

Reachability in Two-Clock Timed Automata is PSPACE-complete

John Fearnley¹ and Marcin Jurdziński²

¹ Department of Computer Science, University of Liverpool, UK

² Department of Computer Science, University of Warwick, UK

Abstract. A recent result has shown that reachability in two-clock timed automata is log-space equivalent to reachability in bounded one-counter automata [6]. We show that reachability in bounded one-counter automata is PSPACE-complete.

1 Introduction

Timed automata [1] are a successful and widely used formalism, which are used in the analysis and verification of real time systems. A timed automaton is a non-deterministic finite automaton that is equipped with a number of real-valued *clocks*, which allow the automaton to measure the passage of time.

Perhaps the most fundamental problem for timed automata is the *reachability* problem: given an initial state, can the automaton perform a sequence of transitions in order to reach a specified target state? In their seminal paper on timed automata [1], Alur and Dill showed that this problem is PSPACE-complete. To show hardness for PSPACE, their proof starts with a linear bounded automaton (LBA), which is a non-deterministic Turing machine with a fixed tape length n . They produce a timed automaton with $2n + 1$ clocks, and showed that the timed automaton can reach a specified state if and only if the LBA halts.

However, the work of Alur and Dill did not address the case where the number of clocks is small. This was rectified by Courcoubetis and Yannakakis [4], who showed that reachability in timed automata with only three clocks is still PSPACE-complete. Their proof cleverly encodes the tape of an LBA in a single clock, and then uses the two additional clocks to perform all necessary operations on the encoded tape. In contrast to this, Laroussinie et al. have shown that reachability in one-clock timed automata is complete for NLOGSPACE, and therefore no more difficult than computing reachability in directed graphs [7].

The complexity of reachability in two-clock timed automata has been left open. The best known lower bound was given by Laroussinie et al., who gave a proof that the problem is NP-hard via a very natural reduction from subset-sum [7]. Moreover, the problem lies in PSPACE, because reachability in two-clock timed automata is obviously easier than reachability in three-clock timed automata. However, the PSPACE-hardness proof of Courcoubetis and Yannakakis seems to fundamentally require three clocks, and does not naturally extend to the two-clock case. Naves [8] has shown that several extensions to two-clock

timed automata lead to PSPACE-completeness, but his work does not advance upon the NP-hard lower bound for unextended two-clock timed automata.

In a recent paper, Haase et al. have shown a link between reachability in timed automata and reachability in *bounded counter automata* [6]. A bounded counter automaton is a non-deterministic finite automaton equipped with a set of counters, and the transitions of the automaton may add or subtract arbitrary integer constants to the counters. The state space of each counter is bounded by some natural number b , so the counter may only take values in the range $[0, b]$. Moreover, transitions may only be taken if they do not increase or decrease a counter beyond the allowable bounds. This gives these seemingly simple automata a surprising amount of power, because the bounds can be used to implement inequality tests against the counters.

Haase et al. show that reachability in two-clock timed automata is log-space equivalent to reachability in bounded *one-counter automata*. Reachability in bounded one-counter automata has also been studied in the context of one-clock timed automata with energy constraints [2], where it was shown that the problem lies in PSPACE, and is NP-hard. It has also been shown that the reachability problem for *unbounded* one-counter automata is NP-complete [5], but the NP containment proof does not seem to generalise to bounded one-counter automata.

Our contribution. We show that satisfiability for quantified boolean formulas can be reduced, in polynomial time, to reachability in bounded one-counter automata. Hence, we show that reachability in bounded one-counter automata is PSPACE-complete, and therefore we resolve the complexity of reachability in two-clock timed automata. Our reduction uses two intermediate steps: *subset-sum games* and *bounded counter-stack automata*.

Counter automata are naturally suited for solving subset-sum problems, so our reduction starts with a quantified version of subset-sum, which we call subset-sum games. One interpretation of satisfiability for quantified boolean formulas is to view the problem as a game between an *existential* player, and a *universal* player. The players take in turns to set their propositions to true or false, and the existential player wins if and only if the boolean formula is satisfied. Subset-sum games follow the same pattern, but apply it to subset-sum: the two players alternate in choosing numbers from sets, and the existential player wins if and only if the chosen numbers sum to a given target. Previous work by Travers can be applied to show that subset-sum games are PSPACE-complete [9].

We reduce subset-sum games to reachability in bounded one-counter automata. However, we will not do this directly. Instead, we introduce bounded counter-stack automata, which are able to store multiple counters, but have a stack-like restriction on how these counters may be accessed. These automata are a convenient intermediate step, because having access to multiple counters makes it easier for us to implement subset-sum games. Moreover, the stack based restrictions means that it is relatively straightforward to show that reachability in bounded counter-stack automata is reducible, in polynomial time, to reachability in bounded one-counter automata, which completes our result.

2 Bounded one-counter automata

A bounded one-counter automaton has a single counter that can store values between 0 and some bound $b \in \mathbb{N}$. The automaton may add or subtract values from the counter, so long as the bounds of 0 and b are not overstepped. This can be used to test inequalities against the counter. For example, to test whether the counter is larger than some $n \in \mathbb{N}$, we first attempt to subtract $n + 1$ from the counter, then, if that works, we add $n + 1$ back to the counter. This creates a sequence of two transitions which can be taken if, and only if, the counter is greater than n . A similar construction can be given for less-than tests. For the sake of convenience, we will include explicit inequality testing in our formal definition, with the understanding that this is not actually necessary.

We now give a formal definition. For two integers $a, b \in \mathbb{Z}$ we define $[a, b] = \{n \in \mathbb{Z} : a \leq n \leq b\}$ to be the subset of integers between a and b . A bounded one-counter automaton is defined by a tuple (L, b, Δ, l_0) , where L is a finite set of locations, $b \in \mathbb{N}$ is a global counter bound, Δ specifies the set of transitions, and $l_0 \in L$ is the initial location. Each transition in Δ has the form (l, p, g_1, g_2, l') , where l and l' are locations, $p \in [-b, b]$ specifies how the counter should be modified, and $g_1, g_2 \in [0, b]$ give lower and upper guards for the counter.

