On Transition Minimality of Bideterministic Automata

Hellis Tamm

This talk is based on a paper to be presented at the conference on Developments in Language Theory (DLT 2007), Turku, July 2007, and to be published in LNCS 4588.

State complexity

- In automata theory, descriptional complexity issues have been of interest for decades.
- It is well known that the number of states of the minimal DFA (deterministic state complexity) for a given language can be exponentially larger than the number of states in a minimal NFA (nondeterministic state complexity).
- The minimal DFA is unique but there may be several minimal NFAs.
- Many cases where the maximal blow-up of size when converting an NFA to DFA does not occur.
- Some sufficient conditions have been identified which imply that the deterministic and nondeterministic state complexities are the same (for example, bideterminism).

Transition complexity

- While the state-minimal DFA is also minimal with respect to the number of transitions, this is not necessarily the case with NFAs.
- Even allowing one more state in an NFA can produce a considerable reduction in the number of transitions.
- The number of transitions may be even a better measure for the size of an NFA than the number of states.
- Furthermore, allowing ϵ-transitions in an NFA (ϵ-NFAs) it is possible to have automata with even less transitions than NFAs.

Bideterministic automata: state minimality

- A bideterministic automaton is any deterministic automaton such that its reversal automaton is also deterministic
- A bideterministic automaton is a state-minimal DFA (easy)
- Any bideterministic automaton is a state-minimal NFA (Tamm and Ukkonen 2003)
- What about transition minimality?

Bideterministic automata: transition minimality

The results presented in the current paper:

- A bideterministic automaton is a transition-minimal NFA (preliminary result in my PhD thesis, 2004)
- Transition minimality of bideterministic automata is not unique
- The necessary and sufficient conditions for a bideterministic automaton to be a unique transition-minimal NFA
- More generally: a bideterministic automaton is a transition-minimal ϵ-NFA.

Universal automaton

A universal automaton is a canonical automaton of a given regular language.

Let Σ be a finite alphabet and let $L \subseteq \Sigma^*$.

A *factorization* of L is a maximal couple (with respect to the inclusion) of languages (U, V) such that $UV \subseteq L$.

The *universal automaton* of L is $U_L = (Q, \Sigma, E, I, F)$ where
Q is the set of factorizations of L,
$I = \{(U, V) \in Q \mid \epsilon \in U\}$,
$F = \{(U, V) \in Q \mid U \subseteq L\}$,
$E = \{((U, V), a, (U', V')) \in Q \times a \times Q \mid Ua \subseteq U'\}$.

Fact: universal automaton of the language L is a finite automaton that accepts L.

Universal automaton: the construction

S. Lombardy (2002) has given the following effective method for constructing the universal automaton from the minimal DFA of the given language:

Let $A = (Q, \Sigma, E, \{q_0\}, F)$ be the minimal DFA accepting L and let P be the set of states of the automaton $D(A^R)$.
Let $P_\cap$ be the closure of P under intersection, without the empty set: if $X, Y \in P_\cap$ and $X \cap Y \neq \emptyset$ then $X \cap Y \in P_\cap$.
Then, the universal automaton U_L is isomorphic to $(P_\cap, \Sigma, H, I, J)$ where $H = \{(X, a, Y) \in P_\cap \times \Sigma \times P_\cap \mid X \cdot a \subseteq Y$ and for all $p \in X,\ p \cdot a \neq \emptyset\}$,
$I = \{X \in P_\cap \mid q_0 \in X\}$, and
$J = \{X \in P_\cap \mid X \subseteq F\}$.

Automaton morphism and the universal automaton

Let $A = (Q, \Sigma, E, I, F)$ and $A' = (Q', \Sigma, E', I', F')$ be two NFAs. Then a mapping μ from Q into Q' is a *morphism* of automata if and only if $p \in I$ implies $p\mu \in I'$, $p \in F$ implies $p\mu \in F'$, and $(p, a, q) \in E$ implies $(p\mu, a, q\mu) \in E'$ for all $p, q \in Q$ and $a \in \Sigma$.

Known properties:

- Let A be a trim automaton that accepts L. Then there exists an automaton morphism from A into U_L.
- In particular, U_L contains as a subautomaton every state-minimal NFA accepting L.

Universal automaton of a bideterministic language

Now, let us construct the universal automaton of a bideterministic language L.

Let $A = (Q, \Sigma, E, \{q_0\}, \{q_f\})$ be a trim bideterministic automaton. It is known that A is the minimal DFA. Since the reversal automaton of A is deterministic, $D(A^R) = A^R$ and the set P as well as $P_\cap$ consist of all sets $\{q\}$ such that $q \in Q$.

