

Star height via games

Mikołaj Bojańczyk
(University of Warsaw)

Abstract—This paper proposes a new algorithm deciding the star height problem. As shown by Kirsten, the star height problem reduces to a problem concerning automata with counters, called limitedness. The new contribution is a different algorithm for the limitedness problem, which reduces it to solving a Gale-Stewart game with an ω -regular winning condition.

The star height of a regular expression is its nesting depth of the Kleene star. The star height problem is to compute the star height of a regular language of finite words, i.e. the smallest star height of a regular expression that defines the language. For instance, the expression $(a^*b^*)^*$ has star height two, but its language has star height one, because it is defined by the expression $(a + b)^*$ of star height one, and it cannot be defined by a star-free expression. Here we consider regular expressions without complementation, in which case star-free regular expressions define only finite languages.

The star height problem is notorious for being one of the most technically difficult problems in automata theory. The problem was first stated by Eggan in 1963 and remained open for 25 years until it was solved by Hashiguchi in [Has88]. Hashiguchi’s solution is very difficult. Much work was later devoted to simplifying the proof, resulting in important new ideas like the tropical semiring [Sim88a] or Simon’s Factorisation Forest Theorem [Sim88b]. This work culminated in a simplified proof, with an elementary complexity bound, which was given by Kirsten in [Kir05]. Kirsten introduces an automaton model, called a nested distance desert automaton, and shows that the star height problem can be reduced to a decision problem for this automaton model, called the limitedness problem. The reduction of the star height problem to the limitedness problem is not difficult, and most of the effort in [Kir05] goes into deciding the limitedness problem. Later solutions to the limitedness problem can be found in the work of Colcombet on cost functions [Col09], or in the PhD thesis of Toruńczyk [Tor11]. All of these solutions are based on variants of the limitedness problem.

The contribution of this paper is a conceptually simple algorithm deciding the limitedness problem, and therefore also the star height problem. The algorithm is a reduction of the limitedness problem to the Church Synthesis Problem, i.e. the problem of finding the winner in a Gale-Stewart game with an ω -regular winning condition. The reduction is quite simple, and reduces the star height problem to a problem that has been widely studied and is well understood.

To make the paper self-contained, we show how the star height problem is reduced to the limitedness problem in Section I, and in Section II we present the new contribution, namely the algorithm for limitedness.

I would like to thank Nathanaël Fijalkow, Denis Kuperberg, Christof Löding and the anonymous referees for their helpful comments. This work has been supported by the Polish NCN grant 2014-13/B/ST6/03595.

I. REDUCTION TO LIMITEDNESS OF COST AUTOMATA

This section introduces cost automata, and describes a reduction of the star height problem to a problem for cost automata, called limitedness. The exact reduction used in this paper was shown in [Kir05] and is included to make the paper self-contained. The idea to use counters and limitedness in the study of star height dates back to the original proof of Hashiguchi.

Cost automata and the limitedness problem: The limitedness problem concerns an automaton model with counters, which in the literature appears under the following names: *nested distance desert automaton* in [Kir05]; *hierarchical B-automaton* in [BC06] and [Col09]; and *R-automaton* in [AKY08]. (The last model is slightly more general, being equivalent to the non-hierarchical version of B-automata, but its limitedness problem is no more difficult.) Here, we simply use the name *cost automaton*. Define a cost automaton to be a nondeterministic finite automaton, additionally equipped with a finite set of *counters* which store natural numbers. For each counter c , there is a distinguished subset of transitions that *increment counter c* , and a distinguished subset of transitions that *reset counter c* . The counters are totally ordered, say they are called $0, \dots, n$, and every transition that resets or increments a counter c must also reset all counters $0, \dots, c-1$. At the beginning of a run, all counters have value zero.

