

PUSHING FOR WEIGHTED TREE AUTOMATA *
— DEDICATED TO THE MEMORY OF ZOLTÁN ÉSIK (1951–2016) —

THOMAS HANNEFORTH, ANDREAS MALETTI, AND DANIEL QUERNHEIM

Universität Potsdam, Human Sciences Faculty, Department Linguistik
Karl-Liebknecht-Str. 24–25, 14476 Potsdam, Germany
e-mail address: thomas.hanneforth@uni-potsdam.de

Universität Leipzig, Faculty of Mathematics and Computer Science, Institute of Computer Science
PO box 100 920, 04009 Leipzig, Germany
e-mail address: maletti@informatik.uni-leipzig.de

Universität Stuttgart, Institute for Natural Language Processing
Pfaffenwaldring 5b, 70569 Stuttgart, Germany
e-mail address: daniel.quernheim@ims.uni-stuttgart.de

ABSTRACT. A weight normalization procedure, commonly called pushing, is introduced for weighted tree automata (wta) over commutative semifields. The normalization preserves the recognized weighted tree language even for nondeterministic wta, but it is most useful for bottom-up deterministic wta, where it can be used for minimization and equivalence testing. In both applications a careful selection of the weights to be redistributed followed by normalization allows a reduction of the general problem to the corresponding problem for bottom-up deterministic unweighted tree automata. This approach was already successfully used by MOHRI and EISNER for the minimization of deterministic weighted string automata. Moreover, the new equivalence test for two wta M and M' runs in time $\mathcal{O}((|M| + |M'|) \log(|Q| + |Q'|))$, where Q and Q' are the states of M and M' , respectively, which improves the previously best run-time $\mathcal{O}(|M| \cdot |M'|)$.

1. INTRODUCTION

Weighted tree automata [FV09] have recently found various applications in fields as diverse as natural language and XML processing [KM09], system verification [Jac11], and pattern recognition. Most applications require efficient algorithms for basic manipulations of tree automata such as determinization [BMV10], inference [MKV10], and minimization [HMM09,

2012 ACM CCS: [Theory of computation]: Formal languages and automata theory — Tree languages; [Theory of computation]: Formal languages and automata — Automata extensions — Quantitative automata.

2010 Mathematics Subject Classification: 68Q45, 68Q25.

Key words and phrases: pushing — weighted tree automaton — minimization — equivalence testing.

* This is a revised and extended version of [MALETTI, QUERNHEIM: *Pushing for weighted tree automata*. Proc. 36th Int. Conf. Mathematical Foundations of Computer Science, LNCS 6907, p. 460–471, 2011].

Financially supported by the German Research Foundation (DFG) grant MA / 4959 / 1-1.

HMM07]. For example, in the system verification domain the properties to be verified are typically easily expressed as a formula in a logic. It is well-known [TW68] that tree automata are as expressive as monadic second-order logic with two successors. This celebrated result was recently generalized to the weighted setting for various weight structures [DV06, Man08, DGMM11, VDH16], so quantitative specifications are readily available. However, one of the main insights gained in the development of the MONA toolkit [KM01] (or the SPASS system [WDF⁺09]) was that the transformation of a formula into an equivalent tree automaton heavily relies on the minimization of the constructed deterministic tree automata as the automata otherwise grow far too quickly. Similarly, a major inference setup, also used in the synthesis subfield in system verification, is ANGLUIN’s minimally adequate teacher setup [Ang87]. In this setup, the learner is given access to an oracle that correctly supplies coefficients of trees in the weighted tree language to be learned, which are called coefficient queries, and certificates that the proposed weighted tree automaton indeed represents the weighted tree language to be learned, which are called equivalence queries. In implementations of the oracle the latter queries are typically answered by equivalence tests.

As already mentioned, quantitative models have recently enjoyed a lot of attention. For example, in natural language processing, weighted devices are often used to model probabilities, cost functions, or other features. In this contribution, we consider pushing [Moh97, Eis03] for weighted tree automata [BR82, FV09] over commutative semifields [HW98, Gol99]. Roughly speaking, pushing moves transition weights along a path. If the weights are properly selected, then pushing can be used to canonicalize a (bottom-up) deterministic weighted tree automaton [Bor05]. The obtained canonical representation has the benefit that it can be minimized using unweighted minimization, in which the weight is treated as a transition label. This strategy has successfully been employed in [Moh97, Eis03] for deterministic weighted (finite-state) string automata, and similar approaches have been used to minimize sequential transducers [Cho03] and bottom-up tree transducers [FSM11]. Here we adapt the strategy for tree automata. In particular, we improve the currently best minimization algorithm [Mal09] for a deterministic weighted tree automaton M with states Q from $\mathcal{O}(|M| \cdot |Q|)$ to $\mathcal{O}(|M| \log |Q|)$, which coincides with the complexity of minimization in the unweighted case [HMM09]. The improvement is achieved by a careful selection of the signs of life [Mal09]. Intuitively, a sign of life for a state q is a context that takes q into a final state. In [Mal09] the signs of life are computed by a straightforward exploration algorithm, which is very efficient, but does not guarantee that states that are later checked for equivalence receive the same sign of life. During the (pair-wise) equivalence checks in [Mal09] the evaluation of the weight of a state in the sign of life of another state thus becomes unavoidable, which causes the increased complexity. In this contribution, we precompute an equivalence relation, which, in general, is still coarser than the state equivalence to be determined, but equivalent states in this equivalence permit the same sign of life. Then we determine a sign of life for each equivalence class. Later we only refine this equivalence relation to obtain the state equivalence, so each state will only be evaluated in its sign of life and this evaluation can be precomputed. Moreover, the weights obtained in this evaluation, also called pushing weights, allow a proper canonicalization in the sense that equivalent states will have exactly the same weights on corresponding transitions after pushing. This property sets our algorithm apart from Algorithm 1 of [Mal09] and allows us to rely on unweighted minimization [HMM09]. Our pushing procedure, which is defined for general (potentially nondeterministic) weighted tree automata, always preserves the semantics, so it might also be useful in other setups.

Secondly, we apply pushing to the problem of testing equivalence. The currently fastest algorithm [DHM11] for checking equivalence of two deterministic weighted tree automata M and M' runs in time $\mathcal{O}(|M| \cdot |M'|)$. It is well known that two minimal deterministic weighted tree automata M and M' are equivalent if and only if they can be obtained from each other by a pushing operation (with proper pushing weights). In other words, equivalent automata M and M' have the same transition structure, but their transition weights can differ by consistent factors. We extend our approach to minimization also to equivalence testing, so we again carefully determine the pushing weight and the sign of life of each state q of M such that it shares the sign of life with all equivalent states of M but also with all corresponding states in M' . This allows us to minimize both input automata and then treat the obtained automata as unweighted automata and test them for isomorphism. This approach reduces the run-time complexity to $\mathcal{O}((|M| + |M'|) \log(|Q| + |Q'|))$, where Q and Q' are the states of M and M' , respectively.

2. PRELIMINARIES

We write \mathbb{N} for the set of all nonnegative integers and $[1, u]$ for its subset $\{i \mid 1 \leq i \leq u\}$ given $u \in \mathbb{N}$. The k -fold CARTESIAN product of a set Q is written as Q^k , and the empty tuple $() \in Q^0$ is often written as ε . Every finite and nonempty set is also called alphabet, of which the elements are called symbols. A ranked alphabet (Σ, rk) consists of an alphabet Σ and a mapping $\text{rk}: \Sigma \rightarrow \mathbb{N}$, which assigns a rank to each symbol. If the ranking ‘rk’ is obvious from the context, then we simply write Σ for the ranked alphabet. For each $k \in \mathbb{N}$, we let Σ_k be the set $\{\sigma \in \Sigma \mid \text{rk}(\sigma) = k\}$ of k -ary symbols of Σ . Moreover, we let $\Sigma(Q) = \{\sigma w \mid \sigma \in \Sigma, w \in Q^{\text{rk}(\sigma)}\}$. The set $T_\Sigma(Q)$ of all Σ -trees indexed by Q is inductively defined to be the smallest set T such that $Q \subseteq T$ and $\Sigma(T) \subseteq T$. Instead of $T_\Sigma(\emptyset)$ we simply write T_Σ . The size $|t|$ of a tree $t \in T_\Sigma(Q)$ is inductively defined by $|q| = 1$ for every $q \in Q$ and $|\sigma(t_1, \dots, t_k)| = 1 + \sum_{i=1}^k |t_i|$ for every $k \in \mathbb{N}$, $\sigma \in \Sigma_k$, and $t_1, \dots, t_k \in T_\Sigma(Q)$. To increase readability, we often omit quantifications like “for all $k \in \mathbb{N}$ ” if they are obvious from the context.