Each state of the automaton consists of a location $l \in L$ along with a counter value c . Thus, we define the set of states to be $L \times [0, b]$. A transition exists between a state $(l, c) \in S$, and a state $(l', c') \in S$ if there is a transition $(l, p, g_1, g_2, l') \in \Delta$, where $g_1 \leq c \leq g_2$, and $c' = c + p$.

The reachability problem for bounded one-counter automaton is: starting at the state $(l_0, 0)$, can the automaton reach a specified target state (l_t, c_t) ? It has been shown that the reachability problem for bounded one-counter automata is equivalent to the reachability problem for two-clock timed automata.

Theorem 1 ([6]). *Reachability in bounded one-counter automata is log-space equivalent to reachability in two-clock timed automata.*

3 Subset-sum games

A subset-sum game is played between an *existential* player and a *universal* player. The game is specified by a pair (ψ, T) , where $T \in \mathbb{N}$, and ψ is a list:

$$\forall \{A_1, B_1\} \exists \{E_1, F_1\} \dots \forall \{A_n, B_n\} \exists \{E_n, F_n\},$$

where A_i, B_i, E_i , and F_i , are all natural numbers.

The game is played in rounds. In the first round, the universal player chooses an element from $\{A_1, B_1\}$, and the existential player responds by choosing an element from $\{E_1, F_1\}$. In the second round, the universal player chooses an element from $\{A_2, B_2\}$, and existential player responds by choosing an element from $\{E_2, F_2\}$. This pattern repeats for rounds 3 through n . Thus, at the end of the game, the players will have constructed a sequence of numbers, and the existential player wins if and only if the sum of these numbers is T .

Formally, the set of *plays* of the game is the set:

$$\mathcal{P} = \prod_{1 \leq j \leq n} \{A_j, B_j\} \times \{E_j, F_j\}.$$

A play $P \in \mathcal{P}$ is winning for the existential player if and only if $\sum P = T$.

A strategy for the existential player is a list of functions $\mathbf{s} = (s_1, s_2, \dots, s_n)$, where each function s_i dictates how the existential player should play in the i th round of the game. Thus, each function s_i is of the form:

$$s_i : \prod_{1 \leq j \leq i} \{A_j, B_j\} \rightarrow \{E_i, F_i\}.$$

This means that the function s_i maps the first i moves of the universal player to a decision for the existential player in the i th round.

A play P conforms to a strategy \mathbf{s} if the decisions made by the existential player in P always agree with \mathbf{s} . More formally, for each i in the range $1 \leq i \leq n$, we define $F_i = P \cap \prod_{1 \leq j \leq i} \{A_j, B_j\}$ to be the first i moves made by the universal player. The play P conforms to a strategy $\mathbf{s} = (s_1, s_2, \dots, s_k)$ if $s_i(F_i) \in P$, for all i . Given a strategy \mathbf{s} , we define the set of conforming plays to be $\text{Plays}(\mathbf{s})$. Note that, since the universal player makes exactly n choices, the set $\text{Plays}(\mathbf{s})$ contains exactly 2^n different plays.

A strategy \mathbf{s} is *winning* if every play $P \in \text{Plays}(\mathbf{s})$ is winning for the existential player. The *subset-sum game problem* is to decide, for a given SSG instance (ψ, T) , whether the existential player has a winning strategy for (ψ, T) .

The SSG problem clearly lies in PSPACE, because it can be solved on a polynomial time alternating Turing machine. A quantified version of subset-sum has been shown to be PSPACE-hard, via a reduction from quantified boolean formulas [9]. Since SSGs are essentially a quantified version of subset-sum, the proof of PSPACE-hardness easily carries over. See Appendix A for further details.

Lemma 2. *The subset-sum game problem is PSPACE-complete.*

4 Counter-Stack Automata

Outline. In this section we ask: can we use a bounded one-counter automaton to store multiple counters? The answer is yes, but doing so forces an interesting set of restrictions on the way in which the counters are accessed. By the end of this section, we will have formalised these restrictions as *counter-stack* automata.

Suppose that we have a bounded-one counter automaton with counter c and bound $b = 15$. Hence, the width of the counter is 4 bits. Now suppose that we wish to store two 2-bit counters c_1 and c_2 in c . We can do this as follows:

$$\mathbf{c} = \begin{array}{cc} \boxed{1} & \boxed{0} \\ \boxed{0} & \boxed{1} \end{array} \begin{array}{c} c_2 \\ c_1 \end{array}$$

We allocate the top two bits of c to store c_2 , and the bottom two bits to store c_1 . We can easily write to both counters: if we want to increment c_2 then we add 4 to c , and if we want to increment c_1 then we add 1 to c .

However, if we want to test equality, then things become more interesting. It is easy to test equality against c_2 : if we want to test whether $c_2 = 2$, then we test whether $8 \leq c \leq 11$ holds. But, we cannot easily test whether $c_1 = 2$ because we would have to test whether c is 2, 6, 10, or 14, and this list grows exponentially as the counters get wider. However, if we know that $c_2 = 1$, then we only need to test whether $c = 6$. Thus, we arrive at the following guiding principal: if you want to test equality against c_i , then you must know the values of c_j for all $j > i$. Counter-stack automata are a formalisation of this principal.

Counter-stack automata. A counter-stack automaton has a set of k distinct counters, which are referred to as c_1 through c_k . For our initial definitions, we will allow the counters to take all values from \mathbb{N} , but we will later refine this by defining *bounded* counter-stack automata. The defining feature of a counter-stack automaton is that the counters are arranged in a stack-like fashion:

- All counters may be increased at any time.
- c_i may only be tested for equality if the values of c_{i+1} through c_k are known.
- c_i may only be reset if the values of c_i through c_k are known.

When the automaton increases a counter, it adds a specified number $n \in \mathbb{N}$ to that counter. The automaton has the ability to perform equality tests against a counter, but the stack-based restrictions must be respected. An example of a valid equality test would be $c_k = 3 \wedge c_{k-1} = 10$, because $c_{k-1} = 10$ only needs to be tested in the case where $c_k = 3$ is known to hold. Conversely, the test $c_{k-1} = 10$ by itself is invalid, because it places no restrictions on the value of c_k .

The automaton may also reset a counter, but the stack-based restrictions apply. Counter c_i may only be reset by a transition, if that transition tests equality against the values of c_i through c_k . For example, c_{k-1} may only be reset if the transition is guarded by a test of the form $c_{k-1} = n_1 \wedge c_{k-2} = n_2$.

