

BÜCHI COMPLEMENTATION AND SIZE-CHANGE TERMINATION *

SETH FOGARTY ^a AND MOSHE Y. VARDI ^b

^{a,b} Department of Computer Science, Rice University, Houston, TX
e-mail address: sfogarty@gmail.com, vardi@cs.rice.edu

ABSTRACT. We compare tools for complementing nondeterministic Büchi automata with a recent termination-analysis algorithm. Complementation of Büchi automata is a key step in program verification. Early constructions using a Ramsey-based argument have been supplanted by rank-based constructions with exponentially better bounds. In 2001 Lee et al. presented the size-change termination (SCT) problem, along with both a reduction to Büchi automata and a Ramsey-based algorithm. The Ramsey-based algorithm was presented as a more practical alternative to the automata-theoretic approach, but strongly resembles the initial complementation constructions for Büchi automata.

We prove that the SCT algorithm is a specialized realization of the Ramsey-based complementation construction. To do so, we extend the Ramsey-based complementation construction to provide a containment-testing algorithm. Surprisingly, empirical analysis suggests that despite the massive gap in worst-case complexity, Ramsey-based approaches are superior over the domain of SCT problems. Upon further analysis we discover an interesting property of the problem space that both explains this result and provides a chance to improve rank-based tools. With these improvements, we show that theoretical gains in efficiency of the rank-based approach are mirrored in empirical performance.

1. INTRODUCTION

The automata-theoretic approach to formal program verification reduces questions about program adherence to a specification to questions about language containment. Representing liveness, fairness, or termination properties requires finite automata that operate on infinite words. One automaton, \mathcal{A} , encodes the behavior of the program, while another automaton, \mathcal{B} , encodes the formal specification. To ensure adherence, verify that the intersection of \mathcal{A} with the complement of \mathcal{B} is empty. Finite automata on infinite words are classified by their acceptance condition and transition structure. We consider here nondeterministic Büchi automata, in which a run is accepting when it visits at least one accepting

1998 ACM Subject Classification: D.2.4.

Key words and phrases: Büchi Complementation, Model Checking, Formal Verification, Automata, Büchi Automata.

* Earlier version appeared in TACAS09.

^a Work supported in part by NSF grants CCR-0124077, CCR-0311326, CCF-0613889, ANI-0216467, and CCF-0728882, by BSF grant 9800096, and by a gift from Intel.

state infinitely often. For these automata, the complementation problem is known to involve an exponential blowup [25]. Thus the most difficult step in checking containment is constructing the complementary automata \overline{B} .

The first complementation constructions for nondeterministic Büchi automata employed a Ramsey-based combinatorial argument to partition the set of all infinite words into a finite set of omega-regular languages. Proposed by Büchi in 1962 [5], this construction was shown in 1987 by Sistla, Vardi, and Wolper to be implementable with a blow-up of $2^{O(n^2)}$ [29]. This brought the complementation problem into singly-exponential blow-up, but left a gap with the $2^{\Omega(n \log n)}$ lower bound proved by Michel [25].

The gap was tightened in 1988, when Safra described a $2^{O(n \log n)}$ construction [26]. Work since then has focused on improving the practicality of $2^{O(n \log n)}$ constructions, either by providing simpler constructions, further tightening the bound [27], or improving the derived algorithms. In 2001, Kupferman and Vardi employed a rank-based analysis of Büchi automata to simplify complementation [23]. Recently, Doyen and Raskin have demonstrated the utility of using a subsumption technique in the rank-based approach, providing a direct universality checker that scales to automata several orders of magnitude larger than previous tools [10].

Separately, in the context of program termination analysis, Lee, Jones, and Ben-Amram presented the size-change termination (SCT) principle in 2001 [24]. This principle states that, for domains with well-founded values, if every infinite computation contains an infinitely decreasing value sequence, then no infinite computation is possible. Lee et al. describe a method of size-change termination analysis and reduce this problem to the containment of two Büchi automata. Stating the lack of efficient Büchi containment solvers, they also propose a Ramsey-based combinatorial solution that captures all possible call sequences in a finite set of graphs. The Lee, Jones, and Ben-Amram (LJB) algorithm was provided as a practical alternative to reducing the verification problem to Büchi containment, but bears a striking resemblance to the 1987 Ramsey-based complementation construction [29].

In this paper we show that the LJB algorithm for deciding SCT [24] is a specialized implementation of the 1987 Ramsey-based complementation construction [29]. Section 2 presents the background and notation for the paper. Section 3 expands the Ramsey-based complementation construction into a containment algorithm, and then presents the proof that the LJB algorithm is a specialized realization of this Ramsey-based containment algorithm. In Section 4, we empirically explore Lee et al.'s intuition that Ramsey-based algorithms are more practical than Büchi complementation tools on SCT problems. Initial experimentation does suggest that Ramsey-based tools are superior on SCT problems. This is surprising, as the worst-case complexity of the LJB algorithm is significantly worse than that of rank-based tools. Investigating this discovery in Section 5, we note that it is natural for SCT problems to be reverse-deterministic, and that for reverse-deterministic problems the worst-case bound for Ramsey-based algorithms matches that of the rank-based approach. This suggests improving the rank-based approach in the face of reverse determinism. Indeed, we find that reverse-deterministic SCT problems have a maximum rank of 2, collapsing the complexity of rank-based complementation to $2^{O(n)}$. Revisiting our experiments, we discover that with this improvement rank-based tools are superior on the domain of SCT problems. To further explore the phenomena, we generate a set of non-reverse-deterministic SCT problems from monotonicity constraint systems, a more complex termination problem. We conclude with a discussion in Section 6.




Figure 1: An automaton which accepts words with finitely many a 's.

2. PRELIMINARIES

In this section we review the relevant details of the Büchi complementation and size-change termination, introducing along the way the notation used throughout this paper. A *nondeterministic Büchi automaton on infinite words* is a tuple $\mathcal{B} = \langle \Sigma, Q, Q^{in}, \rho, F \rangle$, where Σ is a finite nonempty alphabet, Q a finite nonempty set of states, $Q^{in} \subseteq Q$ a set of initial states, $F \subseteq Q$ a set of accepting states, and $\rho : Q \times \Sigma \rightarrow 2^Q$ a nondeterministic transition function. We lift the ρ function to sets of states and words of arbitrary length as follows. Given a set of states R , define $\rho(R, \sigma)$ to be $\bigcup_{q \in R} \rho(q, \sigma)$. Inductively, given a set of states R : let $\rho(R, \epsilon) = R$, and for every word $w = \sigma_0 \dots \sigma_n$ let $\rho(R, w)$ be defined as $\rho(\rho(R, \sigma_0), \sigma_1 \dots \sigma_n)$.

A *run* of a Büchi automaton \mathcal{B} on a word $w \in \Sigma^\omega$ is a infinite sequence $r = q_0 q_1 \dots \in Q^\omega$ such that $q_0 \in Q^{in}$ and, for every $i \geq 0$, we have $q_{i+1} \in \rho(q_i, w_i)$. A run is *accepting* iff $q_i \in F$ for infinitely many $i \in \mathbb{N}$. A word $w \in \Sigma^\omega$ is accepted by \mathcal{B} if there is an accepting run of \mathcal{B} on w . The words accepted by \mathcal{B} form the language of \mathcal{B} , denoted by $L(\mathcal{B})$. Correspondingly, a *path* in \mathcal{B} from q to r on a word $w \in \Sigma^+$ is a finite sequence $r = q_0 \dots q_n \in Q^+$ such that $q_0 = q$, $q_n = r$, and, for every $i \in \{0 \dots n - 1\}$, we have $q_{i+1} \in \rho(q_i, w_i)$. A path is *accepting* if some state in the path is in F .

Example 2.1. An example automaton is shown in Figure 1, which accepts words with a finite, but non-zero, number of a 's. The automaton waits in q and guesses when it has seen the last a , transitioning on that a to r . If the automaton moves to r prematurely, it can transition to s before it encounters any remaining a 's to continue the run. From s , it can guess once again when it has seen the last a , transitioning this time to t .

A Büchi automaton \mathcal{A} is contained in a Büchi automaton \mathcal{B} iff $L(\mathcal{A}) \subseteq L(\mathcal{B})$, which can be checked by verifying that the intersection of \mathcal{A} with the complement $\bar{\mathcal{B}}$ of \mathcal{B} is empty: $L(\mathcal{A}) \cap L(\bar{\mathcal{B}}) = \emptyset$. We know that the language of an automaton is non-empty iff there are states $q \in Q^{in}$, $r \in F$ such that there is a path from q to r and an accepting path from r to itself. The initial path is called the *prefix*, and the combination of the prefix and cycle is called a *lasso* [31]. Furthermore, the intersection of two automata can be constructed, having a number of states proportional to the product of the number states of the original automata [6]. Thus, the most computationally demanding step is constructing the complement of \mathcal{B} . In the formal verification field, existing empirical work has focused on the simplest form of containment testing, *universality* testing, where \mathcal{A} is the universal automaton [9, 30].




Figure 2: Three graphs in $\tilde{Q}_{\mathcal{B}}$ for the automaton of Figure 1. From left to right, the graph describing the word a , the graph describing the word b , and the graph describing the word ab .

2.1. Ramsey-Based Universality. When Büchi introduced these automata in 1962, he described a complementation construction involving a Ramsey-based combinatorial argument [5]. We describe an optimized implementation presented in 1987 [29]. To construct the complement of $\mathcal{B} = \langle \Sigma, Q, Q^{in}, \rho, F \rangle$, where $Q = \{q_0, \dots, q_{n-1}\}$, we construct a set $\tilde{Q}_{\mathcal{B}}$ whose elements capture the essential behavior of \mathcal{B} . Each element corresponds to an answer to the following question:

Given a finite nonempty word w , for every two states $q, r \in Q$:

- (1) Is there a path in \mathcal{B} from q to r over w ?
- (2) If so, is some such path accepting?

Define $Q' = Q \times \{0, 1\} \times Q$, and $\tilde{Q}_{\mathcal{B}}$ to be the subset of $2^{Q'}$ whose elements, for every $q, r \in Q$, do not contain both $\langle q, 0, r \rangle$ and $\langle q, 1, r \rangle$. Each element of $\tilde{Q}_{\mathcal{B}}$ is a $\{0, 1\}$ -arc-labeled graph on Q . An arc represents a path in \mathcal{B} , and the label is 1 if the path is accepting. Note that there are 3^{n^2} such graphs. With each graph $\tilde{g} \in \tilde{Q}_{\mathcal{B}}$ we associate a language $L(\tilde{g})$, the set of words for which the answer to the posed question is the graph encoded by \tilde{g} .

Definition 2.2. Let $\tilde{g} \in \tilde{Q}_{\mathcal{B}}$ and $w \in \Sigma^+$. Then $w \in L(\tilde{g})$ iff, for all pairs of states $q, r \in Q$:

- (1) $\langle q, a, r \rangle \in \tilde{g}$, $a \in \{0, 1\}$, iff there is a path in \mathcal{B} from q to r over w .
- (2) $\langle q, 1, r \rangle \in \tilde{g}$ iff there is an accepting path in \mathcal{B} from q to r over w .

Example 2.3. Three graphs from $\tilde{Q}_{\mathcal{B}}$ are shown in Figure 1. All graphs have a non-empty language. The word a is in the language of the first graph, the word b is in the language of the second graph, and the word ab is in the language of the third graph.

Lemma 2.4. [5, 29]

- (1) $\{L(\tilde{g}) \mid \tilde{g} \in \tilde{Q}_{\mathcal{B}}\}$ is a partition of Σ^+
- (2) If $u \in L(\tilde{g})$, $v \in L(\tilde{h})$, and $uv \in L(\tilde{k})$, then $L(\tilde{g}) \cdot L(\tilde{h}) \subseteq L(\tilde{k})$

The languages $L(\tilde{g})$, for the graphs $\tilde{g} \in \tilde{Q}_{\mathcal{B}}$, form a partition of Σ^+ . With this partition of Σ^+ we can devise a finite family of ω -languages that cover Σ^ω . For every $\tilde{g}, \tilde{h} \in \tilde{Q}_{\mathcal{B}}$, let $Y(\tilde{g}, \tilde{h})$ be the ω -language $L(\tilde{g}) \cdot L(\tilde{h})^\omega$. Say that a language $Y(\tilde{g}, \tilde{h})$ is *proper* if $Y(\tilde{g}, \tilde{h})$ is non-empty, $L(\tilde{g}) \cdot L(\tilde{h}) \subseteq L(\tilde{g})$, and $L(\tilde{h}) \cdot L(\tilde{h}) \subseteq L(\tilde{h})$. There are a finite, if exponential, number of such languages. A Ramsey-based argument shows that every infinite string

belongs to a language of this form, and that $\overline{L(\mathcal{B})}$ can be expressed as the union of languages of this form.

Lemma 2.5. [5, 29] $\Sigma^\omega = \bigcup \{Y(\tilde{g}, \tilde{h}) \mid Y(\tilde{g}, \tilde{h}) \text{ is proper}\}$

Proof: The proof is based on Ramsey's Theorem. Consider an infinite word $w = \sigma_0\sigma_1\dots$. By Lemma 2.4, every prefix of the word w is in the language of a unique graph \tilde{g}_i . Let $k = 3^{n^2}$ be the number of graphs. Thus w defines a partition of \mathbb{N} into k sets D_1, \dots, D_k such that $i \in D_l$ iff $\sigma_0\dots\sigma_{i-1} \in L(\tilde{g}_l)$. Clearly there is some m such that D_m is infinite.

Similarly, by Lemma 2.4 we can use the word w to define a partition of all *pairs* of elements (i, j) from D_m , where $i < j$. This partition consists of k sets C_1, \dots, C_k , such that $\langle i, j \rangle \in C_l$ iff $\sigma_i\dots\sigma_{j-1} \in L(\tilde{g}_l)$. Ramsey's Theorem tells us that, given such a partition, there exists an infinite subset $\{i_1, i_2, \dots\}$ of D_m and a C_n such that for all pairs of distinct elements i_j, i_k , it holds that $\langle i_j, i_k \rangle \in C_n$.

This implies that the word w can be partitioned into

$$w_1 = \sigma_0\dots\sigma_{i_1-1}, \quad w_2 = \sigma_{i_1}\dots\sigma_{i_2-1}, \quad w_3 = \sigma_{i_2}\dots\sigma_{i_3-1}, \quad \dots,$$

where $w_1 \in L(\tilde{g}_m)$ and $w_i \in L(\tilde{g}_n)$ for $i > 1$. By construction, $\sigma_0\dots\sigma_{i_j-1} \in L(\tilde{g}_m)$ for every i_j , and thus we have that $w_1w_2 \in L(\tilde{g}_m)$. In addition, as $\sigma_{i_j}\dots\sigma_{i_k-1} \in L(\tilde{g}_n)$ for every pair i_j, i_k , we have that $w_2w_3 \in L(\tilde{g}_n)$. By Lemma 2.4, it follows that $L(\tilde{g}_m) \cdot L(\tilde{g}_n) \subseteq L(\tilde{g}_m)$, and that $L(\tilde{g}_n) \cdot L(\tilde{g}_n) \subseteq L(\tilde{g}_n)$, and thus $Y(\tilde{g}_m, \tilde{g}_n)$ is proper. \square

Furthermore, each proper language is entirely contained or entirely disjoint from $L(\mathcal{B})$. This provides a way to construct the complement of $L(\mathcal{B})$: take the union every proper language that is disjoint from $L(\mathcal{B})$.