We reserve the use of a special symbol \square that is not an element in any considered alphabet. Its function is to mark a designated position in certain trees called contexts. Formally, the set $C_\Sigma(Q)$ of all Σ -contexts indexed by Q is defined as the smallest set C such that $\square \in C$ and $\sigma(t_1, \dots, t_{i-1}, c, t_{i+1}, \dots, t_k) \in C$ for every $\sigma \in \Sigma_k$, $t_1, \dots, t_k \in T_\Sigma(Q)$, $i \in [1, k]$, and $c \in C$. As before, we simplify $C_\Sigma(\emptyset)$ to C_Σ . In simple words, a context is a tree, in which the special symbol \square occurs exactly once and at a leaf position. Note that $C_\Sigma(Q) \cap T_\Sigma(Q) = \emptyset$, but $C_\Sigma(Q) \subseteq T_\Sigma(Q \cup \{\square\})$, which allows us to treat contexts like trees. Given $c \in C_\Sigma(Q)$ and $t \in T_\Sigma(Q \cup \{\square\})$, the tree $c[t]$ is obtained from c by replacing the unique occurrence of \square in c by t . In particular, $c[c'] \in C_\Sigma(Q)$ given that $c, c' \in C_\Sigma(Q)$.

A commutative semiring [HW98, Gol99] is a tuple $(S, +, \cdot, 0, 1)$ such that $(S, +, 0)$ and $(S, \cdot, 1)$ are commutative monoids and $s \cdot 0 = 0$ and $s \cdot (s_1 + s_2) = (s \cdot s_1) + (s \cdot s_2)$ for all $s, s_1, s_2 \in S$ (*i.e.*, \cdot distributes over $+$). It is a commutative semifield if $(S \setminus \{0\}, \cdot, 1)$ is a commutative group (*i.e.*, in addition, for every $s \in S \setminus \{0\}$ there exists $s^{-1} \in S$ such that $s \cdot s^{-1} = 1$). Typical commutative semifields include

- the BOOLEAN semifield $\mathbb{B} = (\{0, 1\}, \max, \min, 0, 1)$,
- the field $(\mathbb{Q}, +, \cdot, 0, 1)$ of rational numbers, and
- the VITERBI semifield $(\mathbb{Q}_{\geq 0}, \max, \cdot, 0, 1)$, where $\mathbb{Q}_{\geq 0} = \{q \in \mathbb{Q} \mid q \geq 0\}$.

Given a mapping $f: A \rightarrow S$, we write $\text{supp}(f)$ for the set $\{a \in A \mid f(a) \neq 0\}$ of elements that are mapped via f to a non-zero semiring element.

For the rest of the paper, let $(S, +, \cdot, 0, 1)$ be a commutative semifield.¹

A weighted tree automaton [BLB83, Boz99, Kui98, BV03, Bor05, FV09] (for short: wta) is a tuple $M = (Q, \Sigma, \mu, F)$, in which

- Q is an alphabet of states,
- Σ is a ranked alphabet of input symbols,
- $\mu: \Sigma(Q) \times Q \rightarrow S$ assigns a weight to each transition, and
- $F \subseteq Q$ is a set of final states.

We often write elements of $T_\Sigma(Q) \times Q$ as $t \rightarrow q$ instead of (t, q) . The size $|M|$ of the wta M is

$$|M| = \sum_{t \rightarrow q \in \text{supp}(\mu)} (|t| + 1) .$$

We extend the transition weight assignment μ to a mapping $h_\mu: T_\Sigma(Q) \times Q \rightarrow S$ by

$$h_\mu(p \rightarrow q) = \begin{cases} 1 & \text{if } p = q \\ 0 & \text{otherwise} \end{cases}$$

$$h_\mu(\sigma(t_1, \dots, t_k) \rightarrow q) = \sum_{q_1, \dots, q_k \in Q} \mu(\sigma(q_1, \dots, q_k) \rightarrow q) \cdot \prod_{i=1}^k h_\mu(t_i \rightarrow q_i)$$

for all $p, q \in Q$, $\sigma \in \Sigma_k$, and $t_1, \dots, t_k \in T_\Sigma(Q)$. The wta M recognizes the weighted tree language $M: T_\Sigma \rightarrow S$ such that $M(t) = \sum_{q \in F} h_\mu(t \rightarrow q)$ for every $t \in T_\Sigma$. Two wta M and M' are equivalent if their recognized weighted tree languages coincide. The unweighted (finite-state) tree automaton [GS84, GS97, CDG⁺07] (for short: fta) corresponding to M is $\text{unw}(M) = (Q, \Sigma, \text{supp}(\mu), F)$.² We note that $\text{supp}(M) \subseteq L(\text{unw}(M))$, where $L(\text{unw}(M))$ is the tree language recognized by the fta $\text{unw}(M)$.

The wta $M = (Q, \Sigma, \mu, F)$ is (*bottom-up deterministic* (or a dwta) if for every $t \in \Sigma(Q)$ there exists at most one $q \in Q$ such that $t \rightarrow q \in \text{supp}(\mu)$). In other words, a wta M is deterministic if and only if $\text{unw}(M)$ is bottom-up deterministic. In a dwta we can (without loss of information) treat μ and h_μ as partial mappings $\mu: \Sigma(Q) \dashrightarrow Q \times S$ and $h_\mu: T_\Sigma(Q) \dashrightarrow Q \times S$. We use $\mu^{(1)}$ and $\mu^{(2)}$ as well as $h_\mu^{(1)}$ and $h_\mu^{(2)}$ for the corresponding projections to the first and second output component, respectively (e.g., $\mu^{(1)}: \Sigma(Q) \dashrightarrow Q$ and $\mu^{(2)}: \Sigma(Q) \dashrightarrow S$). To avoid complicated distinctions, we treat undefinedness like a value (i.e., it is equal to itself, but different from every other value). We observe that $\text{supp}(M) = L(\text{unw}(M))$ for each dwta M .³ Moreover, the restriction to final states instead of final weights in the definition of a wta does not restrict the expressive power [Bor05, Lemma 6.1.4], which applies to both wta and dwta. In addition, the transformation of a wta with final weights into an equivalent wta with final states does not add additional states, so all our results also apply to wta with final weights.

¹Clearly, weighted tree automata can also be defined for semirings or even more general weight structures, but already minimization for deterministic finite-state string automata becomes NP-hard for simple semirings that are not semifields (see [Eis03, Section 3]).

²An fta computes in the same manner as a wta over the BOOLEAN semifield \mathbb{B} .

³The statement holds because each commutative semifield is zero-divisor free [Bor03, Lemma 1].

An equivalence relation \equiv on a set A is a reflexive, symmetric, and transitive subset of A^2 . The equivalence class (or block) $[a]_{\equiv}$ of the element $a \in A$ is $\{a' \in A \mid a \equiv a'\}$, and we let $(A'/\equiv) = \{[a']_{\equiv} \mid a' \in A'\}$ for every $A' \subseteq A$. Whenever \equiv is obvious from the context, we simply omit it. The equivalence \equiv respects a set $A' \subseteq A$ if $[a] \subseteq A'$ or $[a] \subseteq A \setminus A'$ for every $a \in A$ (i.e., each equivalence class is either completely in A' or completely outside A').

Let $M = (Q, \Sigma, \mu, F)$ be a dwta. An equivalence relation $\equiv \subseteq Q^2$ is a congruence (of M) if $\mu^{(1)}(\sigma(q_1, \dots, q_k)) \equiv \mu^{(1)}(\sigma(q'_1, \dots, q'_k))$ for every $\sigma \in \Sigma_k$ and all equivalent states $q_1 \equiv q'_1, \dots, q_k \equiv q'_k$. Note that this definition of congruence completely disregards the weights, which yields that \equiv is a congruence for M if and only if \equiv is a congruence for $\text{unw}(M)$. Two states $q_1, q_2 \in Q$ are weakly equivalent, written as $q_1 \sim_M q_2$, if $h_{\mu}^{(1)}(c[q_1]) \in F$ if and only if $h_{\mu}^{(1)}(c[q_2]) \in F$ for all contexts $c \in C_{\Sigma}(Q)$. In other words, weak equivalence coincides with classical equivalence [GS84, Definition II.6.8] for $\text{unw}(M)$. Consequently, the weak equivalence relation \sim_M is actually a congruence of M that respects F [GS84, Theorem II.6.10]. The weak equivalence relation \sim_M can be computed in time $\mathcal{O}(|M| \log |Q|)$ [HMM09]. Finally, two states are (strongly) equivalent, written as $q_1 \equiv_M q_2$ if there exists a factor $s \in S \setminus \{0\}$ such that for all $c \in C_{\Sigma}(Q)$ we have

$$h_{\mu}^{(2)}(c[q_1]) \cdot \chi_F(h_{\mu}^{(1)}(c[q_1])) = s \cdot h_{\mu}^{(2)}(c[q_2]) \cdot \chi_F(h_{\mu}^{(1)}(c[q_2])) ,$$

where $\chi_F: Q \rightarrow \{0, 1\}$ is the characteristic function of F ; i.e., $F(q) = 1$ if and only if $q \in F$ for all $q \in Q$. The equivalence relation \equiv_M is called the MYHILL-NERODE equivalence relation [Mal09, Definition 3]. It is also a congruence that respects F [Mal09, Lemma 4]. If M is clear from the context, then we just write \equiv instead of \equiv_M .

