Lecture 6a: Alternating Automata: Direct arguments and a different formulation

Given an NFA $A = (Q, \Sigma, \delta, s, F)$ and word $w = a_1 a_2 \ldots a_n$, we can characterize when A accepts the word w as follows : Construct a tree of height $n+1$. The root of the tree is labelled by s. If a node at level i, $1 \leq i \leq n$, is labelled by a state q and $\delta(q, a_i) = X$ then it has $|X|$ children labelled by the elements of X. Now treat each internal node in this tree as the logical $\vee$ operator to obtain a propositional formula over the set of propositions Q. This formula evaluates to true on a valuation that assigns true to the elements of F (and false to all other states) iff the word w is accepted. We shall write X to stand for the function σ_X which assigns true to each element of X and false to elements of $Q \setminus X$.

Formally, for each state q and each word $w \in \Sigma^*$ we define a formula $\mathcal{F}(q, w)$ as follows:

1. $\mathcal{F}(q, \epsilon) = q$
2. $\mathcal{F}(q, aw) = \bigvee_{p \in \delta(q,a)} \mathcal{F}(p, w)$

Then, we have the following proposition which is easy to prove by induction on the length of the word w.

Proposition 1 *There is an accepting run from the state q on the input w iff $F \models \mathcal{F}(q, w)$.*

It is easy to generalize this idea to alternating automata. We construct the tree exactly as above, but in turning it into a formula, we replace each state in an internal node by $\vee$ if it belongs to $Q_\exists$ and by $\wedge$ if it belongs to $Q_\forall$. Consider the following alternating automaton from Lecture 6.

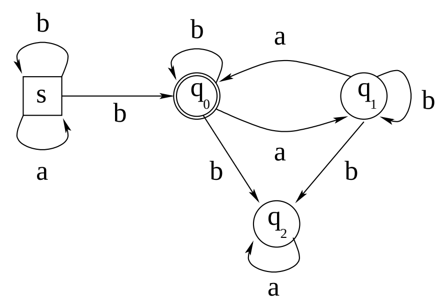

The tree and the formula corresponding to this automaton's behaviour on the input word $babaa$ are given below. As you can verify, the formula is not satisfied by a valuation that assigns true to s and q_0 and false to the other states confirming that this word is not accepted by this automaton.

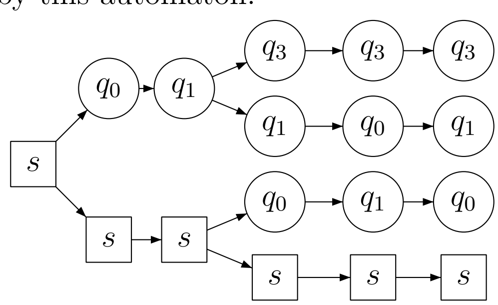

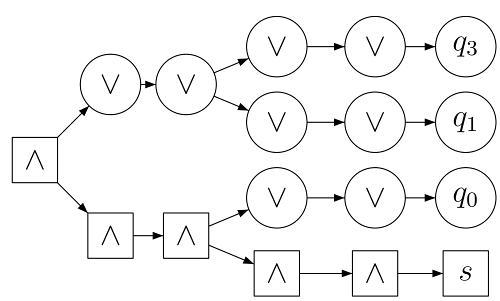

This translation of acceptance to satisfiability of formulas an be formalised as follows: For each state q and each word $w \in \Sigma^*$ we define a formula $\mathcal{F}(q, w)$ as follows:

1. $\mathcal{F}(q, a) = q$
2. $\mathcal{F}(q, aw) = \bigvee_{p \in \delta(q,a)} \mathcal{F}(p, w)$ if $q \in Q_\exists$
3. $\mathcal{F}(q, aw) = \bigwedge_{p \in \delta(q,a)} \mathcal{F}(p, w)$ if $q \in Q_\forall$

The following proposition follows from an easy induction on the length of the word w.