Formal definition. A counter-stack automaton is a tuple (L, C, Δ, l_0) , where L is a finite set of locations, $C = [1, k]$ is a set of counter indexes, $l_0 \in L$ is an initial state, and Δ specifies the transition relation. Each transition in Δ has the form (l, E, I, R, l') where:

- $l, l' \in L$ is a pair of locations,
- E is a partial function from C to \mathbb{N} which specifies the equality tests. If $E(i)$ is defined for some i , then $E(j)$ must be defined for all $j \in C$ with $j > i$.
- $I \in \mathbb{N}^k$ specifies the how the counters must be increased,
- $R \subseteq C$ specifies the set of counters that must be reset. It is required that $E(r)$ is defined for every $r \in R$.

Each state of the automaton is a location annotated with values for each of the k counters. That is, the state space of the automaton is $L \times \mathbb{N}^k$. A state $(l, c_1, c_2, \dots, c_k)$ can transition to a state $(l', c'_1, c'_2, \dots, c'_k)$ if, and only if, there exists a transition $(l, E, I, R, l') \in \Delta$, where the following conditions hold:

- For every i for which $E(i)$ is defined, we must have $c_i = E(i)$.
- For every $i \in R$, we must have $c'_i = 0$.
- For every $i \notin R$, we must have $c'_i = c_i + I_i$.

A *run* is a sequence of states s_0, s_1, \dots, s_n , where each s_i can transition to s_{i+1} . To solve the reachability problem for counter-stack automata, we must decide whether there is a run from $(l_0, 0, 0, \dots, 0)$ to a target state $(l_t, t_1, t_2, \dots, t_k)$.

A counter-stack automaton is *b-bounded*, for some $b \in \mathbb{N}$, if it is impossible for the automaton to increase a counter beyond b . Formally, this condition requires that, for every state $(l, c_1, c_2, \dots, c_k)$ that can be reached by a run from $(l_0, 0, 0, \dots, 0)$, we have $c_i \leq b$ for all i . We say that a counter-stack automaton is bounded, if it is b -bounded for some $b \in \mathbb{N}$.

Simulation by a bounded one-counter automaton. A bounded counter-stack automaton is designed to be simulated by a bounded one-counter automaton. To do this, we follow the construction outlined at the start of this section: we split the bits of the counter c into k chunks, where each chunk represents one of the counters c_i . Note that the boundedness assumption is crucial, because otherwise incrementing c_i may overflow the allotted space, and inadvertently modify the value of c_{i+1} . See Appendix B for more details of the construction.

Lemma 3. *Reachability in bounded counter-stack automata is polynomial-time reducible to reachability in bounded one-counter automata.*

5 Outline Of The Construction

Our goal is to show that reachability in bounded counter-stack automata is PSPACE-hard. To do this, we will show that subset-sum games can be solved by bounded counter-stack automata. In this section, we give an overview of our construction using the following two-round QSS game.

$$(\forall \{A_1, B_1\} \exists \{E_1, F_1\} \forall \{A_2, B_2\} \exists \{E_2, F_2\}, T).$$

For brevity, we will refer to this instance as (ψ, T) for the rest of this section. The construction is split into two parts: the *play* gadget, and the *reset* gadget.




Fig. 1. The play gadget

The play gadget. The play gadget is shown in Figure 1. The construction uses 9 counters. The locations are represented by circles, and the transitions are

represented by edges. The annotations on the transitions describe the increments, resets, and equality tests: the notation $c_i + n$ indicates that n is added to counter i , the notation $R(c_i)$ indicates that counter i is reset to 0, and the notation $c_i = n$ indicates that the transition may only be taken when $c_i = n$ is satisfied.

This gadget allows the automaton to implement a play of the SSG. The locations u_1 and u_2 allow the automaton to choose the first and second moves of the universal player, while the locations e_1 and e_2 allow the automaton to choose the first and second moves for the existential player. As the play is constructed, a running total is stored in c_9 , which is the top counter on the stack. The final transition between w_1 and w_2 checks whether the existential player wins the play, and then resets c_9 . Thus, the set of runs between u_1 and w_2 corresponds precisely to the set of plays won by the existential player in the SSG.

In addition to this, each outgoing transition from u_i or e_i comes equipped with its own counter. This counter is incremented if and only if the corresponding edge is used during the play, and this allows us to check precisely which play was chosen. These counters will be used by the reset gadget. The idea behind our construction is to force the automaton to pass through the play gadget multiple times. Each time we pass through the play gadget, we will check a different play, and our goal is to check a set of plays that verify whether the existential player has a winning strategy for the SSG.

Which plays should be checked? In our example, we must check four plays. The format of these plays is shown in Table 1.

Play	u_1	e_1	u_2	e_2
1	A_1	E_1 or F_1	A_2	E_2 or F_2
2	A_1	Unchanged	B_2	E_2 or F_2
3	B_1	E_1 or F_1	A_2	E_2 or F_2
4	B_1	Unchanged	B_2	E_2 or F_2

Table 1. The set of plays that the automaton will check.

The table shows four different plays, which cover every possible strategy choice of the universal player. Clearly, if the existential player does have a winning strategy, then that strategy should be able to win against all strategy choices of the universal player. The plays are given in a very particular order: the first two plays contain A_1 , while the second two plays contain B_1 . Moreover, we always check A_2 , before moving on to B_2 .

We want to force the decisions made at e_1 and e_2 to form a coherent strategy for the existential player. In this game, a strategy for the existential player is a pair $\mathbf{s} = (s_1, s_2)$, where s_i describes the move that should be made at e_i . It is critical to note that s_1 only knows whether A_1 or B_1 was chosen at u_1 . This restriction is shown in the table: the automaton may choose freely between E_1 and F_1 in the first play. However, in the second play, the automaton must make the same choice as it did in the first play. The same relationship holds between the third and fourth plays. These restrictions ensure that the plays shown in Table 1 are a description of a strategy for the existential player.