Lemma 2.6. [5, 29]

- (1) For $\tilde{g}, \tilde{h} \in \tilde{Q}_{\mathcal{B}}$, either $Y(\tilde{g}, \tilde{h}) \cap L(\mathcal{B}) = \emptyset$ or $Y(\tilde{g}, \tilde{h}) \subseteq L(\mathcal{B})$.
- (2) $\overline{L(\mathcal{B})} = \bigcup \{Y(\tilde{g}, \tilde{h}) \mid Y(\tilde{g}, \tilde{h}) \text{ is proper and } Y(\tilde{g}, \tilde{h}) \cap L(\mathcal{B}) = \emptyset\}$

To obtain the complementary Büchi automaton $\overline{\mathcal{B}}$, Sistla et al. construct, for each $\tilde{g} \in \tilde{Q}_{\mathcal{B}}$, a deterministic automata on finite words, \mathcal{B}_g , that accepts exactly $L(\tilde{g})$. Using the automata \mathcal{B}_g , one can then construct the complementary automaton $\overline{\mathcal{B}}$ [29]. We can then use a lasso-finding algorithm on $\overline{\mathcal{B}}$ to prove the emptiness of $\overline{\mathcal{B}}$, and thus the universality of \mathcal{B} . However, we can avoid an explicit lasso search by employing the rich structure of the graphs in $\tilde{Q}_{\mathcal{B}}$. For every two graphs $\tilde{g}, \tilde{h} \in \tilde{Q}_{\mathcal{B}}$, determine if $Y(\tilde{g}, \tilde{h})$ is proper. If $Y(\tilde{g}, \tilde{h})$ is proper, test if it is contained in $L(\mathcal{B})$ by looking for a lasso with a prefix in \tilde{g} and a cycle in \tilde{h} . \mathcal{B} is universal if every proper $Y(\tilde{g}, \tilde{h})$ is so contained.

Lemma 2.7. [29] *Given an Büchi automaton \mathcal{B} and the set of graphs $\tilde{Q}_{\mathcal{B}}$,*

- (1) \mathcal{B} is universal iff for every proper $Y(\tilde{g}, \tilde{h})$, it holds that $Y(\tilde{g}, \tilde{h}) \subseteq L(\mathcal{B})$.
- (2) Let $\tilde{g}, \tilde{h} \in \tilde{Q}_{\mathcal{B}}$ be two graphs where $Y(\tilde{g}, \tilde{h})$ is proper. $Y(\tilde{g}, \tilde{h}) \subseteq L(\mathcal{B})$ iff there exists $q \in Q^{in}$, $r \in Q$, $a \in \{0, 1\}$ where $\langle q, a, r \rangle \in \tilde{g}$ and $\langle r, 1, r \rangle \in \tilde{h}$.

Lemma 2.7 yields a PSPACE algorithm to determine universality [29]. Simply check each $\tilde{g}, \tilde{h} \in \tilde{Q}_{\mathcal{B}}$. If $Y(\tilde{g}, \tilde{h})$ is both proper and not contained in $L(\mathcal{B})$, then the pair (\tilde{g}, \tilde{h}) provide a counterexample to the universality of \mathcal{B} . If no such pair exists, the automaton must be universal.

2.2. Rank-Based Complementation. While our focus is mainly on the Ramsey-based approach, in Section 5 we look at the rank-based construction described here. If a Büchi automaton \mathcal{B} does not accept a word w , then every run of \mathcal{B} on w must eventually cease visiting accepting states. The rank-based construction, foreshadowed in [20] and first introduced in [21], uses a notion of ranks to track the progress of each possible run towards this point. Consider a Büchi automaton $\mathcal{B} = \langle \Sigma, Q, Q_\infty, \rho, Q_f \rangle$ and an infinite word $w = \omega_0\omega_1\dots$. The runs of \mathcal{B} on w can be arranged in an infinite DAG (directed acyclic graph), $G_w = \langle V, E \rangle$, where

- $V \subseteq Q \times \mathbb{N}$ is such that $\langle q, l \rangle \in V$ iff some run r of \mathcal{B} on w has $r(l) = q$.
- $E \subseteq \bigcup_{l \geq 0} (Q \times \{l\}) \times (Q \times \{l+1\})$ is $E(\langle q, l \rangle, \langle q', l+1 \rangle)$ iff $\langle q, l \rangle \in V$ and $q' \in \rho(q, \omega_l)$.

G_w , called the *run DAG* of \mathcal{B} on w , exactly embodies all possible runs of \mathcal{B} on w . We define a run DAG G_w to be *accepting* when there exists a path in G_w with infinitely many states in F . This path corresponds to an accepting run of \mathcal{B} on w . When G_w is not accepting, we say it is a *rejecting* run DAG. Say that a node $\langle q, i \rangle$ of a graph is *finite* if it has only finitely many descendants, that it is *accepting* if $q \in Q_f$, and that it is *F-free* if it is not accepting and does not have accepting descendants.

Given a run DAG G_w , we inductively define a sequence of subgraphs by eliminating nodes that cannot be part of accepting runs. A node that is finite can clearly not be part of an infinite run, much less an accepting infinite run. Similarly, a node that is *F-free* may be part of an infinite run, but this infinite run can not visit an infinite number accepting states.

- $G_w(0) = G_w$
- $G_w(2i+1) = G_w(2i) \setminus \{\langle q, l \rangle \mid \langle q, l \rangle \text{ is finite in } G_w(2i)\}$
- $G_w(2i+2) = G_w(2i+1) \setminus \{\langle q, l \rangle \mid \langle q, l \rangle \text{ is } F\text{-free in } G_w(2i)\}$

If the final graph a node appears in is $G_w(i)$, we say that node is of rank i . If the rank of every node is at most i , we say G_w has a rank of i . \mathcal{B} has a maximum rank of i when every rejecting run DAG of \mathcal{B} has a maximum rank of i . Note that nodes with odd ranks are removed because they are *F-free*. Therefore no accepting state can have an odd rank. Kupferman and Vardi prove that the maximum rank of a rejecting run DAG for every automaton is bounded by $2|Q| - 2$ [21]. This allows us to create an automaton that guesses the ranking of rejecting run DAG of \mathcal{B} as it proceeds along the word.

A *level ranking* for an automaton \mathcal{B} with n states is a function $f : Q \rightarrow \{0 \dots 2n - 2, \perp\}$, such that if $q \in F$ then $f(q)$ is even or \perp . Let a be a letter in Σ and f, f' be two level rankings f . Say that f *covers* f' under a when for all q and every $q' \in \rho(q, a)$, if $f(q) \neq \perp$ then $f'(q') \neq \perp$ and $f'(q') \leq f(q)$; i.e. no transition between f and f' on a increases in rank. Let F_r be the set of all level rankings.

Definition 2.8. If $\mathcal{B} = \langle \Sigma, Q, Q^{in}, \rho, F \rangle$ is a Büchi automaton, define $KV(\mathcal{B})$ to be the automaton $\langle \Sigma, F_r \times 2^Q, \langle f_{in}, \emptyset \rangle, \rho', F_r \times \{\emptyset\} \rangle$, where

- $f_{in}(q) = 2n - 2$ for each $q \in Q^{in}$, \perp otherwise.
- Define $\rho' : \langle F_r \times 2^Q \rangle \times \sigma \rightarrow 2^{(F_r \times 2^Q)}$ to be
 - If $o \neq \emptyset$ then $\rho'(\langle f, o \rangle, \sigma) = \{\langle f', o' \setminus d \rangle \mid f \text{ covers } f' \text{ under } \sigma, o' = \rho(o, \sigma), d = \{q \mid f'(q) \text{ odd}\}\}$.
 - If $o = \emptyset$ then $\rho'(\langle f, o \rangle, \sigma) = \{\langle f', o' \rangle \mid f \text{ covers } f' \text{ under } a, o' = \{q \mid f'(q) \text{ even}\}\}$.

Lemma 2.9. [22] *For every Büchi automaton \mathcal{B} , $L(KV(\mathcal{B})) = \overline{L(\mathcal{B})}$.*

This automaton tracks the progress of \mathcal{B} along a word $w = \sigma_0\sigma_1\dots$ by attempting to find an infinite series $f_0f_1\dots$ of level rankings. We start with the most general possible level ranking, and ensure that every rank f_i covers f_{i+1} under σ_i . Every run has a non-increasing rank, and so must eventually become trapped in some rank. To accept a word, the automaton requires that each run visit an odd rank infinitely often. Recall that accepting states cannot be assigned an odd rank. Thus, for a word rejected by \mathcal{B} , every run can eventually become trapped in an odd rank. Conversely, if there is an accepting run that visits an accepting node infinitely often, that run cannot visit an odd rank infinitely often and the complementary automaton rejects it.

An algorithm seeking to refute the universality of \mathcal{B} can look for a lasso in the state-space of the rank-based complement of \mathcal{B} . A classical approach is Emerson-Lei backward-traversal nested fixpoint $\nu Y.\mu X.(Pre(X) \cup (Pre(Y) \cap F))$ [11]. This nested fixpoint employs the observation that a state in a lasso can reach an arbitrary number of accepting states. The outer fixpoint iteratively computes sets Y_0, Y_1, \dots such that Y_i contains all states with a path visiting i accepting states. Universality is checked by testing if Y_∞ , the set of all states with a path visiting arbitrarily many accepting states, intersects Q^{in} . The strongest algorithm implementing this approach, from Doyen and Raskin, takes advantage of the presence of a subsumption relation in the rank-based construction: one state $\langle f, o \rangle$ subsumes another $\langle f', o' \rangle$ iff: $f'(x) \leq f(x)$ for every $x \in Q$; $o' \subseteq o$; and $o = \emptyset$ iff $o' = \emptyset$. When computing sets in the Emerson-Lei approach, it is sufficient to store only the maximal elements under this relation. Furthermore, the predecessor operation for a single state and letter results in at most two incomparable elements. This algorithm has scaled to automata an order of magnitude larger than other approaches [9].

2.3. Size-Change Termination. In [24] Lee et al. proposed the size-change termination (SCT) principle for programs: “If every infinite computation would give rise to an infinitely decreasing value sequence, then no infinite computation is possible.” The original presentation concerned a first-order pure functional language, where every infinite computation arises from an infinite call sequence and values are always passed through a sequence of parameters.

Proving that a program is size-change terminating is done in two phases. The first extracts from a program a set of size-change graphs, \mathcal{G} , containing guarantees about the relative size of values at each function call site. The second phase, and the phase we focus on, analyzes these graphs to determine if every infinite call sequence has a value that descends infinitely along a well-ordered set. For an excellent discussion of the abstraction of functional language semantics, refer to [19]. We consider here a set H of functions, and denote the parameters of a function f by $P(f)$.

Definition 2.10. A *size-change graph* (SCG) from function f_1 to function f_2 , written $G : f_1 \rightarrow f_2$, is a bipartite $\{0, 1\}$ -arc-labeled graph from the parameters of f_1 to the parameters of f_2 , where $G \subseteq P(f_1) \times \{0, 1\} \times P(f_2)$ does not contain both $x \xrightarrow{1} y$ and $x \xrightarrow{0} y$.

Size-change graphs capture information about a function call. An arc $x \xrightarrow{1} y$ indicates that the value of x in the function f_1 is strictly greater than the value passed as y to function f_2 . An arc $x \xrightarrow{0} y$ indicates that x 's value is greater than or equal to the value given to




Figure 3: Size-Change Graphs: A size-change problem with two functions, f and g , and three call sites: a call a to g occurring in the body of f , and two recursive calls, b and c , from g to itself.




Figure 4: The dotted line forms a prefix of a late-start thread in the call sequence $abcbbc\cdots$.

y . We assume that all call sites in a program are reachable from the entry points of the program¹.

A *size-change termination* (SCT) problem is a tuple $L = \langle H, P, C, \mathcal{G} \rangle$, where H is a set of functions, P a mapping from each function to its parameters, C a set of call sites between these functions, and \mathcal{G} a set of SCGs for C . A call site is written $c : f_1 \rightarrow f_2$ for a call to function f_2 occurring in the body of f_1 . The size-change graph for a call site $c : f_1 \rightarrow f_2$ is written as G_c . Given a SCT problem L , a *call sequence* in L is a infinite sequence $cs = c_0, c_1, \dots \in C^\omega$, such that there exists a sequence of functions f_0, f_1, \dots where $c_0 : f_0 \rightarrow f_1, c_1 : f_1 \rightarrow f_2, \dots$. A *thread* in a call sequence c_0, c_1, \dots is a connected sequence of arcs, $x \xrightarrow{a} y, y \xrightarrow{b} z, \dots$, beginning in some call c_i such that $x \xrightarrow{a} y \in G_{c_i}, y \xrightarrow{b} z \in G_{c_{i+1}}, \dots$. We say that L is *size-change terminating* if every call sequence contains a thread with infinitely many 1-labeled arcs. Note that a thread need not begin at the start of a call sequence. A sequence must terminate if a well-founded value decreases infinitely often, regardless of when this decrease begins. Therefore threads can begin in arbitrary function calls, in arbitrary parameters. We call this the *late-start property* of SCT problems. We revisit this property in Section 3.3.

Example 2.11. Three size-change graphs, which will provide a running example for this paper, are presented in Figure 3. The represented problem is size-change terminating. The call sequence $abcbc\cdots$ is displayed in Figure 4, where a thread of infinite descent exists, starting in the second graph. This late-start thread proves the sequence terminating.

Every call sequence can be represented as a word in C^ω , and a SCT problem reduced to the containment of two ω -languages. The first language $Flow(L) = \{cs \in C^\omega \mid cs \text{ is a call sequence}\}$, contains all call sequences. The second language, $Desc(L) = \{cs \in Flow(L) \mid \text{some thread in } cs \text{ has infinitely many 1-labeled arcs}\}$, contains only call

¹The implementation provided by Lee et al. [24] also make this assumption, and in the presence of unreachable functions size-change termination may be undetectable.




Figure 5: $\mathcal{A}_{Flow(L)}$ and $\mathcal{A}_{Desc(L)}$: the automata resulting from applying Definition 2.12 to the SCT problem of Figure 3.

sequences that guarantee termination. A SCT problem L is size-change terminating if and only if $Flow(L) \subseteq Desc(L)$.

Lee et al. [24] describe two Büchi automata, $\mathcal{A}_{Flow(L)}$ and $\mathcal{A}_{Desc(L)}$, that accept these languages. $\mathcal{A}_{Flow(L)}$ is simply the call graph of the program. $\mathcal{A}_{Desc(L)}$ waits in a copy of the call graph and nondeterministically chooses the beginning point of a descending thread. From there it ensures that a 1-labeled arc is taken infinitely often. To do so, it keeps two copies of each parameter, and transitions to the accepting copy only on a 1-labeled arc. Lee et al. prove that $L(\mathcal{A}_{Flow(L)}) = Flow(L)$, and $L(\mathcal{A}_{Desc(L)}) = Desc(L)$. The automata for our running example are provided in Figure 5.

Definition 2.12.²

$\mathcal{A}_{Flow(L)} = \langle C, H, H, \rho_F, H \rangle$, where

- $\rho_F(f_1, c) = \{f_2 \mid c : f_1 \rightarrow f_2\}$

$\mathcal{A}_{Desc(L)} = \langle C, Q_p \cup H, Q_p \cup H, \rho_D, F \rangle$, where

- $Q_p = \{\langle x, r \rangle \mid f \in H, x \in P(f), r \in \{1, 0\}\}$,
- $\rho_D(f_1, c) = \{f_2 \mid c : f_1 \rightarrow f_2\} \cup \{\langle x, 0 \rangle \mid c : f_1 \rightarrow f_2, x \in P(f_2)\}$
- $\rho_D(\langle x, r \rangle, c) = \{\langle x', r' \rangle \mid x \xrightarrow{r'} x' \in \mathcal{G}_c\}$,
- $F = \{\langle x, 1 \rangle \mid f \in H, x \in P(f)\}$

Using the complementation constructions of either Section 2.1 or 2.2 and a lasso-finding algorithm, we can determine the containment of $\mathcal{A}_{Flow(L)}$ in $\mathcal{A}_{Desc(L)}$. Lee et al. propose an alternative graph-theoretic algorithm, employing SCGs to encode descent information about entire call sequences. A notion of composition is used, where a call sequence $c_0 \dots c_{n-1}$ has a thread from x to y if and only if the composition of the SCGs for each call, $G_{c_0}; \dots; G_{c_{n-1}}$, contains the arc $x \xrightarrow{a} y$. The closure of \mathcal{G} under the composition operation, called S , is

²The original LJB construction [24] restricted starting states in $\mathcal{A}_{Desc(L)}$ to functions. This was changed to simplify Section 3.4. The modification does not change the accepted language.




