



I|S|T AUSTRIA

Institute of Science and Technology

The Target Discounted-Sum Problem

Udi Boker and Thomas A Henzinger and Jan Otop

Technical Report No. IST-2015-335-v1+1
Deposited at 18 May 2015 07:50
<http://repository.ist.ac.at/335/1/report.pdf>

IST Austria (Institute of Science and Technology Austria)
Am Campus 1
A-3400 Klosterneuburg, Austria

Copyright © 2012, by the author(s).

All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

The Target Discounted-Sum Problem *

Udi Boker¹, Thomas A. Henzinger², and Jan Otop²

¹The Interdisciplinary Center (IDC) Herzliya

²IST Austria

May 13, 2015

Abstract

The *target discounted-sum* problem is the following: Given a rational discount factor $0 < \lambda < 1$ and three rational values a, b , and t , does there exist a finite or an infinite sequence $w \in \{a, b\}^*$ or $w \in \{a, b\}^\omega$, such that $\sum_{i=0}^{|w|} w(i)\lambda^i$ equals t ?

The problem turns out to relate to many fields of mathematics and computer science, and its decidability question is surprisingly hard to solve.

We solve the finite version of the problem, and show the hardness of the infinite version, linking it to various areas and open problems in mathematics and computer science: β -expansions, discounted-sum automata, piecewise affine maps, and generalizations of the Cantor set. We provide some partial results to the infinite version, among which are solutions to its restriction to eventually-periodic sequences and to the cases that $\lambda \geq \frac{1}{2}$ or $\lambda = \frac{1}{n}$, for every $n \in \mathbb{N}$.

We use our results for solving some open problems on discounted-sum automata, among which are the exact-value problem for nondeterministic automata over finite words and the universality and inclusion problems for functional automata.

1 Introduction

Discounting the influence of future events takes place in many natural processes, such as temperature change, capacitor charge, and effective interest rate, for which reason it is a key paradigm in economics and it is widely studied in game theory, Markov decision processes, and automata theory [1, 4, 5, 8, 10, 11, 16, 17, 20]. Yet, the decidability question of basic problems with regard to these models are still open. For example, the universality and inclusion problems of discounted-sum automata (see Section 6).

It turns out that the following simple *target discounted-sum problem* constitutes the main difficulty in many of these open problems. It was first raised, to the best of our knowledge, by Jean-François Raskin in the course of analyzing discounted-sum automata.

Definition 1. *Given a rational discount factor $0 < \lambda < 1$, a target rational value t , and rational weights a and b , the target discounted-sum problem is the question of whether there exists a finite, resp. an infinite, sequence (solution) $w \in \{a, b\}^*$, resp. $w \in \{a, b\}^\omega$, such that $\sum_{i=0}^{|w|} w(i)\lambda^i$ equals t .*

For distinguishing between the question about a finite and an infinite sequence, we will denote the former problem TDS^F and the latter TDS.

Despite its simple statement, resolving the decidability of TDS appears to be challenging and relates to many open questions in mathematics and computer science.

This problem is a natural milestone for open problems that involve discounting, such as problems on discounted-sum automata [5, 7, 16, 17], discounted-sum two-player games [6, 23], and multi-objective discounted-sum reachability [8]. In particular, TDS reduces to the universality problem of discounted-sum automata over infinite words (Theorem 29), whereas the exact-value problem of discounted-sum automata over finite words reduces to a generalized version of TDS^F (Theorem 26). Using our solution

*This research was supported in part by the European Research Council (ERC) under grant 267989 (QUAREM) and by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE) and Z211-N23 (Wittgenstein Award).

to TDS^F , we solve some of these open problems in Section 6. We further describe them at the end of this section.

Less intuitive is the connection between TDS and piecewise affine maps (Section 7.1). Its reachability problem asks, given a piecewise affine map f , and points s and t in a space of some dimension d , whether there exists $n \in \mathbb{N}$, such that $f^n(s) = t$. The problem is known to be undecidable for 2 or more dimensions [25, 29, 30], and it is open for one dimension [2, 3, 26]. (The reachability problem of an affine map, which is not “piecewise”, is known as the “orbit problem”, and it is decidable for all dimensions over the rationals [24, 35].) The orbit problem also relates to some decidable questions on linear recurrence sequences [15, 31].) We show that TDS reduces to the reachability problem of one-dimensional piecewise affine maps (Theorem 32).