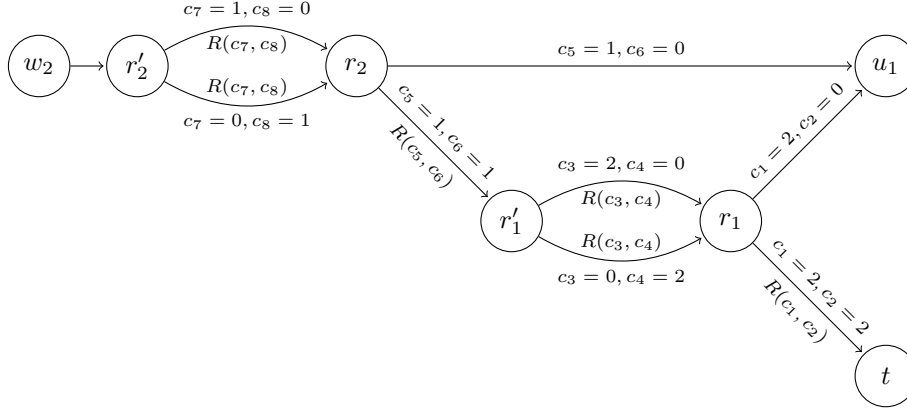


Fig. 2. The reset gadget

The reset gadget. The reset gadget, shown in Figure 2, enforces the constraints shown in Table 1. The locations w_2 and u_1 represent the same locations as they did in Figure 1. To simplify the diagram, we have only included meaningful equality tests. Whenever we omit a required equality test, it should be assumed that the counter is 0. For example, the outgoing transitions from r_2 implicitly include the requirement that c_7 , c_8 , and c_9 are all 0.

We consider the following reachability problem: can $(t, 0, 0, \dots, 0)$ be reached from $(u_1, 0, 0, \dots, 0)$? The structure of the reset gadget places restrictions on the runs that reach t . All such runs pass through the reset gadget exactly four times, and the following table describes each pass:

Pass	Path
1	$w_2 \rightarrow r'_2 \rightarrow r_2 \rightarrow u_1$
2	$w_2 \rightarrow r'_2 \rightarrow r_2 \rightarrow r'_1 \rightarrow r_1 \rightarrow u_1$
3	$w_2 \rightarrow r'_2 \rightarrow r_2 \rightarrow u_1$
4	$w_2 \rightarrow r'_2 \rightarrow r_2 \rightarrow r'_1 \rightarrow r_1 \rightarrow t$

To see why these paths must be taken, observe that, for every $i \in \{1, 3\}$, each pass through the play gadget increments either c_i or c_{i+1} , but not both. This means that the first time that we arrive at r_2 , we must take the transition directly to u_1 , because the guard on the transition to r'_1 cannot possibly be satisfied after a single pass through the play gadget. When we arrive at r_2 on the second pass, we are forced to take the transition to r'_1 , because we cannot have $c_5 = 1$ and $c_6 = 0$ after two passes through the play gadget. This transition resets both c_5 and c_6 , so the pattern can repeat again on the third and fourth visits to r_2 . The location r_1 behaves in the same way as r_2 , but the equality tests are scaled up, because r_1 is only visited on every second pass through the reset gadget.

We can now see that all strategies of the universal player must be considered. The transition between r_2 and u_1 forces the play gadget to increment c_5 , and

therefore the first and third plays must include A_2 . Similarly, the transition between r_2 and r'_1 forces the second and fourth plays to include B_2 . Meanwhile, the transition between r_1 and u_1 forces the first and second plays to include A_1 , and the transition between r_1 and t forces the third and fourth plays to include B_1 . Thus, we select the universal player strategies exactly as Table 1 prescribes.

The transitions between r'_1 and r_1 check that the existential player is playing a coherent strategy. When the automaton arrives at r'_1 during the second pass, it verifies that either E_1 was included in the first and second plays, or that F_1 was included in the first and second plays. If this is not the case, then the automaton gets stuck. The counters c_3 and c_4 are reset when moving to r_1 , which allows the same check to occur during the fourth pass. For the sake of completeness, we have included the transitions between r'_2 and r_2 , which perform the same check for E_2 and F_2 . However, since the existential player is allowed to change this decision on every pass, the automaton can never get stuck at r'_2 .

The end result is that location t can be reached if and only if the existential player has a winning strategy for (ψ, T) . As we will show in the next section, the construction extends to arbitrarily large SSGs, which then leads to a proof that reachability in counter-stack automata is PSPACE-hard. Note that all counters in this construction are bounded: c_9 is clearly bounded by the maximum value that can be achieved by a play of the SSG, and reset gadget ensures that no other counter may exceed 4. Thus, we will have completed our proof of PSPACE-hardness for bounded one-counter automata and two-clock timed automata.

6 Formal Definition and Proof

Sequential strategies for SSGs. We start by formalising the ideas behind Table 1. Recall that the table gives a strategy for the existential player in the form of a list of plays. Moreover, the table gave a very specific ordering in which these plays must appear. We now formalise this ordering.

We start by dividing the integers in the interval $[1, 2^n]$ into i -blocks. The 1-blocks partition the interval into two equally sized blocks. The first 1-block consists of the range $[1, 2^{n-1}]$, and the second 1-block consists of the range $[2^{n-1} + 1, 2^n]$. There are four 2-blocks, which partition the 1-blocks into two equally sized sub-ranges. This pattern continues until we reach the n -blocks.

Formally, for each $i \in \{1, 2, \dots, n\}$, then there are 2^i distinct i -blocks. The set of i -blocks can be generated by considering the intervals $[k + 1, k + 2^{n-i}]$ for the first 2^i numbers $k \geq 0$ that satisfy $k \bmod 2^{n-i} = 0$. An i -block is *even* if k is an even multiple of 2^{n-i} , and it is *odd* if k is an odd multiple of 2^{n-i} .

The ordering of the plays in Table 1 can be described using blocks. There are four 2-blocks, and A_2 appears only in even 2-blocks, while B_2 only appears in odd 2-blocks. Similarly, A_1 only appears in the even 1-block, while B_1 only appears in the odd 1-block. The restrictions on the existential player can also be described using blocks: the existential player's strategy may not change between E_i and F_i during a i -block. We generalise this idea in the following definition.

Definition 4 (Sequential strategy). A sequential strategy for the existential player in (ψ, T) is a list of 2^n plays $\mathcal{S} = P_1, P_2, \dots, P_{2^n}$, where for every i -block L we have:

- If L is an even i -block, then P_j must contain A_i for all $j \in L$.
- If L is an odd i -block, then P_j must contain B_i for all $j \in L$.
- We either have $E_i \in P_j$ for all $j \in L$, or we have $F_i \in P_j$ for all $j \in L$.