Figure 6: The composition of the SCGs for b and c , from Figure 3. The resulting size-change graph is idempotent, contains the arc $y \xrightarrow{1} y$, and describes the call sequence of Figure 4.

then searched for a counterexample describing an infinite call sequence with no infinitely descending thread.

Definition 2.13. Let $G : f_1 \rightarrow f_2$ and $G' : f_2 \rightarrow f_3$ be two SCGs. Their composition $G;G'$ is defined as $G'' : f_1 \rightarrow f_3$ where:

$$\begin{aligned}
 G'' &= \{x \xrightarrow{1} z \mid x \xrightarrow{a} y \in G, y \xrightarrow{b} z \in G', y \in P(f_2), a = 1 \text{ or } b = 1\} \\
 &\cup \{x \xrightarrow{0} z \mid x \xrightarrow{0} y \in G, y \xrightarrow{0} z \in G', y \in P(f_2), \text{ and} \\
 &\quad \text{for all } y', a, b \text{ if } x \xrightarrow{a} y' \in G \text{ and } y' \xrightarrow{b} z \in G' \text{ then } a = b = 0\}
 \end{aligned}$$

Using composition, we can focus on a subset of graphs. Say that a graph $G : f \rightarrow f$ is *idempotent* when $G = G;G$. Each idempotent graph describes a cycle in the call graph, and a Ramsey-based argument shows that each cycle in the call graph can be accounted for by at least one idempotent graph. The composition of two graphs is shown in Figure 6, which describes the call sequence in Figure 4.

Algorithm LJB searches for a counterexample to size-change termination. First, it iteratively build the closure set S : initialize S as \mathcal{G} ; and for every $G : f_1 \rightarrow f_2$ and $G' : f_2 \rightarrow f_3$ in S , include the composition $G;G'$ in S . Second, the algorithm check every $G : f_1 \rightarrow f_1 \in S$ to ensure that if G is idempotent, then G has an associated thread with infinitely many 1-labeled arcs. This thread is represented by an arc of the form $x \xrightarrow{1} x$. There are pathological SCT problems for which the complexity of Algorithm LJB is $2^{O((n/2)^2)}$.

The next theorem, whose proof uses a Ramsey-based argument, demonstrates the correctness of Algorithm LJB in determining the size-change termination of an SCT problem $L = \langle H, P, C, \mathcal{G} \rangle$.

Theorem 2.1. [24] A SCT problem $L = \langle H, P, C, \mathcal{G} \rangle$ is *not* size-change terminating iff S , the closure of \mathcal{G} under composition, contains an idempotent SCG graph $G : f \rightarrow f$ that does *not* contain an arc of the form $x \xrightarrow{1} x$.

3. SIZE-CHANGE TERMINATION AND RAMSEY-BASED CONTAINMENT

The Ramsey-based test of Section 2.1 and the LJB algorithm of Section 2.3 bear a remarkable similarity. In this section we bridge the gap between the Ramsey-based universality test and the LJB algorithm, by demonstrating that the LJB algorithm is a specialized realization of the Ramsey-based containment test. This first requires developing a Ramsey-based framework for Büchi containment testing.

Algorithm 1: LJB($\langle H, P, C, \mathcal{G} \rangle$)

Data: A size-change termination problem $\langle H, P, C, \mathcal{G} \rangle$.
Result: Whether or not the problem is size-change terminating.
Initialize $\mathbf{S} \leftarrow \mathcal{G}$
repeat
 for all pairs $G : f \rightarrow g, G' : g \rightarrow h$ **in** \mathbf{S} **do**
 $G'' : f \rightarrow h \leftarrow G; G'$
 Add G'' to \mathbf{S}
 if $f = h$ **and** $G''; G'' = G''$ **then**
 if there does not exist an arc of the form $x \xrightarrow{1} x$ **in** G'' **then**
 return *Not Terminating*
until \mathbf{S} reaches closure
return *Terminating*

3.1. Ramsey-Based Containment with Supergraphs. To test the containment of a Büchi automaton \mathcal{A} in a Büchi automaton \mathcal{B} , we could construct the complement of \mathcal{B} using either the Ramsey-based or rank-based construction, compute the intersection automaton of \mathcal{A} and $\overline{\mathcal{B}}$, and search this intersection automaton for a lasso. With universality, however, we avoided directly constructing $\overline{\mathcal{B}}$ by exploiting the structure of states in the Ramsey-based construction (see Lemma 2.7). We demonstrate a similar test for containment.

Consider two automata, $\mathcal{A} = \langle \Sigma, Q_{\mathcal{A}}, Q_{\mathcal{A}}^{in}, \rho_{\mathcal{A}}, F_{\mathcal{A}} \rangle$ and $\mathcal{B} = \langle \Sigma, Q_{\mathcal{B}}, Q_{\mathcal{B}}^{in}, \rho_{\mathcal{B}}, F_{\mathcal{B}} \rangle$. When testing the universality of \mathcal{B} , any word not in $L(\mathcal{B})$ is a sufficient counterexample. To test $L(\mathcal{A}) \subseteq L(\mathcal{B})$ we must restrict our search to the subset of Σ^ω accepted by \mathcal{A} . In Section 2.1, we defined a set $\tilde{Q}_{\mathcal{B}}$ of 0-1 arc-labeled graphs, whose elements provide a family of ω -languages that covers Σ^ω (see Lemma 2.5). We now define a set, $\hat{Q}_{\mathcal{A}, \mathcal{B}}$, which provides a family of ω -languages covering $L(\mathcal{A})$.

We first define $\bar{Q}_{\mathcal{A}} = Q_{\mathcal{A}} \times Q_{\mathcal{A}}$ to capture the connectivity in $Q_{\mathcal{A}}$. An element $\bar{g} = \langle q, r \rangle \in \bar{Q}_{\mathcal{A}}$ is a single arc asserting the existence of a path in \mathcal{A} from q to r . With each arc we associate a language, $L(\bar{g})$.

Definition 3.1. Given $w \in \Sigma^+$, say that $w \in L(\langle q, r \rangle)$ iff there is a path in \mathcal{A} from q to r over w .

Define $\hat{Q}_{\mathcal{A}, \mathcal{B}}$ as $\bar{Q}_{\mathcal{A}} \times \tilde{Q}_{\mathcal{B}}$. The elements of $\hat{Q}_{\mathcal{A}, \mathcal{B}}$, called *supergraphs*, are pairs consisting of an arc from $\bar{Q}_{\mathcal{A}}$ and a graph from $\tilde{Q}_{\mathcal{B}}$. Each element simultaneously captures all paths in \mathcal{B} and a single path in \mathcal{A} . The language $L(\langle \bar{g}, \tilde{g} \rangle)$ is then $L(\bar{g}) \cap L(\tilde{g})$. For convenience, we implicitly take $\hat{g} = \langle \bar{g}, \tilde{g} \rangle$, and say $\langle q, a, r \rangle \in \hat{g}$ when $\langle q, a, r \rangle \in \tilde{g}$. Since the language of each graph consists of finite words, we employ the concatenation of languages to characterize infinite runs. To do so, we first prove Lemma 3.2, which simplifies the concatenation of entire languages by demonstrating an equivalence to the concatenation of arbitrary words from these languages.

Lemma 3.2. *If $u \in L(\hat{g})$, $v \in L(\hat{h})$, $uv \in L(\hat{k})$, and $L(\bar{g}) \cdot L(\bar{h}) \subseteq L(\bar{k})$, then $L(\hat{g}) \cdot L(\hat{h}) \subseteq L(\hat{k})$*

Proof: Assume we have such an u and v . We demonstrate every word $w \in L(\hat{g}) \cdot L(\hat{h})$ must be in $L(\hat{k})$. If we expand the premise, we obtain $w \in (L(\bar{g}) \cap L(\tilde{g})) \cdot (L(\bar{h}) \cap L(\tilde{h}))$. This

implies w must be in $L(\bar{g}) \cdot L(\bar{h})$ and in $L(\tilde{g}) \cdot L(\tilde{h})$. Next, we know that $u \in L(\tilde{g})$, $v \in L(\tilde{h})$, and $w \in L(\tilde{k})$. Thus by Lemma 2.4, $L(\tilde{g}) \cdot L(\tilde{h}) \subseteq L(\tilde{k})$, and $w \in L(\tilde{k})$. Along with the premise $L(\bar{g}) \cdot L(\bar{h}) \subseteq L(\bar{k})$, we can now conclude $w \in L(\bar{k}) \cap L(\tilde{k})$, which is $L(\hat{k})$. \square

The languages $L(\hat{g})$, $\hat{g} \in \hat{Q}_{\mathcal{A},\mathcal{B}}$, cover all finite subwords of $L(\mathcal{A})$. A subword of $L(\mathcal{A})$ has at least one path between two states in $Q_{\mathcal{A}}$, and thus is in the language of an arc in $\bar{Q}_{\mathcal{A}}$. Furthermore, by Lemma 2.4 this word is described by some graph, and the pair of the arc and the graph makes a supergraph. Unlike the case of graphs and Σ^+ , the languages of supergraphs do not form a partition of $L(\mathcal{A})$: a word might have multiple paths between states in \mathcal{A} , and so be described by more than one arc in $\bar{Q}_{\mathcal{A}}$. With them we construct the finite family of ω -languages that cover $L(\mathcal{A})$. Given $\hat{g}, \hat{h} \in \hat{Q}_{\mathcal{A},\mathcal{B}}$, let $Z(\hat{g}, \hat{h})$ be the ω -language $L(\hat{g}) \cdot L(\hat{h})^\omega$. In analogy to Section 2.1, call $Z(\hat{g}, \hat{h})$ *proper* if: (1) $Z(\hat{g}, \hat{h})$ is non-empty; (2) $\bar{g} = \langle q, r \rangle$ and $\bar{h} = \langle r, r \rangle$ where $q \in Q_{\mathcal{A}}^{in}$ and $r \in F_{\mathcal{A}}$; (3) $L(\hat{g}) \cdot L(\hat{h}) \subseteq L(\bar{g})$ and $L(\hat{h}) \cdot L(\hat{h}) \subseteq L(\bar{h})$. Call a pair of supergraphs $\langle \hat{g}, \hat{h} \rangle$ *proper* if $Z(\hat{g}, \hat{h})$ is proper. We note that $Z(\hat{g}, \hat{h})$ is non-empty if $L(\hat{g})$ and $L(\hat{h})$ are non-empty, and that, by the second condition, every proper $Z(\hat{g}, \hat{h})$ is contained in $L(\mathcal{A})$.

Lemma 3.3. *Let \mathcal{A} and \mathcal{B} be two Büchi automata, and $\hat{Q}_{\mathcal{A},\mathcal{B}}$ the corresponding set of supergraphs. $L(\mathcal{A}) = \bigcup \{Z(\hat{g}, \hat{h}) \mid \hat{g}, \hat{h} \in \hat{Q}_{\mathcal{A},\mathcal{B}}, Z(\hat{g}, \hat{h}) \text{ is proper}\}$.*

Proof: We extend the Ramsey argument of Lemma 2.5 to supergraphs.

Consider an infinite word $w = \sigma_0\sigma_1\dots$ with an accepting run $p = p_0p_1\dots$ in \mathcal{A} . As p is accepting, we know that $p_0 \in Q_{\mathcal{A}}^{in}$ and $p_i \in F_{\mathcal{A}}$ for infinitely many i . Since $F_{\mathcal{A}}$ is finite, at least one accepting state q must appear infinitely often. Let $D \subseteq \mathbb{N}$ be the set of indexes i such that $p_i = q$.

We pause to observe that, by the definition of the languages of arcs, for every $i \in D$ the word $\sigma_0\dots\sigma_{i-1}$ is in $L(\langle p_0, q \rangle)$, and for every $i, j \in D$, $i < j$, the word $\sigma_i\dots\sigma_{j-1} \in L(\langle q, q \rangle)$. Every language $Z(\hat{g}, \hat{h})$ where $\bar{g} = \langle p_0, q \rangle$ and $\bar{h} = \langle q, q \rangle$ thus satisfies the second requirement of properness.

In addition to restricting our attention the subset of nodes where $p_i = q$, we further partition D into $k = 3^{n^2}$ sets D_1, \dots, D_k based on the prefix of w until that point, where there is a D_l associated with each possible graph \tilde{g}_l . By Lemma 2.4, every finite word is in the language of some graph \tilde{g} . Say that $i \in D_l$ iff $\sigma_0\dots\sigma_{i-1} \in L(\tilde{g}_l)$. As k is finite, for some m D_m must be infinite. Let $\tilde{g} = \tilde{g}_m$.

Similarly, by Lemma 2.4 we can use the word w to define a partition of all unordered *pairs* of elements from D_m . This partition consists of k sets C_1, \dots, C_k , such that $(i, j) \in C_l$ iff $\sigma_i\dots\sigma_{j-1} \in L(\tilde{g}_l)$. Without loss of generality, for $(i, j) \in C_l$, assume $i < j$. Ramsey's Theorem tells us that, given such a partition, there exists an infinite subset $\{i_1, i_2, \dots\}$ of D_m and a C_n such that $(i_j, i_k) \in C_n$ for all pairs of distinct elements i_j, i_k .

This is precisely to say there is a graph \tilde{h} so that, for every $(i_j, i_k) \in C_n$, it holds that $\sigma_{i_j}\dots\sigma_{i_k-1} \in L(\tilde{h})$. C_n thus partitions the word w into

$$w_1 = \sigma_0\dots\sigma_{i_1-1}, w_2 = \sigma_{i_1}\dots\sigma_{i_2-1}, w_3 = \sigma_{i_2}\dots\sigma_{i_3-1}, \dots,$$

such that $w_1 \in L(\tilde{g})$ and $w_i \in L(\tilde{h})$ for $i > 1$. Let $\hat{g} = \langle \langle p_0, q \rangle, \tilde{g} \rangle$ and let $\hat{h} = \langle \langle q, q \rangle, \tilde{h} \rangle$. By the above partition of w , we know that $w \in Z(\hat{g}, \hat{h})$.