Another interesting connection is between TDS and generalized Cantor sets (Section 7.2). The standard Cantor set is obtained by starting with the interval $[0, 1]$, and removing, at each iteration, the middle third of the remaining intervals. A natural generalization is to remove the middle k th (for example, the middle fifth) of the remaining intervals at each iteration [12, 13, 18]. While the membership question of a given number in the standard Cantor set is easily resolved, this is not the case with the general “middle k th Cantor set”. The reason for the difference between removing the middle third and removing, for example, the middle fifth, lies in the fractal behavior of these removals: In the former case, each of the remaining intervals constitutes a third of the original interval, and should thus be multiplied by 3 so as to view it as the original problem. In the latter case, each of the remaining intervals constitutes $\frac{2}{5}$ of the original interval, and should thus be multiplied by $\frac{5}{2}$ in order to view it as the original problem. This is closely related to number representation in an integral base, such as 3, which is very simple, as opposed to representation in a nonintegral base, such as $\frac{5}{2}$, which is significantly more complicated (Section 3). We show that the membership problem in the middle k th Cantor set reduces to TDS (Theorem 34).

Analyzing TDS, we show that it can be reduced to its restricted version, denoted TDS_{01} , in which the weights are fixed to be 0 and 1. The importance of this reduction is that it allows to view TDS as a question within the well studied area of β -expansions (β -representations), which deals with the representation of numbers in a nonintegral base [33]. Namely, a TDS_{01} instance with a discount factor λ and a target value t has a solution iff t has a representation in base $\frac{1}{\lambda}$, using only the 0 and 1 digits.

Unfortunately, though, the structure of β -expansions is still largely a mystery. Nikita Sidorov wrote [36], for example: “Usually the (greedy) expansions in bases like $\frac{3}{2}$ are considered virtually impossible to work with. For instance, if you expand $\frac{2}{5}$, say, in base $\frac{3}{2}$, then virtually nothing is known about this expansion.”

Nevertheless, analyzing TDS through β -expansions leads to some partial results to the problem. An immediate corollary is the decidability for discount factors equal to or bigger than half. This is because every number has a representation in base $1 < \beta \leq 2$, using only the 0 and 1 digits [33]. Other straightforward results are the uniqueness of the solution for a discount factor smaller than half, when a solution exists, and the co-recursively-enumerability of TDS.

A key tool in our analysis of β -representations is the notion of “gaps”: One can explore the representation of a given number in base β using the “greedy exploration” [33] – going from left to right, and adding at each position the maximal possible digit. In this process, every step produces a “remainder”, which should be treated in the next position. A “gap” is intuitively a “normalized remainder”; It is equal, at the n th step of the exploration, to the multiplication of the remainder and β^n . By normalizing the remainders into gaps, the exploration process has a fractal nature, which gives the intuition to its close connection to general Cantor sets and to piecewise affine maps. It is also related to generalizations of Mahler’s $\frac{3}{2}$ problem [28] and Collatz’s problem [9, 27], though more loosely.

Analyzing the possible gaps in the exploration process allows us to solve, in PSPACE, the restriction of TDS to eventually-periodic sequences. As a special case, we get decidability for every discount factor of the form $\frac{1}{n}$, where $n \in \mathbb{N}$. One might be tempted to conclude that another special case is a solution to finite sequences, namely to TDS^F . This is, however, not the case, as an instance of TDS^F cannot be reduced, in general, to a TDS instance with only 0’s and 1’s.

For leveraging our result on eventually-periodic sequences into a solution to TDS^F and into new results about discounted-sum automata, we consider its following three natural extensions. The first is the generalization of TDS to have arbitrarily many weights, denoted GTDS; the second is adding a parameter to the problem, constraining the sequence of weights by an ω -regular expression, denoted CTDS; and the third is their combination, denoted CGTDS. The corresponding generalizations of TDS^F are denoted GTDS^F , CTDS^F (getting a regular expression), and CGTDS^F .

Fortunately, we are able to generalize our results to all of these extended versions. The main challenge here is that, as opposed to TDS, we no longer have the “dichotomy” property, stating that each discount factor either guarantees a solution or guarantees that the solution, if exists, is unique. The underlying reason is that when allowing enough digits, a number might not have, or may have many, and even infinitely many, different representations in a nonintegral base [14, 21]. Nevertheless, we derive the generalized results using nondeterministic explorations, rather than greedy explorations, König’s Lemma, and a few other observations.

Using our results, we solve some open problems on discounted-sum automata over finite words, among which are the exact-value problem for nondeterministic automata (Theorem 26) and the universality and inclusion problems for functional automata (Theorem 27.)