3. SIGNS OF LIFE

First, we demonstrate how to efficiently compute signs of life (Definition 3.1), which are evidence that a final state can be reached. Together with these signs of life we also compute a pushing weight for each state (Section 4). Our Algorithm 1 is a straightforward extension of [Mal09, Algorithm 1] that computes on equivalence classes of states (with respect to a congruence that respects finality) instead of states.⁴ This change guarantees that equivalent states receive the same sign of life, which is an essential requirement for the algorithms in Sections 5 and 6.

Before we start we need to recall the definition of a sign of life [Mal09]. In addition, we recall the relevant properties that we use in our algorithm. For the rest of this section, let $M = (Q, \Sigma, \mu, F)$ be a dwta.

Definition 3.1 ([Mal09, Section 2]). A context $c \in C_{\Sigma}(Q)$ is a *sign of life* for the state $q \in Q$ if $h_{\mu}^{(1)}(c[q]) \in F$. Any state that has a sign of life is *live*; otherwise it is *dead*.

The following lemma justifies that we can compute signs of life for equivalence classes of congruences that respect F instead of individual states since all states of such an equivalence class share the same signs of life.

Lemma 3.2 (see [Mal09, Lemma 9]). *We have $\cong \subseteq \sim_M$ for every congruence \cong that respects F . In particular, $\equiv_M \subseteq \sim_M$. Moreover, every sign of life for $q \in Q$ is also a sign of life for every $q' \in [q]_{\cong}$.*

⁴Note that our algorithm is not simply the previous algorithm executed on the quotient dwta with respect to the congruence. The original dwta is used essentially in the computation of the pushing weights.

Alg. 1 COMPUTESOL: Compute a sign of life and its weight for each state.

Require: dwta $M = (Q, \Sigma, \mu, F)$ and congruence $\cong \subseteq Q^2$ of M that respects F

Ensure: return live state partition (L/\cong) and the mappings $\text{sol}: (L/\cong) \rightarrow C_\Sigma(Q)$ and

$\lambda: L \rightarrow S \setminus \{0\}$ such that $\lambda(q) = h_\mu^{(2)}(\text{sol}([q]_{\cong})[q])$ for every $q \in L$

```

1:  $L \leftarrow (F/\cong)$  // final states are trivially live ...
2: for all  $B \in L$  do
    $\text{sol}(B) \leftarrow \square$  // ... with the trivial context as sign of life...
4:  $\lambda(q) \leftarrow 1$  for all  $q \in B$  // ... and trivial pushing weight
    $U \leftarrow L$  // start from the final states
6: while  $U \neq \emptyset$  do
   take  $B \in U$  and  $U \leftarrow U \setminus \{B\}$  // get an unexplored class
8: for all  $\sigma(q_1, \dots, q_k) \in \Sigma(Q)$  such that  $\mu^{(1)}(\sigma(q_1, \dots, q_k)) \in B$  do
   for all  $i \in [1, k]$  such that  $[q_i]_{\cong} \notin L$  do
10:  $c \leftarrow \sigma(q_1, \dots, q_{i-1}, \square, q_{i+1}, \dots, q_k)$  // prepare context
    $L \leftarrow L \cup \{[q_i]_{\cong}\}; U \leftarrow U \cup \{[q_i]_{\cong}\}$  // add class to  $L$  and  $U$ 
12:  $\text{sol}([q_i]_{\cong}) \leftarrow \text{sol}(B)[c]$  // add transition to target block's sign of life
    $\lambda(q) \leftarrow \lambda(\mu^{(1)}(c[q])) \cdot \mu^{(2)}(c[q])$  for all  $q \in [q_i]_{\cong}$  // multiply transition weight
14: return  $(L, \text{sol}, \lambda)$ 

```

Proof. It is known that \sim_M is the coarsest congruence that respects F [GS84, Theorem II.6.10].⁵ Consequently, $\cong \subseteq \sim_M$ and $\equiv_M \subseteq \sim_M$ since we already remarked that \equiv_M is also a congruence that respects F . Based on the definition of \sim_M it is trivial to see that all elements of an equivalence class of \sim_M share the same signs of life [Mal09, Lemma 9]. Since $[q]_{\cong} \subseteq [q]_{\sim_M}$ we obtain the desired statement. \square

Algorithm 1 simply attempts to reach all states from the final states computing a context that takes the state to a final state (*i.e.*, a sign of life) as well as its weight in the process. Due to Lemma 3.2 the signs of life are computed for equivalence classes (or blocks) instead of individual states. Now let us explain Algorithm 1 in detail. Every final state $q \in F$ is trivially live as evidenced by the trivial sign of life \square . Since the congruence \cong respects F , the set (F/\cong) contains equivalence classes that contain only final states. We set the sign of life for each class to \square [see Line 3], and for each involved state q we set its pushing weight to 1 [see Line 4]. Overall, this initialization takes time $\mathcal{O}(|F|)$. Next, we add all those blocks to the live states L and to the blocks U yet to be explored. As long as there are still unexplored blocks, we select a block B from U and remove it from U . Then we consider all transitions that end in a state that belongs to the block B and check whether it contains a source state that is not yet present in L . For each such source state q_i , we add its equivalence class $[q_i]_{\cong}$ to both L and U . Then we set the sign of life for this class to the sign of life for B extended by the considered transition [see Line 12]. Finally, we select each state q from $[q_i]_{\cong}$ and compute a pushing weight by multiplying the weight of the currently considered transition with q_i replaced by q to the already computed pushing weight for the target state reached by the modified transition [see Line 13].

⁵Mind that \sim_M coincides with classical equivalence on $\text{unw}(M)$ and that our notion of congruence completely disregards the weights.

Theorem 3.3. *Algorithm 1 is correct and runs in time $\mathcal{O}(|M| + |Q|)$.*

Proof. Since our algorithm is similar to the one of [Mal09], our proof closely resembles the proofs of [Mal09, Lemma 10 and Theorem 11] adjusted to equivalence classes. We already argued that the initialization runs in time $\mathcal{O}(|F|) \subseteq \mathcal{O}(|Q|)$. It is easy to see that $U \subseteq L$ at all times in the main loop [Line 6–13] of the algorithm. Consequently, each block can be added at most once to U since it is added at the same time to L and only blocks not in L can be added to U . This yields that the main loop executes at most $|(Q/\cong)| \leq |Q|$ times. The inner loop [Line 9–13] can execute at most $|M|$ times since each transition is considered at most once in the middle loop and at most once for each source state of the transition. The statements in the inner loop all execute in constant time except for Line 13, which can be executed once for each state $q \in Q$. Overall, we thus obtain the running time $\mathcal{O}(|M| + |Q|)$.

Now let us prove the post-conditions. By Lemma 3.2 we know that signs of life are shared between elements in an equivalence class of \cong . The remaining statements are proved by induction along the outer main loop. Initially, we set

$$\lambda(q) = 1 = h_\mu^{(2)}(q) = h_\mu^{(2)}(\square[q])$$

by Lines 3–4, which proves the post-condition because $\text{sol}([q]_{\cong}) = \square$. In the main loop, we set $\lambda(q) = \lambda(\mu^{(1)}(c[q])) \cdot \mu^{(2)}(c[q])$ in Line 13. The equivalence class of $q' = \mu^{(1)}(c[q])$ has already been explored in a previous iteration because $q \cong q_i$, which by the congruence property yields $\mu^{(1)}(c[q]) \cong \mu^{(1)}(c[q_i])$ and the latter was in the explored equivalence class B , which in turn yields that the former is in B . Consequently, we can employ the induction hypothesis and obtain $\lambda(q') = h_\mu^{(2)}(\text{sol}(B)[q'])$. In addition,

$$\begin{aligned} \lambda(q) &= \lambda(q') \cdot \mu^{(2)}(c[q]) = h_\mu^{(2)}(\text{sol}(B)[q']) \cdot \mu^{(2)}(c[q]) \\ &= h_\mu^{(2)}(\text{sol}(B)[c[q]]) = h_\mu^{(2)}((\text{sol}(B)[c])[q]) , \end{aligned}$$

which proves the post-condition because $\text{sol}([q]_{\cong}) = \text{sol}([q_i]_{\cong}) = \text{sol}(B)[c]$ by Line 12. Clearly, $\text{sol}([q]_{\cong})$ is a sign of life for q , which proves that q is live. Finally, suppose that there is a live state $q \in Q$ such that $[q]_{\cong} \notin L$ (*i.e.*, we assume a live state that is not classified as such by Algorithm 1). Since it is live, it has a sign of life $c \in C_\Sigma(Q)$. By induction on c we can prove that, when processing $c[q]$, there exists a transition that uses a source state q_i such that $[q_i]_{\cong} \notin L$, whereas the target state q' is such that $[q']_{\cong} \in L$.⁶ However, since $[q']_{\cong}$ was explored, the considered transition was considered in the algorithm, which means that the equivalence class $[q_i]_{\cong}$ was added to L . This contradicts the assumption, which shows that all states that are not represented in L are indeed dead. \square

