi$ to mean "for all LTS T, $\|\varphi\|^T_\varnothing = \|\psi\|^T_\varnothing$".

4.4 Syntactic Approximants Using Infinitary Syntax

The fixpoint formulas of the modal mu-calculus denote fixpoints of monotone functions in a complete lattice. When reasoning about fixpoints, it is convenient to introduce a notation for approximants of these fixpoints. To this end, we introduce an infinitary syntax.

Let λ range over limit ordinals

$$\begin{aligned} \mu^0 Z.\varphi(Z) &:= \mathsf{f} \\ \mu^{\alpha+1} Z.\varphi(Z) &:= \varphi(\mu^\alpha Z.\varphi(Z)) \\ \mu^\lambda Z.\varphi(Z) &:= \bigvee_{\alpha<\lambda} \mu^\alpha Z.\varphi(Z) \end{aligned}$$

The semantics of these infinitary terms are defined as:

$$\begin{aligned} \|\mu^{\alpha+1}Z.\varphi(Z)\|^T_V &:= \|\varphi(Z)\|^T_{V[Z\mapsto\|\mu^\alpha Z.\varphi(Z)\|^T_V]} \\ \|\mu^{\lambda}Z.\varphi(Z)\|^T_V &:= \bigvee_{\alpha<\lambda}\|\mu^\alpha Z.\varphi(Z)\|^T_V \end{aligned}$$

Thus $\|\mu Z.\varphi(Z)\|^T_V = \|\mu^\kappa Z.\varphi(Z)\|^T_V$ where κ is the closure ordinal. If $\|\mu^\alpha Z.\varphi(Z)\|^T_V = \|\mu^{\alpha+1}Z.\varphi(Z)\|^T_V$ then $\|\mu^\alpha Z.\varphi(Z)\|^T_V = \|\mu Z.\varphi(Z)\|^T_V$.

Similarly for approximants of the greatest fixpoints:

$$\begin{aligned} \nu^0 Z.\varphi(Z) &:= \mathsf{t} \\ \nu^{\alpha+1}Z.\varphi(Z) &:= \varphi(\nu^\alpha Z.\varphi(Z)) \\ \nu^{\lambda}Z.\varphi(Z) &:= \bigwedge_{\alpha<\lambda}\nu^\alpha Z.\varphi(Z) \end{aligned}$$

and

$$\begin{aligned} \|\nu^{\alpha+1}Z.\varphi(Z)\|^T_V &:= \|\varphi(Z)\|^T_{V[Z\mapsto\|\nu^\alpha Z.\varphi(Z)\|^T_V]} \\ \|\nu^{\lambda}Z.\varphi(Z)\|^T_V &:= \bigwedge_{\alpha<\lambda}\|\nu^\alpha Z.\varphi(Z)\|^T_V \end{aligned}$$

Note that the infinitary terms $\nu^{\alpha+1}Z.\varphi(Z), \nu^\lambda Z.\varphi(Z), \mu^{\alpha+1}Z.\varphi(Z)$, etc. are *not* part of the modal mu-calculus. They are notations denoting subsets of the state-set that are useful for calculations.

Lemma 4.4 (Approximation)**.** *Let $T = \langle S, \to, \rho\rangle$ be an LTS. For any $s \in S$, we have*

(i) $s \in \|\mu Z.\varphi(Z)\|^T_V$ iff $s \in \|\mu^\alpha Z.\varphi(Z)\|^T_V$ for some α.

(ii) $s \in \|\nu Z.\varphi(Z)\|^T_V$ iff $s \in \|\nu^\alpha Z.\varphi(Z)\|^T_V$ for all α.

Proof. (i) We have $\|\mu Z.\varphi(Z)\|^T_V = lfp(f_{\varphi,Z,V})$ by definition. Note that for each ordinal α, $\|\mu^\alpha Z.\varphi(Z)\|^T_V$ coincides with $f^\alpha_{\varphi,Z,V}$ i.e. the element of $\mathcal{P}(S)$ in the chain $f^0_{\varphi,Z,V}, f^1_{\varphi,Z,V}, \cdots$ indexed by α. Hence, by Knaster-Taski

$$\|\mu Z.\varphi(Z)\|^T_V = \bigvee_{\alpha\in\mathbf{Ord}} f^\alpha_{\varphi,Z,V} = \bigcup_{\alpha\in\mathbf{Ord}} \|\mu^\alpha Z.\varphi(Z)\|^T_V.$$

(ii) Exercise

□

Lemma 4.5. *Let $T = \langle S, \to, \rho\rangle$ be an LTS. For any $s \in S$, we have*

(i) If $s \in \|\mu Z.\varphi(Z)\|^T_V$ then there is a least ordinal α such that $s \in \|\mu^\alpha Z.\varphi(Z)\|^T_V$ but $s \notin \|\mu^\beta Z.\varphi(Z)\|^T_V$ for all $\beta < \alpha$.

(ii) If $s \notin \|\nu Z.\varphi(Z)\|^T_V$ then there is a least ordinal α such that $s \notin \|\nu^\alpha Z.\varphi(Z)\|^T_V$ but $s \in \|\mu^\alpha Z.\varphi(Z)\|^T_V$ for all $\beta < \alpha$.

Proof. The sequence

$$\|\mu^0 Z.\varphi\|^T_V \subseteq \|\mu^1 Z.\varphi\|^T_V \subseteq \|\mu^2 Z.\varphi\|^T_V \subseteq \cdots$$

is an increasing chain in $\mathcal{P}(S)$, whose supremum is $\|\mu Z.\varphi\|^T_V$ by Knaster-Tarski. Since $\|\mu Z.\varphi\|^T_V = \bigcup_\gamma \|\mu^\gamma Z.\varphi\|^T_V$, if $s \in \|\mu Z.\varphi\|^T_V$, the set $\{\,\beta \in \mathbf{Ord} : s \in \|\mu^\beta Z.\varphi\|^T_V\,\}$ is non-empty. Since the class of ordinals is *well-founded*[1], the set has a least element α (say). It follows that $s \in \|\mu^\alpha Z.\varphi\|^T_V$ and $s \notin \|\mu^\beta Z.\varphi\|^T_V$ for all $\beta < \alpha$. □

[1] A binary relation, R, is *well-founded* on a class X just if every non-empty subset of X has a minimal element with respect to R.

4.5 Intuitions from Examples

Correctness properties of reactive systems are often classified into safety or liveness properties. Intuitively

- *safety properties* say that "something bad will never happen"
- *liveness properties* say that "something good will eventually happen".

Useful Slogans

1. "ν is looping, whereas μ is *finite* looping"
2. "μ is liveness and ν is safety".

Example 4.3 ("ν is looping"). (i) $\nu Z.P \wedge [a]Z$ relativized 'always' formula. "P is true along every a-path".

(ii) $\nu Z.Q \vee (P \wedge [a]Z)$ relativized 'while' formula. "On every a-path, P holds while Q fails". The formula is true if either Q holds, or P holds and wherever we go next (via a), the formula is true, and .... In particular, if P is always true, and Q never holds, the formula is true. *Cf.* $\mu Z.Q \vee (P \wedge \langle a\rangle Z)$—next slide.

Mu-formulas require something to happen (i.e. to exit the loop), and are thus liveness properties.

Example 4.4 ("μ is finite looping"). (i) $\mu Z.P \vee [a]Z$: "On all a-paths, P eventually holds".

(ii) $\mu Z.Q \vee (P \wedge \langle a\rangle Z)$: "On some a-path, P holds until Q holds (and Q *must* eventually hold)." I.e. we are not allowed to repeat the unfolding forever, so we must eventually "bottom out" through the Q disjunct.

Example 4.5. Consider the simple transition system

$$T_1 = \quad s \circlearrowright a$$

with state-set $S = \{ s \}$ and a single transition $s \xrightarrow{a} s$. We have $s \vDash \nu Z.(\langle a\rangle Z \vee [a]\mathsf{f})$. *Intuitively* this is because from s it is always possible to do an a-transition and then an a-transition and then an a-transition ... *ad infinitum.*

Example 4.6. Take T_1 as before. We show that $s \nvDash \mu Z.(\langle a\rangle Z \vee [a]\mathsf{f})$. If s satisfies $\mu Z.(\langle a\rangle Z \vee [a]\mathsf{f})$ to hold, then it must be possible for s to do a finite[2] number of a-transitions and then satisfy $[a]\mathsf{f}$, but the latter is impossible since there is always an a-transition from s. I.e. s can never satisfy $\underbrace{\langle a\rangle(\cdots(\langle a\rangle}_{\text{finite times}}([a]\mathsf{f})))$. Now we calculate. Note that $\mu^1 Z.(\langle a\rangle Z \vee [a]\mathsf{f}) = (\langle a\rangle\mathsf{f} \vee [a]\mathsf{f})$. Now $\|\langle a\rangle\mathsf{f}\| = \varnothing$ and $\|[a]\mathsf{f}\| = \varnothing$. So $\|\mu^1 Z.(\langle a\rangle Z \vee [a]\mathsf{f})\| = \varnothing$. Similarly $\|\mu^2 Z.(\langle a\rangle Z \vee [a]\mathsf{f})\| = \varnothing$. Hence $\|\mu Z.(\langle a\rangle Z \vee [a]\mathsf{f})\| = \varnothing$

Alternatively consider the function

$$\begin{aligned} f \,:\, U &\mapsto \|\langle a\rangle Z \vee [a]\mathsf{f}\|^T_{\varnothing[Z\mapsto U]} \\ &= \|\langle a\rangle Z\|^T_{\varnothing[Z\mapsto U]} \cup \|[a]\mathsf{f}\| \\ &= U \cup \varnothing \end{aligned}$$

which is the identity map on $\mathcal{P}(S)$. Hence $\|\mu Z.(\langle a\rangle Z \vee [a]\mathsf{f})\| = lfp(f) = \varnothing$.

[2] Because S is finite, the closure ordinal of any monotone function on $\mathcal{P}(S)$ is finite.

Example 4.7. Take the transition system

$$T_2 \;=\; 0 \underset{a}{\overset{a}{\rightleftarrows}} 1 \xrightarrow{b} 2$$

with state-set $S = \{0, 1, 2\}$. We show that $0 \vDash \nu Z.\mu Y.[a]((\langle b\rangle \mathtt{t} \wedge Z) \vee Y)$. First note $\|\langle b\rangle \mathtt{t}\| = \{1\}$. Set

$$F(U) \;=\; \|\mu Y.[a]((\langle b\rangle \mathtt{t} \wedge Z) \vee Y)\|_{[Z\mapsto U]}.$$

W.t.p. $0 \in \mathit{gfp}(F)$. By abuse of notation, we have

$$F\{0, 1, 2\} = F\{0, 1\} = F\{1\} = \|\mu Y.[a](\{1\} \vee Y)\|.$$

Intuitively $s \in \|\mu Y.[a](\{1\} \vee Y)\|$ if and only if every a-labelled path from s reaches 1. What is $\|\mu Y.[a](\{1\} \vee Y)\|$? We have

$$\begin{aligned}
\|\mu^1 Y.[a](\{1\} \vee Y)\| &= \{0, 2\} \\
\|\mu^2 Y.[a](\{1\} \vee Y)\| &= [a](\{1\} \vee \{0, 2\}) = \{0, 1, 2\} \\
\|\mu^3 Y.[a](\{1\} \vee Y)\| &= [a](\{1\} \vee \{0, 1, 2\}) = \{0, 1, 2\}.
\end{aligned}$$

I.e. $F\{0, 1, 2\} = \{0, 1, 2\}$. Hence $\mathit{gfp}(F) = \{0, 1, 2\}$.

4.6 Alternation Depth Hierarchy

Consider the following modal mu-calculus formulas.

- $\mu X.\varphi \vee [-]X$: all paths eventually satisfy φ
- $\nu Y.[\mu Z.P \vee \langle -\rangle Z] \wedge \langle -\rangle Y$: there exists an infinite path along which P is always reachable
- $\nu Y \mu Z.(P \vee \langle -\rangle Z) \wedge \langle -\rangle Y$: there exists a path along which P holds infinitly often

The fixed point quantifiers provide great expressive power. Liveness and safety can be expressed readily and by allowing more nested fixpoint quantifiers (depth) we can express complex fairness constraints. Thus it natural to define a measure of expressiveness in this term.

The *alternation depth* of a formula is the maximum number of μ/ν alternations in a chain of nested fixed points. The *simple hierarchy* counts the syntactic alternation. The *Niwiński hierarchy* considers genuine dependency between the fixed points.

Formally the Niwiński hierarchy is defined as follows. A formula φ is in the classes Π_0^μ and Σ_0^μ if it contains no fixpoint operators i.e. it is a formula of modal logic. The class Σ_{n+1}^μ is the closure of $\Sigma_n^\mu \cup \Pi_n^\mu$ under the following rules.

- If $\varphi, \psi \in \Sigma_{n+1}^\mu$, then $\varphi \wedge \psi, \varphi \vee \psi, [a]\varphi, \langle a\rangle\varphi \in \Sigma_{n+1}^\mu$.
- If $\varphi \in \Sigma_{n+1}^\mu$ and X positive in φ, then $\mu X.\varphi \in \Sigma_{n+1}^\mu$.
- If $\varphi(X), \psi \in \Sigma_{n+1}^\mu$, then $\varphi(\psi) \in \Sigma_{n+1}^\mu$, provided no free variable of ψ becomes bound by a fixpoint quantifier in φ.

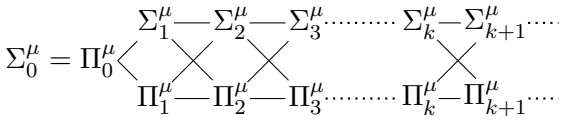

Figure 4.1: Alternation-depth hierarchy.

The class Π_{n+1}^μ is defined analogously.

It is known that indeed arbitrary alternation is necessary to capture all expressible properties, in other words the modal mu-calculus alternation hierarchy is strict (Bradfield 1996).

There is a price to pay for this expressive power and that is complexity. The most notable open problem in modal mu-calculus is the complexity of model checking.

Model Checking Problem: *Given an LTS T, a state s, and a closed modal mu-calculus formula φ, does $s \in \|\varphi\|_\varnothing^T$ hold?*

The best algorithms to date are all essentially exponential in the depth of the formula φ and the best known bound is $\mathbf{NP} \cap \mathbf{co\text{-}NP}$. The problem is conjectured to be in $\mathbf{P}$. The harder problem of satisfiability is known to be **EXPTIME**-complete.

Satisfiability Problem: *Given a closed modal mu-calculus formula φ, is it satisfiable?* I.e. is there an LTS T with a state s such that $s \in \|\varphi\|_\varnothing^T$?

Although modal mu-calculus is a decidable logic, it is too expressive for real use. Its importance comes mainly from providing a meta-language to establish results about other temporal logics. Most temporal logics such as LTL and CTL and their extension can be interpreted into modal mu-calculus using two nested quantifiers.

4.7 An Interlude: Computational Tree Logic (CTL)

$$\varphi \ ::= \ P \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid [K]\varphi \mid \mathbf{A}(\varphi \,\mathbf{U}\, \psi) \mid \mathbf{E}(\varphi \,\mathbf{U}\, \psi)$$

where $P \in \mathit{Prop}$ and $K \subseteq \mathcal{L}$.

CTL is a *branching time* logic; LTL is a linear-time logic.

Semantics of CTL We interpret CTL formulas in the same structures as modal mu-calculus. Let $T = \langle S, \longrightarrow \subseteq S \times \mathcal{L} \times S, \rho : \mathit{Prop} \longrightarrow 2^S \rangle$ be an LTS.

$$\begin{array}{rcl}
s \vDash P & := & s \in \rho(P) \\
s \vDash \neg\varphi & := & s \nvDash \varphi \\
s \vDash \varphi_1 \wedge \varphi_2 & := & s \vDash \varphi_1 \text{ and } s \vDash \varphi_2 \\
s \vDash [K]\varphi & := & \text{for all finite or infinite runs } s \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \cdots, \\
 & & \text{if } a_1 \in K, \text{ then } s_1 \vDash \varphi. \\
s \vDash \mathbf{A}(\varphi\,\mathbf{U}\,\psi) & := & \text{for all finite or infinite runs } s = s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \cdots, \\
 & & \text{there is } i \geq 0 \text{ with } s_i \vDash \psi \text{ and for all } 0 \leq j < i,\ s_j \vDash \varphi \\
s \vDash \mathbf{E}(\varphi\,\mathbf{U}\,\psi) & := & \text{for some finite or infinite runs } s = s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} s_2 \cdots, \\
 & & \text{there is } i \geq 0 \text{ with } s_i \vDash \psi \text{ and for all } 0 \leq j < i,\ s_j \vDash \varphi
\end{array}$$

In LTL, we write $\mathbf{F}\,\varphi := \mathsf{t}\,\mathbf{U}\,\varphi$ and $\mathbf{G}\,\varphi := \neg(\mathbf{F}\,\neg\varphi)$. In CTL, we write $\mathbf{A}\,\mathbf{F}\,\varphi := \mathbf{A}(\mathsf{t}\,\mathbf{U}\,\varphi)$ and $\mathbf{A}\,\mathbf{G}\,\varphi := \neg\mathbf{E}(\mathsf{t}\,\mathbf{U}\,\neg\varphi)$; similarly for $\mathbf{E}\,\mathbf{F}\,\varphi$ and $\mathbf{E}\,\mathbf{G}\,\varphi$.