Proposition 2 *Let A be an alternating automaton and let w be a word over its alphabet. There is an accepting run from the state q on the input w iff $F \models \mathcal{F}(q, w)$.*

One consequence of this characterization is the complementation construction for the alternating automata. Let $\mathcal{F}^d(q, w)$ be the dual formula obtained by replacing every $\vee$ by a $\wedge$ and vice versa. Then,

$$
\begin{array}{rcll}
A \text{ rejects } w & \iff & F \not\models \mathcal{F}(s, w) & \text{By Proposition above} \\
 & \iff & F \models \neg\mathcal{F}(s, w) & \text{Definition of } \models \\
 & \iff & F \models \mathcal{F}^d(s, w)[q \mapsto \neg q] & \text{DeMorgan's Laws} \\
 & \iff & (Q \setminus F) \models \mathcal{F}^d(s, w) &
\end{array}
$$

Using $\mathcal{F}^d(s, w)$ for $\mathcal{F}(s, w)$ corresponds to interchanging the sets $Q_\exists$ and $Q_\forall$ and using $Q \setminus F$ in place of F corresponds to switching the accepting and non-accepting states. This gives the following theorem

Theorem 3 *Let $A = (Q_\exists, Q_\forall, \Sigma, \delta, s, F)$ be an alternating automaton. Then, the automaton $\overline{A} = (P_\exists = Q_\forall, P_\forall = Q_\exists, \Sigma, \delta, s, Q \setminus F)$ accepts the complement of the language accepted by A.*

A better formulation of Alternating Automata

The translation of transitions to boolean functions $\wedge$ and $\vee$ above, suggests an obvious generalization to the definition of alternating automata, one that allows not just $\wedge$ and $\vee$ in the definition of transitions but any expression involving these operators (and hence eliminating the distinction between $Q_\exists$ and $Q_\forall$).

We now define an alternating automaton to be a tuple $(Q, \Sigma, \delta, s, F)$ where $\delta : Q \times \Sigma \to \mathcal{B}^+(Q)$, where $\mathcal{B}^+(Q)$ is the set of *positive boolean formulas over Q* i.e. expressions constructed from Q using the operators $\wedge$ and $\vee$.

A run of an automaton of this form is the natural generalization of the run defined earlier. A run starting at state q on the word ϵ is just the tree with the single node labelled q. On a word $a_1 a_2 \ldots a_n$ it is a tree with $n + 1$ levels where the root is labelled by q. Further, if a node at level i, $1 \leq i \leq n$, is labelled by a state p and X is the set of labels of its children then $X \models \delta(q, a_i)$. The run is accepting if all the leaves are labelled by elements of F. A word is accepted if there is an accepting run on that word starting at the state s.

Next, we show that acceptance can be reduced to the truth of a propositional formula for this model as well. For each state q in the automaton $A = (Q, \Sigma, \delta, s, F)$ and word $w \in \Sigma^*$ we define a formula $\mathcal{F}^A(q, a)$ as follows (where we omit the A):

1. $\mathcal{F}(q, \epsilon) = q$
2. $\mathcal{F}(q, aw) = \delta(q, a)[p \mapsto \mathcal{F}(p, w) \mid p \in Q]$

where $\delta(q, a)[p \mapsto \mathcal{F}(p, w) \mid p \in Q]$ stands for the formula obtained by replacing the proposition p by $\mathcal{F}(p, w)$, for each state p, in the formula $\delta(q, a)$. We can prove that

Proposition 4 *A has an accepting run starting at state q on the word w iff* $F \models \mathcal{F}(q, w)$.

Proof: The proof proceeds by induction on the length of w. If $w = \epsilon$ then w is accepted iff $q \in F$ iff $F \models q$ and $\mathcal{F}(q, \epsilon) = q$.

Suppose, A accepts aw. Then there is a set $X \subseteq Q$, such that $X \models \delta(q, a)$ and for each $p \in X$, there is an accepting on the word w (by the definition of accepting runs). By the induction hypothesis, for each $p \in X$, $F \models \mathcal{F}(p, w)$. But then

$$\begin{array}{rcll} F \models \mathcal{F}(q, aw) & \iff & F \models \delta(q, a)[p \mapsto \mathcal{F}(p, w) \mid p \in Q] & \text{by definition of } \mathcal{F}(q, aw) \\ & \iff & \{p \mid F \models \mathcal{F}(p, w)\} \models \delta(q, a) & \text{Definition of } \models \\ & if & X \models \delta(q, a) & X \subseteq \{p \mid F \models \mathcal{F}(p, w)\}, \text{ monotonicity} \end{array}$$