Define the *value* of a run to be the smallest m such that all counters have value at most m at all times during the run. In other words, this is the smallest m such that for every counter c , there can be at most m transitions in the run that increment counter c and which are not separated by a transition which resets counter c . A cost automaton is said to be *limited* over a language L if there exists a bound m such that for every word in L there exists a run of the cost automaton which is both accepting (i.e. begins in an initial state and ends in a final state) and has value at most m . The *limitedness problem* is to decide, given a cost automaton and a regular language L , whether the automaton is limited over L . The rest of Section I is devoted to proving the reduction stated in the following theorem. This is the same reduction that was used in Kirsten’s proof from [Kir05], a detailed proof can be found in Proposition 6.8 of [Kir06].

Theorem 1. *The star height problem reduces to the limitedness problem.*

In the reduction we consider a normal form of regular expressions, which are called *string expressions* following [Coh70]. The main idea is that in a string expression, a concatenation of unions is converted into a union of concatenations. String expressions have two parameters, called (star) height and degree, and are defined by induction on height. A string expression of height 0 and degree at most m is defined to be any regular expression describing a finite set of words of length at most m . For $h \geq 1$, a string expression of height at most h and degree at most m is defined to be any finite union of expressions of the form

$$w_1 f_1^* w_2 f_2^* \cdots w_i f_i^*$$

where i is at most the degree m , w_1, \dots, w_i are words of length bounded by the degree m , and f_1, \dots, f_i are string expressions of height strictly smaller than h and same degree m . A straightforward induction proof, see Lemma 3.1 in [Coh70], shows that every regular expression can be converted into a string expression without increasing star height. Therefore, a language has star height at most h if and only if it can be defined by a string expression of height at most h and some degree.

The main observation in the reduction of the star height problem to the limitedness problem is the following proposition, which says that for every star height h and every regular language L , there is a cost automaton which maps an input word to the smallest degree of a string expression that contains the input word, has star height at most h , and is contained in L .

Proposition 1. *For every regular language $L \subseteq A^*$ and $h \in \mathbb{N}$ one can compute a cost automaton such that for every word $w \in A^*$ and $m \in \mathbb{N}$, the following conditions are equivalent:*

- 1) *there is a string expression e of height at most h and degree at most m such that $w \in e \subseteq L$.*
- 2) *the automaton has a run on w with value at most m .*

Proof. The proposition is proved by induction on h . The induction base of $h = 0$ is straightforward: the cost automaton has one counter, has an accepting run if and only if the word belongs to L , and the value of the counter in this run is the length of the input word.

Let us now do the induction step. Suppose that we have proved the proposition for $h - 1$, and we want to prove it for h . Take L as in the statement of the proposition. Since L is regular, it is recognised by some homomorphism

$$\alpha : A^* \rightarrow M$$

into a finite monoid. For a subset $N \subseteq M$ define

$$[N]_h^m \subseteq A^*$$

the union of all string expressions that have degree at most m , height at most h , and which are contained in $\alpha^{-1}(N)$. By induction assumption we know that there is a cost automaton, call it \mathcal{A}_N , such that a word admits a run of value at most m if and only if it belongs to $[N]_{h-1}^m$.

In the lemma below, we use the following notation: for $w_1, \dots, w_i \in A^*$ and $N_1, \dots, N_i \subseteq M$, we write

$$w_1 N_1^* \cdots w_i N_i^* \subseteq M$$

to be all possible elements of the monoid M that can be obtained by taking a product $\alpha(w_1)n_1 \cdots \alpha(w_i)n_i$ where each n_j is in the submonoid of M generated by N_j .

Lemma 1. *Let $h, m \in \mathbb{N}$. For every input word w , condition 1 in the statement of the proposition is equivalent to:*