Example 4.8. (i) $s \vDash \mathbf{E}\,\mathbf{F}\,P$. There is a path from s on which P is eventually true.

(ii) $s \vDash \mathbf{A}\,\mathbf{F}\,\mathbf{E}\,\mathbf{G}\,P$. On every path from s, there is a state from which there is a path where P holds everywhere.

(iii) $s \vDash \mathbf{E}\,\mathbf{G}\,\mathbf{A}\,\mathbf{F}\,P$. There is a path from s such that every state t on it satisfies the property that on every path from t, P is eventually true.

Example 4.9. What does the mu-formula $\mu Z.[\mathcal{L}]Z$ mean? We use approximants to analyse the formula.

- $\mu^0 Z.[\mathcal{L}]Z := \varnothing$
- $\mu^1 Z.[\mathcal{L}]Z := [\mathcal{L}]\varnothing$, which is the set of terminal or "deadlocked" states.
- $\cdots$
- $\mu^{i+1} Z.[\mathcal{L}]Z := [\mathcal{L}]\cdots[\mathcal{L}]\varnothing$ ($i+1$ nested boxes), which is the set of states s such that every path from s has length at most i and necessarily ends at a "deadlocked" state.

Hence $\mu Z.[\mathcal{L}]Z$ describes the set of states s of an LTS such that every path from s necessarily ends in a deadlocked state. In fact we have $\mu Z.[\mathcal{L}]Z \equiv \mathbf{A}\,\mathbf{F}\,([\mathcal{L}]\mathsf{f})$.

Example 4.10. *CTL formula* $\forall\,\boldsymbol{G}\,\varphi$: "for every path (or run), φ always holds on it". It is the property X *such that* φ *is true now, and for every successor,* X *remains true.* I.e. X satisfies the modal equation:

$$X \;=\; \varphi \wedge [-]X$$

where $[-]X$ means "X is true at every successor".

Which fixpoint? If a state satisfies any solution of the equation then surely it satisfies $\forall\,\boldsymbol{G}\,\varphi$. So the "equation" should be

$$X \;\Rightarrow\; \varphi \wedge [-]X$$

or more precisely, $\|X\| \;\subseteq\; \|\varphi \wedge [-]X\|$. Thus the meaning is the greatest postfixed point $\nu X.\varphi \wedge [-]X$.

Example 4.11. *CTL formula* $\exists \boldsymbol{F} \varphi$: "there exists a path on which φ eventually holds". It is the property Y *such that either* φ *holds now, or there is some successor on which* Y *is true.* I.e. Y satisfies the modal equation:

$$Y \;=\; \varphi \vee \langle - \rangle Y$$

where $\langle - \rangle Y$ means "Y is true at some successor".

We can argue: if a state satisfies $\exists \boldsymbol{F} \varphi$, then it surely satisfies any solution of the equation; and so, we want the least solution of

$$Y \;\Leftarrow\; \varphi \vee \langle - \rangle Y$$

or more precisely $\|Y\| \;\supseteq\; \|\varphi \vee \langle - \rangle Y\|$. I.e. the meaning is the least prefixed point $\mu Y.\varphi \vee \langle - \rangle Y$.

Problems

4.1 Let φ be a monotone function on the powerset of S. Let $U \subseteq S$. Prove

$$U \subseteq \mathit{gfp}(\varphi) \quad \Leftrightarrow \quad U \subseteq \varphi(\mathit{gfp}(\varphi^U))$$

where $\varphi^U \colon \mathcal{P}(S) \to \mathcal{P}(S)$ is the function given by $V \mapsto U \cup \varphi(V)$.

4.2 Prove the following:

(a) If $s \in \|\mu Z.\varphi\|_V^T$ then there is a least ordinal α such that $s \in \|\mu^\alpha Z.\varphi\|_V^T$, and for all $\beta < \alpha$, $s \notin \|\mu^\beta Z.\varphi\|_V^T$.

(b) If $s \notin \|\nu Z.\varphi\|_V^T$ then there is a least ordinal α such that $s \notin \|\nu^\alpha Z.\varphi\|_V^T$, and for all $\beta < \alpha$, $s \in \|\nu^\beta Z.\varphi\|_V^T$.

4.3 Prove that if the state-set S has n elements, then for any $s \in S$

(a) $s \vDash_V \nu Z.\varphi$ iff $s \in \|\nu^n Z.\varphi\|_V^T$

(b) $s \vDash_V \mu Z.\varphi$ iff $s \in \|\mu^n Z.\varphi\|_V^T$

4.4 We say that φ and ψ are *equivalent* just if for any T and for any V, we have $\|\varphi\|_V^T = \|\psi\|_V^T$. Prove

(a) $\mu Z.\varphi(Z)$ and $\varphi(\mu Z.\varphi)$ are equivalent (similarly $\nu Z.\varphi$ and $\varphi(\nu Z.\varphi)$ are equivalent).

(b) $\mu Z.\varphi(Z)$ and $\neg\nu Z.\neg\varphi(\neg Z)$ are equivalent.

4.5 Is there a labelled transition system T and a valuation V such that for every modal mu-calculus formula φ, $\|\mu Z.\varphi\|_V^T = \|\nu Z.\varphi\|_V^T$?

4.6

(a) Show that $s_0 \in \|\mu Z.[a]Z\|_V^T$ iff there is no infinite transition sequence

$$s_0 \xrightarrow{a} s_1 \xrightarrow{a} s_2 \xrightarrow{a} \cdots$$

(b) What is the meaning of $\nu Z.[a]Z$?

4.7 Consider the transition system

$$T_2 \;=\; 0 \underset{a}{\overset{a}{\rightleftarrows}} 1 \xrightarrow{b} 2$$

with state-set $S = \{\, 0, 1, 2 \,\}$.

(a) Compute $\|\mu Y.\nu Z.[a]((\langle b\rangle \mathtt{t} \vee Y) \wedge Z)\|$.

(b) Prove or disprove the following by drawing the game graph and argue by the existence (or not) of a winning strategy.

i. $0 \vDash \mu Z.\nu Y.[a]((\langle b\rangle \mathtt{t} \vee Y) \wedge Z)$.

ii. $2 \vDash \mu Z.\nu Y.[a]((\langle b\rangle \mathtt{t} \vee Y) \wedge Z)$.

4.8 We say that variable Z in a formula φ is *guarded* if every occurrence of Z in φ is in the scope of some modal operator $[a]-$ or $\langle a\rangle -$. We say that a formula is *guarded* if for every subformula $\sigma Z.\psi$ of φ, Z is guarded in ψ (where $\sigma = \mu, \nu$). Prove *Kozen's Lemma*: Every modal mu-calculus formula is equivalent to some positive guarded formula.

4.9 What properties are expressed by the following modal mu-calculus formulas?

(a) $\mu Z.[\mathcal{L}]Z$

(b) $\nu Z.\langle a\rangle Z \wedge \langle b\rangle Z$

4.10 Fix a set $\mathcal{L}$ of labels. Express the following in modal mu-calculus formulas.

(a) Eventually either a happens or P becomes true.

(b) Along some path P is always true

(c) There is a path along which P holds continuously and Q holds infinitely often.

Chapter 5

Games and Tableaux for Modal Mu-Calculus

[*We gratefully acknowledge Bahareh Afshari's contribution to the section on tableaux for modal mu-calculus.*]

Synopsis

Modal mu-calculus model checking games. Streett and Emerson's Fundamental Semantic Theorem. 2-person infinite games. Memoryless winning strategies. Signature and signature decrease lemma. Tableaux. Tree model property. Finite model property. Parity games. Computing winning regions. PARITY is in **NP** ∩ co-**NP**. Determinacy for finite parity games.

References (Streett and Emerson, 1989; Bradfield and Stirling, 2007; Stirling, 1997, 2001; Walukiewicz, 1995; Niwinski and Walukiewicz, 1997)

5.1 Game Characterisation of Model Checking

Modal Mu-Calculus Model Checking Problem Given a state s_0 of a labelled transition system T, a valuation V, and a mu-calculus formula φ, does $s_0 \models^T_V \varphi$ hold?.

We aim to give a game characterisation of the problem. We shall consider games played by players V (Verifier) and R (Refuter) on directed graphs of a certain kind. Given a state s_0 of a transition system T, a valuation V and a modal mu-calculus formula φ, we define a game between V and R, $\mathcal{G}^T_V(s_0, \varphi)$, such that $s_0 \models^T_V \varphi$ if, and only if, there is a winning (memoryless) strategy for V in $\mathcal{G}^T_V(s_0, \varphi)$.

The characterisation may be viewed as a version of the Fundamental Semantic Theorem of Streett and Emerson (Streett and Emerson, 1989).

Some preliminaries on syntax We assume that φ is *positive normal* i.e.

(i) Negation is only applied to atomic propositions.

(ii) If $\sigma_1 Z_1$ and $\sigma_2 Z_2$ are two different occurrences of binders in φ, then $Z_1 \neq Z_2$.

(iii) Free variables are disjoint from bound variables.

Every formula can be converted into positive normal form using *de Morgan laws* and by renaming bound variables. For example

$$\mu Z.\langle J\rangle Y \vee (\langle b\rangle Z \wedge \mu Y.\nu Z.([b]Y \wedge [K]Z))$$

can be rewritten

$$\mu Z.\langle J\rangle Y \vee (\langle b\rangle Z \wedge \mu X.\nu U.([b]X \wedge [K]U))$$

Let $Sub(\varphi)$ be the set of subformulas of φ. For example, $Sub(\mu X.\nu Y.([b]X \wedge [K]Y))$ is the set

$$\{\ \mu X.\nu Y.([b]X \wedge [K]Y),\ \nu Y.[b]X \wedge [K]Y,\ [b]X \wedge [K]Y,\ [b]X,\ [K]Y,\ X,\ Y\ \}$$

If φ is (positive) normal and $\sigma Z.\psi \in Sub(\varphi)$, then the (bound) variable Z can be used to identify this subformula.

For $\sigma_1 X_1.\psi_1, \sigma_2 X_2.\psi_2 \in Sub(\varphi)$, we say that X_1 *subsumes* X_2, written $X_1 \succcurlyeq X_2$, just if $\sigma_2 X_2.\psi_2 \in Sub(\sigma_1 X_1.\psi_1)$. In the example above, X subsumes Y (but not vice versa).

Lemma 5.1. *Assume a positive normal φ.*

(i) If X subsumes Y and Y subsumes Z, then X subsumes Z.

(ii) If X subsumes Y and $X \neq Y$, then it is not the case that Y subsumes X.

Definition 5.1 (Model Checking Game $\mathcal{G}_V^T(s,\varphi)$)**.** Fix a transition system $T = \langle\, S,\ \to \subseteq S \times \mathcal{L} \times S,\ \rho : Prop \to \mathcal{P}(S)\,\rangle$, a state $s_0 \in S$, a valuation V, and a positive normal φ. The two players are Refuter (R) and Verifier (V).

- R attempts to show $s_0 \not\models_V^T \varphi$, whereas
- V attempts to show $s_0 \models_V^T \varphi$.

We define the graph underlying $\mathcal{G}_V^T(s_0,\varphi)$. The vertices (also called *positions*, or *moves*) are pairs (t,ψ) where $t \in S$ and $\psi \in Sub(\varphi)$. The initial vertex is (s_0,φ). Edges are organised into three groups according to the shape of the subformula.

- *Boolean subformulas.* For each $\psi_1 \vee \psi_2, \psi_1 \wedge \psi_2 \in Sub(\varphi)$, $s \in S$, $i \in \{\,1,2\,\}$, the following are edges:
$$(s,\psi_1 \vee \psi_2) \to (s,\psi_i), \qquad (s,\psi_1 \wedge \psi_2) \to (s,\psi_i)$$
- *Modal subformulas.* For each $K \subseteq \mathcal{L}$, $a \in K$, and each state t where $s \xrightarrow{a} t$, the following are edges:
$$(s,[K]\psi) \to (t,\psi), \qquad (s,\langle K\rangle\psi) \to (t,\psi)$$
- *Fixpoint subformulas.* For each $\sigma Z.\psi \in Sub(\varphi)$, each $s \in S$, the following are edges:
$$(s,\sigma Z.\psi) \to (s,Z), \qquad (s,Z) \to (s,\psi)$$

Observe that the size of ψ decreases in $(s,\psi) \to (s',\psi')$ in all cases except when ψ is a fixpoint variable.

Plays There are no out-going edges from (s,ψ) where ψ is an atomic proposition (P or $\neg P$), or *free* variable Z (i.e. does not identify any fixpoint subformula). Certain vertices are

"owned" by one of the two players; as for the rest of the variables (all of which have out-degree 1), it is unimportant who owns them.

V	R	Ownership is irrelevant (say, V)
$(s, \psi_1 \vee \psi_2)$	$(s, \psi_1 \wedge \psi_2)$	
$(s, \langle K \rangle \psi)$	$(s, [K]\psi)$	(s, Z) for bound Z
(s, P), $s \notin \rho(P)$	(s, P), $s \in \rho(P)$	$(s, \sigma Z.\psi)$
(s, Z), $s \notin V(Z)$	(s, Z), $s \in V(Z)$	

A *play* of $\mathcal{G}_V^T(s_0, \varphi)$ is a path in the game graph that starts from the initial vertex (s_0, φ), namely, $(s_0, \varphi), (s_1, \varphi_1), (s_2, \varphi_2), \cdots$, such that for each i, if (s_i, φ_i) has label V (resp. R), then V (resp. R) chooses (s_{i+1}, φ_{i+1}).

Example 5.1. Take the transition system with state-set $S = \{ s \}$ with $\mathcal{L} = \{ a, b \}$, and a transition $s \xrightarrow{a} s$.

$T_1 =$ 

The game $\mathcal{G}_{\varnothing}^{T_1}(s, \nu Z.[b]\mathsf{f} \wedge \langle a \rangle Z)$ has the following game graph:

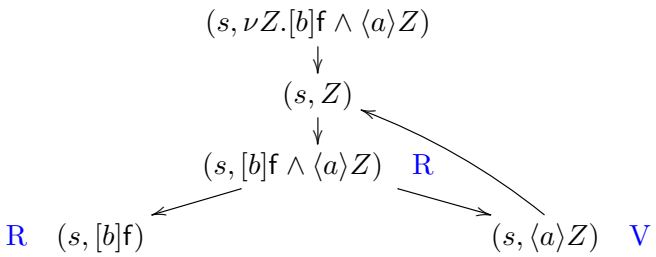

Winning conditions R wins a play just if

(i) The play is $(s_0, \varphi_0), (s_1, \varphi_1), \cdots, (s_n, \varphi_n)$ and

  a. $\varphi_n = P$ and $s_n \notin \rho(P)$ or

  b. $\varphi_n = Z$ and Z is free in φ_0 and $s_n \notin V(Z)$, or

  c. $\varphi_n = \langle K \rangle \psi$ and $\{ t : s_n \xrightarrow{a} t \text{ and } a \in K \} = \varnothing$

(ii) The play $(s_0, \varphi_0), (s_1, \varphi_1), \cdots, (s_n, \varphi_n), \cdots$ is infinite, and the *unique* fixed point variable X, which occurs infinitely often and which subsumes all other variables occuring infinitely often, identifies a least fixpoint subformula of φ.

V wins a play just if

(i) The play is $(s_0, \varphi_0), (s_1, \varphi_1), \cdots, (s_n, \varphi_n)$ and

  a. $\varphi_n = P$ and $s_n \in \rho(P)$, or

  b. $\varphi_n = Z$ and Z is free in φ_0 and $s_n \in V(Z)$, or

  c. $\varphi_n = [K]\psi$ and $\{ t : s_n \xrightarrow{a} t \text{ and } a \in K \} = \varnothing$

(ii) The play $(s_0, \varphi_0), (s_1, \varphi_1), \cdots, (s_n, \varphi_n), \cdots$ is infinite, and the *unique* fixed point variable X, which occurs infinitely often and which subsumes all other variables occuring infinitely often, identifies a greatest fixpoint subformula of φ.