Thus $F \models \mathcal{F}(q, aw)$. The converse is as easy and the reasoning is almost identical

$$\begin{array}{rcll} F \models \mathcal{F}(q, aw) & \iff & F \models \delta(q, a)[p \mapsto \mathcal{F}(p, w) \mid p \in Q] & \text{by definition of } \mathcal{F}(q, aw) \\ & \iff & \{p \mid F \models \mathcal{F}(p, w)\} \models \delta(q, a) & \text{Definition of } \models \end{array}$$

By induction hypothesis, there are accepting runs on w for each p such that $F \models \mathcal{F}(p, w)$. The accepting run on aw is constructed by taking the set $\{p \mid F \models \mathcal{F}(p, w)\}$ as the set of children of the root labelled q (which satisfies $\delta(q, a)$ by above) and then following the accepting runs available from the induction hypothesis. ■

Let α^d denote the dual of α for any positive boolean expression $\alpha \in \mathcal{B}^+(Q)$, obtained by interchanging the $\wedge$ and $\vee$ operators. Then, the following Lemma is an easy consequence of the proposition above.

Lemma 5 *Let* $A = (Q, \Sigma, \delta, s, F)$ *be an alternating automaton. Then,* $\overline{A} = (Q, \Sigma, \delta^d, s, Q \setminus F)$ *where* $\delta^d(q, a) = (\delta(q, a))^d$ *accepts the complement of the language accepted by A.*

Proof: First observe that $\mathcal{F}^{A^d}(q, w) = (\mathcal{F}^A(q, w))^d$. Thus,

$$\begin{array}{rcll} A \text{ rejects } w & \iff & F \not\models \mathcal{F}^A(s, w) & \text{By Proposition above} \\ & \iff & F \models \neg \mathcal{F}^A(s, w) & \text{definition of } \models \\ & \iff & F \models (\mathcal{F}^A(s, w))^d[q \mapsto \neg q] & \text{DeMorgan's Laws} \\ & \iff & (Q \setminus F) \models \mathcal{F}^{A^d}(s, w) & \text{Observation above} \\ & \iff & \overline{A} \text{ accepts } w & \end{array}$$

■

Further, for this version of alternating automata, closure under other boolean operations is also easy.

Lemma 6 *Let $A_1 = (Q_1, \Sigma, \delta_1, s_1, F_1)$ and $A_2 = (Q_2, \Sigma, \delta_2, s_2, F_2)$ be alternating automata accepting the languages L_1 and L_2. Then, $B_\cup = (Q_1 \cup Q_2 \cup \{s\}, \Sigma, \delta_\cup, s, F_1 \cup F_2)$ with $\delta_\cup = \delta_1 \cup \delta_2 \cup [(s, a) \mapsto \delta(s_1, a) \vee \delta(s_2, a)]$ accepts $L_1 \cup L_2$. Further, $B_\cap = (Q_1 \cup Q_2 \cup \{s\}, \Sigma, \delta_\cap, s, F_1 \cup F_2)$ with $\delta_\cap = \delta_1 \cup \delta_2 \cup [(s, a) \mapsto \delta(s_1, a) \wedge \delta(s_2, a)]$ accepts $L_1 \cap L_2$.*

The proof of this lemma is left as an easy exercise. For the formulation of alternating automata in Lecture 6, it is not easy to carry out these constructions (For $L_1 \cup L_2$, how do we deal with the case when s_1 or s_2 or both are $\forall$ states?).

Converting this version of alternating automata to nondeterministic automata is easy and we leave the proof of the following lemma as an exercise as well.

Lemma 7 *Let $A = (Q, \Sigma, \delta, s, F)$ be an alternating automaton. Then the nondeterministic automaton $(2^Q, \Sigma, \delta_n, \{s\}, 2^F)$ with $\delta_n(X, a) = \{Y | Y \models \delta(q, a), q \in X\}$ accepts the same language as A.*

Connection to Games

In the game graph construction in Lecture 6, we had one copy of each state in each of the $n+1$ levels (where the word under consideration is $a_1 a_2 \ldots a_n$). State q at level i is connected to the states from $\delta(q, a_i)$ at level $i+1$. Any node labelled by $q \in Q_\forall$, belongs to the pathfinder and captures the fact that all successors should have accepting runs on $a_{i+1} \ldots a_n$ while a node labelled by a $q \in Q_\exists$ belongs to the automaton is required to pick one successor with an accepting run on $a_{i+1} \ldots a_n$.