(*) *there exists a number $i \leq m$ and subsets*

$$N_1, \dots, N_i \subseteq M$$

such that w admits a factorisation

$$w = w_1 v_1 w_2 v_2 \cdots w_i v_i$$

such that the following conditions hold

$$w_1 N_1^* \cdots w_i N_i^* \subseteq \alpha(L) \quad (1)$$

$$|w_j| \leq m \quad \text{for } j \in \{1, \dots, i\} \quad (2)$$

$$v_j \in ([N_j]_{h-1}^m)^* \quad \text{for } j \in \{1, \dots, i\}. \quad (3)$$

The proof of the above lemma is omitted, as it follows rather simply from the definition of string expressions. To complete the proof of the proposition, we will show that there exists a cost automaton, which admits a run of value $m \in \mathbb{N}$ if and only if condition (*) in the above lemma is satisfied. When reading an input word w , the automaton proceeds as follows. It uses nondeterminism to guess a factorisation

$$w = w_1 v_1 w_2 v_2 \cdots w_i v_i,$$

as in the statement of the claim, and it also uses nondeterminism to guess the subsets N_1, \dots, N_i . The subset N_j is only stored in the memory of the automaton while processing the block v_j . Condition (1) is verified using the finite control of the automaton: when the automaton is processing the block $w_j v_j$, its finite memory stores N_j as well as

$$w_1 N_1^* \cdots w_j N_j^* \subseteq M.$$

The remaining conditions are verified using counters. The most important counter is used to the length i of the factorisation, i.e. it is incremented whenever a block of the form $w_j v_j$ is processed. A less important counter is used to bound the length of the words w_1, \dots, w_i , as required by condition (2), i.e. it is incremented whenever reading a single letter of w_i , and reset after finishing w_i . Finally, for condition (3), when processing w_j we use the inductively obtained cost automaton \mathcal{A}_{N_j} , which works for the smaller height $h - 1$. This inductively obtained cost automaton is run in a loop, corresponding to the Kleene star in condition (3), and after each iteration of this loop all of its counters are reset. \square

Theorem 1, which reduces the star height problem to the limitedness problem, is a straightforward corollary of the above proposition. Suppose that we want to decide if a regular language L has star height h . Apply Proposition 1 to L and

h , yielding a cost automaton. We claim that L has star height h if and only if the cost automaton is limited over L .

If the cost automaton is limited over L , say every word in L admits a run with value at most m , then by the proposition, we know that every word $w \in L$ belongs to some string expression e which has height at most h and degree at most m . There are finitely many string expressions of given height and degree, and therefore one can take the union of all of these string expressions, yielding a regular expression (in fact, a string expression) of star height at most h which defines L .

If L has star height h , then it is defined by a regular expression of star height h , and therefore also by a string expression of the same height (and some degree), and the cost automaton is limited by the degree.

This completes the reduction of the star height problem to the limitedness problem.

II. THE LIMITEDNESS PROBLEM IS DECIDABLE

The reduction in the previous section was used already in the proof of Kirsten, and is included here only to make this paper self-contained. This section contains the original contribution of this paper, namely an alternative decidability argument for the limitedness problem. We show that the limitedness problem reduces to solving Gale-Stewart games with ω -regular winning conditions, a problem which was shown decidable by Büchi and Landweber [BL69]. Also in this section, most of the material is not new: apart from using the the Büchi-Landweber theorem as a target of the reduction, the reduction itself is mainly based on a technique from [CL08].

Gale-Stewart Games: A *Gale-Stewart game* is a game of perfect information and infinite duration played by two players, call them A and B. Each player has their own finite alphabet, call these alphabets A and B . The game is played in rounds; in each round player A picks a letter from his alphabet, and player B responds with a letter from his alphabet. After ω rounds have been played, the game results in an ω -word of the form $a_1b_1a_2b_2\cdots$ with $a_1, a_2, \dots \in A$ and $b_1, b_2, \dots \in B$. The objective of player B is to ensure that the resulting sequence belongs to a given set of ω -words, called the *winning condition for player B*, while the objective of player A is to avoid this.

A special case of Gale-Stewart games that is important in computer science is when the winning condition is ω -regular, i.e. it is recognised by a nondeterministic Büchi automaton, or any other equivalent model. A problem known as the *Church synthesis problem* is to decide, given a Gale-Stewart game with an ω -regular winning condition, which of the players, if any, has a winning strategy. This problem was solved by Büchi and Landweber, who proved the following theorem.