It follows from the winning condition that given a maximal play (either finite and terminal, or infinite), exactly one of V and R wins.

Proposition 9. *If* $(s_0, \varphi_0), (s_1, \varphi_1), \cdots, (s_n, \varphi_n), \cdots$ *is an infinite play of the game* $\mathcal{G}_V^T(s_0, \varphi)$*, then there is a unique variable* X *that*

(i) occurs infinitely often (i.e. $\varphi_j = X$ *for infinitely many* j*), and*

(ii) if Y *also occurs infinitely often, then* X *subsumes* Y*.*

Proof. Because φ_i decreases in size except when it is a fixpoint variable, there are infinite variable occurrences in an infinite play. Suppose, for a contradiction, X and Y are maximal among the infinitely occurring variables and none subsumes the other. It follows that:

(i) The respective fixpoint subformulas named by X and Y occur in different branches of a conjunction or disjunction.

(ii) Further the conjunction (say) is in the scope of some fixpoint subformula named by Z (say), for otherwise, at most one of X and Y can occur infinitely often.

Hence Z subsumes both X and Y and occurs infinitely often in the play, which is a contradiction. □

Example 5.2 (T_1 revisited)**.**

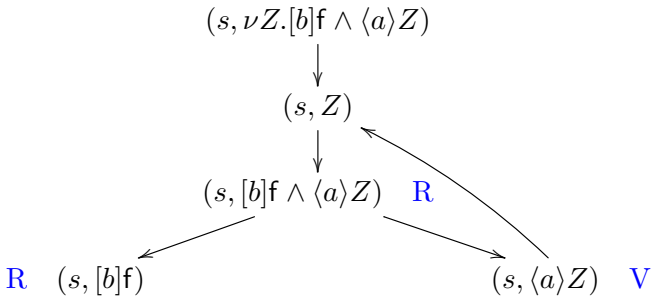

Who wins the game? Two cases according to R's move:

- R chooses the left branch: V wins.
- R chooses the right branch: If an infinite play should arise, since Z identifies a ν-fixpoint, V wins.

Hence V wins.

Is it always the case that exactly one of R and V has a winning strategy?

Example 5.3 (T_2 revisited)**.** $\mathcal{G}_{\varnothing}^{T_2}(0, \mu Y.\nu Z.[a]((\langle b \rangle \mathsf{t} \vee Y) \wedge Z))$ where

$$T_2 \;=\; 0 \underset{a}{\overset{a}{\rightleftarrows}} 1 \xrightarrow{b} 2$$

with state-set $S = \{0, 1, 2\}$. *Who wins the game?* (See the picture on the next page)

R has a winning strategy:

- At $(1, (\langle b \rangle \mathsf{t} \vee Y) \wedge Z)$: choose right branch
- At $(0, (\langle b \rangle \mathsf{t} \vee Y) \wedge Z)$: choose right branch

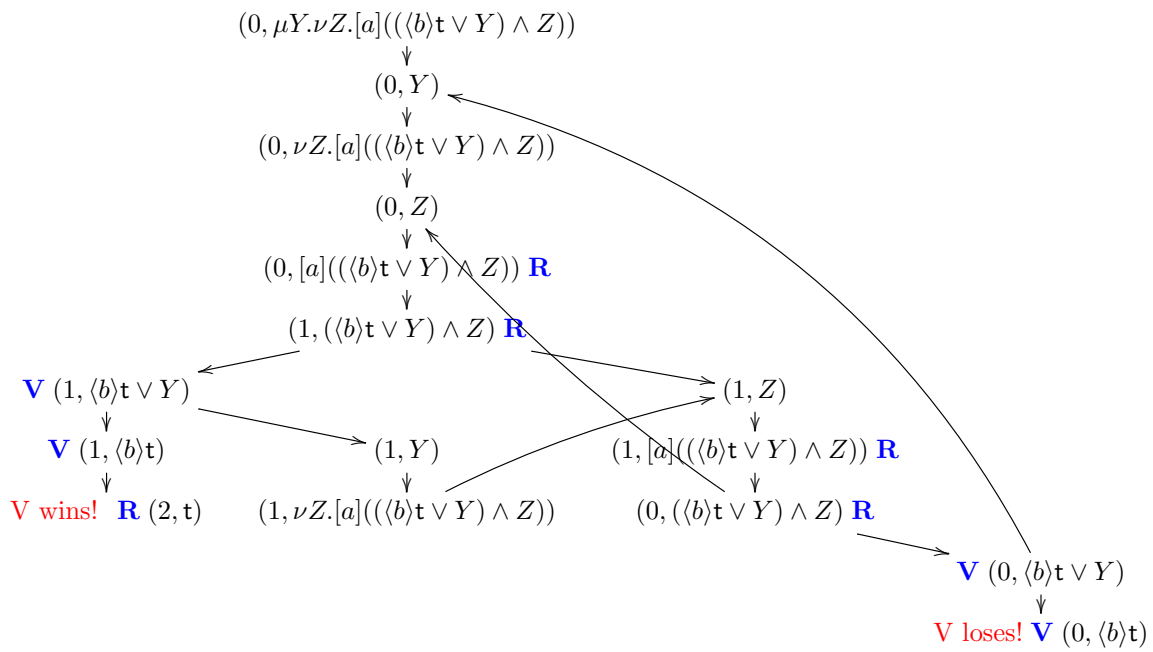

Figure 5.1: Game graph of T_2

At $(0, \langle b\rangle \mathtt{t} \vee Y)$:

- V loses at once if he chooses the left disjunct.
- V also loses if he chooses the right disjunct all the time, since the infinitely occurring and subsuming variable identifies a μ-fixpoint.

Hence R has a winning strategy.

Example 5.4. We present the game graph $\mathcal{G}^{T_2}_{\varnothing}(0, \mu Y.\nu Z.[a](((\langle b\rangle \mathtt{t} \vee Y) \wedge Z))$ in Figure 5.1.

A *strategy* for a player is a set of rules telling the player how to move. A player *uses strategy* π in a play provided all his moves in the play obey the rules in π. *Memoryless strategies* (often called *history-free strategies* in the literature) are strategies that depend only on the last move or position of the play. It follows that for player R, rules have the form:

- at position $(s, \varphi_1 \wedge \varphi_2)$ choose (s, φ_i) where $i = 1$ or $i = 2$
- at position $(s, [K]\varphi)$ such that $\{\, t : s \xrightarrow{a} t \text{ and } a \in K \,\} \neq \varnothing$, choose (t, φ) where $s \xrightarrow{a} t$ and $a \in K$.

Similarly for player V, rules have the form:

- at position $(s, \varphi_1 \vee \varphi_2)$ choose (s, φ_i) where $i = 1$ or $i = 2$
- at position $(s, \langle K\rangle\varphi)$ such that $\{\, t : s \xrightarrow{a} t \text{ and } a \in K \,\} \neq \varnothing$, choose (t, φ) where $s \xrightarrow{a} t$ and $a \in K$.

A strategy π for a player is *winning* if the player wins every play in which he uses π (regardless of how his opponent plays). We aim to prove:

Theorem 5.1 (Fundamental Semantic Theorem)**.** $s_0 \models^T_V \varphi$ *if and only if player V has a memoryless winning strategy for* $\mathcal{G}^T_V(s_0, \varphi)$.

We shall establish the Theorem by proving the following:

(i) If $s_0 \models^T_V \varphi$ then player V has a memoryless winning strategy for $\mathcal{G}^T_V(s_0, \varphi)$.

(ii) If $s_0 \not\models^T_V \varphi$ then player R has a (memoryless) winning strategy for $\mathcal{G}^T_V(s_0, \varphi)$.

It follows from the definition that if one player has a winning strategy, then the other does *not*. I.e. *at most* one player has a winning strategy for a given game $\mathcal{G}^T_V(s_0, \varphi)$. Thus it follows from statements (i) and (ii) in the preceding that *at least* one player has a winning strategy for a given game.

A class of games is said to be *determined* if for every game, there is a winning strategy for exactly one of the players. Thus model checking games of the kind, $\mathcal{G}^T_V(s_0, \varphi)$, are determined. We shall see shortly that these games are *parity games*. In 1975, Donald A. Martin proved that all *Borel games* (which include parity games) are determined.

5.2 Proof of the Fundamental Semantic Theorem

Assume $s_0 \models^T_V \varphi$. We aim to show that V is always able to preserve the "truth of game positions" by making judicious choices, so winning every play.

We list all the fixpoint subformulas of φ, in decreasing order of size, as follows:

$$\sigma_1 Z_1.\psi_1,\ \sigma_2 Z_2.\psi_2,\ \cdots,\ \sigma_n Z_n.\psi_n$$

It follows that:

(i) if $i < j$ then it is impossible for Z_j to subsume Z_i

(ii) if Z_i subsumes Z_j then $i \leq j$.

A *position* in the game $\mathcal{G}^T_V(s, \varphi_0)$ has the form (t, ψ) where ψ may contain free variables in $\{ Z_1, \cdots, Z_n \}$. Our strategy is to show that V can play in such a way that if it reaches (s', ψ) then $s' \models^T_V \psi$. An immediate problem is that since some of the Z_is may occur free in ψ, $\|\psi\|^T_V$ may not be defined!

True positions and valuations We first define valuations $V_0, \cdots, V_n$ by induction on n:

$$\begin{aligned} V_0 &:= V \\ V_{i+1} &:= V_i[Z_{i+1} \mapsto \|\sigma_{i+1} Z_{i+1}.\psi_{i+1}\|^T_{V_i}] \end{aligned}$$

The idea is V_i maps $Z_1, \cdots, Z_i$ to their respective (correct) denotations. E.g. $V_1 : Z_1 \mapsto \|\sigma_1 Z_1.\psi_1\|^T_{V_0}$. Note that $V = V_0$ is not well-defined on $Z_1, Z_2, \cdots, Z_n$, but this is ok since $\sigma_1 Z_1.\psi_1$—as Z_1 is the most subsuming—does not contain any *free* occurrence of $Z_2, \cdots, Z_n$. Similarly, $\|\sigma_2 Z_2.\psi_2\|^T_{V_1}$ is well-defined since it does not contain free occurrences of $Z_3, \cdots, Z_n$, and so on. Thus V_n captures the semantics of all bound variables Z_is i.e. $\|\psi\|^T_{V_n}$ is well-defined for every $\psi \in Sub(\varphi)$.

We say that (t, ψ) is a *true position* just if $t \models^T_{V_n} \psi$. Note that (s_0, φ_0) is a true position, by assumption.

Next we give a more refined valuation that identifies the smallest *least fixpoint approximant* making a given position true. Let

$$\mu Y_1.\chi_1,\ \mu Y_2.\chi_2,\ \cdots,\ \mu Y_m.\chi_m$$

be the set of least fixpoint subformulas of φ, again in decreasing order of size.

A *signature* is just an m-long vector of ordinals. We say that signatures $r < r'$ if r *lexicographically precedes*[1] r'.

Given a signature $r = \alpha_1 \cdots \alpha_m$ and a valuation V, we define valuations $V_0^r, \cdots, V_n^r$ by induction:

$$\begin{aligned} V_0^r &:= V \\ V_{i+1}^r &:= V_i^r \left[Z_{i+1} \mapsto \begin{cases} \|\sigma_{i+1} Z_{i+1}.\psi_{i+1}\|^T_{V_i^r} & \text{if } \sigma_{i+1} = \nu \\ \|\mu Y_j^{\alpha_j}.\chi_j\|^T_{V_i^r} & \text{else if } \sigma_{i+1} Z_{i+1} = \mu Y_j \end{cases} \right] \end{aligned}$$

Thus we use a signature $r = \alpha_1 \cdots \alpha_m$ to interpret the *least* fixpoint subformulas in φ so that the j-th such is interpreted by its α_j-th approximant.

μ-signature of a true position

Lemma 5.2. *If $t \models^T_{V_n} \psi$ then there is a smallest signature r such that $t \models^T_{V_n^r} \psi$.*

Thus given a true position (t, ψ), we define its *μ-signature* (or simply *signature*), written $sig^\mu(t, \psi)$, to be the least r such that $t \models^T_{V_n^r} \psi$ holds.

Notation Let r and r' be signatures. We write

- $r(k)$ to mean the kth component of the signature r, and
- $r =_k r'$ to mean that the first k components of the signatures r and r' are identical.

The μ-signature is *unchanged or decreases* when passing through boolean, modal or ν-variable dependencies / edges, and when passing through $(-, Y_j)$, it strictly decreases in the jth-component and is unchanged in each of the $1, \cdots, j-1$ component.

Lemma 5.3 (Signature Decrease)**.** *Whenever the LHS of the relation in question is defined, we have*

(i) $sig^\mu(s, \varphi_1 \vee \varphi_2) = sig^\mu(s, \varphi_i)$, *for some* $i \in \{1, 2\}$.

(ii) $sig^\mu(s, \varphi_1 \wedge \varphi_2) \geq \max(sig^\mu(s, \varphi_1), sig^\mu(s, \varphi_2))$

(iii) $sig^\mu(s, \langle a \rangle \varphi_1) = sig^\mu(t, \varphi_1)$, *for some* t *such that* $s \xrightarrow{a} t$

(iv) $sig^\mu(s, [a]\varphi_1) \geq sig^\mu(t, \varphi_1)$, *for all* t *such that* $s \xrightarrow{a} t$

(v) Assuming Z_i *is a* ν*-variable,* $sig^\mu(s, \nu Z_i.\psi_i) = sig^\mu(s, Z_i) = sig^\mu(s, \psi_i)$.

(vi) Assuming $Z_i = Y_j$ *is a* μ*-variable*

  (a) $sig^\mu(s, \mu Y_j.\chi_j) =_{j-1} sig^\mu(s, Y_j)$

  (b) $sig^\mu(s, Y_j)(j) > sig^\mu(s, \chi_j)(j)$ *and* $sig^\mu(s, Y_j) =_{j-1} sig^\mu(s, \chi_j)$.

Proof. Exercise. □

[1] We define $a_1 \cdots a_m < b_1 \cdots b_m$ iff $a_1 < b_1$, or ($a_1 = b_1$ and $a_2 \cdots a_m < b_2 \cdots b_m$).

Winning memoryless stratregy for V Finally we simultaneously give the memoryless strategy for V and prove that it is winning, by appealing to the *Signature Decrease Lemma*.

Suppose (s_m, φ_m) is the current position in the play. Assume that it is a true position, and so, $s_m \models_{V_n^{r_m}} \varphi_m$ where $r_m = sig^{\mu}(s_m, \varphi_m)$. If (s_m, φ_m) is a final position, then V is the winner. Therefore either V wins or the play is not yet complete. In the latter, we show how the play is extended to (s_{m+1}, φ_{m+1}) by a case analysis on φ_m.

- $\varphi_m = \psi_1 \wedge \psi_2$: Then R chooses ψ_i for some $i \in \{1, 2\}$ as the next move, and the next position $(s_{m+1}, \varphi_{m+1}) = (s_m, \psi_i)$ remains true, and $s_{m+1} \models_{V_n^{r_{m+1}}} \varphi_{m+1}$ with $r_{m+1} \leq r_m$, by the Lemma.
- $\varphi_m = [K]\psi$: Then R chooses as the next move $(s_{m+1}, \varphi_{m+1}) = (t, \psi)$, for some t and some $a \in K$ such that $s_m \xrightarrow{a} t$. Since $s_m \in \|[K]\psi\|^T_{V_n^{r_m}}$, we have $s_{m+1} \in \|\psi\|^T_{V_n^{r_m}}$, so (s_{m+1}, φ_{m+1}) is true and $r_{m+1} \leq r_m$ by the Lemma.

Applying the Signature Decrease Lemma

- $\varphi_m = \psi_1 \vee \psi_2$: V chooses one of ψ_1 and ψ_2 (say ψ_1) that holds. The next position $(s_{m+1}, \varphi_{m+1}) = (s_m, \psi_1)$ remain true. By the Lemma, $r_{m+1} = r_m$.
- $\varphi_m = \langle K \rangle \psi$: similar to above.
- $\varphi_m = \sigma_i Z_i.\psi_i$. Then $(s_{m+1}, \varphi_{m+1}) = (s_m, Z_i)$. Two cases:
  - $\sigma_i = \nu$: (s_{m+1}, φ_{m+1}) is a true position, and $r_{m+1} = r_m$.
  - $\sigma_i = \mu$: (s_{m+1}, φ_{m+1}) is a true position, but $r_{m+1} =_{i-1} r_m$.
- $\varphi_m = Z_i$. Then $(s_{m+1}, \varphi_{m+1}) = (s_m, \psi_i)$. Two cases:
  - $\sigma_i = \nu$: (s_{m+1}, φ_{m+1}) is a true position, and $r_{m+1} = r_m$.
  - $\sigma_i = \mu$: (s_{m+1}, φ_{m+1}) is a true position, with $r_{m+1} < r_m$ and $r_{m+1} =_{i-1} r_m$.