Example 3.4. Our example dwta $N = (Q, \Sigma, \mu, F)$ is depicted left in Figure 1. For any transition (small circle, the annotation specifies the input symbol and the weight separated by a colon), the arrow leads to the target state and the source states q_1, \dots, q_k have been arranged in a counter-clockwise fashion starting from the target arrow. For example, the bottom center transition labeled $\sigma : 4$ in the left dwta of Figure 1 corresponds to $\mu(\sigma(q_b, q_1) \rightarrow q_2) = 4$; *i.e.*, its target state is q_2 , its symbol is σ , its source states are $q_b q_1$ (in this order), and its weight is 4. As usual, final states are doubly circled. The graphical representation of wta is explained in detail in [Bor05]. The coarsest congruence \cong respecting $F = \{q_1, q_f\}$ is represented by the set $\{\{q_1, q_f\}, \{q_2, q_b\}\}$ of equivalence classes (*i.e.*, partition). We use this congruence in Algorithm 1. First, the block F of final states is marked as live and added to U .

⁶Such a switch must exist because all the final states are represented in L .

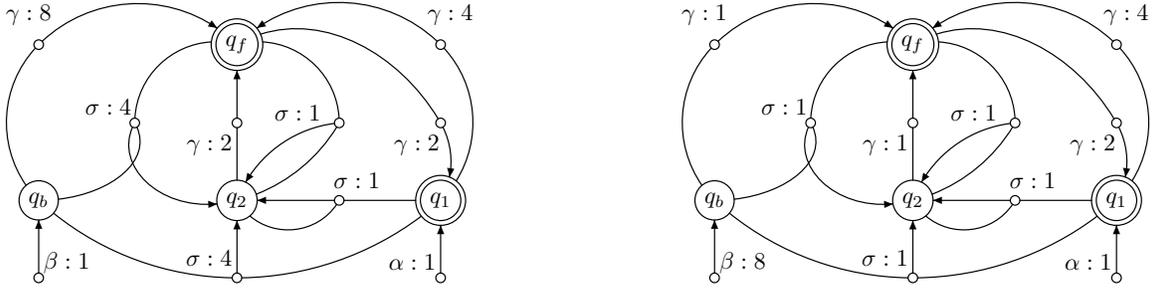


Figure 1: Dwta over the rational numbers before (left) and after (right) pushing.

It is assigned the trivial context \square as sign of life and each final state is assigned the trivial weight 1. Clearly, we can only select one equivalence class $B = F$ in the main loop. Let us consider the transition $\gamma(q_b) \rightarrow q_f$, whose target state q_f is in B . Since $[q_b]_{\cong} = \{q_2, q_b\}$ has not yet been marked as live, we add it to both L and U . In addition, we set its sign of life to $\gamma(\square)$. Finally, we set the pushing weights to $\lambda(q_b) = \lambda(q_f) \cdot \mu^{(2)}(\gamma(q_b) \rightarrow q_f) = 8$ and $\lambda(q_2) = \lambda(q_f) \cdot \mu^{(2)}(\gamma(q_2) \rightarrow q_f) = 2$. Now all states are live, so the loops will terminate. Consequently, we have computed all signs of life and the pushing weights

$$\lambda(q_1) = \lambda(q_f) = 1 \quad \lambda(q_2) = 2 \quad \text{and} \quad \lambda(q_b) = 8 .$$

4. PUSHING

The MYHILL-NERODE congruence requires that there is a unique scaling factor for every pair (q, q') of equivalent states. Thus, any fixed sign of life c for both q and q' [for which $\chi_F(h_\mu^{(1)}(c[q])) = 1 = \chi_F(h_\mu^{(1)}(c[q']))$] yields non-zero weights $h_\mu^{(2)}(c[q])$ and $h_\mu^{(2)}(c[q'])$, which can be used to determine this unique scaling factor between q and q' . In fact, we already computed those weights $\lambda(q)$ and $\lambda(q')$ in Algorithm 1. By Lemma 3.2, states that are not weakly equivalent (and thus might not have the same sign of life after executing Algorithm 1 with \sim_M) also cannot be equivalent. For the remaining pairs of live states, we computed a sign of life $\text{sol}([q]_{\sim_M})$ for the equivalence class $[q]_{\sim_M}$ of q in the previous section. In addition, we computed pushing weights $\lambda(q)$ and $\lambda(q')$. Now, we will use these weights to normalize the wta by *pushing* [Moh97, Eis03, PG09]. Intuitively, pushing cancels the scaling factor for equivalent states, which we will prove in the next section. In general, it just redistributes weights along the transitions. In weighted (finite-state) string automata [Sak09], pushing is performed from the final states towards the initial states [Moh97]. Since we work with bottom-up wta [Bor05] (*i.e.*, our notion of determinism is bottom-up), this works analogously here by moving weights from the root towards the leaves. However, we introduce our notion of pushing for arbitrary, not necessarily deterministic wta. To this end, we lift the corresponding definition [Moh97, page 296] from string to tree automata.

In this section, let $M = (Q, \Sigma, \mu, F)$ be an arbitrary wta and $\lambda: Q \rightarrow S \setminus \{0\}$ be an arbitrary mapping such that $\lambda(q) = 1$ for every $q \in F$.

Definition 4.1. The *pushed* wta $\text{push}_\lambda(M)$ is (Q, Σ, μ', F) such that for every $\sigma \in \Sigma_k$ and $q, q_1, \dots, q_k \in Q$

$$\mu'(\sigma(q_1, \dots, q_k) \rightarrow q) = \lambda(q) \cdot \mu(\sigma(q_1, \dots, q_k) \rightarrow q) \cdot \prod_{i=1}^k \lambda(q_i)^{-1} .$$

The mapping λ indicates the pushed weights. It is non-zero everywhere and has to be 1 for final states because our model does not have final weights.⁷ In the pushed wta $\text{push}_\lambda(M)$, the weight of every transition leading to the state $q \in Q$ is obtained from the weight of the corresponding transition in M by multiplying the weight $\lambda(q)$. To compensate, the weight of every transition leaving the state q will cancel the weight $\lambda(q)$ by multiplying with $\lambda(q)^{-1}$. Thus, we expect an equivalent wta after pushing, which we confirm by showing that M and $\text{push}_\lambda(M)$ are indeed equivalent. The corresponding statement for string automata is [Moh97, Lemma 4].

Proposition 4.2. *The wta M and $\text{push}_\lambda(M)$ are equivalent. Moreover, if M is deterministic, then so is $\text{push}_\lambda(M)$.*

Proof. Let $\text{push}_\lambda(M) = M' = (Q, \Sigma, \mu', F)$. The preservation of determinism is obvious because $\text{supp}(\mu') \subseteq \text{supp}(\mu)$.⁸ We prove that $h_{\mu'}(t \rightarrow q) = \lambda(q) \cdot h_\mu(t \rightarrow q)$ for every $t \in T_\Sigma$ and $q \in Q$ by induction on t . Let $t = \sigma(t_1, \dots, t_k)$ for some $\sigma \in \Sigma_k$ and $t_1, \dots, t_k \in T_\Sigma$. By the induction hypothesis, we have $h_{\mu'}(t_i \rightarrow q_i) = \lambda(q_i) \cdot h_\mu(t_i \rightarrow q_i)$ for every $i \in [1, k]$ and $q_i \in Q$. Consequently,

$$\begin{aligned} h_{\mu'}(t \rightarrow q) &= \sum_{q_1, \dots, q_k \in Q} \mu'(\sigma(q_1, \dots, q_k) \rightarrow q) \cdot \prod_{i=1}^k h_{\mu'}(t_i \rightarrow q_i) \\ &= \sum_{q_1, \dots, q_k \in Q} \lambda(q) \cdot \mu(\sigma(q_1, \dots, q_k) \rightarrow q) \cdot \prod_{i=1}^k \lambda(q_i)^{-1} \cdot \prod_{i=1}^k (\lambda(q_i) \cdot h_\mu(t_i \rightarrow q_i)) \\ &= \lambda(q) \cdot h_\mu(t \rightarrow q) . \end{aligned}$$