We say that \mathcal{S} is winning for the existential player if $\sum P_j = T$ for every $P_j \in \mathcal{S}$. Since a sequential strategy is simply a strategy written in the form of a list, we have the following lemma. See Appendix C for further details.

Lemma 5. *The existential player has a winning strategy if and only if the existential player has a sequential winning strategy.*

The base automaton. We describe the construction in two steps. Recall, from Figures 1 and 2, that the top counter is used by the play gadget to store the value of the play, and to test whether the play is winning. We begin by constructing a version of the automaton that omits the top counter. That is, if c_k is the top counter, we modify the play gadget by removing all increases to c_k , and the equality test for c_k between w_1 and w_2 . We call this the *base* automaton. Later, we will add the constraints for c_k back in, to construct the *full* automaton.

We now give a formal definition of the base automaton. Throughout this definition, we keep consistency with the location and counter names used in Figures 1 and 2. For each natural number n , we define a counter-stack automaton \mathcal{A}_n as follows. The automaton has the following set of locations

- For each $i \in [1, n]$ we have a location u_i and a location e_i .
- We have two check states w_1 and w_2 .
- For each $i \in [1, n]$ we have two reset locations r_i and r'_i .
- We have a goal location t .

The automaton uses $k = 2n + 1$ counters. The top counter c_k is reserved for the full automaton, and will not be used in this construction. We will identify counters 1 through $2n$ using the following shorthands. For each integer i , we define $a_i = c_{4(i-1)+1}$, we define $b_i = c_{4(i-1)+2}$, we define $e_i = c_{4(i-1)+3}$, and we define $f_i = c_{4(i-1)+4}$. Note that, in Figure 1, we have $a_1 = c_1$ and $a_2 = c_5$, and these are precisely the counters associated with A_1 and A_2 , respectively. The same relationship holds between b_i and B_i , and so on.

The transitions of the automaton are defined as follows. Whenever we omit a required equality test against a counter c_i , it should be assumed that the transition includes the test $c_i = 0$.

- Each location u_i has two transitions to e_i .
 - A transition that adds 1 to a_i .
 - A transition that adds 1 to b_i .
- We define u_{n+1} to be a shorthand for w_1 . Each location e_i has two transitions to u_{i+1} .

- A transition that adds 1 to e_i .
- A transition that adds 1 to f_i .
- Location w_1 has a transition to w_2 , and w_2 has a transition to r'_n . These transitions do not increase any counter, and do not test any equalities.
- Each location r'_i has two outgoing transitions to r_i .
 - A transition that tests $e_i = 2^{n-i}$ and $f_i = 0$.
 - A transition that tests $e_i = 0$ and $f_i = 2^{n-i}$.
- We define r'_0 to be shorthand for location t . Each location r_i has two outgoing transitions.
 - A transition to u_1 that tests $a_i = 2^{n-i}$ and $b_i = 0$.
 - A transition to r'_{i-1} that tests $a_i = 2^{n-i}$ and $b_i = 2^{n-i}$.

Runs in the base automaton. We now describe the set of runs are possible in the base automaton. We decompose every run of the automaton into segments, such that each segment contains a single pass through the play gadget. More formally, we decompose R into segments R_1, R_2, \dots , where each segment R_i starts at u_1 , and ends at the next visit to u_1 . We say that a run gets *stuck* if the run does not end at $(t, 0, 0, \dots, 0)$, and if the final state of the run has no outgoing transitions. We say that a run R gets stuck during an i -block L if there exists a $j \in L$ such that R_j gets stuck. The following lemma gives a characterisation of the runs in \mathcal{A}_n . See Appendix D for further details.

Lemma 6. *Let R be a run in \mathcal{A}_n . R does not get stuck if and only if, for every i -block L , all of the following hold.*

- If L is an even i -block, then R_j must increment a_i for every $j \in L$.
- If L is an odd i -block, then R_j must increment b_i for every $j \in L$.
- Either R_j increments e_i for every $j \in L$, or R_j increments f_i for every $j \in L$.

We say that a run is *successful* if it eventually reaches $(t, 0, 0, \dots, 0)$. By definition, a run is successful if and only if it never gets stuck. Also, the transition from r_1 to t ensures that every successful run must have exactly 2^n segments. With these facts in mind, if we compare Lemma 6 with Definition 4, then we can see that the set of successful runs in \mathcal{A}_n corresponds exactly to the set of sequential strategies for the existential player in the SSG.

Since we eventually want to implement \mathcal{A}_n as a bounded one-counter automaton, it is important to prove the \mathcal{A}_n is bounded. We do this in the following Lemma. See Appendix E for full details.

Lemma 7. *Along every run of \mathcal{A}_n we have that:*

- a_i and b_i are bounded by 2^{n-i+1} , and
- e_i and f_i are bounded by 2^{n-i} .

The full automaton. Let (ψ, T) be an SSG instance, where ψ is:

$$\forall \{A_1, B_1\} \exists \{E_1, F_1\} \dots \forall \{A_n, B_n\} \exists \{E_n, F_n\}.$$

We will construct a counter-stack automaton \mathcal{A}_ψ from \mathcal{A}_n . Recall that the top counter c_k is unused in \mathcal{A}_n . We modify the transitions of \mathcal{A}_n as follows. Let δ be a transition. If δ increments a_i then it also adds A_i to c_k , if δ increments b_i then it also adds B_i to c_k , if δ increments e_i then it also adds E_i to c_k , and if δ increments f_i then it also adds F_i to c_k . We also modify the transition between w_1 and w_2 , so that it checks whether $c_k = T$, and resets c_k .

Since we only add extra constraints to \mathcal{A}_n , the set of successful runs in \mathcal{A}_ψ is contained in the set of successful runs of \mathcal{A}_n . Recall that the set of successful runs in \mathcal{A}_n encodes the set of sequential strategies for the existential player in (ψ, T) . In \mathcal{A}_ψ , we simply check whether each play in the sequential strategy is winning for the existential player. Thus, we have shown the following lemma.