We now show that $Z(\hat{g}, \hat{h})$ is proper. First, as $w \in Z(\hat{g}, \hat{h})$, we know $Z(\hat{g}, \hat{h})$ is non-empty. Second, as noted above, the second requirement is satisfied by the arcs $\langle p_0, q \rangle$

and $\langle q, q \rangle$. Finally, we demonstrate the third condition holds. As $\sigma_0 \dots \sigma_{i-1} \in L(\tilde{g})$ for every $i \in C_n$, we have that $w_1 w_2 \in L(\tilde{g})$. Both w_1 and $w_1 w_2$ are in $L(\langle p_0, q \rangle)$ and so $w_1, w_1 w_2 \in L(\hat{g})$. By the definition of the language of arcs, $L(\langle p_0, q \rangle) \cdot L(\langle q, q \rangle) \subseteq L(\langle p_0, q \rangle)$. Thus by Lemma 3.2, we can conclude that $L(\hat{g}) \cdot L(\hat{h}) \subseteq L(\hat{g})$. Next observe that as $\sigma_i \dots \sigma_{j-1} \in L(\tilde{h})$ for every pair $i, j \in C_n$, we have that $w_2 w_3 \in L(\tilde{h})$. As $w_2, w_2 w_3$ are both in $\langle q, q \rangle$, it holds that $w_2, w_2 w_3 \in L(\hat{h})$. By the definition of the language of arcs, $L(\langle q, q \rangle) \cdot L(\langle q, q \rangle) \subseteq L(\langle q, q \rangle)$. By Lemma 3.2 we can now conclude $L(\hat{h}) \cdot L(\hat{h}) \subseteq L(\hat{h})$. Therefore $Z(\hat{g}, \hat{h})$ is a proper language containing w . \square

Lemma 3.4. *Let \mathcal{A} and \mathcal{B} be two Büchi automata, and $\hat{Q}_{\mathcal{A}, \mathcal{B}}$ the corresponding set of supergraphs.*

- (1) *For all proper $Z(\hat{g}, \hat{h})$, either $Z(\hat{g}, \hat{h}) \cap L(\mathcal{B}) = \emptyset$ or $Z(\hat{g}, \hat{h}) \subseteq L(\mathcal{B})$.*
- (2) *$L(\mathcal{A}) \subseteq L(\mathcal{B})$ iff every proper language $Z(\hat{g}, \hat{h}) \subseteq L(\mathcal{B})$.*
- (3) *Let \hat{g}, \hat{h} be two supergraphs such that $Z(\hat{g}, \hat{h})$ is proper. $Z(\hat{g}, \hat{h}) \subseteq L(\mathcal{B})$ iff there exists $q \in Q_{\mathcal{B}}^{in}$, $r \in Q_{\mathcal{B}}$, $a \in \{0, 1\}$ such that $\langle q, a, r \rangle \in \hat{g}$ and $\langle r, 1, r \rangle \in \hat{h}$.*

Proof: Given two supergraphs $\hat{g} = \langle \tilde{g}, \tilde{g} \rangle$ and $\hat{h} = \langle \tilde{h}, \tilde{h} \rangle$, recall that $Y(\tilde{g}, \tilde{h})$ is the ω -language $L(\tilde{g}) \cdot L(\tilde{h})^\omega$. Further note that $L(\hat{g}) \subseteq L(\tilde{g})$ and $L(\hat{h}) \subseteq L(\tilde{h})$, and therefore $Z(\hat{g}, \hat{h}) \subseteq Y(\tilde{g}, \tilde{h})$.

1: Consider two supergraphs \hat{g}, \hat{h} . By Lemma 2.6 either $Y(\tilde{g}, \tilde{h}) \cap L(\mathcal{B}) = \emptyset$ or $Y(\tilde{g}, \tilde{h}) \subseteq L(\mathcal{B})$. Since $Z(\hat{g}, \hat{h}) \subseteq Y(\tilde{g}, \tilde{h})$, it holds that $Z(\hat{g}, \hat{h}) \cap L(\mathcal{B}) = \emptyset$ or $Z(\hat{g}, \hat{h}) \subseteq L(\mathcal{B})$.

2: Immediate from Lemma 3.3 and clause 1.

3: By Lemma 2.6 either $Y(\tilde{g}, \tilde{h}) \subseteq L(\mathcal{B})$ or $Y(\tilde{g}, \tilde{h}) \cap L(\mathcal{B}) = \emptyset$. By Lemma 2.7 $Y(\tilde{g}, \tilde{h}) \subseteq L(\mathcal{B})$ iff a q, r and a exist such that $\langle q, a, r \rangle \in \tilde{g}$ and $\langle r, 1, r \rangle \in \tilde{h}$. Since $Z(\hat{g}, \hat{h}) \subseteq Y(\tilde{g}, \tilde{h})$, $Z(\hat{g}, \hat{h}) \subseteq L(\mathcal{B})$ iff such a q, r and a exist. \square

In an analogous fashion to Section 2.1, we can use supergraphs to test the containment of two automata, \mathcal{A} and \mathcal{B} . Search all pairs of supergraphs, $\hat{g}, \hat{h} \in \hat{Q}_{\mathcal{A}, \mathcal{B}}$ for a pair that is both proper and for which there does not exist a $q \in Q_{\mathcal{B}}^{in}$, $r \in Q_{\mathcal{B}}$, $a \in \{0, 1\}$ such that $\langle q, a, r \rangle \in \hat{g}$ and $\langle r, 1, r \rangle \in \hat{h}$. Such a pair is a counterexample to containment. If no such pair exists, then $L(\mathcal{A}) \subseteq L(\mathcal{B})$.

3.2. Composition of Supergraphs. Employing supergraphs to test containment faces difficulty on two fronts. First, the number of supergraphs is very large. Second, verifying properness requires checking language nonemptiness and containment: PSPACE-hard problems. To address these problems we construct only supergraphs with non-empty languages. Borrowing the notion of composition from Section 2.3 allows us to use exponential space to compute exactly the needed supergraphs. Along the way we develop a polynomial-time test for the containment of supergraph languages. Our plan is to start with graphs corresponding to single letters and compose them until we reach closure. The resulting subset of $\hat{Q}_{\mathcal{A}, \mathcal{B}}$, written $\hat{Q}_{\mathcal{A}, \mathcal{B}}^f$, contains exactly the supergraphs with non-empty languages. In addition to removing the need to check for emptiness, composition allows us to test the sole




Figure 7: The composition of a graph with itself.

remaining aspect of properness, language containment, in time polynomial in the size of the supergraphs. We begin by defining the composition of simple graphs.

Definition 3.5. Given two graphs \tilde{g} and \tilde{h} define their composition, written as $\tilde{g};\tilde{h}$, as the graph

$$\begin{aligned} & \{ \langle q, 1, r \rangle \mid q, r, s \in Q_{\mathcal{B}}, \langle q, b, s \rangle \in \tilde{g}, \langle s, c, r \rangle \in \tilde{h}, b = 1 \text{ or } c = 1 \} \\ \cup & \{ \langle q, 0, r \rangle \mid q, r, s \in Q_{\mathcal{B}}, \langle q, 0, s \rangle \in \tilde{g}, \langle s, 0, r \rangle \in \tilde{h}, \text{ and} \\ & \text{for all } t \in Q_{\mathcal{B}}, b, c \in \{0, 1\} \text{ if } \langle q, a, t \rangle \in \tilde{g} \text{ and } \langle t, b, r \rangle \in \tilde{h} \text{ then } a = b = 0 \} \end{aligned}$$

Example 3.6. Figure 7 shows the composition of a simple graph with itself. Figure 2 is also illustrative, as the third graph is the composition of the first two.

We can then define the composition of two supergraphs $\hat{g} = \langle \langle q, r \rangle, \tilde{g} \rangle$ and $\hat{h} = \langle \langle r, s \rangle, \tilde{h} \rangle$, written $\hat{g};\hat{h}$, as the supergraph $\langle \langle q, s \rangle, \tilde{g};\tilde{h} \rangle$. To generate exactly the set of supergraphs with non-empty languages, we start with supergraphs describing single letters. For a containment problem $L(\mathcal{A}) \subseteq L(\mathcal{B})$, define the subset of $\hat{Q}_{\mathcal{A},\mathcal{B}}$ corresponding to single letters to be $\hat{Q}_{\mathcal{A},\mathcal{B}}^1 = \{ \hat{g} \mid \hat{g} \in \hat{Q}_{\mathcal{A},\mathcal{B}}, a \in \Sigma, a \in L(\hat{g}) \}$. For completeness, we present a constructive definition of $\hat{Q}_{\mathcal{A},\mathcal{B}}^1$.

Definition 3.7.

$$\begin{aligned} \hat{Q}_{\mathcal{A},\mathcal{B}}^1 = & \{ \langle \langle q, r \rangle, \tilde{g} \rangle \mid q \in Q_{\mathcal{A}}, r \in \rho_{\mathcal{A}}(q, a), a \in \Sigma, \\ & \tilde{g} = \{ \langle q', 0, r' \rangle \mid q' \in Q_{\mathcal{B}} \setminus F_{\mathcal{B}}, r' \in (\rho_{\mathcal{B}}(q', a) \setminus F_{\mathcal{B}}) \} \cup \\ & \{ \langle q', 1, r' \rangle \mid q' \in Q_{\mathcal{B}}, r' \in \rho_{\mathcal{B}}(q', a), q' \text{ or } r' \in F_{\mathcal{B}} \} \} \end{aligned}$$

We then define $\hat{Q}_{\mathcal{A},\mathcal{B}}^f$ to be the closure of $\hat{Q}_{\mathcal{A},\mathcal{B}}^1$ under composition. Algorithm `DoubleGraphSearch`, which we prove correct below, employs composition to check the containment of two automata. It first generates the set of initial supergraphs, and then computes the closure of this set under composition. Along the way it tests properness by using composition. Every time it encounters a proper pair of supergraphs, it either verifies that a satisfying pair of arcs exist, or halts with a counterexample to containment. We call this search the *double-graph search*.

To begin proving our algorithm correct, we link composition and the concatenation of languages, first for simple graphs and then for supergraphs.

Lemma 3.8. *For every two graphs \tilde{g} and \tilde{h} , it holds that $L(\tilde{g}) \cdot L(\tilde{h}) \subseteq L(\tilde{g};\tilde{h})$.*

Algorithm 2: DoubleGraphSearch(\mathcal{A}, \mathcal{B})**Data:** Two Büchi automata, \mathcal{A} and \mathcal{B} .**Result:** Whether $L(\mathcal{A})$ is contained in $L(\mathcal{B})$.Initialize $\widehat{Q}^f \leftarrow \widehat{Q}_{\mathcal{A}, \mathcal{B}}^1$ **repeat** **for all** pairs $\widehat{g}, \widehat{h} \in \widehat{Q}^f$ where $\bar{g} = \langle q, r \rangle$ and $\bar{h} = \langle r, s \rangle$ **do** Add $\widehat{g}; \widehat{h}$ to \widehat{Q}^f **if** $q \in Q_{\mathcal{A}}^{in}$, $r \in F_{\mathcal{A}}$, $s = r$, $\widetilde{g}; \widetilde{h} = \widetilde{g}$ and $\widetilde{h}; \widetilde{h} = \widetilde{h}$ **then** **if** there do not exist $\langle q, a, r \rangle \in \widehat{g}$ and $\langle r, 1, r \rangle \in \widehat{h}$ where $q \in Q_{\mathcal{B}}^{in}$ **then**
 return *Not Contained***until** \widehat{Q}^f reaches closure**return** *Contained*

Proof: Consider two words $w_1 \in L(\widetilde{g})$, $w_2 \in L(\widetilde{h})$. By Definition 2.2, to prove $w_1w_2 \in L(\widetilde{g}; \widetilde{h})$ we must show that for every $q, r \in Q$: both (1) $\langle q, a, r \rangle \in \widetilde{g}; \widetilde{h}$ iff there is a path from q to r over w_1w_2 , and (2) that $a = 1$ iff there is an accepting path.

If an arc $\langle q, a, r \rangle \in \widetilde{g}; \widetilde{h}$ exists, then there is an $s \in Q$ such that $\langle q, b, s \rangle \in \widetilde{g}$ and $\langle s, c, r \rangle \in \widetilde{h}$. By Definition 2.2, this implies the existence of a path x_1s from q to s over w_1 , and a path sx_2 from s to r over w_2 . Thus x_1sx_2 is a path from q to r over w_1w_2 .

If a is 1, then either b or c must be 1. By Definition 2.2, b (resp., c) is 1 iff there is an accepting path x'_1s (resp., sx'_2) over w_1 (resp., w_2) from q to s (resp., s to r). In this case $x'_1sx'_2$ (resp., $x_1sx'_2$) is an accepting path in \mathcal{B} from q to r over w_1w_2 .

Symmetrically, if there is a path x from q to r over w_1w_2 , then after reading w_1 we are in some state s and have split x into x_1sx_2 , so that x_1s is a path from q to s and sx_2 a path from s to r . Thus by Definition 2.2 $\langle q, b, s \rangle \in \widetilde{g}$, $\langle s, c, r \rangle \in \widetilde{h}$, and $\langle q, a, s \rangle \in \widetilde{g}; \widetilde{h}$.

Furthermore, if there is an accepting path from q to r over w_1w_2 , then after reading w_1 we are in some state s and have split the path into x_1sx_2 , so that x_1s is a path from q to s , and sx_2 a path from s to r . Either x_1s or sx_2 must be accepting, and thus by Definition 2.2 $\langle q, b, s \rangle \in \widetilde{g}$, $\langle s, c, r \rangle \in \widetilde{h}$, and either b or c must be 1. Therefore a must be 1. \square

Lemma 3.9. Let $\widehat{g}, \widehat{h}, \widehat{k}$ be supergraphs in $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^f$ such that $\bar{g} = \langle q, r \rangle$, $\bar{h} = \langle r, s \rangle$, and $\bar{k} = \langle q, s \rangle$. Then $\widehat{g}; \widehat{h} = \widehat{k}$ iff $L(\widehat{g}) \cdot L(\widehat{h}) \subseteq L(\widehat{k})$.

Proof: Assume $\widehat{g}; \widehat{h} = \widehat{k}$ as a premise. This implies $\bar{k} = \langle \langle q, s \rangle, \widetilde{g}; \widetilde{h} \rangle$. If either $L(\widehat{g})$ or $L(\widehat{h})$ are empty, then $L(\widehat{g}) \cdot L(\widehat{h})$ is empty and this direction holds trivially. Otherwise, take two words $u \in L(\widehat{g})$, $v \in L(\widehat{h})$. By construction, $u \in L(\langle q, r \rangle)$ and $v \in L(\langle r, s \rangle)$. The definition of the languages of arcs therefore implies the existence of a path from q to r over u and a path from r to s over v . Thus $uv \in L(\langle q, s \rangle)$. Similarly, $u \in L(\widetilde{g})$, $v \in L(\widetilde{h})$, and Lemma 3.8 implies that $uv \in L(\widetilde{k})$. Thus uv is in $L(\langle \langle q, r \rangle, \widetilde{k} \rangle)$. and by Lemma 3.2 $L(\widehat{g}) \cdot L(\widehat{h}) \subseteq L(\widehat{k})$.

In the other direction, if $L(\widehat{g}) \cdot L(\widehat{h}) \subseteq L(\widehat{k})$, we show that $\widehat{g}; \widehat{h} = \widehat{k}$. By definition, $\widehat{g}; \widehat{h}$ is $\langle \langle q, s \rangle, \widetilde{g}; \widetilde{h} \rangle$. As $\widehat{g}, \widehat{h} \in \widehat{Q}_{\mathcal{A}, \mathcal{B}}^f$, they are the composition of a finite number of graphs from $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^1$. The above direction then demonstrates that they are non-empty, and there is

a word $w \in L(\widehat{g}) \cdot L(\widehat{h})$. This expands to $w \in (L(\bar{g}) \cap L(\widetilde{g})) \cdot (L(\bar{h}) \cap L(\widetilde{h}))$, which implies $w \in L(\widetilde{g}) \cdot L(\widetilde{h})$. By Lemma 3.8, w is then in $L(\widetilde{g}; \widetilde{h})$. Since, by Lemma 2.4, w is the language of exactly one graph, we have that $\widetilde{g}; \widetilde{h} = \widehat{k}$, which proves $\widehat{g}; \widehat{h} = \widehat{k}$. \square

Lemma 3.9 provides the polynomial time test for properness employed in Algorithm `DoubleGraphSearch`. Namely, given two supergraphs $\widehat{g} = \langle \langle q, r \rangle, \widetilde{g} \rangle$ and $\widehat{h} = \langle \langle r, r \rangle, \widetilde{h} \rangle$ from $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^f$, the pair $\langle \widehat{g}, \widehat{h} \rangle$ is proper exactly when $q \in Q_{\mathcal{A}}^{in}$, $r \in F_{\mathcal{A}}$, $\widetilde{g}; \widetilde{h} = \widetilde{g}$ and $\widetilde{h}; \widetilde{h} = \widetilde{h}$. We now provide the final piece of our puzzle: proving that the closure of $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^1$ under composition contains every non-empty supergraph.