The case of infinite plays Take an infinite play $(s_0, \varphi_0), (s_1, \varphi_1), \cdots$ in which after (say) (s_k, φ_k), every occurrence of a fixpoint variable is subsumed by Z_i. Suppose, for a contradiction, $Z_i = Y_j$ which identifies a least fixpoint subformula $\mu Y_j.\chi_j$.

Let $k_1, k_2, \cdots$ be the positions in the play where Y_j occurs. The move from (s_{k_i}, Y_j) to (s_{k_i+1}, χ_j) causes a strict decrease of the μ-signature. We show that the remaining moves between k_i and k_{i+1} cannot cancel this decrease out.

We only need to worry about $(s, \mu Y_l.\chi_l) \to (s, Y_l)$, as this is the only case which may cause an increase in the μ-signature. By assumption Y_l is subsumed by Y_j; it follows that $l > j$. Note that the transition $(t, \mu Y_j.\chi_j)$ to (t, Y_j) is not possible in the segment from k_i to k_{i+1} because Y_j is assumed to be subsuming. Hence the increase in μ-signature occurs after the l-th component, and so, there is still an overall decrease in the respective j-prefix of r_z as z progresses from k_i to k_{i+1}. This contradicts the well-foundedness of (fixed-length vectors of) ordinals.

A duality We use a "symmetric" argument to establish the other direction: *If* $s \not\models^T_V \varphi$ *then player R has a memoryless winning strategy for* $\mathcal{G}^T_V(s, \varphi)$.

Specifically

- For a position (t, ψ) that is *false* (i.e. $s \notin \|\psi\|^T_V$), define its *ν-signature*, written $sig^{\nu}(t, \psi)$, to be the least r such that $t \not\models^T_{V_n^r} \psi$ holds.

- Establish a corresponding *Signature Decrease Lemma* for ν-signatures.
- Define a memoryless winning strategy for R.

Hence V has a memoryless winning strategy for $\mathcal{G}_V^T(s,\varphi)$ iff $s \models_V^T \varphi$.

Lemma 5.4. *A player does not have a memoryless winning strategy for $\mathcal{G}_V^T(s,\varphi)$ if and only if he has a memoryless strategy for $\mathcal{G}_V^T(s,\widetilde{\varphi})$, where $\widetilde{\varphi}$ is the positive normal form of $\neg\varphi$.*

Proof. We consider V; the situation for R is exactly dual.

V has no memoryless winning strategy for $\mathcal{G}_V^T(s,\varphi)$
iff $s \not\models_V^T \varphi$
iff $s \models_V^T \neg\varphi$
iff V has a memoryless winning strategy for $\mathcal{G}_V^T(s,\widetilde{\varphi})$

□

5.3 Tableaux for modal mu-calculus

Preliminaries We will not distinguish between different labels namely only consider formulae $[\mathcal{L}]\varphi$ and $\langle\mathcal{L}\rangle\varphi$ which we will denote with $\Box\varphi$ and $\Diamond\varphi$. Note that this is not a necessary condition for the following but makes for concise presentation.

From now on, all formulae considered are assumed to be closed, in positive normal form and guarded. A formula φ is *guarded* if for every $\sigma Z.\psi \in Sub(\varphi)$, every occurrence of the bound variable Z in ψ appears in the scope of a modal operator. Every formula can be converted into guarded form by replacing unguarded μ- and ν-variables by $P \wedge \neg P$ and $P \vee \neg P$ respectively where P is any atomic proposition. For example $\mu Z.Z \vee \varphi$ is equivalent to $\mu Z.\varphi$ and $\nu Y.(Q \wedge (\mu Z.Y \wedge \Diamond Z)) \vee \Diamond Y$ is equivalent to $\nu Y.(Q \wedge \mu Z.\Diamond Z) \vee \Diamond Y$.

Satisfaction and Models We say a transition system $T = (S, \to, \lambda)$ with a distinguished node s_0 *satisfies* the formula φ of modal mu-calculus, if $s_0 \models^T \varphi$ (i.e. $s_0 \in \|\varphi\|^T$). In this case we call T a *model* of φ and say that φ is *satisfiable*.

Tableaux: an informal view Consider the formula $\varphi = \nu Y \mu X.(\Box\Box P \wedge \Diamond X) \vee (\neg P \wedge \Diamond Y)$. We want to ask if φ is satisfiable. We aim to find a model of φ namely a transition system

$T = (S, \to, \lambda)$ and $s_0 \in S$ such that $s_0 \models^T \varphi$. We can proceed as follows.

$$
\begin{array}{ll}
 & s_0 \models \nu Y \mu X.(\Box\Box P \wedge \Diamond X) \vee (\neg P \wedge \Diamond Y) \\
 & s_0 \models (\Box\Box P \wedge \Diamond X) \vee (\neg P \wedge \Diamond Y) \\
 & s_0 \models \Box\Box P \wedge \Diamond X \\
 & s_0 \models \{\Box\Box P, \Diamond X\} \\
\exists s_1 : s_0 \to s_1 & s_1 \models \{\Box P, X\} \\
 & s_1 \models \{\Box P, (\Box\Box P \wedge \Diamond X) \vee (\neg P \wedge \Diamond Y)\} \\
 & s_1 \models \{\Box P, \neg P \wedge \Diamond Y\} \\
 & s_1 \models \{\Box P, \neg P, \Diamond Y\} \\
\exists s_2 : s_1 \to s_2 & s_2 \models \{P, Y\} \\
 & s_2 \models \{P, (\Box\Box P \wedge \Diamond X) \vee (\neg P \wedge \Diamond Y)\} \\
 & s_2 \models \{P, \Box\Box P \wedge \Diamond X\} \\
 & s_2 \models \{P, \Box\Box P, \Diamond X\} \\
\exists s_3 : s_2 \to s_3 & s_3 \models \{\Box P, X\} \\
 & \vdots
\end{array}
$$

This "model search" naturally gives rise to the transition system illustrated in Figure 5.2. This search for satisfiability (or soundness) can be formalised using tableaux. The idea is to encapsulate all possible plays on an arbitrary transition system as a tree where each node is labelled by a subset of $Sub(\varphi)$.

$$s_0 \text{ —— } s_1 \text{ —— } s_2 \text{ —— } s_3 \text{ —— } \cdots \qquad s_i \in \lambda(P) \text{ if and only if } i > 0 \text{ is even}$$

Figure 5.2: A model of $\nu Y \mu X.(\Box\Box P \wedge \Diamond X) \vee (\neg P \wedge \Diamond Y)$.

Trees and Paths

Definition 5.2. A tree is a tuple $T = (S, \to, \lambda)$ with a distinguished node ρ which satisfies the following conditions.

- $(S, \to)$ is a connected, directed graph.
- λ is a function from S to a set of labels.
- The are no transitions into ρ.
- For every $s \in S \setminus \{\rho\}$ there is exactly one $s_0 \in S$ such that $s_0 \to s$.

The node ρ is referred to as the *root* of the tree; any node without outgoing transitions is a *leaf* and λ is the *labelling function* of the tree T. We use subscripts to emphasise which tree they refer to.

If $T = (S, \to, \lambda)$ is a tree then a *path* through T is an enumerable set $\mathbb{P} \subseteq S$ such that $\rho_T \in \mathbb{P}$, if $s_0 \to s \in \mathbb{P}$ then $s_0 \in \mathbb{P}$, and for every $s \in \mathbb{P}$ either s is a leaf or there exists exactly one $s' \in S$ such that $s \to s'$ and $s' \in \mathbb{P}$. Hence a path can be visualised as a sequence $\rho_T = s_0 \to s_1 \to s_2 \to \ldots \to s_n \to \ldots$ and we write $\mathbb{P}(n)$ to denote s_n.

Recall $Prop$ is a (possibly infinite) set of propositions from which the syntax of $\mathcal{L}_\mu$ is defined. Let $Prop^\neg = \{\neg P \colon P \in Prop\}$. Uppercase Greek letters such as Γ and Δ denote *sequents*, finite sets of formulae. $\Box\Gamma$ abbreviates the set $\{\Box\varphi : \varphi \in \Gamma\}$ and $\Diamond\Gamma$ is defined analogously. We write Γ, φ to mean $\Gamma \cup \{\varphi\}$, and Γ, Δ to denote $\Gamma \cup \Delta$.

$$(\wedge)\ \frac{\Gamma, \varphi_0, \varphi_1}{\Gamma, \varphi_0 \wedge \varphi_1} \qquad (\vee)\ \frac{\Gamma, \varphi_i}{\Gamma, \varphi_0 \vee \varphi_1}\ i = 0 \text{ or } 1$$

$$(\sigma)\ \frac{\Gamma, Z}{\Gamma, \sigma Z.\varphi}\ \sigma \in \{\mu, \nu\} \qquad (Z)\ \frac{\Gamma, \varphi}{\Gamma, Z}\ Z \text{ identifies } \sigma Z.\varphi$$

$$(\text{mod})\ \frac{\Gamma, \delta \text{ for each } \delta \in \Delta}{\Box\Gamma, \Diamond\Delta, \Theta}\ \Theta \subseteq Prop \cup Prop^\neg \text{ consistent}$$

Table 5.1: Rules for generating pre-tableaux.

Definition 5.3. A *pre-tableau* for Γ is any tree generated by the rules in Table 5.1 from the sequent Γ in which any finite branch terminates by reaching a sequent of the form $\Box\Xi, \Diamond\Delta, \Theta$ where Δ is empty. Note that this tree is finitely branching and branching only occurs at a (mod)-rule.

For each rule in Table 5.1 the distinguished formulae in the lower and upper sequents are called respectively the *principal* and *residue* formulae of the rule. In case of a (mod)-rule all the formulae occurring are considered distinguished including those in Θ.

Note that our trees grow upwards. Thus the tableau rules are presented such that they read bottom-up. For example, in the $(\vee)$-rule from the sequent $\Gamma, \varphi_0 \vee \varphi_1$ we may proceed to either Γ, φ_0 or Γ, φ_1.

Exercise 5.1. *Write down a (mod)-rule for $\mathcal{L}_\mu$-formulae with labels that satisfies the informal soundness property described above.*

Proposition 10. *Every infinite path in a pre-tableau passes through a (mod)-rule infinitely often. Moreover, the (mod)-rule can be applied only if no other rule is applicable.*

Proof. Exercise 5.3. □

Fix a pre-tableau $t = (V, \to, \lambda)$ for Γ and a path $\mathbb{P}$ through t. A finite *thread* (through $\mathbb{P}$) is a sequence of formulae $\varphi_0, \varphi_1, \ldots, \varphi_n$ such that

- $\varphi_i \in \lambda(\mathbb{P}(i))$ for each $i \leq n$;
- $\varphi_{i+1} = \varphi_i$ if φ_i is not principal in the rule applied at node $\mathbb{P}(i)$, otherwise φ_{i+1} is the residual *subformula* of φ_i occurring in the label of $\mathbb{P}(i+1)$.

An infinite sequence of formulae $\varphi_0, \varphi_1, \ldots$ is a thread if every finite initial sequence is.

Lemma 5.5. *For each infinite thread there exists a variable that appears infinitely often in the thread and subsumes all other infinitely occurring variables.*

Proof. Exercise. □

In each infinite thread the unique variable identified by Lemma 5.5 will be referred to as the *most significant variable* of the thread. We call an infinite thread a *μ-thread* if its most significant variable is a μ-variable; otherwise it is a *ν-thread*.

Definition 5.4. A pre-tableau is a *tableau* if

(i) the sequent at the leaf of each finite branch is of the form $\Box\Gamma, \Theta$ where Θ is a consistent set of propositions, and

(ii) every infinite thread is a ν-thread.

Lemma 5.6. *Let $t = (V, \to, \lambda)$ be a tableau. Then for every $v \in V$ we have $\lambda(v) \cap (Prop \cup Prop^{\neg})$ is consistent.*

Proof. Exercise 5.3. □

Example 5.5. Of the four pre-tableaux given below for the formula $\mu Z.Q \vee (P \wedge \Diamond Z)$ only the first tree are tableaux. The fourth one which is assumed to always pick the right disjunct in $Q \vee (P \wedge \Diamond Z)$ is not a tableau.

$$\dfrac{\dfrac{\dfrac{\dfrac{Q}{Q \vee (P \wedge \Diamond Z)}(\vee)}{Z}(Z)}{\mu Z.Q \vee (P \wedge \Diamond Z)}(\mu)}$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{Q}{Q \vee (P \wedge \Diamond Z)}(\vee)}{Z}(Z)}{P, \Diamond Z}(\text{mod})}{P \wedge \Diamond Z}(\wedge)}{Q \vee (P \wedge \Diamond Z)}(\vee)}{\mu Z.Q \vee (P \wedge \Diamond Z)}\begin{matrix}Z\\(Z)\end{matrix}\\(\mu)$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{Q}{Q \vee (P \wedge \Diamond Z)}(\vee)}{Z}(Z)}{P, \Diamond Z}(\text{mod})}{P \wedge \Diamond Z}(\wedge)}{Q \vee (P \wedge \Diamond Z)}(\vee)}{Z}(Z)}{\mu Z.Q \vee (P \wedge \Diamond Z)}(\mu)$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\vdots}{Z}}{P \wedge \Diamond Z}(\vee)}{Q \vee (P \wedge \Diamond Z)}(Z)}{Z}(\text{mod})}{P, \Diamond Z}(\wedge)}{P \wedge \Diamond Z}(\vee)}{Q \vee (P \wedge \Diamond Z)}(Z)}{Z}(\mu)$$

Soundness and Completeness

Lemma 5.7 (Soundness)**.** *If φ has a tableau then φ is satisfiable.*

Proof. Let $t = (V, \to, \lambda)$ be a tableau for φ. Define a tree $T = (S, \to_T, \lambda_T)$ and a map $\tau\colon V \to S$ such that

(i) $\tau(\rho) = \rho_T$

(ii) If $v \to u$ and the tableau rule applied at v is (mod) then $\tau(v) \to_T \tau(u)$, otherwise $\tau(u) = \tau(v)$.

(iii) $P \in \lambda_T(s)$ if and only if there exists $v \in V$ such that $\tau(v) = s$ and $P \in \lambda(v)$.

We will use t to define a winning strategy for Verifier in the model checking game $\mathcal{G}^T(\rho_T, \varphi)$, whence the Fundamental Semantic Theorem yields $\rho_T \models^T \varphi$.

Suppose we have a play (s_0, φ_0), (s_1, φ_1), …, (s_n, φ_n) and a corresponding thread φ'_0, φ'_1, …, φ'_m in the tableau t such that

- there exists $i_0 = 0 < i_1 < \cdots < i_n < i_{n+1} = m + 1$ so that for every $j \leq n$ and every $k \in [i_j, i_{j+1})$, $\varphi'_k = \varphi_j$, and
- φ'_m is principal in the node v of t associated to it and $\tau(v) = s_n$.

Suppose $\varphi_n \notin Prop \cup Prop^{\neg}$. We show how the play and the thread can be extended.

Case I. It is Verifier's move then either $\varphi_n = \psi_0 \vee \psi_1$ or $\varphi_n = \Diamond\psi_0$. As t is a tableau and $\varphi'_m = \varphi_n$ the thread can be extended by some $\psi \in \{\psi_0, \psi_1\}$. If u denotes the node associated to ψ in this thread set Verifier's choice to be $(s_{n+1}, \varphi_{n+1}) = (\tau(u), \psi)$.

Case II. It is Refuter's move. Then for any any choice of (s_{n+1}, φ_{n+1}) by Refuter the thread can be extended accordingly.