We complete the proof as follows.

$$M'(t) = \sum_{q \in F} h_{\mu'}(t \rightarrow q) = \sum_{q \in F} \lambda(q) \cdot h_\mu(t \rightarrow q) = \sum_{q \in F} h_\mu(t \rightarrow q) = M(t)$$

because $\lambda(q) = 1$ for every $q \in F$. □

Theorem 4.3. *The wta $\text{push}_\lambda(M)$ is equivalent to M and can be obtained in time $\mathcal{O}(|M|)$.*

Example 4.4. Let us return to our example dwta N left in Figure 1 and perform pushing. The pushing weights λ are given in Example 3.4. We consider the transition $\sigma(q_b, q_f) \rightarrow q_2$, which has weight 4 in N . In $\text{push}_\lambda(N)$ this transition has the weight

$$\lambda(q_2) \cdot \mu(\sigma(q_b, q_f) \rightarrow q_2) \cdot \lambda(q_b)^{-1} \cdot \lambda(q_f)^{-1} = 2 \cdot 4 \cdot 8^{-1} \cdot 1^{-1} = 1 .$$

The dwta $\text{push}_\lambda(N)$ is presented right in Figure 1. With a little effort, we can confirm that q_2 and q_b are equivalent in $\text{push}_\lambda(N)$, whereas q_1 and q_f are not.

⁷As already mentioned, the restriction to final states is a convenience and not an essential restriction.

⁸In fact, $\text{supp}(\mu') = \text{supp}(\mu)$ because semifields are zero-divisor free [Bor03, Lemma 1].

Alg. 2 Overall structure of our minimization algorithm; see [Mal09] for details on the procedure MERGESTATES. Note that the final merging is performed on the input dwta M . The alphabetic dwta N is only needed to compute the equivalence \equiv_M .

Require: a dwta M with states Q

Ensure: return a minimal, equivalent dwta

$\sim_M \leftarrow \text{COMPUTECOARSESTCONGRUENCE}(M, Q \times Q)$	// complexity: $\mathcal{O}(M \log Q)$
2: $(L, \text{sol}, \lambda) \leftarrow \text{COMPUTESOL}(M, \sim_M)$	// complexity: $\mathcal{O}(M)$
$N \leftarrow \text{alph}(\text{push}_\lambda(M))$	// complexity: $\mathcal{O}(M)$
4: $\equiv_M \leftarrow \text{COMPUTECOARSESTCONGRUENCE}(N, \sim_M)$	// complexity: $\mathcal{O}(M \log Q)$
return MERGESTATES(M, \equiv_M, λ)	// complexity: $\mathcal{O}(M)$

5. MINIMIZATION

Our main application of weight pushing is efficient dwta minimization, which we present next. The overall structure of our minimization procedure is presented in Algorithm 2. As mentioned earlier, the coarsest congruence \sim_M for a dwta $M = (Q, \Sigma, \mu, F)$ that respects F can be obtained by minimization [HMM09] of $\text{unw}(M)$. We call this procedure COMPUTECOARSESTCONGRUENCE and supply it with a dwta M and an equivalence relation. It returns the coarsest congruence (of M) that refines the given equivalence relation.

*Let $M = (Q, \Sigma, \mu, F)$ be a dwta (without useless states) and $\lambda: Q \rightarrow S \setminus \{0\}$ be the pushing weights computed by Algorithm 1 when run on M and \sim_M .⁹
In addition, we let $\text{push}_\lambda(M) = M' = (Q, \Sigma, \mu', F)$.*

The dwta M' has the property that $(\mu')^{(2)}(\sigma(q_1, \dots, q_k)) = (\mu')^{(2)}(\sigma(q'_1, \dots, q'_k))$ for all $\sigma \in \Sigma_k$ and states $q_i \equiv_M q'_i$ for every $i \in [1, k]$. We will prove this property (5.1) in Lemma 5.2. It is this property, which, in analogy to the string case [Moh97, Eis03], allows us to compute the equivalence $\equiv_M = \sim_N$ on an unweighted fta N , in which we treat the transition weight as part of the input symbol. For example, the algorithm of [HMM09] can then be used to compute \sim_N . Finally, we merge the equivalent states using the information about the scaling factors contained in the pushing weights λ in the same way as in [Mal09]. Let us start with the formal definitions.¹⁰

Definition 5.1. Let $M = (Q, \Sigma, \mu, F)$ be a dwta, and let $S' = \{\mu(\tau) \mid \tau \in \text{supp}(\mu)\}$ be the finite set of non-zero weights that occur as transition weights in M . The *alphabetic dwta* $\text{alph}(M)$ for M is $(Q, \Sigma \times S', \mu'', F)$, where

- $\text{rk}(\langle \sigma, s \rangle) = \text{rk}(\sigma)$ for every $\sigma \in \Sigma$ and $s \in S'$,
- $\mu''(\tau) = 1$ for every $\tau \in \text{supp}(\mu'')$, and
- for every $\sigma \in \Sigma_k$, $s \in S'$, and $q, q_1, \dots, q_k \in Q$

$$\mu''(\langle \sigma, s \rangle(q_1, \dots, q_k) \rightarrow q) = 1 \quad \iff \quad \mu(\sigma(q_1, \dots, q_k) \rightarrow q) = s .$$

Clearly, the construction of $\text{alph}(M)$ can be performed in time $\mathcal{O}(|M|)$. Next, we show that the equivalence \equiv_M in M coincides with the equivalence $\sim_{\text{alph}(M')}$ in $\text{alph}(M')$, where $M' = \text{push}_\lambda(M)$. We achieve this proof by showing both inclusions.

⁹In a dwta without useless states we have $|Q| \leq |M|$.

¹⁰We avoid a change of the weight structure from our semifield to the BOOLEAN semifield \mathbb{B} since the multiplicative submonoid induced by $\{0, 1\}$ is isomorphic to the multiplicative monoid of \mathbb{B} . Thus, our dwta with weights in $\{0, 1\}$ compute in the same manner as a dwta over \mathbb{B} or equivalently a deterministic fta.

Lemma 5.2. *The congruence \equiv_M of M is a congruence of $\text{alph}(M')$ that respects F .*

Proof. Let $\text{alph}(M') = (Q, \Sigma \times S', \mu'', F)$. Since M and $\text{alph}(M')$ have the same final states F , \equiv_M trivially respects F because it is a congruence of M that respects F . Naturally, \equiv_M is an equivalence, so it remains to prove the congruence property for $\text{alph}(M')$. Let $\sigma \in \Sigma_k$ and $q_i \equiv_M q'_i$ for every $i \in [1, k]$. Then

$$\mu^{(1)}(\sigma(q_1, \dots, q_k)) \equiv_M \mu^{(1)}(\sigma(q'_1, \dots, q'_k))$$

because \equiv_M is a congruence of M . For the moment, let us assume that¹¹

$$(\mu')^{(2)}(\sigma(q_1, \dots, q_k)) = s = (\mu')^{(2)}(\sigma(q'_1, \dots, q'_k)) ,$$

then

$$\begin{aligned} (\mu'')^{(1)}(\langle \sigma, s \rangle(q_1, \dots, q_k)) &= (\mu')^{(1)}(\sigma(q_1, \dots, q_k)) = \mu^{(1)}(\sigma(q_1, \dots, q_k)) \\ \equiv_M \mu^{(1)}(\sigma(q'_1, \dots, q'_k)) &= (\mu')^{(1)}(\sigma(q'_1, \dots, q'_k)) = (\mu'')^{(1)}(\langle \sigma, s \rangle(q'_1, \dots, q'_k)) . \end{aligned}$$

For all the remaining combinations of $\langle \sigma, s' \rangle$ we have that both $(\mu'')^{(1)}(\langle \sigma, s' \rangle(q_1, \dots, q_k))$ and $(\mu'')^{(1)}(\langle \sigma, s' \rangle(q'_1, \dots, q'_k))$ are undefined and thus equal. We have thus proved the congruence property given the assumption. Consequently, it remains to show that the assumption

$$(\mu')^{(2)}(\sigma(q_1, \dots, q_k)) = (\mu')^{(2)}(\sigma(q'_1, \dots, q'_k)) \quad (5.1)$$

is true. By Definition 4.1, we have

$$(\mu')^{(2)}(\sigma(q_1, \dots, q_k)) = \lambda(\mu^{(1)}(\sigma(q_1, \dots, q_k))) \cdot \mu^{(2)}(\sigma(q_1, \dots, q_k)) \cdot \prod_{i=1}^k \lambda(q_i)^{-1} \quad (5.2)$$

$$(\mu')^{(2)}(\sigma(q'_1, \dots, q'_k)) = \lambda(\mu^{(1)}(\sigma(q'_1, \dots, q'_k))) \cdot \mu^{(2)}(\sigma(q'_1, \dots, q'_k)) \cdot \prod_{i=1}^k \lambda(q'_i)^{-1} . \quad (5.3)$$