Lemma 3.10. *For two Büchi automata \mathcal{A} and \mathcal{B} , every $\widehat{h} \in \widehat{Q}_{\mathcal{A}, \mathcal{B}}$, where $L(\widehat{h}) \neq \emptyset$, is in $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^f$.*

Proof: Let $\widehat{h} = \langle \langle q, r \rangle, \widetilde{h} \rangle$ where $L(\widehat{h}) \neq \emptyset$. Then there is at least one word $w = \sigma_0 \dots \sigma_{n-1} \in L(\widehat{h})$, which is to say $w \in L(\langle q, r \rangle) \cap L(\widetilde{h})$. By the definition of the languages of arcs, there is a path $p = p_0 \dots p_n$ in \mathcal{A} over w such that $p_0 = q$ and $p_n = r$.

Define \widetilde{g}_{σ_i} to be the graph in $\widetilde{Q}_{\mathcal{B}}^1$ containing σ_i . Let \widehat{g}_{σ_i} be $\langle \langle p_i, p_{i+1} \rangle, \widetilde{g}_{\sigma_i} \rangle$, and let \widehat{g}_w be $\widehat{g}_{\sigma_0}; \widehat{g}_{\sigma_1}; \dots; \widehat{g}_{\sigma_{n-1}}$. Note that each $\widehat{g}_{\sigma_i} \in \widehat{Q}_{\mathcal{A}, \mathcal{B}}^1$. By Lemma 3.9 $w \in \widetilde{g}_w$. By Lemma 2.4, w is in only one graph and $\widetilde{g}_w = \widetilde{h}$. By construction, $\bar{g}_w = \langle q, r \rangle$. Therefore $\langle \bar{g}_w, \widetilde{g}_w \rangle = \langle \langle q, r \rangle, \widetilde{h} \rangle = \widehat{h}$, and \widehat{h} is in the closure of $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^1$ under composition. \square

We can now show the correctness of Algorithm `DoubleGraphSearch`, using Lemma 3.9 to justify testing properness with composition, and Lemma 3.11 below to justify the correctness and completeness of our search for a counterexample.

Lemma 3.11. *Let \mathcal{A} and \mathcal{B} be two Büchi automata. $L(\mathcal{A})$ is not contained in $L(\mathcal{B})$ iff $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^f$ contains a pair of supergraphs \widehat{g}, \widehat{h} such that $\langle \widehat{g}, \widehat{h} \rangle$ is proper and there do not exist arcs $\langle q, a, r \rangle \in \widehat{g}$ and $\langle r, 1, r \rangle \in \widehat{h}$, $q \in Q_{\mathcal{B}}^{in}$.*

Proof: As all proper graphs are non-empty, this follows from parts (2) and (3) of Lemma 3.4 and Lemma 3.10. \square

Theorem 3.1. For every two Büchi automata \mathcal{A} and \mathcal{B} , it holds that $L(\mathcal{A}) \subseteq L(\mathcal{B})$ iff `DoubleGraphSearch`(\mathcal{A}, \mathcal{B}) returns Contained.

Proof: By Lemma 3.9, testing for composition is equivalent to testing for language containment, and the outer conditional in Algorithm `DoubleGraphSearch` holds only for proper pairs of supergraphs. By Lemma 3.11, the inner conditional checks if a proper pair of supergraphs is a counterexample, and if no such proper pair in $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^f$ is a counterexample then containment must hold. \square

3.3. Strongly Suffix Closed Languages. Algorithm `DoubleGraphSearch` has much the same structure as Algorithm `LJB`. The most noticeable difference is that Algorithm `DoubleGraphSearch` checks pairs of supergraphs, where Algorithm `LJB` checks only single size-change graphs. Indeed, Theorem 2.1 suggests that, for some languages, a cycle implies the existence of a lasso. When proving containment of Büchi automata with such languages, it is sufficient to search for a graph $\widehat{h} \in \widehat{Q}_{\mathcal{B}}$, where $\widehat{h}; \widehat{h} = \widehat{h}$, with no arc $\langle r, 1, r \rangle$. This *single-graph search* reduces the complexity of our algorithm significantly. What enables

this in size-change termination is the late-start property: threads can begin at arbitrary points. We here define the class of automata amenable to this optimization, first presenting the case for universality testing, without proof, for clarity.

In size-change termination, the late-start property asserts that an accepting cycle can start at an arbitrary point. Intuitively, this suggests that an arc $\langle r, 1, r \rangle \in \tilde{h}$ might not need a matching prefix $\langle q, a, r \rangle$ in some \tilde{g} : the cycle can just start at r . In the context of universality, we can apply this method when it is safe to add or remove arbitrary prefixes of a word. To describe these languages we extend the standard notion of *suffix closure*. A language L is suffix closed when, for every $w \in L$, every suffix of w is in L .

Definition 3.12. A language L is *strongly suffix closed* if it is suffix closed and for every $w \in L$, $w_1 \in \Sigma^+$, we have that $w_1w \in L$.

Lemma 3.13. Let \mathcal{B} be an Büchi automaton where every state in Q is reachable and $L(\mathcal{B})$ is strongly suffix closed. \mathcal{B} is not universal iff the set of supergraphs with non-empty languages, $\tilde{Q}_{\mathcal{B}}^f$, contains a graph \tilde{h} such that $\tilde{h}; \tilde{h}$ and \tilde{h} does not contain an arc of the form $\langle r, 1, r \rangle$.

As an intuition for the correctness of Lemma 3.13, note that the existence of an 1-labeled cyclic arc in \tilde{h} implies a loop, that Q being reachable implies a prefix can be prepended to this loop to make a lasso, and that strong suffix closure allows us to swap this prefix for the prefix of every other word that share this cycle.

To extend this notion to handle containment questions $L_1 \subseteq L_2$, we restrict our focus to words in L_1 . Instead of requiring L_2 to be closed under arbitrary prefixes, L_2 need only be closed under prefixes that keep the word in L_1 .

Definition 3.14. A language L_2 is *strongly suffix closed with respect to L_1* when L_2 is suffix closed and, for every $w \in L_1 \cap L_2$, $w_1 \in \Sigma^+$, if $w_1w \in L_1$ then $w_1w \in L_2$.

When checking the containment of \mathcal{A} in \mathcal{B} for the case when $L(\mathcal{B})$ is strongly suffix closed with respect to $L(\mathcal{A})$, we can employ a the simplified algorithm below. As in Algorithm `DoubleGraphSearch`, we search all supergraphs in $\hat{Q}_{\mathcal{A},\mathcal{B}}^f$. Rather than searching for a proper pair of supergraphs, however, Algorithm `SingleGraphSearch` searches for a single supergraph \hat{h} where $\hat{h}; \hat{h} = \hat{h}$ that does not contain an arc of the form $\langle r, 1, r \rangle$. We call this search the *single-graph search*.

We now prove Algorithm `SingleGraphSearch` correct. Theorem 3.2 demonstrates that, under the requirements specified, the presence of a single-graph counterexample refutes containment, and the absence of such a supergraph proves containment.

Theorem 3.2. Let \mathcal{A} and \mathcal{B} be two Büchi automata where $Q_{\mathcal{A}}^{in} = Q_{\mathcal{A}}$, every state in $Q_{\mathcal{B}}$ is reachable, and $L(\mathcal{B})$ is strongly suffix closed with respect to $L(\mathcal{A})$. Then $L(\mathcal{A})$ is *not* contained in $L(\mathcal{B})$ iff $\hat{Q}_{\mathcal{A},\mathcal{B}}^f$ contains a supergraph $\hat{h} = \langle \langle s, s \rangle, \tilde{h} \rangle$ such that $s \in F_{\mathcal{A}}$, $\hat{h}; \hat{h} = \hat{h}$, and \hat{h} does *not* contain an arc $\langle r, 1, r \rangle$.

Proof: In one direction, assume $\hat{Q}_{\mathcal{A},\mathcal{B}}^f$ contains a supergraph $\hat{h} = \langle \langle s, s \rangle, \tilde{h} \rangle$ where $s \in F_{\mathcal{A}}$, $\hat{h}; \hat{h} = \hat{h}$, and there is no arc $\langle r, 1, r \rangle \in \hat{h}$. We show that $Z(\hat{h}, \hat{h})$ is a proper language not contained in $L(\mathcal{B})$. As $\hat{h} \in \hat{Q}_{\mathcal{A},\mathcal{B}}^f$, we know $L(\hat{h})$ is not empty, implying $Z(\hat{h}, \hat{h})$ is non-empty. As $Q_{\mathcal{A}} = Q_{\mathcal{A}}^{in}$, it holds that $s \in Q_{\mathcal{A}}^{in}$. By Lemma 3.9, the premise $\hat{h}; \hat{h} = \hat{h}$ implies $L(\hat{h}) \cdot L(\hat{h}) \subseteq L(\hat{h})$, and $Z(\hat{h}, \hat{h})$ is proper. Finally, as there is no $\langle r, 1, r \rangle \in \hat{h}$, by Theorem 3.1, $Z(\hat{h}, \hat{h}) \notin L(\mathcal{B})$, and $Z(\hat{h}, \hat{h})$ is a counterexample to $L(\mathcal{A}) \subseteq L(\mathcal{B})$.

Algorithm 3: SingleGraphSearch(\mathcal{A}, \mathcal{B})

Data: Two Büchi automata, \mathcal{A} and \mathcal{B} .

Require $Q_{\mathcal{A}}^{in} = Q_{\mathcal{A}}$, $Q_{\mathcal{B}}$ is reachable, and $L(\mathcal{B})$ is strongly suffix closed w.r.t. $L(\mathcal{A})$

Result: Whether $L(\mathcal{A})$ is contained in $L(\mathcal{B})$.

Initialize $\widehat{Q}^f \leftarrow \widehat{Q}_{\mathcal{A}, \mathcal{B}}^1$

repeat

for all pairs $\widehat{g}, \widehat{h} \in \widehat{Q}^f$ where $\bar{g} = \langle q, r \rangle$ and $\bar{h} = \langle r, s \rangle$ **do**

$\widehat{k} \leftarrow \widehat{g}; \widehat{h}$

 Add \widehat{k} to \widehat{Q}^f

if $q \in F_{\mathcal{A}}$, $q = s$, and $\widehat{k}; \widehat{k} = \widehat{k}$ **then**

if there does not exist an arc $\langle r, 1, r \rangle \in \widehat{k}$ **then**

return *Not Contained*

until \widehat{Q}^f reaches closure

return *Contained*

In the opposite direction, assume the premise that $\widehat{Q}_{\mathcal{A}, \mathcal{B}}^f$ does not contain a supergraph $\widehat{h} = \langle \langle s, s \rangle, \widehat{h} \rangle$ where $s \in F_{\mathcal{A}}$, $\widehat{h}; \widehat{h} = \widehat{h}$, and there is no arc $\langle r, 1, r \rangle \in \widehat{h}$. We prove that every word $w \in L(\mathcal{A})$ is also in $L(\mathcal{B})$. Take a word $w \in L(\mathcal{A})$. By Lemma 3.3, w is in some proper language $Z(\widehat{g}, \widehat{h})$ and can be broken into $w_1 w_2$ where $w_1 \in L(\widehat{g})$, $w_2 \in L(\widehat{h})^\omega$.

Because $Z(\widehat{g}, \widehat{h})$ is proper, Lemma 3.9 implies $\widehat{h} = \langle \langle s, s \rangle, \widehat{h} \rangle$ where $s \in F_{\mathcal{A}}$ and $\widehat{h}; \widehat{h} = \widehat{h}$. This, along with our premise, implies \widehat{h} contains an arc $\langle r, 1, r \rangle$. Since all states in $Q_{\mathcal{B}}$ are reachable, there is $q \in Q_{\mathcal{B}}^{in}$ and $u \in \Sigma^+$ with a path in \mathcal{B} from q to r over u . By Lemma 2.7, this implies uw_2 is accepted by \mathcal{B} . For $L(\mathcal{B})$ to strongly suffix closed with respect to $L(\mathcal{A})$, it must be suffix closed. Therefore $w_2 \in L(\mathcal{B})$. Now we move to $L(\mathcal{A})$, and note that the premise $Q_{\mathcal{A}}^{in} = Q_{\mathcal{A}}$ implies $L(\mathcal{A})$ is suffix closed. Thus the fact that $w_1 w_2 \in L(\mathcal{A})$ implies $w_1 \in L(\mathcal{A})$. Since $L(\mathcal{B})$ is strongly suffix closed with respect to $L(\mathcal{A})$, and $w_2 \in L(\mathcal{B})$, it must be that $w_1 w_2 \in L(\mathcal{B})$. \square

3.4. From Ramsey-Based Containment to Size-Change Termination. We now delve into the connection between the LJB algorithm for size-change termination and the single-graph search algorithm for Büchi containment. We will show that Algorithm LJB is a specialized realization of Algorithm **SingleGraphSearch**. Given an SCT problem L , size-change graphs in $\text{LJB}(L)$ are direct analogues of supergraphs in $\text{SingleGraphSearch}(\mathcal{A}_{Flow(L)}, \mathcal{A}_{Desc(L)})$. For convenience, take $L = \langle H, P, C, \mathcal{G} \rangle$, $\mathcal{A}_{Flow(L)} = \langle C, Q_{Fl}, Q_{Fl}^{in}, \rho_{Fl}, F_{Fl} \rangle$, and $\mathcal{A}_{Desc(L)} = \langle C, Q_{Ds}, Q_{Ds}^{in}, \rho_{Ds}, F_{Ds} \rangle$.

We first show that $\mathcal{A}_{Flow(L)}$ and $\mathcal{A}_{Desc(L)}$ satisfy the preconditions of Algorithm **SingleGraphSearch**: that $Q_{Fl}^{in} = Q_{Fl}$; that every state in Q_{Ds} is reachable; and that $Desc(L)$ is strongly suffix closed with respect to $Flow(L)$. For the first and second requirement, it suffices to observe that every state in both $\mathcal{A}_{Flow(L)}$ and $\mathcal{A}_{Desc(L)}$ is initial.³

For the third, strong suffix closure is a direct consequence of the definition of a thread: since a thread can start at arbitrary points, it does not matter what call path we use to

³In the original reduction, 1-labeled parameters may not have been reachable.




Figure 8: Size-Change Graphs vs. Supergraphs: The SCGs for call sites a and b , from Figure 3, and corresponding supergraphs for the characters a and b , from $\hat{Q}_{\mathcal{A}_{Flow(L)}, \mathcal{A}_{Desc(L)}}^1$.

reach that point. Adding a prefix to a call path cannot cause that call path to become non-terminating. Thus the late-start property is precisely $Desc(L)$ being strongly suffix closed with respect to $Flow(L)$, and we can employ the single-graph search.