This is in fact a winning strategy for Verifier. Suppose $(s_0, \varphi_0), (s_1, \varphi_1), \ldots, (s_n, \varphi_n)$ is a finite play using the above strategy. It follows that there exits a node $v \in V$ such that $\tau(v) = s_n$ and $\varphi_n \in \lambda(v)$ is principal at v. If $\varphi_n \in Prop$, this play is winning for Verifier by definition. If $\varphi_n = \neg P$ then since $\lambda(v) \cap (Prop \cup Prop^{\neg})$ is consistent, $P \notin \lambda_T(s_n)$ and Verifier wins this play too. Furthermore, an infinite play corresponds to an infinite thread in the tableau. Since every infinite thread is a ν-thread this implies that the play is winning for Verifier. □

Lemma 5.8 (Completeness)**.** *If φ is satisfiable then φ has a tableau.*

Proof. Let T be a transition system with distinguished node s_0 such that $s_0 \models^T \varphi$. Use Verifier's winning strategy in $\mathcal{G}^T(s_0, \varphi)$ (given by the Fundamental Semantic Theorem) to make your choices in defining a pre-tableau for φ. Then every infinite thread in the tableau corresponds to an infinite play in the game and hence is a ν-thread. Moreover, every finite branch in the pre-tableau will be of the form $\Box\Gamma, \Theta$ where Θ is consistent as otherwise there will be maximal plays that end in a position (s, P) with $s \not\models P$ or $(s, \Diamond\psi)$ which would be losing for Verifier. □

The Soundness and Completeness Lemmata give us a characterisation of satisfaction in terms of the existence of tableaux:

Theorem 5.2. *A formula is satisfiable iff it has a tableau.*

Corollary 5.1 (Tree Model Property)**.** *If a formula φ of modal mu-calculus has a model then it has a tree model.*

Proof. Suppose φ is satisfiable. The Completeness Lemma 5.8 implies that φ has a tableau. The proof of the Soundness Lemma 5.7 then shows how to extract a tree model for φ from this tableau. □

Finite Model Property Tableaux are a powerful tool for decision procedures. They are closely related to games and automata, but provide an insight that is unique. We will demonstrate this by proving the following classic result.

Theorem 5.3 (Finite Model Property)**.** *If a formula φ of modal mu-calculus has a model then it has a finite model.*

Proof. If φ has a model then by Lemma 5.8 it has a tableau, say t. Let S be the collection of finite threads p that are extended by some infinite thread whose most significant quantifier is encountered already in p. S naturally forms a finitely branching tree in which all branches are finite, whence by Köning's Lemma, S is a finite set. The leaves of S define a finite set of nodes $F_0 \subset V_t$ such that every infinite path in t passes through exactly one node of F_0. Similarly for each $n > 1$ there exists a finite set $F_n \subset V_t$ such that every infinite path through t intersects

F_n exactly once, and every infinite thread through t sees its most significant variable at least once between the sets F_{n-1} and F_n.

We can now construct the finite model for φ. Fix $k = 2^{|Sub(\varphi)|}$ and an arbitrary node $v \in F_k$. For $j \leq k$, denote by v_j the unique node in F_j below v. As k is sufficiently large, there exists $i < j \leq k$ such that $\lambda_t(v_i) = \lambda_t(v_j)$. In t, remove the subtree of t at v_j and add an edge from the predecessor of v_j to v_i. Note that the resulting structure is not a tableau but its "unravelling" will give rise to a tableau due to the definition of F_j. Repeat this process for each node of F_k. This yields a finite tree-with-back-edges. From this finite structure we can extract a finite transition system by the same construction as in the soundness lemma which, by an analogous argument, is a model of φ. □

The following stronger result can also be proved using tableaux. However, the proof is beyond the scope of this lecture.

Theorem 5.4 (Small Model Property)**.** *If a formula φ of modal mu-calculus has a model then it has a model of size at most $2^{2^{\mathcal{O}(n)}}$ where n is the size of φ.*

The idea is to show that if φ has a tableau then it has a tableau in which any two applications of the $(\vee)$-rule to the same sequent with the same principal formula yields the same residue. This will allow us to mark a frontier of depth $2^{\mathcal{O}(n)}$ from which we can already make the back-edges as in the proof of the finite model property.

An immediate consequence of the theorem is

Corollary 5.2. *The satisfaction problem for the modal mu-calculus is decidable. Indeed, the problem is EXPTIME (complete).*

5.4 Parity Games

Definition 5.5. A *parity game* is a tuple $G = \langle N, E, v_I, \lambda, \Omega \rangle$ where

- $\langle N, E \subseteq N \times N \rangle$ is a directed graph such that E is *total* (i.e. for each u, there is some v, such that $(u, v) \in E$), and $v_I \in N$ is the *start vertex*
- $\lambda : N \longrightarrow \{ V, R \}$ labels each vertex with R (Refuter) or V (Verifier); $\lambda(v)$ indicates which player is responsible for moving from v
- $\Omega : N \longrightarrow \{ 0, \cdots, p \}$ assigns a *priority* (or *colour*) to each vertex v

A play begins with a token on the start vertex v_I. When the token is on u and $\lambda(u) = P$, player P moves it along an outgoing edge (u, v) to v. A *play* is an infinite path $v_I\, v_1 \cdots v_i \cdots$ in the graph visited by the token.

The winner of a play is determined by the *min-parity* condition: if the least priority that occurs infinitely often in the sequence $\Omega(v_I)\, \Omega(v_1)\, \Omega(v_2) \cdots$ is even then V wins; otherwise R wins. There is an equivalent *max-parity* condition.

Example 5.6 (A parity game)**.** V-moves are circled; R-moves are boxed. Prioities are indicated in red after the colon. Who has a winning strategy starting from a?

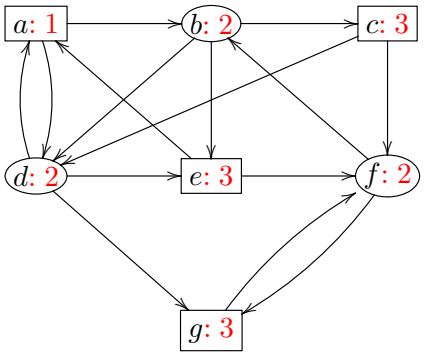

Recall: V wins an infinite play just if the least infinitely occurring priority is even. Ans: V has a winning strategy: $b \mapsto c,\ f \mapsto g,\ d \mapsto g$.

Example 5.7.

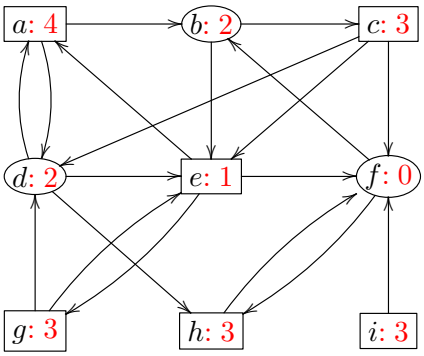

From which vertices does V have a winning strategy? Equivalently, what is the *winning region* of V?

Example 5.8. Parity Game $\mathcal{G}_{\varnothing}^{T_2}[0, \mu Y.\nu Z.[a]((\langle b\rangle \mathtt{t} \vee Y) \wedge Z)]$ where T_2 is defined in Exam-

ple 6.2.

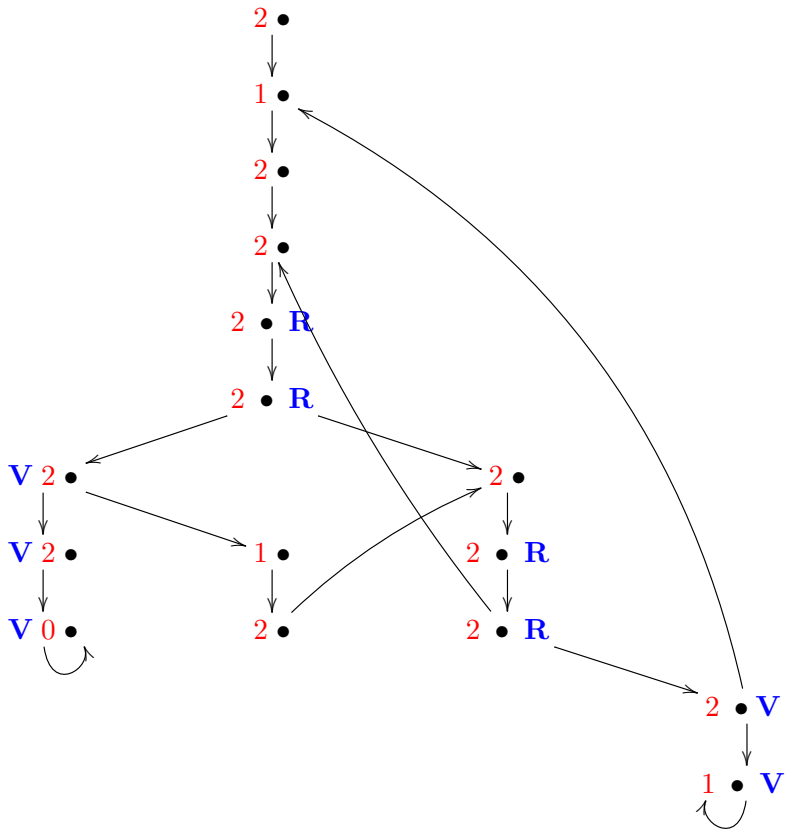

Every modal mu-calculus model checking instance determines a parity game. Precisely,

Proposition 11. *The mu-calculus model checking problem* reduces *to the decision problem* PARITY*: given a parity game, does V have a winning solution?*

5.5 Solvability and Determinacy for Finite Parity Games

Let $G = \langle N, \to, v_0, \lambda, \Omega \rangle$ be a parity game, and $v \in N$. Write $G(v) := \langle N, \to, v, \lambda, \Omega \rangle$. Let $P \in \{ V, R \}$. The *winning region of* P, written W_P, is the subset of G-vertices v such that P has a winning strategy starting from v (i.e. in the game $G(v)$). It follows from the definition that $W_R \cap W_V = \varnothing$. Trivially $W_R \cup W_V \subseteq N$.

Theorem 5.5 (Solvability)**.** *Let $G = \langle N, \to, v_0, \lambda, \Omega \rangle$ be a finite parity game. Both the winning regions, W_V and W_R, and the corresponding memoryless winning strategies for V and R are computable.*

Intuitively $\mathsf{Force}^i_P(X)$ is the set of vertices from which player P can force play to enter the set X of vertices in at most i moves.

Definition 5.6 (Force Sets)**.** Let $X \subseteq N$.

$$
\begin{array}{rcl}
\mathsf{Force}^0_P(X) & := & X \quad \text{for } P \in \{V, R\} \\
\mathsf{Force}^{i+1}_R(X) & := & \mathsf{Force}^i_R(X) \\
& \cup & \{\, j : \lambda(j) = R \wedge \exists k \in \mathsf{Force}^i_R(X).j \to k \,\} \\
& \cup & \{\, j : \lambda(j) = V \wedge \forall k.j \to k \Rightarrow k \in \mathsf{Force}^i_R(X) \,\} \\
\mathsf{Force}^{i+1}_V(X) & := & \mathsf{Force}^i_V(X) \\
& \cup & \{\, j : \lambda(j) = V \wedge \exists k \in \mathsf{Force}^i_V(X).j \to k \,\} \\
& \cup & \{\, j : \lambda(j) = R \wedge \forall k.j \to k \Rightarrow k \in \mathsf{Force}^i_V(X) \,\} \\
\mathsf{Force}_P(X) & := & \bigcup_{i \geq 0} \mathsf{Force}^i_P(X)
\end{array}
$$

Rank and Forcing Strategy The definition of force set provides a method for computing it. As i increases, we calculate $\mathsf{Force}^i_P(X)$ until it is the same as $\mathsf{Force}^{i-1}_P(X)$. Clearly this must hold when $i \leq (|N| - |X|) + 1$.

If $j \in \mathsf{Force}_P(X)$ and current position is j, then P can force play from j into X, regardless of how the opponent moves. (Vertex j itself need not belong to X.) The *rank* of such a vertex j is the least index i such that $j \in \mathsf{Force}^i_P(X)$.

For each $i \in \mathsf{Force}_P(X)$ belonging to P, either $i \in X$ or there is $i \to k$ and $k \in \mathsf{Force}_P(X)$, so the *forcing strategy* for P is to choose a k with the least rank.

Example 5.9 (Force Sets of a Parity Game)**.** V-moves are circled; R-moves are boxed. Prioities are indicated in red.

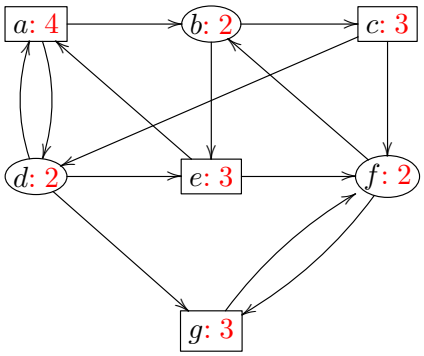

i	$\mathsf{Force}^i_V(\{ f \})$
0	$\{ f \}$
1	$\{ f, g \}$
2	$\{ f, g, d \}$
3	$\{ f, g, d, c \}$
4	$\{ f, g, d, c, b \}$
5	$\{ f, g, d, c, b, a \}$
6	$\{ f, g, d, c, b, a, e \}$

Since $\mathsf{Force}_V(\{ f \})$ is the entire vertex set, and f has the least priority which is even, V has a winning strategy from every vertex. V's forcing strategy: $f \mapsto g,\ d \mapsto g,\ b \mapsto c$

Subgames If $X \subseteq N$, then $G - X$ is the result of removing all vertices in X from G, all edges from vertices in X, all edges into vertices in X.

The subgraph $G - X$ may or may not be a parity game (because not all vertices may be total). If it is, then $G - X$ is a *subgame* of G.

Proposition 12. *If G is a game and X is a subset of vertices, then the subgraph $G - \mathsf{Force}_P(X)$ is a subgame.*

Exercise 5.2. Prove the proposition.

Algorithmic Solution of Finite Parity Games $\mathbf{WReg}(G)$ computes G's winning regions, W_V and W_R. (The respective memoryless winning strategies can be extracted from the correctness proof.) Assume the least priority is even; otherwise swap players in the algorithm.

Algorithm $\mathbf{WReg}(G)$. Output: W_V and W_R

1. Let v be a G-vertex of the least priority (assumed even). Set $X := \mathsf{Force}_V(\{\, v \,\})$.
2. If $X = N$ then return $W_V := N$ and $W_R := \varnothing$.
3. Else run $\mathbf{WReg}(G - X)$, and let W'_R and W'_V be the winning regions.
   (a) If V can guarantee transition from v to $W'_V \cup X$, i.e.,
   $$\left\{ \begin{array}{l} (a)\ \lambda(v) = V \ \wedge\ \exists v' \,.\, v \to v' \ \wedge\ v' \in (W'_V \cup X), \text{ or} \\ (b)\ \lambda(v) = R \ \wedge\ \forall v' \,.\, v \to v' \ \Rightarrow\ v' \in (W'_V \cup X) \end{array} \right\} \text{ then}$$
   return $\left\{ \begin{array}{lcl} W_V & := & W'_V \cup X \\ W_R & := & W'_R. \end{array} \right.$
   (b) Else set $X' := \mathsf{Force}_R(W'_R)$ in G. Run $\mathbf{WReg}(G - X')$, and let W''_V and W''_R be the winning regions. Return $\left\{ \begin{array}{lcl} W_R & := & W''_R \cup X' \\ W_V & := & W''_V. \end{array} \right.$

Example 5.10 (Computing Winning Regions)**.** Parity game G:

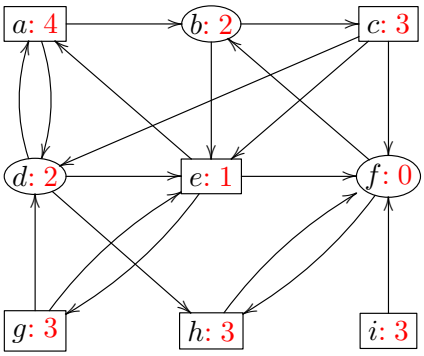

Run $\mathbf{WReg}(G)$: let winning regions be W_V and W_R. Set $X = \mathsf{Force}_V(\{\, f \,\}) = \{\, d, f, h, i \,\}$. $\mathbf{WReg}(G - X)$ returns regions $W'_V = \varnothing$ and $W'_R = (G - X)$.

Since $f \to h$ and $h \in X$, return $W_V = X$ and $W_R = W'_R$.

Parity game $G - X$:

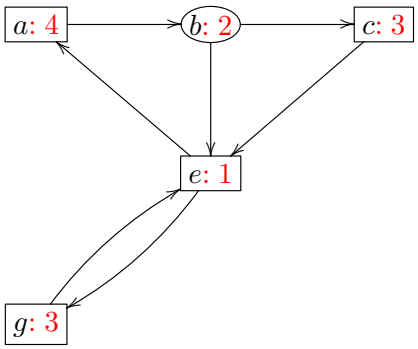

In $G - X$, $\mathsf{Force}_R(\{\, e \,\}) = (G - X)$. Hence $\mathbf{WReg}(G - X)$ returns winning regions $W'_V = \varnothing$ and $W'_R = (G - X)$.