2 Problem Restrictions and Extensions

We analyze below the natural restrictions and extensions of TDS (Definition 1). Most interesting are the restriction to fixed weights of 0 and 1, which will be shown to be equivalent to the original problem, and the generalization to arbitrarily many weights, for which all the positive results will follow.

Restrictions

One may wonder whether some restricted versions of TDS are easier to solve. A natural approach in this direction is to fix one of the four TDS parameters.

Fixing the discount factor may indeed ease the problem. We will show solutions to the cases where $\lambda \geq \frac{1}{2}$ (Theorem 8) and $\lambda = \frac{1}{n}$, for every natural number n (Theorem 15).

Fixing the target value, on the other hand, cannot help: The general problem, with a discount factor λ , a target t , and weights a and b , can be reduced to a restricted problem that only allows a constant target value T , by choosing new weights $a' = a \cdot \frac{T}{t}$ and $b' = b \cdot \frac{T}{t}$.

Fixing the weights also cannot ease the problem. We show below that even when the weights are restricted to be exactly 0 and 1, the problem remains exactly as hard as in the general case. This observation is the key for approaching TDS via β -expansions (Section 3). To fit into the β -expansion setting, we formally define the 0-1 restriction of the problem to start the sequence of summations $\sum_{i=0}^{\infty} w(i)\lambda^i$ with λ^1 rather than with λ^0 , i.e., we put $w(0) = 0$.

Definition 2 (TDS_{01}). *Given a rational discount factor $0 < \lambda < 1$ and a target rational value t , the 0-1 target discounted-sum problem (TDS_{01}) is the question of whether there exists an infinite sequence (solution) $w \in \{0, 1\}^{\omega}$, such that $\sum_{i=1}^{\infty} w(i)\lambda^i$ is equal to t .*

We show below that TDS_{01} is exactly as hard as the general infinite TDS.

Theorem 3. *TDS reduces to TDS_{01} . The reduction preserves eventual periodicity and non-eventual periodicity of the solution.*

Proof. We make the reduction in three steps:

1. We claim that the original TDS, denoted \mathcal{P} , is equivalent to the TDS, denoted \mathcal{P}' , with a target $t' = t - \frac{a}{1-\lambda}$ and weights $a' = 0$ and $b' = b - a$. Indeed, subtracting a from every element in a discounted-sum sequence $w \in \{a, b\}^{\omega}$, provides a discounted-sum sequence $w' \in \{0, b - a\}^{\omega}$, such that $\sum_{i=1}^{\infty} w'(i)\lambda^i = \sum_{i=1}^{\infty} w(i)\lambda^i - \sum_{i=1}^{\infty} a\lambda^i = \sum_{i=1}^{\infty} w(i)\lambda^i - \frac{a}{1-\lambda}$.
2. We claim that the TDS \mathcal{P}' is equivalent to the TDS, denoted \mathcal{P}'' , with a target $t'' = \frac{t'}{b'} = \frac{t - \lambda t - a}{(1-\lambda)(b-a)}$ and weights $a'' = 0$ and $b'' = 1$. Indeed, dividing every element in a discounted-sum sequence $w' \in \{0, b'\}^{\omega}$ by b' , provides a discounted-sum sequence $w'' \in \{0, 1\}^{\omega}$, such that $\sum_{i=1}^{\infty} w''(i)\lambda^i = \frac{1}{b'} \sum_{i=1}^{\infty} w'(i)\lambda^i$.
3. The TDS \mathcal{P}'' already uses the weights 0 and 1. It is obviously equivalent to the TDS_{01} with a target $t''' = \lambda t'' = \frac{\lambda(t - \lambda t - a)}{(1-\lambda)(b-a)}$, starting the summation with λ^1 rather than with λ^0 .

□

Observe that TDS_{01} is more general than the corresponding problem with respect to finite sequences, as every finite sequence w can be considered as the infinite sequence $w0^\omega$. Yet, TDS^F is not subsumed by TDS , and cannot be reduced to only have the 0 and 1 weights. (The first step in the proof of Theorem 3 only holds for infinite sequences.) We shall return to TDS^F at the end of this section.

Extensions

We consider two natural extensions of TDS and TDS^F , as well as their combination.