Consider supergraphs in $\hat{Q}_{\mathcal{A}_{Flow(L)}, \mathcal{A}_{Desc(L)}}$, from here simply denoted by \hat{Q}_L . The state space of $\mathcal{A}_{Flow(L)}$ is the set of functions H , and the state space of $\mathcal{A}_{Desc(L)}$ is the union of H and Q_p , the set of all $\{0, 1\}$ -labeled parameters. A supergraph in \hat{Q}_L thus comprises an arc $\langle q, r \rangle$ in H and a $\{0, 1\}$ -labeled graph \tilde{g} over $H \cup Q_p$. The arc asserts the existence of a call path from q to r , and the graph \tilde{g} captures the relevant information about corresponding paths in $\mathcal{A}_{Desc(L)}$.

These supergraphs are almost the same as SCGs, $G : q \rightarrow r$ (See Figure 8). Aside from notational differences, both contain an arc asserting the existence of a call path between two functions, and a $\{0, 1\}$ -labeled graph. There are nodes in both graphs that correspond to parameters of functions, and arcs between two such nodes describe a thread between the corresponding parameters. The analogy falls short, however, on three points:

- (1) In SCGs, nodes are always parameters of functions. In supergraphs, nodes can be either parameters of functions or function names.
- (2) In SCGs, nodes are unlabeled. In supergraphs, nodes are labeled either 0 or 1.
- (3) In an SCG, only the nodes corresponding to the parameters of two specific functions are present. In a supergraph, nodes corresponding to every parameter of every function exist.

Each difference is an opportunity to specialize the Ramsey-based containment algorithm, Algorithm `SingleGraphSearch`, by simplifying supergraphs. When these specializations are taken together, we have Algorithm `LJB`.

- 1 No functions in H are accepting for $\mathcal{A}_{Desc(L)}$, and once we transition out of H into Q_p we can never return to H . Therefore nodes r corresponding to function names can never be part of a descending arc $\langle r, 1, r \rangle$. Since we only search for arcs of the form $\langle r, 1, r \rangle$, we can simplify supergraphs in \hat{Q}_L by removing all nodes corresponding to functions.
- 2 The labels on parameters are the result of encoding a Büchi edge acceptance condition in a Büchi state acceptance condition automaton, and can be dropped from supergraphs with no loss of information. Consider an arc $\langle \langle f, a \rangle, b, \langle g, c \rangle \rangle$. If b is 1, we know the corresponding thread contains a descending arc. The value of c tells us if the final arc

in the thread is descending, but *which arc* is descending is irrelevant. Thus it is safe to simplify supergraphs in \widehat{Q}_L by removing labels on parameters.

- 3 While all parameters have corresponding states in $\mathcal{A}_{Desc(L)}$, each supergraph describes threads in a call sequence between two particular functions. There are no threads in this call sequence between parameters of other functions, and so no supergraph with a non-empty language has arcs between the parameters of other functions. We can thus simplify supergraphs in \widehat{Q}_L by removing all nodes corresponding to parameters of other functions.

To formalize this notion of simplification, we first define, \widehat{G}_L , the set of simplified supergraphs and show that \widehat{G}_L is in one-to-one correspondence with S , the closure of \mathcal{G} under composition.

Definition 3.15. $\widehat{G}_L = \{\langle \langle f_1, f_2 \rangle, \tilde{k} \rangle \mid f_1, f_2 \in H, \tilde{k} \subseteq 2^{P(f_1) \times \{0,1\} \times P(f_2)}\}$

Say that $\langle r, \tilde{g} \rangle \in \widehat{Q}_L$ *simplifies to* $\langle r, \tilde{k} \rangle \in \widehat{G}_L$ when $\langle q, b, r \rangle \in \tilde{k}$ iff there exists $a, c \in \{0, 1\}$ such that $\langle \langle q, a \rangle, b, \langle r, c \rangle \rangle \in \tilde{g}$. Let \widehat{G}_L^1 be $\{\tilde{k} \mid \widehat{g} \in \widehat{Q}_L^1, \widehat{g} \text{ simplifies to } \tilde{k}\}$, and \widehat{G}_L^f be the closure of \widehat{G}_L^1 under composition.

We can map SCGs directly to elements of \widehat{G}_L . Say $G : f_1 \rightarrow f_2 \equiv \langle \langle f_1, f_2 \rangle, \tilde{g} \rangle$ when $q \xrightarrow{a} r \in G$ iff $\langle q, a, r \rangle \in \tilde{g}$. Note that the composition operations for supergraphs of this form is identical to the composition of SCGs: if $G_1 \equiv \widehat{g}$ and $G_2 \equiv \widehat{h}$, then $G_1; G_2 \equiv \widehat{g}; \widehat{h}$. Therefore every element of \widehat{Q}_L^f simplifies to some element of \widehat{G}_L^f .

We now show that supergraphs whose languages contain single characters are in one-to-one correspondence with \mathcal{G} , and that every idempotent element of \widehat{G}_L^f contains an arc of the form $\langle r, 1, r \rangle$ exactly when the closure of \mathcal{G} under composition does not contain a counterexample graph.

Lemma 3.16. *Let $L = \langle H, P, C, \mathcal{G} \rangle$ be an SCT problem.*

- (1) *The \equiv relation is a one-to-one correspondence between \widehat{G}_L^1 and \mathcal{G}*
- (2) *L is not size-change terminating iff \widehat{G}_L^f contains a supergraph \widehat{k} such that $\widehat{k}; \widehat{k} = \widehat{k}$ and there does not exist an arc of the form $\langle r, 1, r \rangle$ in \widehat{k} .*

Proof:

(1): Given a size-change graph $G \in \mathcal{G}$, we construct a unique supergraph $\widehat{k} \in \widehat{G}_L^1$ such that $\widehat{k} \equiv G$. Every size-change graph $G : f_1 \rightarrow f_2 \in \mathcal{G}$ is the SCG for a call site c from f_1 to f_2 . This is a call sequence of length one. Thus there is a $\widehat{g} \in \widehat{Q}_L^1$ so that $c \in L(\widehat{g})$ and $\widehat{g} = \langle f_1, f_2 \rangle$. We show that the simplification of \widehat{g} is equivalent to G . By the reduction of Definition 2.12 and the definition of graphs in Definition 2.2, the arc $\langle \langle q, b \rangle, a, \langle r, c \rangle \rangle \in \widehat{g}$, for some $b, c \in \{0, 1\}$, exactly when $q \xrightarrow{a} r \in G$. The supergraph \widehat{g} simplifies to some $\widehat{k} \in \widehat{G}_L$. By the definition of simplification, $\langle \langle q, b \rangle, a, \langle r, c \rangle \rangle \in \widehat{g}$ exactly when $\langle q, a, r \rangle \in \widehat{k}$. Thus $\widehat{k} \equiv G : f_1 \rightarrow f_2$.

In the other direction, $\widehat{k} \in \widehat{G}_L^1$ iff there exists a $\widehat{g} = \langle \langle f_1, f_2 \rangle, \tilde{g} \rangle \in \widehat{Q}_L^1$ that simplifies to \widehat{k} . By Definition 3.7, which defines \widehat{Q}_L^1 and Definition 2.12, \widehat{g} exists because there is a call site c . This call site corresponds to a SCG $G : f_1 \rightarrow f_2$. Analogously to the above, by Definition 2.12 the arcs in \tilde{g} between parameters correspond to arcs in G : $\langle \langle q, b \rangle, a, \langle r, c \rangle \rangle \in \tilde{g}$, for some $b, c \in \{0, 1\}$, exactly when $q \xrightarrow{a} r \in G$. These are the only arcs that remain after simplification, during which the labels are removed. Thus $\widehat{k} \equiv G : f_1 \rightarrow f_2$.

(2): By (1), \mathcal{G} is in one-to-one correspondence with \widehat{G}_L^1 under the \equiv relation. Since composition of supergraphs and SCGs is identical, S , the closure of \mathcal{G} under composition, is in one-to-one correspondence with the \widehat{G}_L^f . Claim (2) then follows from Theorem 2.1 and claim (1). \square

In conclusion, we can specialize the Ramsey-based containment algorithm for $L(\mathcal{A}_{Flow(L)}) \subseteq L(\mathcal{A}_{Desc(L)})$ in two ways. First, by Theorem 3.2 we know that $Flow(L) \subseteq Desc(L)$ if and only if \widehat{Q}_L contains an idempotent graph $\widehat{g} = \widehat{g};\widehat{g}$ with no arc of the form $\langle r, 1, r \rangle$. Thus we can employ the single-graph search instead of the double-graph search. Secondly, we can simplify supergraphs in \widehat{Q}_L by removing the labels on nodes and keeping only nodes associated with appropriate parameters for the source and target function. The simplifications of supergraphs whose languages contain single characters are in one-to-one correspondence with \mathcal{G} , the initial set of SCGs. As every state in $Flow(L)$ is accepting, every idempotent supergraph can serve as a counterexample. Therefore $Desc(L) \subseteq Flow(L)$ if and only if the closure of the set of simplified supergraphs, which is in one-to-one correspondence with \mathcal{G} , under composition does not contain an idempotent supergraph with no arc of the form $\langle r, 1, r \rangle$. This is precisely Algorithm LJB.

4. EMPIRICAL ANALYSIS

All the Ramsey-based algorithms presented in Section 2.3 have worst-case running times that are exponentially larger than those of the rank-based algorithms. We now compare existing, Ramsey-based, SCT tools to a rank-based Büchi containment solver on the domain of SCT problems. To facilitate a fair comparison, we briefly describe two improvements to the algorithms presented above.

4.1. Towards an Empirical Comparison. First, in constructing the analogy between SCGs in the LJB algorithm and supergraphs in the Ramsey-based containment algorithm, we noticed that supergraphs contain nodes for every parameter, while SCGs contain only nodes corresponding to parameters of relevant functions. These nodes are states in $\mathcal{A}_{Desc(L)}$. While we can specialize the Ramsey-based test to avoid them, Büchi containment solvers might suffer. These states duplicate information. As we already know which functions each supergraph corresponds to, there is no need for each node to be unique to a specific parameter.

These extra states emerge because $Desc(L)$ only accepts strings that are contained in $Flow(L)$, and in doing so demands that parameters only be reached by appropriate call paths. But the behavior of $\mathcal{A}_{Desc(L)}$ on strings not in $Flow(L)$ is irrelevant to the question of $Flow(L) \subseteq Desc(L)$, and we can replace the names of parameters in $\mathcal{A}_{Desc(L)}$ with their location in the parameter list. Further, we can rely on $Flow(L)$ to verify the sequence of function calls before our accepting thread and make do with a single waiting state. As an example, see Figure 9.




Figure 9: $\mathcal{A}_{Desc(L)}$ (left), from the original reduction of Definition 2.12, and $\mathcal{A}'_{Desc(L)}$ (right) from the reduction of Definition 4.1.

Definition 4.1. Given an SCT problem $L = \langle H, P, C, \mathcal{G} \rangle$ and a projection Ar of all parameters onto their positions $1..n$ in the parameter list, define:

$$\begin{aligned} \mathcal{A}'_{Desc(L)} &= \langle C, S \cup \{q_0\}, S \cup \{q_0\}, \rho_D, F \rangle \\ \text{where } S &= \{1..n\} \times \{1, 0\} \\ \rho_D(q_0, c) &= \{q_0\} \cup \{ \langle Ar(x), 0 \rangle \mid c : f_1 \rightarrow f_2, x \in P(f_2) \} \\ \rho_D(\langle h, a \rangle, c) &= \{ \langle Ar(y), a' \rangle \mid x \xrightarrow{a'} y \in \mathcal{G}_c, h = Ar(x) \} \\ F &= \{1..n\} \times \{1\} \end{aligned}$$

Lemma 4.2. $L(\mathcal{A}_{Flow(L)}) \subseteq L(\mathcal{A}_{Desc(L)})$ iff $L(\mathcal{A}_{Flow(L)}) \subseteq L(\mathcal{A}'_{Desc(L)})$

Proof: The languages of $\mathcal{A}_{Desc(L)}$ and $\mathcal{A}'_{Desc(L)}$ are not the same. What we demonstrate is that for every word in $Flow(L)$, we can convert an accepting run in one of $\mathcal{A}_{Desc(L)}$ or $\mathcal{A}'_{Desc(L)}$ into an accepting run in the other. Recall that the states of $\mathcal{A}_{Flow(L)}$ are functions $f \in H$. States of $\mathcal{A}_{Desc(L)}$ are either elements of H or elements of $Q_p = \bigcup_{f \in H} P(f) \times \{1, 0\}$, the set of labeled parameters. For convenience, given a pair $\langle x, a \rangle \in Q_p$, define $Ar(\langle x, a \rangle)$ to be $\langle Ar(x), a \rangle$.

Consider an accepting run $r = r_0 r_1 \dots$ of $\mathcal{A}_{Desc(L)}$ over a word w . Let $s = s_0 s_1 \dots$ be the sequence of states in $\mathcal{A}'_{Desc(L)}$ such that when $r_i \in H$, $s_i = q_0$, and when $r_i \in Q_p$, $s_i = Ar(r_i)$. By the definition of $\mathcal{A}'_{Desc(L)}$, q_0 always transitions to q_0 and a transition between r_i and r_{i+1} implies a transition between $Ar(r_i)$ and $Ar(r_{i+1})$. Therefore s is a run of $\mathcal{A}'_{Desc(L)}$ over w . Furthermore, if $\langle x, a \rangle$ is an accepting state in $\mathcal{A}_{Desc(L)}$, $a = 1$ and $\langle Ar(x), a \rangle$ is an accepting state in $\mathcal{A}'_{Desc(L)}$. Thus, s is an accepting run of $\mathcal{A}'_{Desc(L)}$ over w .

Conversely, consider a word w with an accepting run $r = r_0 r_1 \dots$ of $\mathcal{A}'_{Desc(L)}$ and an accepting run $s = s_0 s_1 \dots$ of $\mathcal{A}_{Flow(L)}$. We define an accepting run $t = t_0 t_1 \dots$ of $\mathcal{A}_{Desc(L)}$ on w . Each t_i depends on the corresponding r_i and s_i . If $r_i = q_0$ and $s_i = f$, then $t_i = f$. If $r_i = \langle k, a \rangle$ and $s_i = f$, then $t_i = \langle x, a \rangle$ where x is the k th parameter in f 's parameter list.

For a call $c : f_1 \rightarrow f_2$, take two labeled parameters, q a labeled parameter of f_1 and r a labeled parameter of f_2 . If $\langle Ar(q), c, Ar(r) \rangle$ is a transition in $\mathcal{A}'_{Desc(L)}$, then $\langle q, c, r \rangle$ is




Figure 10: Subsumption: two graphs \tilde{g} and \tilde{h} , where $\tilde{g} \preceq \tilde{h}$.

a transition in $\mathcal{A}_{Desc(L)}$. Therefore t is a run of $\mathcal{A}_{Desc(L)}$ on w . Furthermore, note that $\langle x, a \rangle \in F_{\mathcal{A}_{Desc(L)}}$ and $\langle Ar(x), a \rangle \in F_{\mathcal{A}_{Desc(L)}}$ iff $a = 1$. Therefore t is an accepting run. \square

Second, in [4], Ben-Amram and Lee present a polynomial approximation of the LJB algorithm for SCT. To facilitate a fair comparison, they optimize the LJB algorithm for SCT by using subsumption to remove certain SCGs when computing the closure under composition. This suggests that the single-graph search of Algorithm `SingleGraphSearch` can also employ subsumption. When computing the closure of a set of graphs under compositions, we can ignore elements when they are approximated by other elements. Intuitively, a graph \tilde{g} approximates another graph \tilde{h} when it is strictly harder to find a 1-labeled sequence of arcs through \tilde{g} than through \tilde{h} . If we can replace \tilde{h} with \tilde{g} without losing arcs, we do not have to consider \tilde{h} . When the right arc can be found in \tilde{g} , then it also occurs in \tilde{h} . On the other hand, when \tilde{g} does not have a satisfying arc, then we already have a counterexample.