Theorem 2. [BL69] *Consider a Gale-Stewart game with an ω -regular winning condition.*

- *The game is determined, i.e. one of the players has a winning strategy.*
- *One can decide which player has a winning strategy.*
- *The player with a winning strategy has a finite memory winning strategy (see proof of Lemma 2 for a definition).*

This is a non-trivial theorem and its proof is not included here. A modern proof, see e.g. Theorem 6.16 in [Tho96], consists of two results: memoryless determinacy of parity games, and conversion of nondeterministic Büchi automata into deterministic parity automata. Using the Büchi-Landweber theorem, we show below that the limitedness problem is decidable, because it reduces to determining the winner in a Gale-Stewart game with an ω -regular winning condition.

Theorem 3. *The limitedness problem reduces to solving Gale-Stewart games with ω -regular winning conditions.*

Fix a cost automaton \mathcal{A} and a regular language L , both over the same input alphabet. Let us begin by fixing some notation. We view a run of the cost automaton, not necessarily accepting, as a sequence of transitions such that the source state of the first transition is the initial state, and which is consistent in the usual sense: for every transition in the run, its target state is the same as the source state of the following transition. We will also talk about infinite runs. If $\delta_1, \dots, \delta_n$ is a sequence of sets of transitions in the cost automaton, then a run in $\delta_1 \cdots \delta_n$ is defined to be a run where the i -th transition is from the set δ_i ; likewise for infinite sequences of sets of transitions.

For a number $m \in \mathbb{N} \cup \{\infty\}$, consider the following Gale-Stewart game, call it the *limitedness game with bound m* . The alphabet for player A is the input alphabet of the cost automaton. The alphabet for player B is the powerset of the set of transitions in the cost automaton. The idea is that player A constructs an infinite word, and player B constructs a representation of a set of runs that is sufficient to accept prefixes of this word with counter values that are small with respect to m . This idea is formalised by the winning condition, defined as follows. When m is a finite number, then the winning condition for player B is the set of sequences

$$a_1\delta_1a_2\delta_2\cdots$$

such that a_1, a_2, \dots are letters and $\delta_1, \delta_2, \dots$ are sets of transitions, and the following conditions hold for every i :

- 1) every transition in the set δ_i is over letter a_i ; and
- 2) every run in $\delta_1 \cdots \delta_i$ has value at most m ; and
- 3) if $a_1 \cdots a_i \in L$ then some run in $\delta_1 \cdots \delta_i$ is accepting.

In particular, the winning condition implies that if $a_1 \cdots a_i$ belongs to L , then it admits an accepting run of the cost automaton which has value at most m . For $m = \infty$, item 2) is replaced by

- 2') for every infinite run in $\delta_1\delta_2\cdots$ and every counter c , if the run increments c infinitely often, then it resets c infinitely often.

The following lemma is also true for finite values of m , but we only use the case of $m = \infty$ in the reduction.

Lemma 2. *For $m = \infty$ the winning condition is ω -regular.*

Proof. For items 1) and 3) one only needs a deterministic automaton with an acceptance condition, sometimes called the *trivial acceptance condition*, where a run is accepting unless

it does not visit some designated sink state. For item 1), only one additional state except the sink state is needed; for item 3) it suffices to compute the set of states of the cost automaton that are accessible via runs in $\delta_1 \cdots \delta_i$. For item 2'), the full power of ω -automata is used. The complement of item 2') is recognised by a nondeterministic Büchi automaton whose number of states is polynomial in the number of counters; the Büchi automaton nondeterministically guesses a run of the cost automaton which increments a counter infinitely often without resetting it. Therefore by determinisation, e.g. using the Safra construction, the whole winning condition is recognised by a deterministic parity automaton of size exponential in the cost automaton, but with a polynomial number of ranks in the parity condition. See [Tho96] for descriptions of the discussed automata and their determinisation. \square

In particular, thanks to the Büchi-Landweber Theorem, one can decide the winner in the limitedness game with bound ∞ . This can be done in exponential time, because it amounts to solving an exponential size parity game with a polynomial number of ranks [Tho96]. Therefore, Theorem 3 will follow from Lemma 2 below, which reduces the limitedness problem to finding the winner in the limitedness game with bound ∞ .