Proof of Theorem 5.5: Correctness of WReg(G)

Proof. By induction on $n = |N|$. Base case $n = 1$ is trivial. Inductive case. $G - X$ has less than n vertices. By the induction hypothesis $\mathbf{WReg}(G - X)$ computes winning regions W'_V and W'_R.

Case 1. V can guarantee transition from v to $W'_V \cup X$ iff (a) or (b). Claim: (i) $W'_V \cup X \subseteq W_V$; and (ii) $W'_R \subseteq W_R$; this is sufficient since $W'_R \cup W'_V \cup X = N$. Thus the memoryless V-strategy on $W'_V \cup X$ is: 1. On W'_V, play the winning strategy (thanks to the induction hypothesis). 2. On X, play the "forcing" strategy, eventually reaching v. 3. From v, move back to $W'_V \cup X$. For R, use the memoryless winning strategy given by the induction hypothesis Proof of Claim (i): If the play eventually remains in W'_V then V wins by the induction hypothesis; otherwise the play passes through v infinitely often, then V wins since the priority of v, which is the least, is even. (ii) holds because, starting in W'_R, R can guarantee that the play remains in W'_R (because no V-move in $G - X$ can transition to X).

Case 2. R can guarantee transition from v to W'_R. Thus $v \in X'$. By the induction hypothesis $\mathbf{WReg}(G - X')$ computes the winning regions W''_V and W''_R. Claim: (i) $W''_R \cup X' \subseteq W_R$ (ii) $W''_V \subseteq W_V$. Proof of (i): R can move to W'_R from any move in X', and there he can guarantee that the play remains in W'_R. From a move in W''_R, V can choose to move to either W''_R or X'. In both cases, R wins the play. (ii) is clear since V can guarantee that the play remains in W''_V.

In the worst case, **WReg** is called $2^{|N|}$ times; thus running time is exponential in $|N|$. □

Uniformly Winning Memoryless Strategies Let $X \subseteq W_V$. A V-strategy is *uniformly winning on* X just if it is winning for V from every $v \in X$.

Lemma 5.9. *Given a parity game, if $W_V \neq \varnothing$ then V has a memoryless strategy that is uniformly winning on W_V.*

Proof. Exercise □

PARITY is the decision problem: Given a parity game $G = \langle\, N, \to, v_0, \lambda, \Omega \,\rangle$, is $v_0 \in W_V$?

Proposition 13. PARITY $\in$ **NP** $\cap$ **co-NP**

Proof. We first show that PARITY is in **NP**. Guess a uniformly winning V-strategy, which is succinct i.e. its size is $O(|N|)$. It can be verified in time polynomial in $|G|$ whether $v_0 \in W_V$. Claim: $v_0 \in W_V$ iff in the *strategy transition graph* (i.e. one edge from V-vertices and all edges from R-vertices), V cannot enter a loop from v_0 such that the least priority is odd. Suppose (w.l.o.g.) that the least odd and the largest priorities are 1 and p (even) respectively: just verify for the subgraphs over

$$\bigcup_{i=1}^{p} \Omega^{-1}(i), \bigcup_{i=3}^{p} \Omega^{-1}(i), \cdots, \bigcup_{i=p-1}^{p} \Omega^{-1}(i)$$

whether they contain a strongly connected component reachable from v_0 which meets the priorities $1, 3, \cdots, p$ respectively.

The complementary problem $v_0 \in (N - W_V)$ is just $v_0 \in W_R$, which is in **NP** using the same argument but with the players swapped. □

One of the best known open problem in the foundations of verification:

Conjecture 5.1. PARITY $\in$ **P**

Determinacy for Finite Parity Games

Theorem 5.6 (Determinacy)**.** *Let $G = \langle N, \to, v_0, \lambda, \Omega \rangle$ be a finite parity game. Then $N \subseteq W_V \cup W_R$. Further if $v \in W_V$ (resp. W_R) then V (resp. R) has a memoryless winning strategy from v.*

Proof. Suppose the image of Ω is $\{0, 1, \cdots, p-1\}$. Proof is by induction on p. Base case is trivial. Inductive case: if least priority is odd, swap players. Let HF_R be the set of vertices from which R has a memoryless winning strategy, and let ρ be a memoryless R-strategy that is *uniformly winning* on HF_R. It suffices to show that V has a memoryless winning strategy from every vertex in $N - HF_R$.

Case 1. No vertex in $G - HF_R$ has the least priority, 0. Apply the induction hypothesis to the subgame $G - HF_R$, since no vertex in it has priority 0. I.e. $G - HF_R$ can be partitioned into winning regions W'_V and W'_R in which memoryless winning strategies exist for the respective players. Now $W_R = HF_R \cup W'_R$ and $W_V = W'_V$; G is thus partitioned and the respective players have the appropriate memoryless winning strategies.

Case 2. $G - HF_R$ contains a vertex with priority 0. Claim: V can guarantee that, starting from a vertex in $G - HF_R$, the play remains there. Now either the play stays in $(G - HF_R) - \mathsf{Force}_V(\Omega^{-1}(0) - HF_R)$ or it visits $\mathsf{Force}_V(\Omega^{-1}(0) - HF_R)$ infinitely often. In the former case, V wins by the induction hypothesis with a memoryless strategy; in the latter, V wins by infinitely many visits to a vertex with priority 0, also with a memoryless strategy. □

Additional Topics

- Equivalence of alternating parity automata and modal mu-calculus (word languages, ranked trees, and graphs).
- Understanding alternation of fixpoints:

$$\mu Y_1.\nu Z_1.\cdots \mu Y_n.\nu Z_n.\varphi(Y_1, \cdots, Y_n, Z_1, \cdots, Z_n)$$

  In terms of expressive power, the alternation depth hierarchy is *strict*.

- AFMC contains CTL, but not CTL*. Model checking AFMC is **P**-complete.

Problems

5.1 This question proves the basic result: *modal mu-calculus is bisimulation invariant.*

First a definition. We say that a relation $B \subseteq S_1 \times S_2$ between the states of labelled transition systems $\mathcal{T}_1$ and $\mathcal{T}_2$, where $\mathcal{T}_i = \langle\, S_i, \longrightarrow_i, \rho_i \,\rangle$ $(i = 1, 2)$, is a *bisimulation* just if whenever $(s, t) \in B$

- for all $P \in \mathit{Prop}$, $s \in \rho_1(P)$ iff $t \in \rho_2(P)$
- for all $a \in \mathcal{L}$
  - for all $s' \in S_1$, if $s \stackrel{a}{\longrightarrow}_1 s'$ then there exists a state $t' \in S_2$ such that $t \stackrel{a}{\longrightarrow}_2 t'$ and $(s', t') \in B$, and
  - for all $t' \in S_2$, if $t \stackrel{a}{\longrightarrow}_2 t'$ then there exists a state $s' \in S_1$ such that $s \stackrel{a}{\longrightarrow}_1 s'$ and $(s', t') \in B$.

Two states s and t are *bisimulation equivalent*, written $s \sim_{\mathcal{T}_1,\mathcal{T}_2} t$, just if there is a bisimulation relation B such that $(s, t) \in B$.

Prove that if $s \sim_{\mathcal{T}_1,\mathcal{T}_2} t$ then for all modal mu-calculus sentences φ, we have $s \models^{\mathcal{T}_1} \varphi$ iff $t \models^{\mathcal{T}_2} \varphi$.

5.2 Prove the *Signature Decrease Lemma*: Whenever the left-hand side of the relation in question is defined, we have:

(a) $\mathit{sig}^\mu(s, \varphi_1 \vee \varphi_2) = \mathit{sig}^\mu(s, \varphi_i)$, for some $i \in \{\, 1, 2 \,\}$.

(b) $\mathit{sig}^\mu(s, \varphi_1 \wedge \varphi_2) \geq \max(\mathit{sig}^\mu(s, \varphi_1), \mathit{sig}^\mu(s, \varphi_2))$

(c) $\mathit{sig}^\mu(s, \langle a \rangle \varphi_1) = \mathit{sig}^\mu(t, \varphi_1)$, for some t such that $s \stackrel{a}{\to} t$

(d) $\mathit{sig}^\mu(s, [a]\varphi_1) \geq \mathit{sig}^\mu(t, \varphi_1)$, for all t such that $s \stackrel{a}{\to} t$

(e) Assuming Z_i is a ν-variable, $\mathit{sig}^\mu(s, \nu Z_i.\psi_i) = \mathit{sig}^\mu(s, Z_i) = \mathit{sig}^\mu(s, \psi_i)$.

(f) Assuming $Z_i = Y_j$ is a μ-variable

  (a) $\mathit{sig}^\mu(s, \mu Y_j.\chi_j) =_{j-1} \mathit{sig}^\mu(s, Y_j)$

  (b) $\mathit{sig}^\mu(s, Y_j)(j) > \mathit{sig}^\mu(s, \chi_j)(j)$ and $\mathit{sig}^\mu(s, Y_j) =_{j-1} \mathit{sig}^\mu(s, \chi_j)$ (and so $\mathit{sig}^\mu(s, Y_j) > \mathit{sig}^\mu(s, \chi_j)$).

5.3 Prove the following.

(i) Every infinite path in a pre-tableau passes through a (mod)-rule infinitely often.

(ii) Let $t = (V, \to, \lambda)$ be a tableau and suppose $v \in V$. Then for all $P \in \mathit{Prop}$ we have $\{P, \neg P\} \not\subseteq \lambda(v)$.

5.4 Show in the proof of the Soundness Lemma 5.7 there is in fact a memoryless strategy for verifier.

Hint: you need to address the following issue. If between two (mod)-rules a disjunction, say $\psi_0 \vee \psi_1$, is broken down more than once with a different disjunct chosen, what will a successful strategy for verifier pick without referring to the history of the play so far?

5.5 Prove that every mu-calculus model checking game $\mathcal{G}_V^T(s,\varphi)$ determines an equivalent parity game namely define a parity game $\mathcal{G}_V^T[s,\varphi]$ such that *Player P has a memoryless winning strategy for* $\mathcal{G}_V^T(s,\varphi)$ *if and only if Player P has a memoryless winning strategy for* $\mathcal{G}_V^T[s,\varphi]$.

5.6 Assume the notations of the lecture course. Let s_0 be a state of a finite labelled transition system T. WLOG let φ be a closed formula in positive normal form with no occurrences of atomic propositions. Suppose φ has n fixpoint variables; and m least fixpoint variables $Y_1, \cdots, Y_m$, naming subformulas $\mu Y_1.\chi_1, \cdots, \mu Y_m.\chi_m$ in decreasing order of size.

Fix a memoryless V-strategy σ for the game $\mathcal{G}_\varnothing^T(s_0,\varphi)$. A *signature assignment* is an assignment $\mathcal{S}$ of signatures (of length m) to each position (t,ψ). We say that a signature assignment is σ*-consistent* just if for each position u

- if $u = (s, Y_j)$ then $\mathcal{S}((s,\chi_j)) <_j \mathcal{S}(u)$ and $\mathcal{S}((s,\chi_j)) =_{j-1} \mathcal{S}(u)$
- if $u = (s, \mu Y_j.\chi_j)$ then $\mathcal{S}((s,Y_j)) =_{j-1} \mathcal{S}(u)$
- if u is a V-position then $\mathcal{S}(\sigma(u)) \leq \mathcal{S}(u)$
- if u is a R-position then for all successor vertices v we have $\mathcal{S}(v) \leq \mathcal{S}(u)$.

(i) Prove that $s \vDash_\varnothing^T \varphi$ if, and only if, there is a memoryless V-strategy σ and a σ-consistent signature assignment $\mathcal{S}$.

(ii) Hence prove that the modal mu-calculus model checking problem is in **NP** $\cap$ **co-NP**.

5.7 The model checking problem of the modal mu-calculus can be given a game characterization:

(a) $s \vDash_V^T \varphi$ iff player V has a memoryless winning strategy for $\mathcal{G}_V^T(s,\varphi)$.

(b) $s \nvDash_V^T \varphi$ iff player R has a memoryless winning strategy for $\mathcal{G}_V^T(s,\varphi)$.

In the lectures, we proved (i). This question proves (ii).

Henceforth we fix a transition system T and a normal formula φ. Let

$$\nu Y_1.\chi_1,\ \nu Y_2.\chi_2,\ \cdots,\ \nu Y_m.\chi_m$$

be the set of greatest fixpoint formulas in φ, again in decreasing order of size.
We define valuations $V_0, \cdots, V_n$ by induction:

$$\begin{aligned} V_0 &= V \\ V_{i+1} &= V_i[Z_{i+1} \mapsto \|\sigma_{i+1} Z_{i+1}.\psi_{i+1}\|_{V_i}^T] \end{aligned}$$

Thus we can make sense of $\|\psi\|^T_{V_n} \subseteq S$, for any $\psi \in Sub(\varphi)$.

Given a signature $r = \alpha_1, \cdots, \alpha_m$ and a valuation V, we define *ν-valuations*[2] $V^r_0, \cdots, V^r_n$ by induction:

$$\begin{aligned} V^r_0 &= V \\ V^r_{i+1} &= V^r_i[E_{i+1}/Z_{i+1}] \end{aligned}$$

where

$$E_{i+1} = \begin{cases} \|\sigma_{i+1} Z_{i+1}.\psi_{i+1}\|^T_{V^r_i} & \text{if } \sigma_{i+1} = \mu \\ \|\nu Y_j^{\alpha_j}.\chi_j\|^T_{V^r_i} & \text{if } \sigma_{i+1} Z_{i+1} = \nu Y_j \end{cases}$$

Recall that if $s \notin \|\nu Z.\psi\|^T_V$ then there is an ordinal α such that $s \notin \|\nu Z^\alpha.\psi\|^T_V$ and for all $\beta < \alpha$, we have $s \in \|\nu Z^\beta.\psi\|^T_V$.

Let $\psi \in Sub(\varphi)$. If $s \notin \|\psi\|^T_V$, we define the *ν-signature* of (s, ψ), written $sig^\nu(s, \psi)$, to be the $<$-least signature $r = \alpha_1 \cdots \alpha_m$ such that $s \notin \|\psi\|^T_{V^r_n}$.

(a) Prove the *Signature Decrease Lemma*: Whenever the left-hand side of the equation or inequality is defined:

i. $sig^\nu(s, \varphi_1 \wedge \varphi_2) = sig^\nu(s, \varphi_1)$ or $sig^\nu(s, \varphi_1 \wedge \varphi_2) = sig^\nu(s, \varphi_2)$

ii. $sig^\nu(s, \varphi_1 \vee \varphi_2) \geq \max(sig^\nu(s, \varphi_1), sig^\nu(s, \varphi_2))$

iii. $sig^\nu(s, [a]\varphi_1) \geq sig^\nu(t, \varphi_1)$, for some t such that $s \xrightarrow{a} t$

iv. $sig^\nu(s, \langle a\rangle\varphi_1) \geq sig^\nu(t, \varphi_1)$ for all t such that $s \xrightarrow{a} t$, or s does not have an a-transition

v. Assuming Z_i is a μ-variable, $sig^\nu(s, \mu Z_i.\psi_i) = sig^\nu(s, Z_i) = sig^\nu(s, \psi_i)$

vi. Assuming $Z_i = Y_j$ is a ν-variable, $sig^\nu(s, \nu Y_j.\chi_j)$ is the same as $sig^\nu(s, Y)$ on the first $j-1$ components

vii. Assuming $Z_i = Y_j$ is a ν-variable, $sig^\nu(s, Y_j) > sig^\nu(s, \chi_j)$, and the first $j-1$ components of the two signatures are the same but not the j-component.

(b) Hence prove that *if $s \not\Vdash^T_V \varphi$ then R has a memoryless winning strategy.*

5.8 Let φ be a monotonic function on a powerset 2^S. Prove that for every $U \subseteq S$,

$$U \subseteq \nu X.\varphi(X) \quad \Leftrightarrow \quad U \subseteq \varphi(\nu X.(U \cup \varphi(X))).$$

[2] This is dual to, and not to be confused with, the μ-valuations as defined in the lectures.

Chapter 6

Tree Automata, Rabin's Theorems and S2S

Synopsis

Büchi tree automata. Deterministic Muller tree automata are more expressive. Complementation of parity tree automata via determinacy of parity games. Rabin's Basis Theorem. S2S: syntax and semantics. Expressivity of S2S: examples. Rabin Tree Theorem: S2S is decidable.