Formally, given two graphs $\tilde{g}, \tilde{h} \in \tilde{Q}_B$ say that \tilde{g} approximates \tilde{h} , written $\tilde{g} \preceq \tilde{h}$, when for every arc $\langle q, a, r \rangle \in \tilde{g}$ there is an arc $\langle q, a', r \rangle \in \tilde{h}$, $a \leq a'$. An example is provided in Figure 10. Note that approximation is a transitive relation. In order to safely employ approximation as a subsumption relation, Ben-Amram and Lee replace the search for a single arc in idempotent graphs with a search for a strongly connected component in all graphs. This was proven to be safe in [13]: when computing the closure of \tilde{Q}_B^1 under composition, it is sufficient to store only maximal elements under this relation.

4.2. Experimental Results. All experiments were performed on a Dell Optiplex GX620 with a single 1.7Ghz Intel Pentium 4 CPU and 512 MB. Each tool was given 3500 seconds, a little under one hour, to complete each task.

Tools: The formal-verification community has implemented rank-based tools in order to measure the scalability of various approaches. The programming-languages community has implemented several Ramsey-based SCT tools. We use the best-of-breed rank-based tool, **Mh**, developed by Doyen and Raskin [9], that leverages a subsumption relation on ranks. We expanded the Mh tool to handle Büchi containment problems with arbitrary languages, thus implementing the full containment-checking algorithm presented in their paper.

We use two Ramsey-based tools. **SCTP** is a direct implementation of the LJB algorithm of Theorem 2.1, written in Haskell [15]. We have extended SCTP to reduce SCT problems to Büchi containment problems, using either Definition 2.12 or 4.1. **sct/scp** is an optimized

C implementation of the SCT algorithm, which uses the subsumption relation of Section 4.1 [4].

Problem Space: Existing experiments on the practicality of SCT solvers focus on examples extracted from the literature [4]. We combine examples from a variety of sources [1, 4, 15, 17, 24, 28, 32]. The time spent reducing SCT problems to Büchi automata never took longer than 0.1 seconds and was dominated by I/O. Thus this time was not counted. We compared the performance of the rank-based Mh solver on the derived Büchi containment problems to the performance of the existing SCT tools on the original SCT problems. If an SCT problem was solved in all incarnations and by all tools in less than 1 second, the problem was discarded as uninteresting. Unfortunately, of the 242 SCT problems derived from the literature, only 5 prove to be interesting.

Experiment Results: Table 1 compares the performance of the rank-based Mh solver against the performance of the existing SCT tools, displaying which problems each tool could solve, and the time taken to solve them. Of the interesting problems, both Sctp and Mh could only complete 3. On the other hand, sct/scp completed all of them, and had difficulty with only one problem.

Problem	Sctp (s)	Mh (s)	sct/scp (s)
ex04 [4]	1.58	Time Out	1.39
ex05 [4]	Time Out	Time Out	227.7
ms [15]	Time Out	0.1	0.02
gexgcd [15]	0.55	14.98	0.023
graphcolour2 [17]	0.017	3.18	0.014

Table 1: SCT problem completion time by tool.

The small problem space makes it difficult to draw firm conclusions, but it is clear that Ramsey-based tools are comparable to rank-based tools on SCT problems: the only tool able to solve all problems was Ramsey based. This is surprising given the significant difference in worst-case complexity, and motivates further exploration.

5. REVERSE-DETERMINISM

In the previous section, the theoretical gap in performance between Ramsey and rank-based solutions was not reflected in empirical analysis. Upon further investigation, it is revealed that a property of the domain of SCT problems is responsible. Almost all problems, and every difficult problem, in this experiment have SCGs whose nodes have an in-degree of at most 1. This property was first observed by Ben-Amram and Lee in their analysis of SCT complexity [4]. After showing how this property explains the performance of Ramsey-based algorithms, we explore why this property emerges and argue that it is a reasonable property for SCT problems to possess. Finally, we improve the rank-based algorithm for problems with this property.

Algorithm 4: $\text{gcd}(X, Y)$

```

1 if  $Y > X$  then
2    $\lfloor \text{gcd}(X, Y - X)$ 
3 else if  $X > Y$  then
4    $\lfloor \text{gcd}(X - Y, Y)$ 
5 else return  $X$ 

```

As stated above, all interesting SCGs in this experiment have nodes with at most one incoming edge. In analogy to the corresponding property for automata, we call this property of SCGs *reverse-determinism*. Say that a SCG is reverse-deterministic if every parameter of f_2 has at most one incoming edge. Given a set of reverse-deterministic SCGs \mathcal{G} , we observe three consequences. First, a reverse-deterministic SCG can have no more than n arcs: one entering each node. Second, there are only $2^{O(n \log n)}$ possible such combinations of n arcs. Third, the composition of two reverse-deterministic SCGs is also reverse-deterministic. Therefore every element in the closure of \mathcal{G} under composition is also reverse-deterministic. These observations imply that the closure of \mathcal{G} under composition contains at most $2^{O(n \log n)}$ SCGs. This reduces the worst-case complexity of the LJB algorithm to $2^{O(n \log n)}$. In the presence of this property, the massive gap between Ramsey-based algorithms and rank-based algorithms vanishes, helping to explain the surprising strength of the LJB algorithm.

Lemma 5.1. *When operating on reverse-deterministic SCT problems, the LJB algorithm has a worst-case complexity of $2^{O(n \log n)}$.*

Proof: A reverse-deterministic SCT problem contains only reverse-deterministic SCGs. Observe that the composition of two reverse-deterministic SCGs is itself reverse-deterministic. As there are only $2^{O(n \log n)}$ possible reverse-deterministic SCGs, the closure computed in the LJB algorithm cannot become larger than $2^{O(n \log n)}$. The LJB algorithm checks each graph in the closure exactly once, and so has a time complexity of $2^{O(n \log n)}$. \square

It is not a coincidence that all SCT problems considered possess this property. As noted in [4], straightforward analysis of functional programs generates only reverse-deterministic problems. In fact, every tool we examined is only capable of producing reverse-deterministic SCT problems. To illuminate the reason for this, imagine a SCG $G : f \rightarrow g$ where f has two parameters, x and y , and g the single parameter z . If G is not reverse deterministic, this implies both x and y have arcs, labeled with either 0 or 1, to z . This would mean that z 's value is both *always* smaller than or equal to x and *always* smaller than or equal to y .

The program in Algorithm 4 can produce non-reverse-deterministic size-change graphs, and serves to demonstrate the difficult analysis required to do so⁴. Consider the SCG for the call on line 2. It is clear there should be a 0-labeled arc from X to X . To reach this point, however, we must satisfy the inequality on line 1. Therefore we can also assert that $Y > X$, and include a 1-labeled arc from Y to X . This is a kind of analysis is difficult to make, and none of the size-change analyzers we examined were capable of detecting this relation.

⁴This example emerges from the Terminweb experiments by Mike Codish, and was translated into a functional language by Amir Ben-Amram and Chin Soon Lee. The authors are grateful to Amir Ben-Amram for bringing this illustrative example to our attention.

5.1. Reverse Determinism and Rank-Based Containment. Since the Ramsey-based approach benefited so strongly from reverse-determinism, we examine the rank-based approach to see if it can similarly benefit. As a first step, we demonstrate that reverse-deterministic automata have a maximum rank of 2, dramatically lowering the complexity of complementation to $2^{O(n)}$. We note, however, that given a reverse-deterministic SCT problem L , the automaton $\mathcal{A}_{Desc(L)}$ is *not* reverse-deterministic. Thus a separate proof is provided to demonstrate that the rank of the resulting automata is still bounded by 2.

An automaton is *reverse-deterministic* when no state has two incoming arcs labeled with the same character. Formally, an automaton is reverse-deterministic when, for each state q and character a , there is at most one state p such that $q \in \rho(p, a)$. As a corollary to Lemma 5.1, the Ramsey-based complementation construction has a worst-case complexity of $2^{O(n \log n)}$ for reverse deterministic automata. With reverse-deterministic automata, we do not have to worry about multiple paths to a state. As a consequence, a maximum rank of 2, rather than $2n - 2$, suffices to prove termination of every path, and the worst-case bound of the rank-based construction improves to $2^{O(n)}$.

Theorem 5.1. Given a reverse-deterministic Büchi automaton \mathcal{B} with n states, there exists an automaton \mathcal{B}' with $2^{O(n)}$ states such that $L(\mathcal{B}') = \overline{L(\mathcal{B})}$.

Proof: In a run DAG G_w of a reverse-deterministic automaton, all nodes have only one predecessor. This implies the run DAG is a tree, and that the number of infinite paths grows monotonically and at some point stabilizes. Call this point k . If G_w is rejecting, we demonstrate that there is a point $j \geq k$ past which all accepting states are finite in G_w . Observe that each infinite path eventually stops visiting accepting states. Let j be the last such point over all infinite paths, or k , whichever is greater. Past j , consider a branch off this path containing an accepting state. This branch cannot be a new infinite path, as the number of infinite paths is stable. This branch cannot lead to an existing infinite path, because that would violate reverse determinism. Therefore this path must be finite, and the accepting state is finite.

Recall that $G_w(0)$ is G_w , $G_w(1)$ is $G_w(0)$ with all finite nodes removed, $G_w(2)$ is $G_w(1)$ with all F -free nodes removed, and $G_w(3)$ is $G_w(2)$ with all finite nodes removed. Because there are no infinite accepting nodes past j , $G_w(1)$ has no accepting nodes at all past j . Thus every node past j is F -free in $G_w(1)$, and $G_w(2)$ has no nodes past j . Thus $G_w(3)$ is empty, and the DAG has a rank of at most 2. We conclude that the maximum rank of rejecting run DAG is 2, and the state space of the automaton in Definition 2.8 can be restricted to level rankings with no ranking larger than 2. \square

Unfortunately, neither the reduction of Definition 2.12 nor the reduction of Definition 4.1 preserve reverse determinism, which is to say that given a reverse-deterministic SCT problem, they do not produce a reverse-deterministic Büchi containment problem. However, we can show that, given a reverse-deterministic SCT problem, the automata produced by Definition 4.1 does have a maximum rank of 2. A similar claim could be made about Definition 2.12 with minor adjustments.

Formally, we prove that for every reverse-deterministic SCT problem L , $\mathcal{A}'_{Desc(L)}$ has a maximum rank of 2. Let w be an infinite word $c_1 c_2 \dots$ not in $L(\mathcal{A}'_{Desc(L)})$, and G_w the rejecting run DAG of $\mathcal{A}'_{Desc(L)}$ on w . There are two kinds of states in $\mathcal{A}'_{Desc(L)}$. There is a waiting state, q_0 , which always transitions to itself, and there are two states for every variable position $h \in 1..n$, $\langle h, 0 \rangle$ and $\langle h, 1 \rangle$. Every state is an initial state. Consider G_w , the




Figure 11: An overapproximation of the run DAG for $\mathcal{A}'_{Desc(L)}$, where the maximum arity of functions in L is 2. For clarity, this figure includes unreachable states. By definition, however, the DAG has only nodes that can be reached from a node on the first level.

run DAG of $\mathcal{A}'_{Desc(L)}$ on a word $c_0c_1\dots$. Each character c_i represents a function call from some function f_i to another function f_{i+1} . At level i of the run DAG, the waiting state has outgoing edges to itself and the positions of 0-labeled parameters of f_{i+1} . Each variable state only has outgoing edges to a 0 or 1-labeled position. To get an idea of what the run DAG looks like, Figure 11 displays a supergraph of the run DAG that includes all states at all levels, even if they are not reachable.

We now prove that the rejecting run DAG G_w has a maximum rank of 2. To do so, we analyze the structure of G_w by first examining subgraphs, and then extending these observations to G_w . Let G'_w be the subgraph of the run DAG that omits the waiting state q_0 at every level of the run DAG. Every path in G'_w corresponds to a (possibly finite) thread in the call sequence $c_1c_2\dots$

Lemma 5.2. *For a level i , $f \in F$, and $x \in P(f)$, at most one of $(\langle Ar(x), 0 \rangle, i)$ and $(\langle Ar(x), 1 \rangle, i)$ has incoming edges in G'_w .*

Proof: For $i = 0$, this holds trivially. For $i > 0$, take a pair of nodes $(\langle h, 0 \rangle, i)$ and $(\langle h, 1 \rangle, i)$. The edges from level $i - 1$ correspond to transitions in $\mathcal{A}'_{Desc(L)}$ on some call c . As c is a call to a single function, we know there is a unique variable x such that $h = Ar(x)$. Because L is reverse deterministic, we know that there is at most one edge in \mathcal{G}_c leading to x . If there is no edge, then there are no edges entering $(\langle h, 0 \rangle, i)$ or $(\langle h, 1 \rangle, i)$.

Otherwise there is exactly one edge in \mathcal{G}_c , $y \xrightarrow{c} x$, $r \in \{0, 1\}$. In this case, the only nodes in level G'_w with an edge to either $(\langle h, 0 \rangle, i)$ or $(\langle h, 1 \rangle, i)$ are $(\langle Ar(y), 0 \rangle, i - 1)$ and $(\langle Ar(y), 1 \rangle, i - 1)$. By the transition function ρ_D , both of these states transition only to $(\langle h, r \rangle, i)$, and only $(\langle h, r \rangle, i)$ has incoming edges in G'_w . \square

Now, define G_w'' to be the subgraph of G_w' containing only nodes with an incoming edge in G_w' . This removes nodes whose only incoming edge was from q_0 . While this excludes nodes that begin threads, this cannot change the accepting or rejecting nature of a thread.

Lemma 5.3.

- (1) G_w'' is a forest.
- (2) Every infinite path in G_w' appears in G_w'' .
- (3) Every accepting node in G_w is also in G_w'' .

Proof:

(1): Lemma 5.2 implies, for every h and i , that only one of $(\langle h, 1 \rangle, i)$ and $(\langle h, 0 \rangle, i)$ is in G_w'' . Combined with the fact that L is reverse deterministic, every node in G_w'' can have at most one incoming edge, and thus it is a forest.

(2): For every infinite path in G_w' , all nodes past the first have an incoming edge from G_w' . Every node with an incoming edge from G_w' is in G_w'' . Thus for every infinite path in G_w' , a corresponding path, perhaps without the first node, occurs in G_w'' .

(3): Only nodes of the form $(\langle h, 1 \rangle, i)$ are accepting. Let $(\langle h, 1 \rangle, i)$ be an accepting node. By the definition of a run DAG, $(\langle h, 1 \rangle, i)$ must be reachable if it is in G_w . Thus there is an edge from another node in G_w to $(\langle h, 1 \rangle, i)$. By the transition function, the waiting state $(q_0, i - 1)$ only has an edge to $(\langle h, 0 \rangle, i)$. Therefore the only nodes that have an edge to $(\langle h, 1 \rangle, i)$ are nodes of the form $(\langle h', r \rangle, i - 1)$. All nodes of this form are in G_w' , and therefore $(\langle h, 1 \rangle, i)$ has an incoming edge from G_w' , and is in G_w'' . \square

We can now make observations about the rejecting run DAGs of $\mathcal{A}'_{Desc(L)}$ that mirror those made about rejecting run DAGs of reverse deterministic automata.

Lemma 5.4. *There exists $j \in \mathbb{N}$ where for all $i > j$, G_w'' has no infinite accepting nodes at level i .*

Proof: As G_w'' is a forest, the number of infinite nodes in level $i + 1$ cannot be smaller than the number of infinite nodes in level i . Thus at some level k the number of infinite nodes reaches a maximum. Past level k , each infinite node has a unique infinite path through G_w'' . As G_w is rejecting, every infinite path eventually stops visiting accepting nodes at some level. Let j be the last such point over all infinite paths, or k , whichever is greater. Past j , consider an accepting node v that branch off an infinite path. This branch cannot be part of the existing infinite path, as this path has ceased visiting accepting nodes. Likewise, this branch cannot be part of a new infinite path, as the number of infinite paths can not increase. Therefore v must be finite. \square

Lemma 5.5. $G_w(3)$ is empty.