The first generalization allows for arbitrarily many weights:

Definition 4 (GTDS). *Given a rational discount factor $0 < \lambda < 1$, a target rational value t , and rational weights a_1, \dots, a_k , for $k \in \mathbb{N}$, the generalized target discounted-sum problem (GTDS) is the question of whether there exists an infinite sequence (solution) $w \in \{a_1, \dots, a_k\}^\omega$, such that $\sum_{i=0}^{\infty} w(i)\lambda^i$ equals t .*

The second extension adds an ω -regular constraint on the allowed sequences. Such a constraint is particularly relevant for linking between TDS and TDS^F , as well as in the scope of discounted-sum automata (Section 6).

Definition 5 (CTDS). *Given a rational discount factor $0 < \lambda < 1$, a target rational value t , rational weights a and b , and an ω -regular expression e , the constrained target discounted-sum problem (CTDS) is the question of whether there exists an infinite sequence (solution) $w \in \{a, b\}^\omega$, such that $\sum_{i=0}^{\infty} w(i)\lambda^i$ equals t and w belongs to the language of e .*

Note that CTDS is a proper extension of TDS . Indeed, a special variant of CTDS with e defined as $(a + b)^\omega$ is equivalent to TDS .

One may then consider the combination of GTDS and CTDS, denoted CGTDS, allowing arbitrarily many weights, and imposing a regular constraint on the allowed sequences.

All of the above extensions also apply to TDS^F , denoted GTDS^F , CTDS^F (getting a regular expression), and CGTDS^F , respectively. By allowing arbitrarily many weights and a constraint, the finite version is subsumed by the infinite version of the problem, as formalized in the following theorem.

Theorem 6. *CGTDS^F reduces to CGTDS .*

Proof. Consider an instance \mathcal{P} of CGTDS^F with a discount factor λ , weights a_1, \dots, a_k , a target t , and a regular expression e . We define an instance \mathcal{P}' of CGTDS as follows: if one of a_1, \dots, a_k is 0 then \mathcal{P}' is the same as \mathcal{P} , except for having an ω -regular expression $e' := e \cdot 0^\omega$. Otherwise, we define \mathcal{P}' to be the same as \mathcal{P} , except for having an additional weight of value 0, and the ω -regular expression $e' := e \cdot 0^\omega$, where e does not contain the additional weight. Then, \mathcal{P} has a finite solution if and only if \mathcal{P}' has an infinite solution. \square

3 TDS as a Question on β -Expansions

Once reducing TDS to TDS_{01} (Theorem 3), we can address the problem via β -expansions, namely by representing numbers in a nonintegral base.

Nonintegral base

We are used to represent numbers in an integral base (radix) – decimal, binary, hexadecimal, etc. For example, the string “3.56” in decimal base is equal to $3 \cdot 10^0 + 5 \cdot 10^{-1} + 6 \cdot 10^{-2}$. Yet, the representation may be with an arbitrary base $\beta > 1$, in which case the string “3.56” is equal to $3 \cdot \beta^0 + 5 \cdot \beta^{-1} + 6 \cdot \beta^{-2}$. Representation in a nonintegral base is known as β -expansion, a notion introduced by Rényi [33] and first studied in the seminal works of Rényi and of Parry [32].

The representation may be finite, as above, or infinite, as in the cases of representing the number $\frac{1}{3}$ in decimal base by the string “0.33333...”, and representing the number $\frac{10}{21}$ in base $\frac{5}{2}$ by “0.10101010...”.

We denote the value of a representation w in base β by $w_{[\beta]}$ (and when no base is mentioned, it is the decimal base). For example the value of $0.102_{[\frac{5}{2}]}$ is $\frac{66}{125}$ (as it equals to $1 \cdot \frac{2}{5} + 2 \cdot \frac{8}{125}$).

When dealing with a nonintegral base $\beta > 1$, the representations are still required to only contain digits that stand for natural numbers. A well known result [33] is that all real numbers have a β -representation, using the numbers $\{0, 1, 2, \dots, \lceil \beta - 1 \rceil\}$ as digits. In general, there might be several, and even infinitely many, representations to the same number [14]. (For an analysis of the numbers with unique representations see [21].)