In the new setting, $\delta(q, a_i)$ is a formula and the game has to capture the fact that for some set X satisfying $\delta(q, a_i)$ every state in X accepts $a_{i+1} \ldots a_n$. One way to do this as a game is the following: Have $2n+1$ levels where each of the levels $1, 3, \ldots, 2n+1$ consist of a copy each of Q while the levels $2, 4, \ldots, 2n$ consist of copies of 2^Q. The nodes in the odd levels belong to the automaton while those in the even numbered levels belong to the pathfinder. The edges from level $2i+1$ to $2i+2$ connect a node labelled q with a node labelled X if and only if $X \models \delta(q, a_i)$ and an edge from level $2i$ to $2i+1$ connects a node labelled X to a node labelled q iff $q \in X$. Thus, on each letter a_i, the automaton proposes an X such that $X \models \delta(q, a_i)$ as the next move on the accepting run and the pathfinder may choose to challenge this by picking any suspect state from this set.

It is quite easy to prove that these games are determined and further that the existence of winning strategies for the automaton in the game defined by the word w is equivalent to acceptance of the word w. We leave this as an exercise.

However, this game suffers from two deficiencies — firstly, the size of the game graph depends exponentially on the number of states Q and secondly, the role of the two players is not symmetric and hence it is not clear how to translate the existence of winning strategies for the pathfinder into an alternating automaton accepting the complement. The former is not such a serious problem (since the game graph is only used as a tool in the proofs) but the latter is a more serious irritation.

We now provide an alternative game, one that is more natural in this setting and does not suffer from either of these deficiencies. It is based on the following simple idea. Instead of picking the set X and then an element q in X in just two moves, we arrange for more elaborate game on the *graph* defined by the formula $\delta(q, a)$ to achieve the same effect.

Let α be a positive boolean formula over Q. We define a game graph associated with α, G_α, by induction on α. This graph is acyclic and constructed by taking the expression tree of the formula α and identifying all leaves labelled by same states (so that there is at most one leaf per state in Q). For the same of uniformity we also assume that there is one leaf labelled by q for each q in Q. All the leaves as well as nodes labelled by $\vee$ belong to the automaton while all the nodes labelled by $\wedge$ belong to the pathfinder. The following figure describes the game graph for the formula $q \wedge ((p \vee q) \wedge s)$ over $Q = \{p, q, r, s\}$.

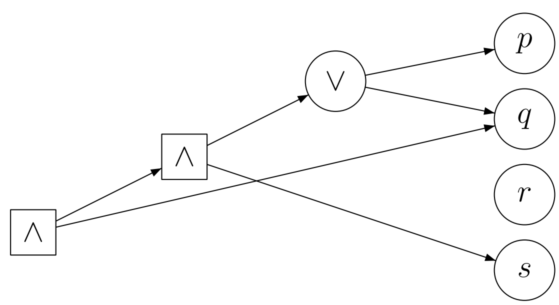

The starting position of the game is the root of the expression tree (the left most $\wedge$ node in the above figure). Here is an useful fact:

Proposition 8 *Let G_α be the game graph as described above. Then, $X \models \alpha$ iff the automaton has a (positional) strategy to ensure that the game ends in a state in X and $X \not\models \alpha$ iff the pathfinder has a (positional) strategy to ensure that the game ends in a state in $Q \setminus X$.*

Proof: The proof as usual is by induction on the size of α.

Suppose $X \models \alpha$. If $\alpha = q$ then the result holds by definition.

If $\alpha = \alpha_1 \vee \alpha_2$ then w.l.o.g. we may assume that $X \models \alpha_1$. Then, at the root of α, the automaton moves to the root of α_1 and then follows the winning strategy on G_{α_1}, available by induction hypothesis, from there on to win this game.