Proof: Let j be the level past which there are no infinite accepting nodes in G_w'' , as per Lemma 5.4. This precisely means that, past j , every accepting node in G_w'' has a finite path. As all accepting nodes in G_w are in G_w'' , past j every reachable accepting node in G_w has a finite path. After level j , $G_w(1)$ contains only non-accepting nodes. This implies that $G_w(2)$ contains no nodes past j , and therefore that $G_w(3)$ is empty. \square

Theorem 5.2. Given a reverse-deterministic SCT problem L with maximum arity n , there is an automaton \mathcal{B}' with at most $2^{O(n)}$ states such that $L(\mathcal{B}') = \overline{L(\mathcal{A}'_{Desc(L)})}$.

Proof: By Lemma 5.5, we know that every rejecting run DAG of $\mathcal{A}'_{Desc(L)}$ has a maximum rank of 2. Therefore it suffices to restrict the rank in Definition 2.8 to 2, replacing all occurrences of $2n - 2$ in the definition with 2. The resulting automata is of size $O(2^{3n})$. \square

5.2. Experiments revisited. In light of this discovery, we revisit the experiments and again compare rank and Ramsey-based approaches on SCT problems. This time we tell Mh, the rank-based solver, that the problems have a maximum rank of 2. Table 2 compares the running time of Mh and sct/scp on the five most difficult problems. As before, time taken to reduce SCT problems to automata containment problems was not counted.

Problem	Mh (s)	sct/scp (s)
ex04	0.01	1.39
ex05	0.13	227.7
ms	0.1	0.02
gexgcd	0.39	0.023
graphcolour2	0.044	0.014

Table 2: SCT problem completion time times by tool, exploiting reverse-determinism.

While our problem space is small, the theoretical worst-case bounds of Ramsey and rank-based approach appears to be reflected in the table. The Ramsey-based sct/scp completes some problems more quickly, but in the worst cases of ex04 and ex05, sct/scp performs significantly more slowly than Mh. It is worth noting, however, that the benefits of reverse-determinism on Ramsey-based approaches emerges automatically, while rank-based approaches must explicitly test for this property in order to exploit it.

5.3. Monotonicity Constraints: Termination Problems Lacking Reverse-Determinism. Monotonicity constraints [8] are a generalization of size-change graphs. While an SCG for a call from f to g is bipartite, with edges only from variables of f to variables of g , monotonicity constraints allow edges between any two variables, even of the same function. In addition, while SCGs only have edges representing *less than* and *less than or equal* relations, monotonicity constraints allow edges representing *equality* relations. A collection of monotonicity constraints is called a monotonicity constraint system (MCS). For a formal presentation, please see [3].

Deciding termination for MCS problems is more involved than for SCT problems, but correctness similarly relies on Ramsey’s Theorem [8]. One method is to reduce a MCS to an SCT problem through *elaboration* [3]. Unfortunately, elaboration is an exponential reduction, and increases the size of the MCS. Alternatively, it is possible to project an individual monotonicity constraint into an SCG in a lossy fashion. To do so, simply remove all edges that are not from a variable of f to a variable of g , and replace equality edges with less-than-or-equal edges. By projecting every monotonicity constraint in an MCS down to a SCG, we obtain a SCT problem. Doing so, however, often removes valuable information that can still be encoded in a size-change graph. To preserve this information, new arcs

that are logically implied by existing arcs can be added to the monotonicity constraint before the constraint is projected to a SCG. The simplest implied arcs are those derived from equality edges: given two arcs $x \xrightarrow{a} y$ and $x \xrightarrow{b} x'$, add $x' \xrightarrow{a} y$ to the monotonicity constraint. Similarly, given an arc $y \xrightarrow{b} y'$, add $x \xrightarrow{a} y'$. More complex implied arcs can be computed by similarly composing other arcs.

We obtained a corpus of 373 monotonicity constraint systems from [7]. In each case, we produced three SCT problems from each MCS: one from directly projecting, one by computing arcs implied by equality before projecting, and one by computing all implied arcs before projecting. We again defined a problem to be interesting if either sct/scp or Mh took more than 1 second to solve the problem. For every interesting problem, there was no difference in result and no significant difference in running time between the two types of implied arcs. Thus we consider only the third, most complex, SCT problem generated from each MC problem, resulting in nine final problems.

None of the interesting SCT problem produced in this fashion were reverse deterministic. Given the complexity of monotonicity constraints, this is perhaps unsurprising. Four of the resulting problems were non-terminating. For these problems, the maximum rank can be computed. To do so, Mh is initially limited to a rank of 1, and the rank is increased until Mh can detect non-termination. Table 3 displays the results for these problems. Despite the lack of reverse-determinism, none of these problems proved difficult for sct/scp: consuming at most 0.4 seconds. However, several were difficult for Mh, including one that took over eight minutes. In cases where we could bound the rank, the running time for Mh often improved dramatically. While we again have only a sparse corpus of interesting problems, these results serve to emphasize the importance of reverse determinism. Perhaps more interestingly, they suggest that, even in cases where reverse determinism does not hold, the Ramsey-based approach performs well.

Problem	rank	Mh (s)	sct/scp (s)
Test3	N/A	4.44	0.047
Test4	N/A	4.65	0.079
Test5	N/A	111.8	0.074
Test6	N/A	482.0	0.097
WorkingSignals	13	1.32 (1.0)	0.098
Gauss	3	1.10 (0.08)	0.146
PartitionList	3	1.38 (0.22)	0.081
Sudoku	5	7.18 (2.42)	0.405

Table 3: MC problem, maximum rank for non-terminating problems, and completion times by tool. Times for Mh in parenthesis are times when given the maximum rank, as if it were precomputed.

6. CONCLUSION

In this paper we demonstrate that the Ramsey-based size-change termination algorithm proposed by Lee, Jones, and Ben-Amram [24] is a specialized realization of the 1987 Ramsey-based complementation construction [5, 29]. With this link established, we compare rank-based and Ramsey-based tools on the domain of SCT problems. Initial experimentation

revealed a surprising competitiveness of the Ramsey-based tools, and led us to further investigation. We discover that SCT problems are naturally reverse-deterministic, reducing the complexity of the Ramsey-based approach. By exploiting reverse determinism, we were able to demonstrate the superiority of the rank-based approach.

Our initial test space of SCT problems was unfortunately small, with only five interesting problems emerging. Despite the very sparse space of problem, they still yielded two interesting observations. First, subsumption appears to be critical to the performance of Büchi complementation tools using both rank and Ramsey-based algorithms. It has already been established that rank-based tools benefit strongly from the use of subsumption [9]. Our results demonstrate that Ramsey-based tools also benefit from subsumption, and in fact experiments with removing subsumption from `sct/scp` seem to limit its scalability. Second, by exploiting reverse determinism, we can dramatically improve the performance of both rank and Ramsey-based approaches to containment checking.

Reverse determinism, however, is not the whole story in comparing the rank and Ramsey based approaches. On a separate corpus of problems derived from Monotonicity Constraints, which are not reverse-deterministic, the Ramsey-based approach outperformed the rank-based approach in every interesting case. It should be noted that, in addition to reverse determinism, there are several ways to achieve a better bound on the maximum rank than $2n - 2$ [16, 18], even for problems that are not known to be non-terminating. The rank-based approach might prove more competitive if such analyses were applied before checking containment. None the less, it is clear that despite the theoretical differences in complexity, we cannot discount the Ramsey-based approach. The competitive performance of Ramsey-based solutions remains intriguing.

In [9, 30], a space of random automata universality problems is used to provide a diverse problem domain. Unfortunately, it is far more complex to similarly generate a space of random SCT problems. First, universality involves a single automaton: SCT problems check the containment of two automata, with a corresponding increase in parameters. Worse, there is no reason to expect that one random automaton will have any probability of containing another random automaton. Sampling this problem space is further complicated by the low transition density of reverse-deterministic problems: in [9, 30] the most interesting problems had a transition density of 2.

On the theoretical side, we have extended the subsumption relation present in `sct/scp`. Recent work has extended the subsumption relation to the double-graph search of Algorithm `DoubleGraphSearch`, and others have improved the relation through the use of simulation [2, 14]. Doing so has enabled us to compare Ramsey and rank-based approaches on the domain of random universality problems [14], with promising results. Future work will investigate how to generate an interesting space of random containment problems, addressing the concerns raised above.

The effects of reverse-determinism on the complementation of automata bear further study. Reverse-determinism is not an obscure property, it is known that automata derived from LTL formula are often reverse-deterministic [12]. As noted above, both rank and Ramsey-based approaches improves exponentially when operating on reverse-deterministic automata. Further, Ben-Amram and Lee have defined SCP, a polynomial-time approximation algorithm for SCT [4]. For a wide subset of SCT problems with restricted in degrees, including the set used in this paper, SCP is exact. In terms of automata, this property is similar, although perhaps not identical, to reverse-determinism. The presence of an exact polynomial algorithm for the SCT case suggests a interesting subset of Büchi containment

problems may be solvable in polynomial time. The first step in this direction would be to determine what properties a containment problem must have to be solved in this fashion.

REFERENCES

- [1] Daedalus. Available on: <http://www.di.ens.fr/~cousot/projects/DAEDALUS/index.shtml>.
- [2] P.A. Abdulla, Y.-F. Chen, L. Holík, R. Mayr, and T. Vojnar. When simulation meets antichains. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science*, pages 158–174. Springer, 2010.
- [3] A.M. Ben-Amram. Size-change termination, monotonicity constraints and ranking functions. *Logical Methods in Computer Science*, 6(3), 2010.
- [4] A.M. Ben-Amram and C.S. Lee. Program termination analysis in polynomial time. *ACM Trans. Program. Lang. Syst.*, 29:5:1–5:37, January 2007.
- [5] J.R. Büchi. On a decision method in restricted second order arithmetic. In *Proc. Int. Congress on Logic, Method, and Philosophy of Science. 1960*, pages 1–12. Stanford University Press, 1962.
- [6] Y. Choueka. Theories of automata on ω -tapes: A simplified approach. *Journal of Computer and Systems Science*, 8:117–141, 1974.
- [7] M. Codish, I. Gonopolskiy, A.M. Ben-Amram, C. Fuhs, and J. Giesl. SAT-based termination analysis using monotonicity constraints over the integers. *Theory and Practice of Logic Programming, 26th Int'l. Conference on Logic Programming (ICLP'11) Special Issue*, 11:503–520, July 2011.
- [8] M. Codish, V. Lagoon, and P.J. Stuckey. Testing for termination with monotonicity constraints. In *Twenty First International Conference on Logic Programming*, volume 3668 of *Lecture Notes in Computer Science*, pages 326–340. Springer-Verlag, October 2005.
- [9] L. Doyen and J.-F. Raskin. Improved algorithms for the automata-based approach to model-checking. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 4424 of *Lecture Notes in Computer Science*, pages 451–465. Springer, 2007.
- [10] L. Doyen and J.-F. Raskin. Antichains for the automata-based approach to model-checking. *Logical Methods in Computer Science*, 5(1), 2009.
- [11] E.A. Emerson and C.-L. Lei. Efficient model checking in fragments of the propositional μ -calculus. In *Proc. 1st IEEE Symp. on Logic in Computer Science*, pages 267–278, 1986.
- [12] E.A. Emerson and A.P. Sistla. Deciding full branching time logics. *Information and Control*, 61(3):175–201, 1984.
- [13] S. Fogarty. Büchi containment and size-change termination. Master's thesis, Rice University, 2008.
- [14] S. Fogarty and M.Y. Vardi. Efficient Büchi universality checking. In *Tools and Algorithms for the Construction and Analysis of Systems*, volume 6015 of *Lecture Notes in Computer Science*, pages 205–220. Springer, 2010.
- [15] C.C. Frederiksen. A simple implementation of the size-change termination principle. Available at: <ftp://ftp.diku.dk/diku/semantics/papers/D-442.ps.gz>, 2001.
- [16] E. Friedgut, O. Kupferman, and M.Y. Vardi. Büchi complementation made tighter. In *2nd Int. Symp. on Automated Technology for Verification and Analysis*, volume 3299 of *Lecture Notes in Computer Science*, pages 64–78. Springer, 2004.
- [17] Arne J. Glenstrup. Terminator II: Stopping partial evaluation of fully recursive programs. Master's thesis, DIKU, University of Copenhagen, June 1999.
- [18] S. Gurumurthy, O. Kupferman, F. Somenzi, and M.Y. Vardi. On complementing nondeterministic Büchi automata. In *Proc. 12th Conf. on Correct Hardware Design and Verification Methods*, volume 2860 of *Lecture Notes in Computer Science*, pages 96–110. Springer, 2003.
- [19] N.D Jones and Nina Bohr. Termination analysis of the untyped λ -calculus. In *Rewriting Techniques and Applications. Proceedings*, volume 3091 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2004.
- [20] N. Klarlund. *Progress Measures and finite arguments for infinite computations*. PhD thesis, Cornell University, 1990.
- [21] O. Kupferman and M.Y. Vardi. Weak alternating automata are not that weak. In *Proc. 5th Israeli Symp. on Theory of Computing and Systems*, pages 147–158. IEEE Computer Society Press, 1997.
- [22] O. Kupferman and M.Y. Vardi. Synthesizing distributed systems. In *Proc. 16th IEEE Symp. on Logic in Computer Science*, pages 389–398, 2001.

- [23] O. Kupferman and M.Y. Vardi. Weak alternating automata are not that weak. *ACM Transactions on Computational Logic*, 2(2):408–429, 2001.
- [24] C.S. Lee, N.D. Jones, and A.M. Ben-Amram. The size-change principle for program termination. In *Proc. 28th ACM Symp. on Principles of Programming Languages*, pages 81–92, 2001.
- [25] M. Michel. Complementation is more difficult with automata on infinite words. CNET, Paris, 1988.
- [26] S. Safra. On the complexity of ω -automata. In *Proc. 29th IEEE Symp. on Foundations of Computer Science*, pages 319–327, 1988.
- [27] S. Schewe. Büchi complementation made tight. In *26th Int. Symp. on Theoretical Aspects of Computer Science*, volume 3, pages 661–672. Schloss Dagstuhl, 2009.
- [28] D. Sereni and N.D. Jones. Termination analysis of higher-order functional programs. In *3rd Asian Symp. on Programming Languages and Systems*, volume 3780 of *Lecture Notes in Computer Science*, pages 281–297. Springer, 2005.
- [29] A.P. Sistla, M.Y. Vardi, and P. Wolper. The complementation problem for Büchi automata with applications to temporal logic. In *Proc. 12th Int. Colloq. on Automata, Languages, and Programming*, volume 194, pages 465–474. Springer, 1985.
- [30] D. Tabakov and M.Y. Vardi. Experimental evaluation of classical automata constructions. In *Proc. 12th Int'l Conf. on Logic for Programming, Artificial Intelligence, and Reasoning*, Lecture Notes in Computer Science 3835, pages 396–411. Springer, 2005.
- [31] M.Y. Vardi. Automata-theoretic model checking revisited. In *Proc. 8th Int. Conf. on Verification, Model Checking, and Abstract Interpretation*, volume 4349 of *Lecture Notes in Computer Science*, pages 137–150. Springer, 2007.
- [32] David Wahlstedt. Detecting termination using size-change in parameter values. Master's thesis, Göteborgs Universitet, 2000.