It is easy to see that the transition relation H of U_L is equal to E, $I = \{q_0\}$, and $J = \{q_f\}$.

Conclusion. *Any bideterministic automaton is the universal automaton for the given language.*

By using algebraic considerations, basically the same fact has been observed by L. Polak (2004).

Let $A = (Q, \Sigma, E, \{q_0\}, \{q_f\})$ be a bideterministic automaton and $A' = (Q', \Sigma, E', I', F')$ be another automaton accepting the same language.

Since $A = U_{L(A)}$, then there exists an automaton morphism μ from A' into A.

Next, we will see that μ defines an automaton transformation.

Proposition. *μ is surjective.*

Proof. Since A is a state-minimal NFA then for each state q of A there exists at least one state q' of A' such that $q'\mu = q$.

Proposition. *There is a transition (p, a, q) of A if and only if there is a transition (p', a, q') of A' such that $p'\mu = p$ and $q'\mu = q$.*

Proof. The "if" part follows from the definition of automaton morphism.

The "only-if" part is proved by contradiction.
Suppose that (p, a, q) is a transition of A but there is no transition (p', a, q') of A' such that $p'\mu = p$ and $q'\mu = q$.
Let $B = (Q, \Sigma, E \setminus \{(p, a, q)\}, \{q_0\}, \{q_f\})$ be a subautomaton of A.
It is clear that μ is an automaton morphism from A' into B.

It is known that for any automaton morphism from X into Y, it holds that $L(X) \subseteq L(Y)$. Therefore, $L(A') \subseteq L(B)$.

Since $L(A) = L(A')$, we also get $L(A) \subseteq L(B)$. But, since A is the unique minimal DFA and B has less transitions than A, it must be that $L(B) \subset L(A)$, a contradiction.

It is not difficult to see that μ defines an automaton transformation from A' to A.

Let $Q = \{q_0, ..., q_{n-1}\}$.

Since μ is surjective, there exists a partition $\Pi = \{Q'_0, ..., Q'_{n-1}\}$ of Q' into $n = |Q|$ disjoint non-empty subsets so that for every $q' \in Q'$ and $i \in \{0, ..., n-1\}$, $q' \in Q'_i$ if and only if $q'\mu = q_i$.

Using Π, A' is transformed into an equivalent automaton A'': for every $i \in \{0, ..., n-1\}$, all states in Q'_i are merged into a single state q''_i of A''.

It is clear that A'' is isomorphic to A.

The number of transitions of A'' is no more than the number of transitions of A'.

Proposition. *Any bideterministic automaton is a transition-minimal NFA.*

Uniqueness of transition minimality

Differently from the state minimality, a bideterministic automaton is not necessarily the only transition-minimal NFA for the corresponding language.

The necessary and sufficient conditions for the unique transition-minimality are given by the following theorem:

Theorem. *A trim bideterministic automaton* $A = (Q, \Sigma, E, \{q_0\}, \{q_f\})$ *is a unique transition-minimal NFA if and only if the following three conditions hold:*

(i) $q_0 \neq q_f$,

(ii) $\textit{indegree}(q_0) > 0$ *or* $\textit{outdegree}(q_0) = 1$,

(iii) $\textit{indegree}(q_f) = 1$ *or* $\textit{outdegree}(q_f) > 0$.

Unambiguous ϵ-NFA

S. John (2003, 2004) has developed a theory to reduce the number of transitions of ϵ-NFAs.

Let A be an ϵ-NFA (Q, Σ, E, I, F) where E is partitioned into two subrelations $E_\Sigma = \{(p, a, q) \mid (p, a, q) \in E, a \in \Sigma\}$ and $E_\epsilon = \{(p, \epsilon, q) \mid (p, \epsilon, q) \in E\}$.

The automaton A is *unambiguous* if and only if for each $w \in L(A)$ there is exactly one path that yields w (without considering ϵ-transitions).

Slices

Let $L \subseteq \Sigma^*$ be a regular language, $U, V \subseteq \Sigma^*$, $a \in \Sigma$.

We call (U, a, V) a *slice* of L if and only if $U \neq \emptyset$, $V \neq \emptyset$ and $UaV \subseteq L$.

Let S be the set of all slices of L.

A partial order on S is defined by:
$(U_1, a, V_1) \leq (U_2, a, V_2)$ if and only if $U_1 \subseteq U_2$ and $V_1 \subseteq V_2$.

The set of *maximal slices* of L is defined by
$S_{max} := \{(U, a, V) \in S \mid$ there is no $(U', a, V') \in S$ with $(U, a, V) < (U', a, V')\}$.