Proposition 2. *The following are equivalent:*

- (a) *The cost automaton \mathcal{A} is limited over language L ;*
- (b) *Player B has a winning strategy in the limitedness game with some finite bound $m < \infty$.*
- (c) *Player B has a winning strategy in the limitedness game with bound ∞ .*

Our reduction gives an exponential time solution to the limitedness problem, which is inferior to Kirsten's optimal polynomial space solution [Kir05]. The rest of this paper is devoted to showing the above proposition. We prove equivalence of (b) and (c), and equivalence of (a) and (b).

Equivalence of (b) and (c): The implication (b) \Rightarrow (c) is immediate. Let us do (c) \Rightarrow (b). Since the limitedness game with bound ∞ is a Gale-Stewart game with an ω -regular winning condition, the Büchi-Landweber theorem implies that if player B has a winning strategy, then player B has a finite memory winning strategy. Let us recall the definition of a finite memory winning strategy for player B. This is a deterministic finite automaton \mathcal{B} whose input alphabet is the alphabet of player A (in our application: the input alphabet of the cost automaton) together with a function f from states of \mathcal{B} to the alphabet of player B (in our application: sets of transitions of the cost automaton) such that when player A chooses letters a_1, \dots, a_n in the first n rounds, then the letter chosen by player B in the n -th round is obtained by applying f to the state of \mathcal{B} after reading the word $a_1 \cdots a_n$.

To prove condition (b) in Proposition 2, consider a strategy defined by a deterministic automaton \mathcal{B} which is winning in the limitedness game with bound ∞ . We will show that when player B uses this strategy, the resulting runs have value at most m , where m is the number of states in the cost automaton times the number of states in the deterministic automaton

defining the strategy. Toward a contradiction, suppose that player A chooses letters a_1, a_2, \dots , player B uses the finite memory strategy to respond with sets of transitions $\delta_1, \delta_2, \dots$, and the winning condition for bound m is violated. This means that for some $n \geq m$, some finite run ρ in $\delta_1 \cdots \delta_n$ has value at least m . By the pigeon hole principle, there must exist i, j with $i < j$ such that: in the run of \mathcal{B} on $a_1 \cdots a_n$ the same state is used in positions i and j ; in the run ρ the same state is used in positions i and j ; and one of the counters is incremented but not reset between positions i and j . Therefore, if player A chooses the infinite word $a_1 \cdots a_{i-1} (a_i \cdots a_j)^\omega$, then the response of player B will contain a run where one of the counters is incremented infinitely often but reset finitely often, a contradiction.

Equivalence of (a) and (b): The implication (b) \Rightarrow (a) is immediate. To prove (a) \Rightarrow (b), we use ideas from [CL08]. Assume that the automaton is limited over L , i.e. there is some $m \in \mathbb{N}$ such that every word in L admits an accepting run with value at most m . For a natural number $m' \geq m$, define a *scoring function with bound m'* to be a function which maps every finite run to a number in $\mathbb{N} \cup \{\infty\}$, called its *score*, subject to the following conditions:

- every run with finite score has value $\leq m'$.
- every run with value $\leq m$ has finite score

A scoring function is called *monotone* if

$$\text{score}(\rho) \leq \text{score}(\rho') \Rightarrow \text{score}(\rho\sigma) \leq \text{score}(\rho'\sigma)$$

holds for every runs ρ, ρ' that end in the same state, and every run σ that begins in this state.

To prove condition (b), Lemma 3 shows that the existence of a monotone scoring function is a sufficient condition for player B winning the limitedness game with a finite bound; Lemma 4 shows that this sufficient condition can be met.