6.1 Trees and Tree automata

We are interested in infinite (full) binary trees whose nodes are labelled by letters of an alphabet Σ. Formally a Σ-*labelled binary tree* is a function $t : \{0,1\}^* \to \Sigma$ i.e. the label of the tree t at node $u \in \{0,1\}^*$ is $t(u)$.

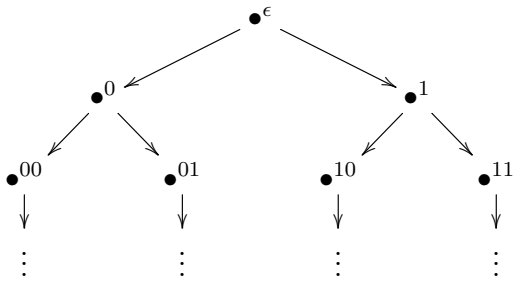

We write $\mathfrak{T}^{\omega}_{\Sigma}$ for the collection of Σ-labelled binary trees. Henceforth by a tree, we mean an element of $\mathfrak{T}^{\omega}_{\Sigma}$. A *tree language* is just a subset $T \subseteq \mathfrak{T}^{\omega}_{\Sigma}$. A *maximal path* (or, simply, *path*) of a tree t is a sequence $\pi = u_0\, u_1\, u_2 \cdots$ of tree nodes whereby $u_0 = \epsilon$ (the root of the tree) and $u_{i+1} = u_i\, 0$ or $u_{i+1} = u_i\, 1$, for every $i \geq 0$. The ω-*word determined by* π is $t(u_0)\; t(u_1)\; t(u_2) \cdots \in \Sigma^{\omega}$.

Definition 6.1. A *tree automaton* A (for Σ-labelled binary trees) is a tuple $(Q, \Sigma, q_0, \Delta, Acc)$, where

- Q is the finite set of states, q_0 is the initial state
- $\Delta \subseteq Q \times \Sigma \times Q \times Q$ is the transition relation, and
- Acc is the acceptance condition (such as Büchi, Muller, Rabin and Parity).

The automaton is *deterministic* just if for every q and a, there is at most one transition (or quadruple) in Δ the first two components of which are q and a.

A *run-tree* of a tree automaton A over a tree t is an assignment of states to tree nodes i.e. a function $\rho : \{0,1\}^* \to Q$ such that

- $\rho(\epsilon) = q_0$, and
- for all $u \in \{0,1\}^*$, $(\rho(u), t(u), \rho(u\,0), \rho(u\,1)) \in \Delta$.

Note that a run-tree is a Q-labelled binary tree.

Let Acc be an acceptance condition for automata over ω-words. A run-tree ρ is accepting w.r.t. condition Acc just if every path of ρ is accepting w.r.t. Acc; we say that a tree automaton $A = (Q, \Sigma, q_0, \Delta, Acc)$ accepts a given tree t just if there is run-tree of A over t which is accepting w.r.t. Acc. Thus we have *Büchi tree automata, Muller tree automata, Rabin tree automata, Parity tree automata*, etc. For example, a Büchi tree automaton $A = (Q, \Sigma, q_0, \Delta, F)$ *accepts* a tree t just if there exists a run-tree ρ of A over t such that in every path of ρ, a final state from F occurs infinitely often.

The *tree language* recognised by the tree automaton A, denoted $L(A)$, is the set of trees accepted by A.

Example 6.1. Consider the language T_1 of $\{a,b\}$-labelled binary trees t such that t has a path with infinitely many a's. The Büchi tree automaton $(\{q_a, q_b, \top\}, \{a,b\}, q_a, \Delta, \{q_a, \top\})$ where

$$\Delta \;:\; \begin{cases} (q_*, a) & \mapsto & \{(q_a, \top), (\top, q_a)\} \\ (q_*, b) & \mapsto & \{(q_b, \top), (\top, q_b)\} \\ (\top, *) & \mapsto & \{(\top, \top)\} \end{cases}$$

recognises the language T_1 (where $*$ mean a or b). Clearly the automaton accepts every tree from the state $\top$. Observe that every run-tree has a single path labelled with states q_a (and possibly q_b), and the rest of the run-tree is labelled with $\top$. There are infinitely many states q_a on this path if and only if there are infinitely many vertices labelled by a on the path of the input tree. Thus the non-deterministic automaton descends the input tree t guessing such a path.

Example 6.2. Consider the language $T_2 := \mathfrak{T}^{\omega}_{\{a,b\}} \setminus T_1$. I.e. T_2 consists of $\{a,b\}$-labelled binary trees t such that every path of t has only finitely many a. Define a deterministic Muller tree automaton $(\{q_a, q_b\}, \{a,b\}, q_a, \Delta, \{\{q_b\}\})$ where

$$\Delta \;:\; \begin{cases} (q_*, a) & \mapsto & \{(q_a, q_a)\} \\ (q_*, b) & \mapsto & \{(q_b, q_b)\} \end{cases}$$

Then, given a tree t, for each path π of t, there are infinitely many occurrences of b (respectively a) on π if and only if the corresponding path of the unique run-tree ρ over t has infinitely many occurrences of q_b (respectively q_a); it follows that $t \in T_2$ if and only if for every path in the unique run-tree, the set of infinitely occurring states is $\{q_b\}$. Hence the automaton recognises T_2.

In constrast to automata over ω-words, deterministic Muller tree automata and non-deterministic Büchi tree automata are not equivalent.

Theorem 6.1. *The language T_2 (of Example 6.2) is not recognisable by any Büchi tree automaton, whether deterministic or not.*

Proof. Assume, for a contradiction, that the Büchi tree automaton $A = (Q, \{a, b\}, q_0, \Delta, F)$ recognises T_2. Let $n = |F| + 1$. Consider the following $\{a, b\}$-labelled binary tree t:

$$t \,:\, u \mapsto \begin{cases} a & u \in (1^+0)^i \text{ for } i \in \{1, \cdots, n\} \\ b & \text{otherwise} \end{cases}$$

Since $t \in T_2$, there is a accepting run-tree ρ of A on t. On the path 1^ω, a final state is visited, say, at $v_0 = 1^{m_0}$. On the path $1^{m_0}\,0\,1^\omega$, final states are visited infinitly often. Suppose a final state is reached, for the first time after 0, at $v_1 = 1^{m_0}\,0\,1^{m_1}$. By repeated the argument, we obtain visits to final states at the nodes $v_0 = 1^{m_1}, v_1 = 1^{m_0}\,0\,1^{m_1}, \cdots, v_n = 1^{m_0}\,0\,1^{m_1}\,0 \cdots 0\,1^{m_n}$. Now, there exist $i < j$ such that the same final state appears at v_i and v_j. It follows from the definition of t that between v_i and v_j, at least one label a occurs (at the node $v_i\,0$).

Construct a new tree t' by copying the (respective labels of the) part between v_i and v_j repeatedly. Similarly we construct the corresponding run-tree ρ' form ρ. Thus A also accepts t' i.e. $t' \in T_2$, but in t', infinitely many a occur on the new path π, which is a contradiction. □

6.2 Parity tree automata

A *parity tree automaton* is a tuple $A = (Q, \Sigma, q_0, \Delta, \Omega)$ with priority map $\Omega : Q \to \{0, \cdots, k\}$. It accepts a tree t just if there is a run-tree ρ of A over t such that for every path of ρ, the least priority that occurs infinitely often is even.

Example 6.3. Consider the parity tree automaton $(\{q_a, q_b\}, \{a, b\}, q_a, \Delta, \Omega)$ where Δ is as defined in Example 6.2, and $\Omega : q_a \mapsto 1; q_b \mapsto 2$. For every path in a given tree, there are only finitely many occurrences of a if, and only if, the there are finitely many occurrences of q_a in the corresponding path of the (unique) run-tree, which is so if, and only if, the least priority that occurs infinitely often is 2. Thus the automaton recognises T_2.

Theorem 6.2. *(i) There is algorithm that, given a parity tree automaton, constructs an equivalent Muller tree automaton.*

(ii) There is algorithm that, given a Muller tree automaton, constructs an equivalent Parity tree automaton.

Proof. Exercise. □

Closure properties of parity tree automata

Lemma 6.1 (Closure under Union)**.** *Given parity tree automata A_1 and A_2, one can construct a parity tree automaton recognising $L(A_1) \cup L(A_2)$.*

Proof. Assume that the state-sets Q_1 and Q_2 of A_1 and A_2 respectively are disjoint, with initial states q_1 and q_2. We define the new automaton with state-set $\{q_0\} \cup Q_1 \cup Q_2$ with new initial state q_0 of prioirty 0 (say). Take all transitions of A_1 and A_2; add, for every transition (q_1, a, r_1, r_2) or (q_2, a, r_1, r_2), the new transition (q_0, a, r_1, r_2). □

Given trees $s \in \mathfrak{T}^\omega_\Gamma, t \in \mathfrak{T}^\omega_\Sigma$, define the tree $s \times t \in \mathfrak{T}^\omega_{\Gamma\times\Sigma}$ by $s \times t(u) = (s(u), t(u))$. The Σ-*projection* of $s \times t$ is the tree t. Given a language T consisting of $(\Gamma \times \Sigma)$-labelled binary trees, define

$$\pi_\Sigma(T) := \{\, t \in \mathfrak{T}^\omega_\Sigma \mid \exists s : s \times t \in T \,\}$$

Lemma 6.2 (Closure under Projection)**.** *Given a parity tree automaton recognising the language T of $(\Gamma \times \Sigma)$-labelled binary trees, one can construct a parity tree automaton recognising $\pi_\Sigma(T)$.*

Proof. Given a parity tree automaton A over $\Gamma \times \Sigma$, define a parity tree automaton B over Σ which does, in any step, the following. For input letter $b \in \Sigma$ guess the Γ-component a and proceed by an A-transition for the input letter (a, b). The state set is not changed. □

Notation Following Vardi and Kupferman, we use acronym XYZ where

- X ranges over automaton modes: deterministic, non-determinisitc and alternating
- Y ranges over acceptance / winning conditions: Büchi, Muller, Rabin, Streett and Parity
- Z ranges over word and tree automaton.

For example, DMW and NPT are shorthand for deterministic Muller word automaton and non-deterministic parity tree automaton respectively.

6.3 Parity Games and Complementation

Acceptance Parity Game Given a NPT $A = (Q, \Sigma, q_I, \Delta, \Omega)$ and a tree t, we define a parity game, called the *acceptance parity game*, $\mathcal{G}_{A,t} = (N, E, (\epsilon, q_I), \lambda, \Omega')$ as follows. Writing $N_V := \lambda^{-1}(V)$ and $N_R := \lambda^{-1}(R)$

- $N_V = \{\, 0, 1 \,\}^* \times Q$
- $N_R = \{\, 0, 1 \,\}^* \times (Q \times Q)$
- for each vertex $(v, q) \in N_V$, for each transition $(q, a, q_0, q_1) \in \Delta$ with $t(v) = a$, we have

$$((v, q), (v, (q_0, q_1)) \in E$$

- for each vertex $(v, (q_0, q_1)) \in N_R$, we have

$$((v, (q_0, q_1)), (v\,0, q_0)), ((v, (q_0, q_1)), (v\,1, q_1)) \in E.$$

- $\Omega' : (v, q) \mapsto \Omega(q)$ and $(v, (q_0, q_1)) \mapsto \max(\Omega(q_0), \Omega(q_1))$.

It follows from the definition of the game $\mathcal{G}_{A,t}$ that there is a one-one correspondence between accepting run-trees of A over t and winning strategies for Verifier in $\mathcal{G}_{A,t}$.

Lemma 6.3 (Run)**.** *Verifier has a winning strategy in $\mathcal{G}_{A,t}$ from vertex (ϵ, q_I) if and only if $t \in L(A)$.*

Given an "input-free" parity tree automaton $A = (Q, q_I, \Delta, \Omega)$ where $\Delta \subseteq Q \times Q \times Q$, we define a simpler parity game $\mathcal{G}_A = (N, E, q_I, \Delta, \Omega')$ in which the tree t and the parameter w in the game positions are suppressed:

- $N_V = Q$
- $N_R = Q \times Q$
- $E = \{\, (q, (q', q'')), ((q', q''), q'), ((q', q''), q'') \mid (q, q', q'') \in \Delta \,\}$
- $\Omega' : q \mapsto \Omega(q)$ and $(q_0, q_1) \mapsto \max(\Omega(q_0), \Omega(q_1))$.

Lemma 6.4 (Run: input-free)**.** *Verifier has a winning strategy in $\mathcal{G}_A$ from vertex q_I if and only if the tree automaton A has an accepting run.*

Theorem 6.3 (Memoryless Determinacy)**.** *Let G be a parity game. From every vertex of G, one of the two players has a memoryless winning strategy.*

Proof. In descriptive set theory, the *Borel determinacy theorem* (Martin, 1975) states that every Gale-Stewart game whose winning set is a Borel set is determined. Parity games lie in the third level of the Borel Hierarchy, hence they are determined. Proofs of memoryless determinacy of parity games can be found in (Emerson and Jutla, 1991; Mostowski, 1991; Zielonka, 1998). □

A major theorem of the chapter is the closure of NPT under complementation. The result was first obtained by Michael Rabin in a landmark paper (Rabin, 1969). Gurevich and Harrington (1982) used games to simplify the proof; they proved a bounded memory theorem for the Rabin condition. A further simplification for the parity condition was proved independently by Emerson and Jutla (1991) and by Mostowski (1991).

Theorem 6.4 (Closure under Complementation)**.** *There is an algorithm that, given a parity tree automaton A, constructs a Muller tree automaton that accepts exactly the trees rejected by A.*

The proof relies on two fundamental results:

(i) memoryless determinacy of parity games (Theorem 6.3), and

(ii) closure under complementation of NPW.

Proof. Let $A = (Q, \Sigma, \Delta, q_0, \Omega)$ be a NPT of Σ-labelled binary trees. Then it follows from Lemma 6.3 that a tree $t \in (\mathfrak{T}^{\omega}_{\Sigma} \setminus L(A))$ if, and only if, Verifier does not have a winning strategy in the acceptance game $\mathcal{G}_{A,t}$. Thanks to the Memoryless Determinacy Theorem for Parity Games (Theorem 6.3), $t \in (\mathfrak{T}^{\omega}_{\Sigma} \setminus L(A))$ if, and only if, Refuter has a memoryless winning strategy in $\mathcal{G}_{A,t}$.

We aim to transform Refuter's memoryless winning strategy to a NMT B that recognises the complement of $L(A)$. A memoryless strategy for Refuter is a function $N_R \to N_V$ that maps a given vertex $(v, (q_0, q_1)) \in N_R$ to one of $(v\,0, q_0)$ and $(v\,1, q_1)$. A key observation is that such a strategy for Refuter can be represented as a *LocStr*-labelled binary tree

$$f : \{0,1\}^* \to \mathit{LocStr}$$

where *LocStr* is the finite set $((Q \times Q) \to \{0,1\})$ of *local strategies*.

Take such a strategy function f and consider the tree $t \times f : \{0,1\}^* \to \Sigma \times \mathit{LocStr}$, in which the node u has label $(t(u), f(u))$. Let $\Sigma' = \Sigma \times \mathit{LocStr} \times \{0,1\}$. Then an ω-word over Σ', $\alpha = (a_0, f_0, d_0)\,(a_1, f_1, d_1)\,\cdots$, such that each $a_i = t(d_0 \cdots d_i)$ and $f_i = f(d_0 \cdots d_i)$, represents a path $\epsilon, d_0, d_0\,d_1, \cdots$ in the tree $t \times f$. We say that a *play over* α is an ω-word $(v_0, q_0)\,(v_0, m_0)\,(v_1, q_1)\,(v_1, m_1)\,\cdots \in (N_V \cdot N_R)^{\omega}$ such that

- $v_0 = \epsilon$
- $(q_i, a_i, m_i) \in \Delta$
- $f_i(m_i) = d_i$ and $q_{i+1} = q^{d_i}$ where $m_i = (q^0, q^1)$.

Thus a play over α is just a play in the acceptance game $\mathcal{G}_{A,t}$ whereby Refuter plays according to strategy f and which determines a path α in the tree $t \times f$.

It is straightforward to construct a NPW C that accepts exactly those words α over Σ' such that there is a play over α that satisfies the parity condition. By McNaughton's Theorem, there is a NMW $\overline{C}$ that accepts the complement of this language. I.e. $\overline{C}$ accepts those words $\alpha \in \Sigma'^{\omega}$ for which all plays over α violate the parity condition.