Transition-minimal unambiguous ϵ-NFA

Let $S' \subseteq S$ be a finite slicing of L. In order to read an automaton $A_{S'}$ out of S', each slice from S' is transformed into a transition of $A_{S'}$, and these transitions are connected via states and ϵ-transitions using a follow-relation $\longrightarrow$ which is defined basically by:
$(U_1, a, V_1) \longrightarrow (U_2, b, V_2)$ if and only if $U_1 a \subseteq U_2$ and $bV_2 \subseteq V_1$

Theorem (S. John). *The three following statements are equivalent for languages $L \subseteq \Sigma^*$ if the slicing S_{max} of L induces an unambiguous ϵ-NFA $A_{S_{max}}$:*
1) *L is accepted by an ϵ-NFA*
2) *$L = L(A_{S'})$ for some finite slicing $S' \subseteq S$*
3) *S_{max} is finite*
Furthermore, $|S_{max}| \leq |S'| \leq |E_\Sigma|$.

Corollary (S. John). *An unambiguous ϵ-NFA $A_{S_{max}}$ has the minimum number of non-ϵ-transitions.*

Transition slice

For each non-ϵ-transition t of an automaton A, we define the *transition slice* of t to be the slice $(U_t, l(t), V_t)$ of $L(A)$ where

– U_t is the set of strings yielded by the paths from an initial state to the source state of t,

– $l(t)$ is the label of t, and

– V_t is the set of strings yielded by the paths from the target state of t to an accepting state.

Using the theory by S. John it is not difficult to prove that a bideterministic automaton is a transition-minimal ϵ-NFA.

Lemma. *For a bideterministic automaton A, let t_1 and t_2 be two different transitions of A, with the same label $a \in \Sigma$ and with the corresponding transition slices (U_{t_1}, a, V_{t_1}) and (U_{t_2}, a, V_{t_2}). Then $U_{t_1} \cap U_{t_2} = \emptyset$ and $V_{t_1} \cap V_{t_2} = \emptyset$.*

Proof. By contradiction. Supposing $U_{t_1} \cap U_{t_2} \neq \emptyset$ implies that A is not deterministic. Similarly $V_{t_1} \cap V_{t_2} = \emptyset$.

Proposition. *Each transition slice of a bideterministic automaton A is maximal.*

Proof. Suppose there is a transition t such that its transition slice (U_t, a, V_t) is not maximal. Then $(U_t, a, V_t) < (U, a, V)$ for some maximal slice (U, a, V). There is a string $uav \in L(A)$ such that $u \in U$ and $v \in V$ but either $u \notin U_t$ or $v \notin V_t$. However, there must be some transition t' with the transition slice $(U_{t'}, a, V_{t'})$ such that $u \in U_{t'}$ and $v \in V_{t'}$ and $(U_{t'}, a, V_{t'}) \leq (U, a, V)$. Now, we know that $U_t \subseteq U$ and $U_{t'} \subseteq U$, and therefore also $U_t \cup U_{t'} \subseteq U$.
In the same way, $V_t \cup V_{t'} \subseteq V$.

Next, we can see that $(U_t \cup U_{t'}, a, V_t \cup V_{t'})$ is a slice of $L(A)$.

Then there is a word $xay \in L(A)$ such that $x \in U_t$ and $y \in V_{t'}$.
Since, by Lemma, there does not exist a transition t'' of A such that $x \in U_{t''}$, $a = l(t'')$ and $y \in V_{t''}$, it can be shown that $xay \notin L(A)$, a contradiction.

Theorem. *A bideterministic automaton A has the minimum number of transitions among all ϵ-NFAs accepting $L(A)$.*

Proof. The set of maximal slices of $L(A)$ is given by
$S_{max} := \{(U_t, l(t), V_t) \mid t \in E\}$, $|S_{max}| = |E|$.
The set S_{max} is used to form the ϵ-NFA $A_{S_{max}}$ by converting every slice from S_{max} into a transition of $A_{S_{max}}$ and connecting these transitions by ϵ-transitions according to the follow-relation.

Since A is bideterministic, A is clearly unambiguous.
There is a one-to-one correspondence between the accepting paths of A and $A_{S_{max}}$. Thus, $A_{S_{max}}$ is also unambiguous.

By Theorem (John) and Corollary (John), $A_{S_{max}}$ has a minimum number of non-ϵ-transitions. Since the number of non-ϵ-transitions of $A_{S_{max}}$ is equal to the number of transitions of A, and there are no ϵ-transitions in A, we conclude that A is transition-minimal among all ϵ-NFAs accepting the given language.