Now we prove that

$$\begin{aligned} &\lambda(\mu^{(1)}(c_j[q_j])) \cdot \mu^{(2)}(c_j[q_j]) \cdot \prod_{i=1}^{j-1} \lambda(q'_i)^{-1} \cdot \prod_{i=j}^k \lambda(q_i)^{-1} \\ &= \lambda(\mu^{(1)}(c_j[q'_j])) \cdot \mu^{(2)}(c_j[q'_j]) \cdot \prod_{i=1}^j \lambda(q'_i)^{-1} \cdot \prod_{i=j+1}^k \lambda(q_i)^{-1} \end{aligned} \quad (5.4)$$

for every $j \in [1, k]$, where $c_j = \sigma(q'_1, \dots, q'_{j-1}, \square, q_{j+1}, \dots, q_k)$. Let $p_j = \mu^{(1)}(c_j[q_j])$ and $p'_j = \mu^{(1)}(c_j[q'_j])$. Since $q_j \equiv_M q'_j$, we also have that $p_j \equiv_M p'_j$ because \equiv_M is a congruence of M . This yields that $p_j \sim_M p'_j$ by Lemma 3.2. Let $c = \text{sol}([p_j]_{\sim_M})$ be a sign of life for both p_j and p'_j . Moreover, we have a constant scaling factor between the equivalent states q_j and q'_j , which yields

$$\frac{\lambda(q_j)}{\lambda(q'_j)} \stackrel{(\dagger)}{=} \frac{h_\mu^{(2)}(c[c_j[q_j]])}{h_\mu^{(2)}(c[c_j[q'_j]])} \stackrel{(\ddagger)}{=} \frac{h_\mu^{(2)}(c[p_j]) \cdot \mu^{(2)}(c_j[q_j])}{h_\mu^{(2)}(c[p'_j]) \cdot \mu^{(2)}(c_j[q'_j])} \quad (5.5)$$

$$\frac{\lambda(p_j)}{\lambda(p'_j)} = \frac{h_\mu^{(2)}(c[p_j])}{h_\mu^{(2)}(c[p'_j])} , \quad (5.6)$$

¹¹Mind that we compare the weights in $M' = \text{push}_\lambda(M)$ here.

where (\dagger) holds because $c[c_j]$ is a sign of life for both q_j and q'_j and (\ddagger) holds essentially by definition. With these equations, let us inspect the main equality.

$$\begin{aligned} & \frac{\lambda(\mu^{(1)}(c_j[q_j])) \cdot \mu^{(2)}(c_j[q_j]) \cdot \prod_{i=1}^{j-1} \lambda(q'_i)^{-1} \cdot \prod_{i=j}^k \lambda(q_i)^{-1}}{\lambda(\mu^{(1)}(c_j[q'_j])) \cdot \mu^{(2)}(c_j[q'_j]) \cdot \prod_{i=1}^j \lambda(q'_i)^{-1} \cdot \prod_{i=j+1}^k \lambda(q_i)^{-1}} \\ &= \frac{\lambda(p_j) \cdot \mu^{(2)}(c_j[q_j]) \cdot \lambda(q_j)^{-1}}{\lambda(p'_j) \cdot \mu^{(2)}(c_j[q'_j]) \cdot \lambda(q'_j)^{-1}} \stackrel{(5.6)}{=} \frac{h_\mu^{(2)}(c[p_j]) \cdot \mu^{(2)}(c_j[q_j])}{h_\mu^{(2)}(c[p'_j]) \cdot \mu^{(2)}(c_j[q'_j])} \cdot \frac{\lambda(q'_j)}{\lambda(q_j)} \stackrel{(5.5)}{=} 1 \end{aligned}$$

Now we are ready to return to the proof obligation expressed in (5.1). We apply (5.4) in total k times to obtain the desired statement.

$$\begin{aligned} (\mu')^{(2)}(\sigma(q_1, \dots, q_k)) &\stackrel{(5.2)}{=} \lambda(\mu^{(1)}(\sigma(q_1, \dots, q_k))) \cdot \mu^{(2)}(\sigma(q_1, \dots, q_k)) \cdot \prod_{i=1}^k \lambda(q_i)^{-1} \\ &= \lambda(\mu^{(1)}(c_1[q_1])) \cdot \mu^{(2)}(c_1[q_1]) \cdot \prod_{i=1}^0 \lambda(q'_i)^{-1} \cdot \prod_{i=1}^k \lambda(q_i)^{-1} \\ &\stackrel{(5.4)}{=} \lambda(\mu^{(1)}(c_1[q'_1])) \cdot \mu^{(2)}(c_1[q'_1]) \cdot \prod_{i=1}^1 \lambda(q'_i)^{-1} \cdot \prod_{i=2}^k \lambda(q_i)^{-1} \\ &= \lambda(\mu^{(1)}(c_2[q_2])) \cdot \mu^{(2)}(c_2[q_2]) \cdot \prod_{i=1}^1 \lambda(q'_i)^{-1} \cdot \prod_{i=2}^k \lambda(q_i)^{-1} \\ &\dots \\ &\stackrel{(5.4)}{=} \lambda(\mu^{(1)}(c_k[q'_k])) \cdot \mu^{(2)}(c_k[q'_k]) \cdot \prod_{i=1}^k \lambda(q'_i)^{-1} \cdot \prod_{i=k+1}^k \lambda(q_i)^{-1} \\ &= \lambda(\mu^{(1)}(\sigma(q'_1, \dots, q'_k))) \cdot \mu^{(2)}(\sigma(q'_1, \dots, q'_k)) \cdot \prod_{i=1}^k \lambda(q'_i)^{-1} \\ &\stackrel{(5.3)}{=} (\mu')^{(2)}(\sigma(q'_1, \dots, q'_k)) , \end{aligned}$$

which completes the proof. \square

Theorem 5.3. *We have $\equiv_M = \sim_N$, where $N = \text{alph}(M')$.*

Proof. Lemma 5.2 shows that \equiv_M is a congruence of N that respects F . Since \sim_N is the coarsest congruence of N that respects F by [GS84, Theorem II.6.10], we obtain that $\equiv_M \subseteq \sim_N$. The converse is simple to prove as states that are weakly equivalent in $\text{alph}(M')$ share exactly the same signs of life with the scaling factor 1. Since the signs of life already indicate the transition weights, we immediately obtain that such weakly equivalent states in $\text{alph}(M')$ have corresponding transitions with equal transition weights in M' , which proves that those states are also equivalent in M' with the scaling factor 1. The latter statement can then be used to prove that they are also equivalent in M (with a scaling factor that is potentially different from 1). \square

The currently fastest dwta minimization algorithm [Mal09] runs in time $\mathcal{O}(|M| \cdot |Q|)$. Our approach, which relies on pushing and is presented in Algorithm 2, achieves the same run-time $\mathcal{O}(|M| \log |Q|)$ as the fastest minimization algorithms for deterministic fta.

Alg. 3 Overall structure of our equivalence test.

Require: accessible dwta $M = (Q, \Sigma, \mu, F)$ and $M' = (Q', \Sigma, \mu', F')$

Ensure: return ‘yes’ if M and M' are equivalent; ‘no’ otherwise

```

   $g \leftarrow \text{COMPUTECORRESPONDENCE}(M, M')$  // complexity:  $\mathcal{O}(|M|)$ 
2:  $\sim_M \leftarrow \text{COMPUTECOARSESTCONGRUENCE}(M, Q \times Q)$  // complexity:  $\mathcal{O}(|M| \log |Q|)$ 
    $\sim_{M'} \leftarrow \text{COMPUTECOARSESTCONGRUENCE}(M', Q' \times Q')$  // complexity:  $\mathcal{O}(|M'| \log |Q'|)$ 
4:  $(L, \text{sol}, \lambda) \leftarrow \text{COMPUTESOL}(M, \sim_M)$  // complexity:  $\mathcal{O}(|M|)$ 
   if  $g$  is not compatible with the congruences  $\sim_M$  and  $\sim_{M'}$  then
6:   return no // see Lemma 6.1; complexity:  $\mathcal{O}(|Q| + |Q'|)$ 
   for all  $q' \in Q'$  do
8:    $\lambda'(q') = h_{\mu'}^{(2)}(c'[q'])$  with  $c' = \text{ren}_g(\text{sol}(\bar{g}^{-1}([q']_{\sim_{M'}})))$  // prepare pushing weights
    $N \leftarrow \text{MINIMIZE}(\text{alph}(\text{push}_\lambda(M)), \sim_M)$  // complexity:  $\mathcal{O}(|M| \log |Q|)$ 
10:  $N' \leftarrow \text{MINIMIZE}(\text{alph}(\text{push}_{\lambda'}(M')), \sim_{M'})$  // complexity:  $\mathcal{O}(|M'| \log |Q'|)$ 
   return  $\text{ISOMORPHIC?}(N, N')$  // complexity:  $\mathcal{O}(|N|)$ 