A strategy function f is not winning for Refuter in $\mathcal{G}_{A,t}$ if, and only if, there is a play which is winning for Verifier when Refuter plays according to f. This play traces out a path of the tree $t \times f$. It follows that automaton C would accept such a path *qua* ω-word over Σ'. But if f is winning then none of the paths is accepted by C. I.e. every such path is accepted by $\overline{C}$. Therefore f is a winning strategy for Refuter if, and only if, each path of the tree $t \times f$ is accepted by $\overline{C}$.

Thus the NMT tree automaton B recognising the complement of A consists of two automata. Given an input tree t, the first constructs a strategy function f by guessing the corresponding local strategy at each node; the second runs $\overline{C}$ along each path in the tree $t \times f$. Thus the automaton B accepts a tree if, and only if, Refuter has a winning memoryless strategy function f. This completes our proof of Theorem 6.4. □

Even though parity tree automata can be complemented, they are not determinisable. The following result says that there is no hope for determinisation.

Lemma 6.5. *Consider the language consisting of $\{a, b\}$-labelled binary trees t such that t has at least one vertex labelled with a. The language is not recognisable by a deterministic tree automaton with any of the acceptance conditions we have considered.*

6.4 The Non-emptiness Problem

Recall that a noneempty regular ω-language contains an ultimately periodic ω-word. We show a corresponding result for non-empty tree languages recognisable by parity tree automata.

Definition 6.2. A tree $t \in \mathfrak{T}_{\Sigma}^{\omega}$ is said to be *regular* just if there is a deterministic finite automaton (DFA) equipped with output, which gives, for every $w \in \{0, 1\}^*$, the label $t(w)$ at node w. The automaton has the form $B = (Q_B, \{0, 1\}, q_0^B, \delta_B, f_B)$ where $\delta_B : Q \times \{0, 1\} \to Q$ is the transition function, and $f_B : Q_B \to \Sigma$ is the output function.

Lemma 6.6. *A tree $t \in \mathfrak{T}_{\Sigma}^{\omega}$ is regular if and only if there is a deterministic input-free tree automaton C with state-set $Q \times \Sigma$ such that the Σ-projection of the unique run-tree of C is the tree t.*

Theorem 6.5 (Rabin Basis Theorem)**.** *(i) The emptiness problem for parity tree automata is decidable.*

(ii) If a parity tree automaton accepts some tree then it accepts a regular tree.

Proof. (i) Let $A = (Q, \Sigma, q_0, \Delta, \Omega)$ be a NPT. We define an input-free NPT $A' = (Q \times \Sigma, \{q_0\} \times \Sigma, -, \Delta', \Omega')$ which non-deterministically generates an input tree t, and processes t like A. Note that A' has possibly several initial states. By Lemma 6.4, A' has an accepting run from (q_0, a) if, and only if, Verifier has a winning strategy from (q_0, a) in the parity game $\mathcal{G}_{A'}$.

(ii) Suppose $L(A) \neq \varnothing$ i.e. A' has an accepting run. It follows that in the parity game $\mathcal{G}_{A'}$, Verifier has a *memoryless* winning strategy from (q_0, a). The memoryless strategy induces a deterministic tree automaton as a "subautomaton" of A', where for each state (q, a), only one transition exists for the continuation of the run. This tree automaton generates a regular tree, which is accepted by A, by construction of A'.

□

6.5 S2S and Rabin's Tree Theorem

The Logical System S2S The logical system $S2S$ (monadic second-order logic of 2 successors) is defined over first-order variables $x, y, \cdots$ ranging over $\{0,1\}^*$ (nodes in the full binary tree) and over second-order variables $X, Y, \cdots$ ranging over $2^{\{0,1\}^*}$ (sets of nodes of the full binary tree).

Terms are built up from first-order variables and ϵ by the two successors, represented as concatenation with 0 and 1 respectively.

Let s and t are terms. The *atomic formulas* are

- $X(s)$ "s is in X"
- $s \leq t$ "s is a prefix of t"
- $s = t$ "s is equal to t".

The formulas of S2S are built up from the atomic formulas using the standard boolean connectives, and closed under first- and second-order quantifiers $\exists$ and $\forall$.

The *structure* of the infinite full binary tree is $\mathbf{t}_2 = (\mathbb{B}^*, \epsilon, S_0, S_1)$ where S_i is the i-th successor function: $S_0(u) = u\,0$ and $S_1(u) = u\,1$ for $u \in \mathbb{B}^*$. The *theory S2S* is the set of S2S-sentences that are true in $\mathbf{t}_2$.

Semantics of S2S S2S-formulas $\varphi(X_1, \cdots, X_n)$, with free 2nd-order variables from $X_1, \cdots, X_n$, are interpreted in expanded structures $\widehat{t} = (\mathbf{t}_2, P_1, \cdots, P_n)$. We write

$$\widehat{t} \;\vDash\; \varphi(X_1, \cdots, X_n)$$

just if $\widehat{t}$ satisfies $\varphi(\overline{X})$. We identify $\widehat{t}$ with the infinite tree $t \in \mathfrak{T}^{\omega}_{\mathbb{B}^n}$ whereby for each $u \in \mathbb{B}^*$, we have

$$t(u) = (b_1, \cdots, b_n) \quad \text{where } b_i = 1 \;\leftrightarrow\; u \in P_i.$$

Given an S2S formula $\varphi(\overline{X})$ the *tree language* defined by $\varphi(\overline{X})$ is the set

$$L(\varphi) := \{\, t \in \mathfrak{T}^{\omega}_{\mathbb{B}^n} \mid \widehat{t} \vDash \varphi \,\}.$$

Theorem 6.6. *A tree language is S2S-definable if and only if it is recognisable by a NPT.*

Theorem 6.7 (Rabin Tree Theorem)**.** *The theory S2S is decidable.*

Bibliography

J. Bradfield and C. P. Stirling. Modal logics and mu-calculi. In A. Ponse and S. Smolka, editors, *Handbook of Process Algebra*, pages 293–332. Springer-Verlag, 2001.

J. Bradfield and C. P. Stirling. Modal mu-calculi. In A. Ponse and S. Smolka, editors, *Handbook of Modal Logic*, pages 721–756. 2007. Studies in Logic and Practical Reasoning Volume 3.

J. R. Büchi. Weak second order arithmetic and finite automata. *Zeitschrift für Maths. Logik und Grundlagen Maths.*, 6:66–92, 1960a.

J. R. Büchi. On a decision method in restricted second order arithmetic. In *Proc. International Congress on Logic, Methodology and Philosophy of Science*, pages 1–11. Stanford Univ. Press, 1960b.

J. R. Büchi. Weak second order arithmetic and finite automata. In *Proc. Int. Congr. Logic, Methodology and Philosophy of Science*, pages 1–12. Stanford University Press, 1962.

Thomas Colcombet and Konrad Zdanowski. A tight lower bound for determinization of transition labeled büchi automata. In *ICALP(2)*, pages 151–162, 2009.

Lawrence C. Eggan. Transition graphs and the star-height of regular events. *Michigan Mathematical Journal*, 10:385–397, 1963.

C. C. Elgot. Decision problems of finite automata design and related arithmetics. *Trans. Amer. Math. Soc.*, 98:21–52, 1961.

E. A. Emerson and C. S. Jutla. Tree automata, mu-calculus and determinacy. In *FOCS*, pages 368–377, 1991.

H. B. Enderton. *Elements of Set Theory.* Academic Press, 1977.

Paul Gastin and Denis Oddoux. Fast LTL to büchi automata translation. In *CAV*, pages 53–65, 2001.

Erich Grädel, Wolfgang Thomas, and Thomas Wilke. *Automata, Logics and Infinite Games.* Springer-Verlag, 2002. LNCS Vol. 2500.

Y. Gurevich and L. Harrington. Tree, automata and games. In *STOC*, pages 60–65, 1982.

Moritz Hammer, Alexander Knapp, and Stephan Merz. Truly on-the-fly LTL model checking. In *TACAS*, pages 191–205, 2005.

Kosaburo Hashiguchi. Algorithms for determining relative star height and star height. *Inf. Comput.*, 78(2):124–169, 1988.

Neil Immerman. *Descriptive Complexity*. Springer, New York, 1999. Graduate Texts in Computer Science.

H. W. Kamp. *The temporal logic of programs*. PhD thesis, University of California, Los Angeles, 1968.

B. Khoussainov and A. Nerode. *Automata Theory and its Applications*, volume 21 of *Progress in Computer Science and Applied Logic*. Birkhäuser, 2001.

Daniel Kirsten. Distance desert automata and the star height problem. *ITA*, 39(3):455–509, 2005.

D. Kozen. *Theory of Computation*. Springer-Verlag, 2006.

Wanwei Liu and Ji Wang. A tighter analysis of piterman's büchi determinization. *Inf. Process. Lett.*, 109:941–945, 2009.

Christof Löding and Wolfgang Thomas. Alternating automata and logics over infinite words. In *IFIP TCS*, pages 521–535, 2000.

D. A. Martin. Borel determinacy. *The Annals of Mathematics*, 102(2):363–371, 1975.

R. McNaughton. Testing and generating infinite sequences by a finite automaton. *Information and Control*, 9:521–530, 1966.

A. R. Meyer and L. J. Stockmeyer. The equivalence problem for regular expressions with squaring requires exponential time. In *Proc. 13th IEEE Symp. on Swithcing and Automata Theory*, pages 125–129, 1972.

A. W. Mostowski. Games with forbidden positions. Technical Report 78, 1991.

D. Muller. Infinite sequences and finite machines. In *Proc. 4th Ann. IEEE Symp. Switching Circuit Theory and Logical Design*, pages 3–16, 1963.

D. Niwinski and I. Walukiewicz. Games for the mu-calculus. *Theoretical Computer Science*, 163:99–116, 1997.

C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

J. P. Pécuchet. On the complementation of büchi automata. *Theoretical Computer Science*, 47:95–98, 1986.

Nir Piterman. From nondeterministic büchi and streett automata to deterministic parity automata. *Logical Methods in Computer Science*, 3, 2007.

Amir Pnueli. The temporal logic of programs. In *Proc. 18th IEEE Symp. Found. of Comp. Sci.*, pages 46–57, 1977.

M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Trans. Amer. Maths. Soc*, 141:1–35, 1969.

A. Rabinovich. A proof of Kamp's theorem. In *CSL*, pages 516–527, 2012.

S. Safra. On the complexity of ω-automaton. In *Proc. 29th IEEE Symp. Foundations of Comp. Sc.*, pages 319–327, 1988.

Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4:177–192, 1970.

Sven Schewe. Tighter bounds for the determinisation of büchi automata. In *FoSSaCS*, pages 167–181, 2009.

Michael Sipser. *Introduction to the Theory of Computation.* Thomson Course Technology, 2005.

A. P. Sistla, M. Y. Vardi, and P. Wolper. The complementation problem for büchi automata with applications to temporal logic. *Theoretical Computer Science*, 49:217–237, 1987.

A. Prasad Sistla and Edmund M. Clarke. The complexity of propositional linear temporal logics. *J. ACM*, 32(3):733–749, 1985.

C. P. Stirling. Bisimulation, model checking and other games. Notes for Mathfit instructional meeting on games and computation, Edinburgh, 1997.

C. P. Stirling. *Modal and Temporal Properties of Processes.* Springer-Verlag, 2001. Texts in Computer Science.

Robert S. Streett and E. Allen Emerson. An automata theoretic decision procedure for the propositional mu-calculus. *Inf. Comput.*, 81(3):249–264, 1989.

R. E. Tarjan. Depth-first search and linear graph algorithms. *SIAM J. Computing*, 1:146–160, 1972.

W. Thomas. Automata on infinite objects. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume B*, pages 134–191. Elsevier, 1990.

W. Thomas. Languages, automata and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3. Springer-Verlag, 1997.

W. Thomas. Solution of Church's problem: A tutorial. In K. Apt and R. van Rooij, editors, *New Perspectives on Games and interaction*, volume 4. Amsterdam University Press, 2008.

M. Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Proceedings of Banff Higher Order Workshop*, pages 238–266. Springer-Verlag, 1996. LNCS Vol. 1043.

M. Y. Vardi. Branching *vs* linear time: final showdown. In *ETAPS 2001.* Springer-Verlag, 2001.

I. Walukiewicz. *Notes on the Propositional Mu-calculus: Completeness and Related Results.* BRICS NS. BRICS, Computer Science Department, University of Aarhus, 1995.

Wieslaw Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theor. Comput. Sci.*, 200(1-2):135–183, 1998.

Appendix A

Ordinals and Transfinite Induction: A Primer

Everybody is familiar with the set $\omega = \{0, 1, 2, \cdots\}$ of *finite ordinals* or *natural numbers*. A standard mathematical tool is the *Principle of Mathematical Induction* on this set. How does one count beyond the finite ordinals? Is there an associated induction principle?

Ordinals are certain sets of sets.

- There are two kinds: *successors* and *limits*.
- They are *well-ordered*.
- There are a lot of them.
- We can do induction on them, using the *Principle of Transfinite Induction*.

A good reference is Enderton's book (Enderton, 1977).

We assume the axioms of Zermelo-Fraenkel set theory. It is impossible to understand ordinals and transfinite induction properly outside the context of set theory.

Definition A.1. (i) A set C of sets is said to be *transitive* just if $A \in C$ whenever $A \in B$ and $B \in C$. Equivalently C is transitive if every element of C is also a subset of C i.e. $C \subseteq 2^C$.

(ii) An *ordinal* is defined to be a set A such that A is transitive, and every element of A is transitive. We use $\alpha, \beta, \gamma, \cdots$ to refer to ordinals.

There are several equivalent definitions of ordinals. It follows that

(i) every element of an ordinal is an ordinal

(ii) every transitive set of ordinals is itself an ordinal.

The collection of all ordinals is not a set, but a proper *class*.

Ordinals are well-ordered A binary relation $<$ on a set A is a *linear order* if it satisfies:

(i) Irreflexivity: $\forall x \in A \,.\, \neg(x < x)$

(i) Transitivity: $\forall x, y, z \in A \,.\, x < y \ \wedge \ y < z \ \rightarrow x < z$

(i) Trichotomy: $\forall x, y \in A \,.\, x < y \ \vee \ y < x \ \vee \ x = y$

A binary relation $<$ on a class S is *well-founded* just if every non-empty subset of S has a $<$-minimal element.

For ordinals α, β, define $\alpha < \beta$ just if $\alpha \in \beta$. It follows that every ordinal is equal to the set of all smaller ordinals i.e. $\alpha = \{\, \beta : \beta < \alpha \,\}$.

Proposition 14. *(i) $<$ is a well-founded linear ordering on the class of ordinals.*

(ii) If α is an ordinal, so is $\alpha \cup \{\, \alpha \,\}$, called the successor ordinal *of α, which is denoted $\alpha + 1$.*

(iii) If A is a set of ordinals then $\bigcup A$ is an ordinal; and it is the supremum of the ordinals in A under $\leq$.

An ordinal is called a *successor ordinal* if it is of the form $\alpha + 1$; otherwise it is a *limit ordinal*. For example, the smallest few ordinals are

$$\begin{aligned} 0 &:= \varnothing \\ 1 &:= \{\, \varnothing \,\} \\ 2 &:= \{\, \varnothing, \{\, \varnothing \,\} \,\} \\ 3 &:= \{\, \varnothing, \{\, \varnothing \,\}, \{\, \varnothing, \{\, \varnothing \,\} \,\} \,\} \\ &\vdots \end{aligned}$$

The first infinite ordinal is $\omega := \{\, 0, 1, 2, \cdots \,\}$. The smallest limit ordinal is 0, the next smallest is ω. In fact there are uncountably many countably infinite ordinals:

$$\begin{aligned} &\omega, \omega + 1, \omega + 2, \cdots, \omega + \omega = \omega \cdot 2, \cdots, \omega \cdot 3, \cdots, \\ &\omega \cdot \omega = \omega^2, \cdots, \omega^3, \cdots, \omega^\omega, \cdots, \omega^{\omega^\omega}, \cdots, \omega^{\omega^{\omega^\omega}}, \cdots, \\ &\epsilon_0 = \omega^{\omega^{\omega^{\omega^{\cdots}}}}, \cdots \end{aligned}$$

The set of all countable ordinals is the first uncountable ordinal, written ω_1.

Principle of Transfinite Induction To prove that "for all ordinals α, the property H_α holds", we establish the following:

(i) *Successor ordinal* $\alpha + 1$: Assuming that H_α holds, then $H_{\alpha+1}$ holds.

(ii) *Limit ordinal* λ: Assuming that H_α holds for all $\alpha < \lambda$, then H_λ holds.

The validity of the principle ultimately rests on the well-foundedness of the relation $\in$.