Lemma 8. *The set of successful runs in \mathcal{A}_ψ corresponds precisely to the set of winning sequential strategies for the existential player in (ψ, T) .*

We also have that \mathcal{A}_ψ is bounded. Counters c_1 through c_{k-1} are bounded due to Lemma 7, and counter c_k is bounded by $\sum\{A_i, B_i, E_i, F_i : 1 \leq i \leq n\}$. This completes the reduction from subset-sum games to bounded counter-stack automata, and gives us our main theorem.

Theorem 9. *Reachability in bounded counter-stack automata is PSPACE-hard.*

Corollary 10. *We have:*

- *Reachability in bounded one-counter automata is PSPACE-complete.*
- *Reachability in 2-clock timed automata is PSPACE-complete.*

References

1. R. Alur and D. L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.
2. P. Bouyer, U. Fahrenberg, K. G. Larsen, N. Markey, and J. Srba. Infinite runs in weighted timed automata with energy constraints. In *Proc. of FORMATS*, pages 33–47, 2008.
3. T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 2009.
4. C. Courcoubetis and M. Yannakakis. Minimum and maximum delay problems in real-time systems. *Formal Methods in System Design*, 1(4):385–415, 1992.
5. C. Haase, S. Kreutzer, J. Ouaknine, and J. Worrell. Reachability in succinct and parametric one-counter automata. In *Proc. of CONCUR*, pages 369–383, 2009.
6. C. Haase, J. Ouaknine, and J. Worrell. On the relationship between reachability problems in timed and counter automata. In *Proc. of RP*, pages 54–65, 2012.
7. F. Laroussinie, N. Markey, and P. Schnoebelen. Model checking timed automata with one or two clocks. In *Proc. of CONCUR*, pages 387–401, 2004.
8. G. Naves. Accessibilité dans les automates temporisés à deux horloges. Rapport de Master, MPRI, Paris, France, 2006.
9. S. Travers. The complexity of membership problems for circuits over sets of integers. *Theoretical Computer Science*, 369(13):211–229, 2006.

A Proof of Lemma 2

Outline. A quantified version of subset-sum has already been shown to be PSPACE-hard [9], and the proof easily carries over for the case of SSGs. For the sake of completeness, we provide a direct proof that SSGs are PSPACE-hard, which closely follows the ideas laid out in [9].

The proof follows the NP-hardness proof for subset-sum, taken from [3][Theorem 34.10]. The key observation is that, if we begin with a quantified version of 3-SAT, then we end up with an SSG.

Subset-sum is NP-hard. We now give a summary of the NP-hardness proof given in [3][Theorem 34.10]. We will describe the reduction using a worked example taken from [3]. Consider the following 3-CNF formula:

$$\begin{aligned}\phi &= C_1 \wedge C_2 \wedge C_3 \wedge C_4 \\ C_1 &= (x_1 \vee \neg x_2 \vee \neg x_3) \\ C_2 &= (\neg x_1 \vee \neg x_2 \vee \neg x_3) \\ C_3 &= (\neg x_1 \vee \neg x_2 \vee x_3) \\ C_4 &= (x_1 \vee x_2 \vee x_3)\end{aligned}$$

This formula has three variables, x_1 , x_2 , and x_3 , and four clauses, C_1 through C_4 . The reduction assumes that there is no clause C_i that contains both x_i and $\neg x_i$, because otherwise C_i would be always be satisfied.

The reduction constructs a subset-sum instance, which is described in the following table:

	x_1	x_2	x_3	C_1	C_2	C_3	C_4
$v_1 =$	1	0	0	1	0	0	1
$v'_1 =$	1	0	0	0	1	1	0
$v_2 =$	0	1	0	0	0	0	1
$v'_2 =$	0	1	0	1	1	1	0
$v_3 =$	0	0	1	0	0	1	1
$v'_3 =$	0	0	1	1	1	0	0
$s_1 =$	0	0	0	1	0	0	0
$s'_1 =$	0	0	0	2	0	0	0
$s_2 =$	0	0	0	0	1	0	0
$s'_2 =$	0	0	0	0	2	0	0
$s_3 =$	0	0	0	0	0	1	0
$s'_3 =$	0	0	0	0	0	2	0
$s_4 =$	0	0	0	0	0	0	1
$s'_4 =$	0	0	0	0	0	0	2
$t =$	1	1	1	4	4	4	4

Each row should be read as a number written in decimal. For example, the first row specifies the number $v_1 = 1001001$. The subset-sum instance asks whether there is a subset of rows v_1 through s'_4 that sums to row t .

The table is constructed according to the following rules. Each column is labelled: the first three columns are labelled by the variables x_1 through x_3 , and the rest of the columns are labelled by the clauses C_1 through C_4 . For each variable x_i we define two rows:

- v_i has a 1 in column x_i , and a 1 in every column C_i that contains x_i .
- v'_i has a 1 in column x_i , and a 1 in every column C_i that contains $\neg x_i$.

In addition to these, for each clause C_i we define two *slack* rows: the row s_i has a 1 in column C_i , and the row s'_i has a 2 in column C_i .

To see that this reduction works, suppose that we know a satisfying assignment of the CNF formula. We can use this to construct a solution to the subset-sum instance. If x_i is true in the satisfying assignment, then we select v_i , and if it is false then we select v'_i . In doing so, we construct a subset with the following properties:

- For each column x_i , we have that the sum of that column is 1, because we never select both v_i and v'_i .
- For each column C_i , we have that the sum of that column is at least 1, because every clause must be satisfied.
- For each column C_i , we have that the sum of that column is at most 3, because each clause contains exactly 3 variables.

These properties ensure that, for each column C_i , we can always select a subset of the slack columns, s_i and s'_i , so that the sum of the column is 4. Thus, every satisfying assignment of the CNF formula corresponds to a solution of the subset-sum instance.

For similar reasons, every solution of the subset-sum instance corresponds to a satisfying assignment of the CNF formula, by simply ignoring the slack rows. Since every column C_i must sum to 4, we know that after removing the slacks, each column must sum to at least 1. This, combined with the fact that v_i and v'_i cannot be selected at the same time, implies that we have a satisfying assignment for the CNF formula.

See [3] for a full proof correctness of the NP-hardness reduction.