Lemma 3. *If there exists a monotone scoring function with bound m' , then player B has a winning strategy in the limitedness game with bound m' .*

Proof. (corresponds to Lemma 9 in [CL08]) Given a scoring function, define a finite run to be *optimal* if its score is finite and minimal among the scores of runs that read the same word and end in the same state. Consider the following strategy for player B: if player A has played a word w , then player B responds with the set of transitions t such that some optimal run over w uses t as the last transition. We show that this strategy is winning in the limitedness game with bound m' . Suppose that in a finite prefix of a play where player B uses this strategy, player A has chosen letters a_1, \dots, a_i and player B has replied with sets of transitions $\delta_1, \dots, \delta_i$.

Claim 1. *The set $\delta_1 \cdots \delta_i$ is equal to the set optimal runs over $a_1 \cdots a_i$.*

Proof. If a run ρ is optimal, then by monotonicity every prefix of ρ is optimal, and therefore every transition used by ρ is the last transition of some optimal run, thus proving membership $\rho \in \delta_1 \cdots \delta_i$. The converse inclusion is proved by induction on

i . Suppose that $\delta_1 \cdots \delta_i$ contains a run of the form ρt , where ρ has length $i - 1$ and t is the last transition used by the run. By induction assumption, ρ is optimal. By definition of δ_i there is some optimal run of the form $\rho' t$. By optimality, ρ has smaller or equal score than ρ' . By monotonicity ρt has smaller or equal score than the optimal run $\rho' t$, and therefore is itself optimal. \square

Using the claim, we verify that the above described strategy of player B is winning in the limitedness game with bound m' . Item 1) in the winning condition is clearly satisfied. Since optimal runs have finite score, and runs with finite score have value at most m' , item 2) is satisfied for bound m' . Since every word $a_1 \cdots a_i \in L$ admits an accepting run with value at most m , and runs of value m have finite score, then $a_1 \cdots a_i$ admits an accepting optimal run, and therefore item 3) is satisfied. \square

To complete the proof of Proposition 2, we show that there always exists a monotone scoring function.

Lemma 4. *There exists a monotone scoring function with some finite bound.*

Proof. (corresponds to Lemma 8 in [CL08]) When there is a single counter, a good scoring function is one which maps runs of value exceeding m to ∞ , and maps other runs to the number of increments after the last reset; in this special case $m = m'$. In particular, the material given so far is sufficient to decide limitedness of distance automata, which are the special case of cost automata with one counter and no resets.

A more fancy scoring function is needed when there are two or more counters. Suppose that the counters are numbered $0, \dots, n$. One natural idea would be to define the score of a run to be ∞ if its value exceeds m , and otherwise to be

$$\sum_{i=0}^n a_i \cdot (m+1)^i$$

assuming that a_i is the number of increments on counter i after its last reset. In other words, a counter valuation would be interpreted as a number in base $m+1$. This scoring function is unfortunately not monotone. Indeed, suppose that there are two counters, and consider two runs:

- ρ does one increment on counter 1 (score $m+1$);
- ρ' does m increments on counter 0 (score m).

Consider a transition t that increments counter 0. Appending t to these runs reverses the order on scores, because ρt has score $m+2$, while $\rho' t$ has score ∞ .

Monotonicity is recovered using the following modification. When an increment to counter i causes its value to reach $m+1$, then instead of setting the whole score to ∞ , the value of counter i is carried over to counter $i+1$, which may trigger further carries. The score becomes ∞ only when the most important counter n reaches $m+1$.

Formally, the score of a run ρ is defined as follows by induction on length. The score of an empty run is 0. Suppose that the score of ρ is already defined and t is a single transition.