If $\alpha = \alpha_1 \wedge \alpha_2$ then if the pathfinder moves to the root of α_i then the automaton follows the w.s. for α_i for $i \in \{1, 2\}$. Thus, the automaton can always force the game into a state in X.

Conversely, suppose the automaton has a strategy to force the game into X on the game G_α. If $\alpha = q$ then $q \in X$ then by definition $X \models q$.

If $\alpha = \alpha_1 \vee \alpha_2$. If the strategy for the automaton picks α_i then by defn. of winning strategy, the automaton also has a winning strategy on the game on G_{α_i} and hence $X \models \alpha_i$ and hence $X \models \alpha$.

If $\alpha = \alpha_1 \wedge \alpha_2$. Since the pathfinder may move the game to root of either of the games G_{α_1} or G_{α_2} it follows that the automaton has a winning strategy in both of these games and hence $X \models \alpha_1$ and $X \models \alpha_2$ and thus $X \models \alpha_1 \wedge \alpha_2$.

We leave the corresponding result for pathfinder winning strategies as an exercise to the interested reader. ■

As a consequence of the above proposition, we can replace the two move subgame defined by the transition $\delta(q, a)$ (automaton picks X satisfying α and then pathfinder picks an element of X) by an entire subgame played on $G_{\delta(q,a)}$ to achieve the same effect but with a clean symmetry between the two players.

We now define the game graph $G_{A,w}$ for an alternating automaton A and word w by induction on the length of w. By construction each such graph has $|w| + 1$ copies of the states from Q, organized as $|w| + 1$ layers (in addition to other nodes) and from a node in layer i only nodes in higher layers are reachable in this graph. Further, the nodes in the 1st layer have no incoming edges and the nodes in the last layer have no outgoing edges.

1. if w is ϵ then the game graph consists of one node each for the states in Q, all of which belong to the player *automaton*

2. The graph $G_{A,aw}$ is constructed from $G_{A,w}$ as follows: Take a new copy of the states of Q. For each state q in this copy, add a copy of the game graph $G_{\delta(q,a)}$ and add an edge from q to the root of $\delta(q, a)$. Identify the nodes labelled q in the leaves of the games graphs $G_{\delta(q,a)}$ with each other as well as the corresonding state in the 1st layer of $G_{A,w}$.

The winning set in the game is the set of states in F in layer $|w| + 1$. We illustrate this definition with an example here.

Let $A = (\{s, p, q\}, a, b, \delta, s, q)$ be the automaton where the transition function is given by the following table:

	a	b
s	$s \wedge (s \vee q)$	s
p	q	$s \wedge q$
q	$q \vee s$	$s \vee p$

The following figure describes the game graph $G_{A,aba}$. The parts added to $G_{A,ba}$ to obtain $G_{A,aba}$ are in green.

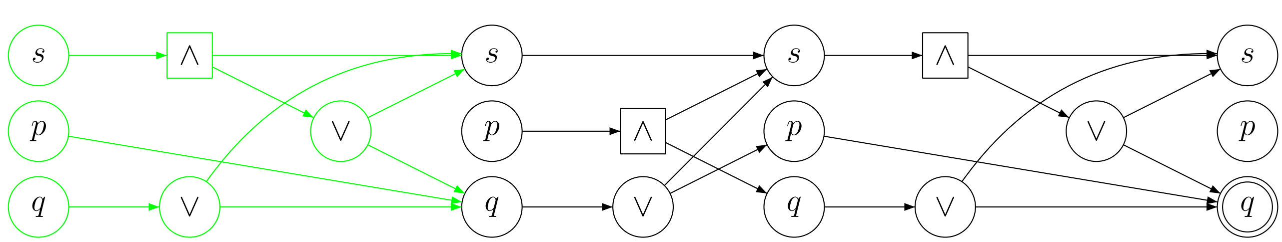

Observation: We observe that the internal nodes in this game labelled with states from Q have only one out going edge and hence it makes no difference as to whether these states belong to the automaton or the pathfinder.

The following theorem states that the winner of the game on $G_{A,w}$ determines whether w is accepted or not.

Lemma 9 *Let $A = (Q, \Sigma, \delta, s, F)$ be an alternating automaton and let $w \in \Sigma^*$. If the automaton has a winning strategy from the level 1 state labelled q in the game $G_{A,w}$ then there is an accepting run on w starting in state q.*

Proof: The proof is by induction on the length of w and the case when $w = \epsilon$ is trivial.