```

Corollary 5.4 (see Algorithm 2). *For every dwta $M = (Q, \Sigma, \mu, F)$, we can compute an equivalent minimal dwta in time $\mathcal{O}(|M| \log |Q|)$.*

6. TESTING EQUIVALENCE

In this final section, we want to decide whether two given dwta are equivalent. To this end, let $M = (Q, \Sigma, \mu, F)$ and $M' = (Q', \Sigma, \mu', F')$ be dwta. The overall approach is presented in Alg. 3. First, we compute a correspondence $g: Q \rightarrow Q'$ between states. For every $q \in Q$, we compute a tree $t \in T_\Sigma$, which is also called *access tree for q* , such that $h_\mu^{(1)}(t) = q$. If no access tree exists, then q is not reachable and can be deleted. A dwta, in which all states are reachable, is called *accessible*. To avoid these details, let us assume that M and M' are accessible, which can always be achieved in time $\mathcal{O}(|M| + |M'|)$. In this case, we can compute an access tree $a(q) \in T_\Sigma$ for every state $q \in Q$ in time $\mathcal{O}(|M|)$ using standard breadth-first search, in which we unfold each state (*i.e.*, explore all transitions leading to it) at most once. To keep the representation efficient, we store the access trees in the format $\Sigma(Q)$, where each state $q \in Q$ refers to its access tree $a(q)$. To obtain the corresponding state $g(q)$, we compute the state of Q' that is reached when processing the access tree $a(q)$. Formally, $g(q) = h_{\mu'}^{(1)}(a(q))$ for every $q \in Q$. This computation can also be achieved in time $\mathcal{O}(|M|)$ since we can reuse the results for the subtrees. Consequently, we have that $h_\mu^{(1)}(a(q)) = q$ and $h_{\mu'}^{(1)}(a(q)) = g(q)$ for every $q \in Q$. Clearly, the computation of the access trees $a: Q \rightarrow T_\Sigma$ and the correspondence $g: Q \rightarrow Q'$ can be performed in time $\mathcal{O}(|M|)$. Next, we compute the coarsest congruences \sim_M and $\sim_{M'}$ for M and M' that respect F and F' , respectively, and the signs of life for M .

Lemma 6.1. *Let M and M' be equivalent. The correspondence $g: Q \rightarrow Q'$ is compatible with the congruences \sim_M and $\sim_{M'}$; *i.e.*, $g(q) \sim_{M'} g(p)$ if and only if $q \sim_M p$ for all $q, p \in Q$. Moreover, for every reachable $q' \in Q'$ there exists $q \in Q$ such that $g(q) \in [q']_{\sim_{M'}}$. Consequently, g induces a bijection $\bar{g}: (Q/\sim_M) \rightarrow (Q'/\sim_{M'})$ on the equivalence classes.*

Proof. Let $q, p \in Q$, and let $t = a(q)$ and $u = a(p)$ be the corresponding access trees. Then

$$\begin{aligned}
& q \sim_M p \\
& \iff \{c \in C_\Sigma(Q) \mid h_\mu^{(1)}(c[q]) \in F\} = \{c \in C_\Sigma(Q) \mid h_\mu^{(1)}(c[p]) \in F\} \\
& \iff \{c \in C_\Sigma \mid h_\mu^{(1)}(c[q]) \in F\} = \{c \in C_\Sigma \mid h_\mu^{(1)}(c[p]) \in F\} \quad (\star) \\
& \iff \{c \in C_\Sigma \mid c[t] \in \text{supp}(M)\} = \{c \in C_\Sigma \mid c[u] \in \text{supp}(M)\} \quad (\dagger) \\
& \iff \{c \in C_\Sigma \mid c[t] \in \text{supp}(M')\} = \{c \in C_\Sigma \mid c[u] \in \text{supp}(M')\} \quad (\text{since } M = M') \\
& \iff \{c \in C_\Sigma \mid h_{\mu'}^{(1)}(c[g(q)]) \in F'\} = \{c \in C_\Sigma \mid h_{\mu'}^{(1)}(c[g(p)]) \in F'\} \quad (\dagger) \\
& \iff \{c \in C_\Sigma(Q) \mid h_{\mu'}^{(1)}(c[g(q)]) \in F'\} = \{c \in C_\Sigma(Q) \mid h_{\mu'}^{(1)}(c[g(p)]) \in F'\} \quad (\star) \\
& \iff g(q) \sim_{M'} g(p) \text{ ,}
\end{aligned}$$

where (\star) follows from [Mal08, Lemma 4] and (\dagger) follows from the easy fact that $h_\mu^{(1)}(c[q]) \in F$ if and only if $c[t] \in \text{supp}(M)$ for all $q \in Q$ and $t \in T_\Sigma$ such that $h_\mu^{(1)}(t) = q$.

For the second statement, let $q' \in Q'$ be a reachable state, and let $t \in T_\Sigma$ be such that $h_{\mu'}^{(1)}(t) = q'$. Clearly, we have

$$\begin{aligned}
& \{c \in C_\Sigma \mid h_{\mu'}^{(1)}(c[q']) \in F'\} \stackrel{(\dagger)}{=} \{c \in C_\Sigma \mid c[t] \in \text{supp}(M')\} = \{c \in C_\Sigma \mid c[t] \in \text{supp}(M)\} \\
& \stackrel{(\dagger)}{=} \{c \in C_\Sigma \mid h_\mu^{(1)}(c[q]) \in F\} \stackrel{(\dagger)}{=} \{c \in C_\Sigma \mid c[a(q)] \in \text{supp}(M)\} \\
& = \{c \in C_\Sigma \mid c[a(q)] \in \text{supp}(M')\} \stackrel{(\dagger)}{=} \{c \in C_\Sigma \mid h_{\mu'}^{(1)}(c[g(q)]) \in F'\}
\end{aligned}$$

where $q = h_\mu^{(1)}(t)$. Consequently, using (\star) we obtain $q' \sim_{M'} g(q)$. \square

We just demonstrated that for equivalent dwta the correspondence g always yields a bijection $\bar{g}: (Q/\sim_M) \rightarrow (Q'/\sim_{M'})$. We can test the compatibility in time $\mathcal{O}(|Q| + |Q'|)$. Next we transfer the signs of life via \bar{g} to the equivalence classes of $\sim_{M'}$ and calculate the corresponding pushing weights for all states $q' \in Q'$. Since the signs of life can contain states of Q , we need to rename them using the correspondence g , so we use the function $\text{ren}_g: T_\Sigma(Q \cup \{\square\}) \rightarrow T_\Sigma(Q' \cup \{\square\})$, which is defined by $\text{ren}_g(\square) = \square$, $\text{ren}_g(q) = g(q)$ for all $q \in Q$, and $\text{ren}_g(\sigma(t_1, \dots, t_k)) = \sigma(\text{ren}_g(t_1), \dots, \text{ren}_g(t_k))$ for all $\sigma \in \Sigma_k$ and trees $t_1, \dots, t_k \in T_\Sigma(Q \cup \{\square\})$. We note that $\text{ren}_g(c) \in C_\Sigma(Q')$ for all $c \in C_\Sigma(Q)$.

Using this approach corresponding equivalence classes receive the same sign of life (modulo the renaming ren_g of the states). We then minimize M and M' using the method of Section 5 (i.e., we perform pushing followed by unweighted minimization). Finally, we test the obtained deterministic fta for isomorphism.