If the score of ρ is ∞ , then also the score of ρt is ∞ . Otherwise, let

$$\sum_{i=0}^n a_i \cdot (m+1)^i$$

be the base $m+1$ representation of the score of ρ . If the transition t does not do any counter operations, then the score of ρt is the same as the score of ρ . If the transition t resets some counter $k \in \{0, \dots, n\}$, then the score of ρt is obtained by resetting digits on positions $0, \dots, k$, i.e. the score of ρt is

$$\sum_{i=k+1}^n a_i \cdot (m+1)^i$$

Finally, if the transition t increments some counter $k \in \{0, \dots, n\}$, then the score of ρt is obtained by resetting digits on positions $0, \dots, k-1$ and incrementing the digit on position k , i.e. the score of ρt is

$$(a_k + 1) \cdot (m+1)^k + \sum_{i=k+1}^n a_i \cdot (m+1)^i$$

If, as a result of the increment, the score reaches the maximal number, call it m' , which can be stored in base $m+1$ representation using digits a_0, \dots, a_n , then the score of ρt is defined to be ∞ . We show below that the function thus defined is a monotone scoring function with the bound being m' defined above.

- Monotonicity. We need to show

$$\text{score}(\rho) \leq \text{score}(\rho') \quad \Rightarrow \quad \text{score}(\rho t) \leq \text{score}(\rho' t).$$

The most interesting case is when the transition t increments some counter k . In this case, the score is obtained by first resetting digits $0, \dots, k-1$, which is monotone, and then adding $(m+1)^k$, which is also monotone.

- Runs with value at most m have finite score. By induction on the length of the run, one shows that if a run has value at most m , then for every $i \in \{0, \dots, n\}$, the i -th digit in the base $m+1$ representation of its score is equal to the number of increments on counter i after its last reset.
- Runs with finite score have value at most m' . Toward a contradiction, suppose that a run with finite score has value strictly bigger than m' . This means that for some counter i , the run does more than m' increments without any reset of counter i . During these increments, the score will carry over so that the value of counter n exceeds m , and therefore the score will become infinite.

\square

This completes the proof of equivalence of (a) and (b) in Proposition 2, and therefore also the proof of Theorem 1.

REFERENCES

- [AKY08] Parosh Aziz Abdulla, Pavel Krcal and Wang Yi. R-Automata. In *CONCUR*, pages 67–81, 2008.
- [BL69] J. Richard Büchi and Lawrence H. Landweber. Solving Sequential Conditions by Finite-State Strategies. *Transactions of the American Mathematical Society*, Volume 138, pages 295–311, 1969.

- [BC06] Mikołaj Bojańczyk and Thomas Colcombet. Bounds in ω -regularity. In *LICS*, pages 285–296, 2006.
- [Coh70] Rina S. Cohen. Star height of certain families of regular events. *Journal of Computer and System Sciences*, Volume 4, Issue 3, pages 281–297, 1970.
- [CL08] Thomas Colcombet and Christof Löding. The nesting depth of disjunctive μ -calculus for tree languages and the limitedness problem. In *CSL*, pages 416–430, 2008.
- [Col09] Thomas Colcombet. The theory of stabilisation monoids and regular cost functions. In *ICALP (2)*, pages 139–150, 2009.
- [Has88] Kosaburo Hashiguchi. Algorithms for determining relative star height and star height. *Inf. Comput.*, 78(2):124–169, 1988.
- [Kir05] Daniel Kirsten. Distance desert automata and the star height problem. *Informatique Théorique et Applications*, 39(3):455–509, 2005.
- [Kir06] Daniel Kirsten. Distance Desert Automata and Star Height Substitutions. Habilitation thesis, Universität Leipzig.
- [Sim88a] Imre Simon. Recognizable sets with multiplicities in the tropical semiring. In *Mathematical Foundations of Computer Science*, pages 107–120, 1988.
- [Sim88b] Imre Simon. Properties of factorization forests. *Formal Properties of Finite Automata and Applications*, pages 65–72, 1988.
- [Tho96] Wolfgang Thomas. Languages, automata and logic. Technical Report 9607, Institut für Informatik und Praktische Mathematik, Christian-Albrechts-Universität zu Kiel, 1996.
- [Tor11] Szymon Toruńczyk. Languages of profinite words and the limitedness problem. *PhD Thesis, University of Warsaw*, 2011.