For the inductive case, suppose the automaton has a winning strategy in the game $G_{A,aw}$ from q. Let X be the set of nodes in the second level of nodes labelled by Q that is reachable starting at q in this game while the automaton is playing this winning strategy. Then, by definition of winning strategy, the automaton has a winning strategy from each of these nodes. But these are the level 1 nodes in the game $G_{A,w}$ and by induction hypothesis, for each of these states $p \in X$, the automaton has an accepting run on w starting at state p. Also, $X \models \delta(q, a)$ (by Proposition 8 above). Thus there is an accepting run from q on aw. ■

Lemma 10 *Let $A = (Q, \Sigma, \delta, s, F)$ be an alternating automaton and let $w \in \Sigma^*$. Let W be the set of states from which A has an accepting run on the word w. The automaton has a postional strategy in the game $G_{A,w}$ that wins from every position at level 1 that is labelled by a state from W.*

A remark is in order before we present the proof. The statment of the lemma is stronger than required in that it claims that a *single* winning strategy works for all the states q from where there is an accepting run. This is needed because in the inductive step, the automaton can only force the game to reach one of a good set of states X on the first letter, from each of which this is accepting run/winning strategy. If different states in X have different winning strategies then we need to combine the these into a single strategy, something that can be avoided with this stronger induction hypothesis.

Proof: The proof is by induction on the length of w and the case of $w = \epsilon$ is trivial.

Suppose the set of states from where the automaton has an accepting run on aw is W. Pick an accepting run starting at each $q \in W$ and let X_q be the set of states labelling the children of q in these accepting runs and let $X = \bigcup_{q \in W} X_q$. Clearly, there is an accepting run on w from each state in X. Thus, by the induction hypothesis there is a strategy for the automaton in the game $G_{A,w}$ that wins from all the positions in X (and perhaps more).

In $G_{A,aw}$ the automaton simply has to play so that when it enters the part of the game graph coming from $G_{A,w}$ it does so in one of the states in X.

Let $q \in W$. In the game graph $G_{\delta(q,a)}$, since the automaton plays the positional strategy that will guarantee that the play ends (i.e. enters the game $G_{A,w}$) in a state in X_q. Such a positional strategy exists by Proposition 8. The game graph $G_{A,aw}$ decomposes into disjoint parts contributed by $G_{\delta(q,a)}$ for each $q \in Q$, followed by the game graph $G_{A,w}$. Thus, the positional strategies for each $G_{\delta(q,a)}$, $q \in W$ that guarantees that the game ends in X_q combined with the strategy for $G_{A,w}$ that wins from all the positions in X works as a winning strategy for the automaton in $G_{A,aw}$. ■

Theorem 11 *Let $A = (Q, \Sigma, \delta, s, F)$ be an alternating automaton and let $w \in \Sigma^*$. The game $G_{A,w}$ is determined with the winner having positional strategies. Further, the automaton has a winning strategy from a level 1 node labelled q iff there is an accepting run starting at q (and the pathfinder has a winning strategy from a level 1 node labelled q iff there is no accepting run starting at state q.)*

Proof: The positional determinacy is proved exactly as for the similar game described in Lecture 6 and the details are left as an exercise. The rest is a consequence of the two lemmas 10 and 9. ■

We can use this result to give an alternate proof of the complementation construction. Observe that the game graph $G_{A,w}$ and $G_{\overline{A},w}$ are duals of each other (i.e. one is obtained from the other by just interchanging the ownership of nodes between the two players, hence interchanging their *roles*). Well, almost, but not entirely. This is because the internal nodes labelled by states from Q are automaton owned in both games. But, as observed earlier, this is irrelevant as those nodes have a single outgoing edge. Thus, the automaton cannot win the game $G_{A,w}$ iff the pathfinder has a winning strategy in the game $G_{A,w}$ to force the game to end in a state in $Q \setminus F$, which by the duality is possible if and only if the automaton has a winning strategy in the game $G_{\overline{A},w}$ to force the game in $Q \setminus F$. Thus $L(\overline{A}) = \overline{L(A)}$.