Changing the format. Our definition of an SSG requires a very specific format for the input instance. In particular, each quantifier is associated with exactly two natural numbers. However, the reduction that we have described can be written down very naturally as a one-player SSG, in which only the existential player is allowed to move. For our example, the instance is $(VS_1S_2S_3S_4, t)$, where:

$$V = \exists \{v_1, v'_1\} \exists \{v_2, v'_2\} \exists \{v_3, v'_3\},$$

$$S_i = \exists \{s_i, 0\} \exists \{s'_i, 0\}.$$

Note that it is valid to force the choice between v_i and v'_i , because no solution of the subset-sum instance can contain both of these numbers.

Subset-sum games are PSPACE-complete. It is now easy to reduce a quantified boolean formula to an SSG. We simply follow the existing reduction, but if variable x_i is universally quantified, then we use $\forall\{v_i, v'_i\}$ rather than $\exists\{v_i, v'_i\}$. For example, if we consider the quantified boolean formula $\forall x_1 \exists x_2 \forall x_3 \phi$, where ϕ is defined as before, then we produce the quantified subset-sum instance $(V'S_1S_2S_3S_4, t)$, where:

$$V' = \forall\{v_1, v'_1\} \exists\{v_2, v'_2\} \forall\{v_3, v'_3\},$$

and S_i is defined as before. The final step is to ensure a strict alternation of quantifiers, which the definition of an SSG requires. This can easily be achieved by inserting “dummy” quantifiers, where necessary. That is, we can insert $\exists\{0, 0\}$ between two consecutive \forall quantifiers, and we can insert $\forall\{0, 0\}$ between two consecutive \exists quantifiers. This change obviously cannot affect the winner of the SSG.

B Proof of Lemma 3

Let $\mathcal{S} = (L, C, \Delta, l_0)$ be a b -bounded counter-stack automaton. Without loss of generality, we will assume that $b = 2^n - 1$, which means that each counter in \mathcal{S} is n bits wide. We will construct a bounded one-counter automaton $\mathcal{B} = (L', b', \Delta', l'_0)$ that simulates \mathcal{S} . We will refer to the counters of \mathcal{S} as c_1 through c_k , and the counter of \mathcal{B} as c .

We will follow the approach laid out at the start of Section 4. That is, we will set the bound $b' = 2^{k \cdot n} - 1$ so that c is $k \cdot n$ bits wide. We then partition these bits in order to implement the counters c_1 through c_k . The counter c_k will use the n most significant bits, the counter c_{k-1} will use the next n most significant bits, and so on.

We introduce some notation to formalise this encoding. Let $x \in [0, b]$ be a counter value for counter c_i . We define $\text{Enc}(x, i) = x \cdot 2^{(i-1) \cdot n}$. To understand this definition, note that for $i = 1$, we have $\text{Enc}(x, i) = x$. Then, for $i = 2$, we have that $\text{Enc}(x, i)$ is the value of x bit-shifted to the left n times. Thus, this definition simply translates x to the correct position in c .

We can now define the translation. We will set $L' = L$ and $l'_0 = l_0$, which means that both automata have the same set of locations, and the same start location. We will use the transitions in Δ' to simulate \mathcal{S} . For each transition $t = (l, E, I, R, l') \in \Delta$, we construct a transition $t' = (l, p, g_1, g_2, l') \in \Delta'$ between the same pair of locations. We want to have the following property: transition t can be used from a state $(l, c_1, c_2, \dots, c_k)$ in \mathcal{S} if and only if transition t' can be used from the state $(l, \sum_i \text{Enc}(c_i, i))$ in \mathcal{B} .

We begin by defining p . We set:

$$p = \sum_{i \notin R} \text{Enc}(I_i, i) - \sum_{i \in R} \text{Enc}(E(i), i).$$

In other words, for each counter $i \notin R$ that is not to be reset, we add $\text{Enc}(I_i, i)$ to c , which correctly adds I_i to c_i . Note that the boundedness assumption on

\mathcal{S} implies that the counters can never overflow due to this operation. For the counters $i \in R$, we subtract $E(i)$ from c_i . Recall that $E(i)$ must always be defined for the indices $i \in R$. Furthermore, the transition may only be taken if $c_i = E(i)$. Thus, subtracting $E(i)$ from c_i will correctly set it to 0.

Next we define the inequality tests. Let j be the smallest index for which $E(j)$ is defined. Our guards are:

$$g_1 = \sum_{i \geq j} \text{Enc}(E(i), i),$$

$$g_2 = \sum_{i \geq j} \text{Enc}(E(i), i) + \text{Enc}(1, j) - 1.$$

It is straightforward to show that, in our encoding scheme, we have $c_i = E(i)$ for all $i \geq j$ if and only if $g_1 \leq c \leq g_2$.

If we are given a target state $s = (t, c_1, c_2, \dots, c_k)$ for \mathcal{S} , then we can translate it into a target state $s' = (t, \sum_i \text{Enc}(c_i, i))$ for \mathcal{B} . The equivalence between the transitions in Δ , and the transitions in Δ' implies that s can be reached from $(l_0, 0, 0, \dots, 0)$ if and only if s' can be reached from $(l'_0, 0)$. This completes the proof of Lemma 3.

C Proof of Lemma 5

Let $\mathbf{s} = (s_1, s_2, \dots, s_n)$ be a winning strategy for the existential player. We define a sequential winning strategy as follows. Recall that $\text{Plays}(\mathbf{s})$ contains exactly 2^n plays. We argue that these plays can be ordered so that they form a sequential strategy. We give an iterative procedure that achieves this task: the first step of the procedure will ensure that the 1-blocks contain the correct plays, the second step will ensure that the 2-blocks contain the correct plays, and so on. In the first step, we observe that exactly 2^{n-1} of the plays contain A_1 , while exactly 2^{n-1} of the plays contain B_1 , so we can order the plays so that the even 1-block contains all plays containing A_1 . Now suppose that we have found the i -blocks. We observe that each i -block L has exactly $2^{n-(i+1)}$ plays that contain A_{i+1} . Therefore, for each i -block L , we can order the plays in L so that the even $(i+1)$ -block has all plays that contain A_{i+1} , and the odd $(i+1)$ -block has all plays that contain B_{i+1} . At the end of this procedure, we will have a list of plays $\mathcal{S} = P_1, P_2, \dots, P_{2^n}$ where:

- P_j contains A_i whenever j is in an even i -block.
- P_j contains B_i whenever j is in an odd i -block.