Lemma 6.2. *We use the symbols of Algorithm 3. Given a compatible correspondence g , the dwta M and M' are equivalent if and only if the deterministic unweighted fta $\text{alph}(\text{push}_\lambda(M))$ and $\text{alph}(\text{push}_{\lambda'}(M'))$ are equivalent.*

Proof. Clearly, if the deterministic fta $\text{alph}(\text{push}_\lambda(M))$ and $\text{alph}(\text{push}_{\lambda'}(M'))$ are equivalent, then also $\text{push}_\lambda(M)$ and $\text{push}_{\lambda'}(M')$ are equivalent since the weights are annotated on the symbols of the former devices. Moreover, since pushing preserves the semantics (see Proposition 4.2), also the dwta M and M' are equivalent, which concludes one direction. For the other direction, let M and M' be equivalent. Then also $\text{push}_\lambda(M)$ and $\text{push}_{\lambda'}(M')$ are equivalent due to Proposition 4.2. An easy adaptation of the proof (of the equality (5.1)

of the transition weights) of Lemma 5.2 can be used to show that the transition weights of corresponding transitions are equal and hence $\text{alph}(\text{push}_\lambda(M))$ and $\text{alph}(\text{push}_{\lambda'}(M'))$ are equivalent. \square

Lemma 6.2 proves the correctness of Algorithm 3 because the minimal deterministic fta for a given tree language is unique (up to isomorphism) [GS84, Theorem 2.11.12]. The run-time of our algorithm should be compared to the previously (asymptotically) fastest equivalence test for dwta of [DHM11], which runs in time $\mathcal{O}(|M| \cdot |M'|)$.

Theorem 6.3. *We can test equivalence of M and M' in time $\mathcal{O}((|M|+|M'|) \log(|Q| + |Q'|))$.*

ACKNOWLEDGMENTS

The authors gratefully acknowledge the insight and suggestions provided by the reviewers of the conference and the current version.

REFERENCES

- [Ang87] Dana Angluin. Learning regular sets from queries and counterexamples. *Information and Computation*, 75(2):87–106, 1987.
- [BLB83] Symeon Bozapalidis and Olympia Louscou-Bozapalidou. The rank of a formal tree power series. *Theoretical Computer Science*, 27(1–2):211–215, 1983.
- [BMV10] Matthias Büchse, Jonathan May, and Heiko Vogler. Determinization of weighted tree automata using factorizations. *Journal of Automata, Languages and Combinatorics*, 15(3–4):229–254, 2010.
- [Bor03] Björn Borchardt. The Myhill-Nerode theorem for recognizable tree series. In *Proc. 7th Int. Conf. Developments in Language Theory*, volume 2710 of *Lecture Notes in Computer Science*, pages 146–158. Springer, 2003.
- [Bor05] Björn Borchardt. *The Theory of Recognizable Tree Series*. PhD thesis, TU Dresden, 2005.
- [Boz99] Symeon Bozapalidis. Equational elements in additive algebras. *Theory of Computing Systems*, 32(1):1–33, 1999.
- [BR82] Jean Berstel and Christophe Reutenauer. Recognizable formal power series on trees. *Theoretical Computer Science*, 18(2):115–148, 1982.
- [BV03] Björn Borchardt and Heiko Vogler. Determinization of finite state weighted tree automata. *Journal of Automata, Languages and Combinatorics*, 8(3):417–463, 2003.
- [CDG⁺07] Hubert Comon, Max Dauchet, Rémi Gilleron, Florent Jacquemard, Denis Lugiez, Christof Löding, Sophie Tison, and Marc Tommasi. Tree automata techniques and applications. Available on: <http://tata.gforge.inria.fr/>, 2007. version 2008-11-18.
- [Cho03] Christian Choffrut. Minimizing subsequential transducers: a survey. *Theoretical Computer Science*, 292(1):131–143, 2003.
- [DGMM11] Manfred Droste, Doreen Götze, Steffen Märcker, and Ingmar Meinecke. Weighted tree automata over valuation monoids and their characterization by weighted logics. In *Bozapalidis Festschrift*, volume 7020 of *Lecture Notes in Computer Science*, pages 30–55. Springer, 2011.
- [DHM11] Frank Drewes, Johanna Högberg, and Andreas Maletti. MAT learners for tree series — an abstract data type and two realizations. *Acta Informatica*, 48(3):165–189, 2011.
- [DV06] Manfred Droste and Heiko Vogler. Weighted tree automata and weighted logics. *Theoretical Computer Science*, 366(3):228–247, 2006.
- [Eis03] Jason Eisner. Simpler and more general minimization for weighted finite-state automata. In *Proc. 2003 Human Language Technology Conf.*, pages 64–71. Association for Computational Linguistics, 2003.
- [FSM11] Sylvia Friese, Helmut Seidl, and Sebastian Maneth. Earliest normal form and minimization for bottom-up tree transducers. *International Journal of Foundations of Computer Science*, 22(7):1607–1623, 2011.

- [FV09] Zoltán Fülöp and Heiko Vogler. Weighted tree automata and tree transducers. In *Handbook of Weighted Automata*, EATCS Monographs in Theoretical Computer Science, chapter 9, pages 313–403. Springer, 2009.
- [Gol99] Jonathan S. Golan. *Semirings and their Applications*. Kluwer Academic Publishers, 1999.
- [GS84] Ferenc Gécseg and Magnus Steinby. *Tree Automata*. Akadémiai Kiadó, Budapest, 1984. 2nd edition available at: <https://arxiv.org/abs/1509.06233>.
- [GS97] Ferenc Gécseg and Magnus Steinby. Tree languages. In *Handbook of Formal Languages*, volume 3, chapter 1, pages 1–68. Springer, 1997.
- [HMM07] Johanna Högberg, Andreas Maletti, and Jonathan May. Bisimulation minimisation for weighted tree automata. In *Proc. 11th Int. Conf. Developments in Language Theory*, volume 4588 of *Lecture Notes in Computer Science*, pages 229–241. Springer, 2007.
- [HMM09] Johanna Högberg, Andreas Maletti, and Jonathan May. Backward and forward bisimulation minimization of tree automata. *Theoretical Computer Science*, 410(37):3539–3552, 2009.
- [HW98] Uwe Hebisch and Hanns J. Weinert. *Semirings — Algebraic Theory and Applications in Computer Science*. World Scientific, 1998.
- [Jac11] Florent Jacquemard. *Extended Tree Automata Models for the Verification of Infinite State Systems*. Mémoire d’habilitation à diriger des recherches, Laboratoire Spécification et Vérification, ENS de Cachan, 2011.
- [KM01] Nils Klarlund and Anders Møller. *MONA Version 1.4 User Manual*. BRICS, Department of Computer Science, Aarhus University, 2001. Notes Series NS-01-1.
- [KM09] Kevin Knight and Jonathan May. Applications of weighted automata in natural language processing. In *Handbook of Weighted Automata*, EATCS Monographs in Theoretical Computer Science, chapter 14, pages 571–596. Springer, 2009.
- [Kui98] Werner Kuich. Formal power series over trees. In *Proc. 3rd Int. Conf. Developments in Language Theory*, pages 61–101. Aristotle University of Thessaloniki, 1998.
- [Mal08] Andreas Maletti. Minimizing deterministic weighted tree automata. In *Proc. 2nd Int. Conf. Language and Automata: Theory and Applications*, volume 5196 of *Lecture Notes in Computer Science*, pages 357–372. Springer, 2008.
- [Mal09] Andreas Maletti. Minimizing deterministic weighted tree automata. *Information and Computation*, 207(11):1284–1299, 2009.
- [Man08] Eleni G. Mandrali. Weighted tree automata with discounting. Master’s thesis, Aristotle University of Thessaloniki, 2008.
- [MKV10] Jonathan May, Kevin Knight, and Heiko Vogler. Efficient inference through cascades of weighted tree transducers. In *Proc. 48th Annual Meeting of the ACL*, pages 1058–1066. Association for Computational Linguistics, 2010.
- [Moh97] Mehryar Mohri. Finite-state transducers in language and speech processing. *Computational Linguistics*, 23(2):269–311, 1997.
- [PG09] Matt Post and Daniel Gildea. Weight pushing and binarization for fixed-grammar parsing. In *Proc. 11th Int. Workshop on Parsing Technologies*, pages 89–98. Association for Computational Linguistics, 2009.
- [Sak09] Jacques Sakarovitch. Rational and recognisable power series. In *Handbook of Weighted Automata*, EATCS Monographs in Theoretical Computer Science, chapter 4, pages 105–174. Springer, 2009.
- [TW68] James W. Thatcher and Jesse B. Wright. Generalized finite automata theory with an application to a decision problem of second-order logic. *Mathematical Systems Theory*, 2(1):57–81, 1968.
- [VDH16] Heiko Vogler, Manfred Droste, and Luisa Herrmann. A weighted MSO logic with storage behaviour and its Büchi-Elgot-Trakhtenbrot theorem. In *Proc. 10th Int. Conf. Language and Automata Theory and Applications*, volume 9618 of *Lecture Notes in Computer Science*, pages 127–139. Springer, 2016.
- [WDF⁺09] Christoph Weidenbach, Dilyana Dimova, Arnaud Fietzke, Rohit Kumar, Martin Suda, and Patrick Wischniewski. SPASS version 3.5. In *Proc. 22nd Int. Conf. Automated Deduction*, volume 5663 of *Lecture Notes in Computer Science*, pages 140–145. Springer, 2009.