So \mathcal{S} satisfies the first two conditions of Definition 4. We argue that \mathcal{S} also satisfies the third condition. Let L be an i -block. By definition, for every $j < i$, there is a unique j -block that contains L . These blocks define a play prefix $F \in \Pi_{1 \leq j \leq i} \{A_i, B_i\}$, and, for each play P_j with $j \in L$, we have $F \subseteq P_j$. Since \mathcal{S} is a reordering of $\text{Plays}(\mathbf{s})$, we must have $s_i(F) \in P_j$ for every $j \in L$. Hence,

\mathcal{S} satisfies Definition 4. Moreover, since \mathbf{s} is winning, we have that every play in $\text{Plays}(\mathbf{s})$ is winning, and therefore \mathcal{S} is a sequential winning strategy.

Now let $\mathcal{S} = P_1, P_2, \dots, P_{2^n}$ be a winning sequential strategy. We give a high level description of a winning strategy for the SSG. At the start of the strategy we set $L_0 = [1, 2^n]$. In each round i of the game, let $D_i \in \{A_i, B_i\}$ be the decision made by the universal player. We select L_i to be the unique i -block in L_{i-1} such that $D_i \in P_j$ for all $j \in L_i$. We play E_i if $E_i \in P_j$ for all $j \in L_i$, and we play F_i if $F_i \in P_j$ for all $j \in L_i$. It is straightforward to encode this strategy in the form $\mathbf{s} = (s_1, s_2, \dots, s_n)$. By construction, when we play \mathbf{s} , the outcome of the game will be some play P_j from \mathcal{S} . Since every play P_j in \mathcal{S} is winning for the existential player, we have that \mathbf{s} is a winning strategy.

D Proof of Lemma 6

Let R be a run in \mathcal{A}_n . The following lemma describes the set of reset states that each segment of R must pass through.

Lemma 11. *Let R be a run in \mathcal{A}_n . Either:*

- R_j visits precisely the reset locations $\{r'_i, r_i\}$ for which $j \bmod 2^{n-i} = 0$, or
- R_j gets stuck.

Proof. We will prove this lemma by induction over i . The base case, where $i = n$, is trivial because $j \bmod 2^{n-n}$ is always equal to 0, and it is clear from the construction that every segment R_j must always visit both r'_n and r_n .

For the inductive step, suppose that the lemma has been shown for $i + 1$, and will show that the lemma holds for i . We know that, in order to reach r'_i or r_i , a segment must first visit r'_{i+1} . By the inductive hypothesis, we know that only segments R_j with $j \bmod 2^{n-(i+1)}$ visit r_{i+1} . At the start of R , we have $a_i = b_i = 0$. On the first visit to r_{i+1} , we clearly cannot take the transition to r'_i , because we have $a_i + b_i = 2^{n-(i+1)}$, and the transition to r'_i requires $a_i + b_i = 2^{n-i}$. Thus, we either have to take the transition to u_1 , or we get stuck. On the second visit to r_{i+1} , we cannot take the transition to u_1 , because we have $a_i + b_i = 2^{n-i}$, and the transition to u_1 requires $a_i + b_i = 2^{n-(i+1)}$. Thus, either we get stuck, or we take the transition to r'_i . The transition between r_{i+1} and r'_i resets a_i and b_i . Thus, we can repeat the argument, and conclude that locations r'_i and r_i are only visited by segments R_j where $j \bmod 2^{n-i} = 0$. \square

Having shown Lemma 11 it is now easy to prove Lemma 6. Let R be a run of \mathcal{A}_n . For the counters a_i and b_i , we have the following facts:

- At the start of the first i -block, we have $a_i = b_i = 0$.
- Each i -block contains exactly 2^{n-i} segments. Each segment must increment one of a_i or b_i , but not both.
- At the end of each odd i -block, we must take the transition from r_i to u_1 to avoid getting stuck. This transition requires $a_i = 2^{n-i}$ and $b_i = 0$.

- At the end of each even i -block, we must take the transition from r_i to r'_{i-1} to avoid getting stuck. This transition requires $a_i = 2^{n-i}$ and $b_i = 2^{n-i}$, and resets a_i and b_i to 0.

These facts imply that a_i must be incremented during every run in an odd i -block to prevent the automaton getting stuck, and b_i must be incremented during every run in an even i -block to prevent the automaton getting stuck. It can also be verified that, if a_i is incremented during every run in an odd i -block, and b_i is incremented during every run in an even i -block, then the automaton will never get stuck at r_i .

Similarly, for the counters e_i and f_i we have the following facts.

- At the start of the first i -block, we have $e_i = f_i = 0$.
- Each i -block contains exactly 2^{n-i} runs. Each run must increment one of e_i or f_i , but not both.
- At the end of each i -block, we must take one of the two transitions from r'_i to r_i to avoid getting stuck. These transitions require that $e_i = 2^{n-i}$ and $f_i = 0$, or $e_i = 0$ and $f_i = 2^{n-i}$.

These facts imply that either e_i is incremented during every run in an i -block, or f_i is incremented during every run in an i -block, or the automaton will get stuck when moving from r'_i to r_i at end of the i -block. It can also be verified that, if the automaton increases e_i during every run in an i -block, then the automaton will not get stuck moving from r'_i to r_i , and if the automaton increases f_i during every run in an i -block, then the automaton will not get stuck moving from r'_i to r_i .

Note that, in \mathcal{A}_n , it is only possible for R to get stuck at the locations r'_i and r_i . Therefore, we have shown that R does not get stuck if and only if the three conditions of Lemma 6 hold for R .

E Proof of Lemma 7

This lemma follows from Lemma 11. Let R be a run. Lemma 11 implies that the transition from r_i to r'_{i-1} is taken in every segment R_j such that $j \bmod 2^{n-(i-1)}$. This transition resets both a_i and b_i to 0. Therefore, neither of these counters may exceed $2^{n-(i-1)}$. Similarly, Lemma 11 implies that every segment R_j such that $j \bmod 2^{n-i} = 0$ must move from r'_i to r_i . Both of the transitions from between r'_i and r_i reset e_i and f_i , and therefore neither of these counters may exceed 2^{n-i} .