Greedy and lazy explorations

The most common scheme for generating a representation for a given number is the *greedy exploration* (*greedy expansion*) [33], going from left to right, and adding at each position the maximal possible digit. For example, when representing $\frac{5}{8}$ in binary, we start with the digit 1 for the value $\frac{1}{2}$, getting a *remainder* of $\frac{1}{8}$. We cannot continue with another 1, since $\frac{1}{4} > \frac{1}{8}$, so we put 0 (and the remainder is still $\frac{1}{8}$). We then put 1 for $\frac{1}{8}$, and we are done, as the remainder is 0, getting the representation $.101_{[2]}$. Note that the greedy exploration provides the largest possible representation, lexicographic wise. Another common scheme is the *lazy exploration* (*lazy expansion*), adding at each position the minimal possible digit, providing the minimal possible representation, lexicographic wise. Finally, a unifying scheme for greedy and lazy explorations is the *nondeterministic exploration*, nondeterministically choosing at each step an eligible digit, i.e., a digit such that the remainder is non-negative and is smaller than the maximal number that can be represented starting from the current position.

TDS₀₁ – reformulated

TDS₀₁ can be naturally written as a question about representing a number in a nonintegral base:

Proposition 7. *A TDS₀₁ instance with a discount factor λ and a target value t has a solution iff t has a representation in base $\frac{1}{\lambda}$, using only the 0 and 1 digits.*

Motivated by TDS₀₁, we only consider rational bases. In addition, as it is trivial to decide the TDS₀₁ problem for a discount factor $\lambda \geq \frac{1}{2}$ (Theorem 8), which relates to a base $\beta \leq 2$, we will mostly consider a base $\beta = \frac{p}{q} > 2$, and assume that p and q are co-prime.

4 TDS – Analysis and Results

We handle in this section the target discounted-sum problem (TDS) and its constrained version (CTDS).

We start with an immediate corollary of viewing TDS as a question of β -expansions, providing a solution to the case that the discount factor is bigger than half. We continue with defining “gaps” – an alternative notion to the remainders that are maintained in exploring the representation of a number. The gaps will be a key tool in our analysis of the representations.

Unlike the general case of representing a rational number in a nonintegral base, we show that once the representations only use the 0 and 1 digits, the representation, if exists, is unique for every rational in every base $\beta > 2$. A direct corollary is that TDS is co-recursively enumerable. A more delicate analysis of the possible representations allows us to decide whether a given rational has an eventually-periodic representation. As a special case, we get that TDS is decidable for every discount factor in the form of $\frac{1}{n}$, where $n \in \mathbb{N}$.

All of the above results, except for the case that $\lambda \geq \frac{1}{2}$, also hold for CTDS.

We conclude with looking into an opposite question – given a non-eventually-periodic representation, does it represent a rational number? We show that in some cases, for example when the growth rate of blocks of identical digits is more than exponential, the answer is negative.

The case of $\lambda \geq \frac{1}{2}$

An immediate benefit of reducing TDS to a question on β -representations is the solution for λ 's greater than $\frac{1}{2}$.

Theorem 8. *TDS is decidable for every discount factor $\lambda \geq \frac{1}{2}$.*

Proof. By Theorem 3 and Section 3, we view the problem as asking whether a target number t has a representation in base $\beta = \frac{1}{\lambda}$ with only 0's and 1's.

Thus, there are two simple cases:

- The target t is bigger than $\frac{1}{\beta-1}$, implying that t does not have a β -representation, since $0.1_{[\beta]}^{\omega} = \frac{1}{\beta-1} < t$.
- The target t is equal to or smaller than $\frac{1}{\beta-1}$, implying that t has a β -representation: Every number has a β -representation with the digits $\{0, 1, 2, \dots, \lceil \beta - 1 \rceil\}$ [33], and since $\lambda \geq \frac{1}{2}$ and $\beta = \frac{1}{\lambda}$, it follows that $\beta \leq 2$, meaning that the 0 and 1 digit suffice.

□

The simplicity of solving TDS for $\lambda \geq \frac{1}{2}$ suggests to check whether it also holds when slightly extending TDS with a regular constraint; i.e., CTDS. Indeed, a slight extension, such as requiring a 0 in every odd position, breaks the solution, as it is analogous to TDS over λ^2 .

Example 9. Consider a TDS_{01} instance with $\lambda = \frac{2}{3}$. One can apply Theorem 8 for solving it, as $\frac{2}{3} \geq \frac{1}{2}$. However, Theorem 8 cannot hold for a CTDS with the same parameters and an additional constraint that the sequence of 0's and 1's is in the language of $(0(0+1))^{\omega}$. Indeed, $\sum_{i=0}^{\infty} w(i)\lambda^i = \sum_{i=0}^{\infty} w(2i)\lambda^{2i}$, which is a TDS_{01} problem for $\lambda = (\frac{2}{3})^2 = \frac{4}{9} < \frac{1}{2}$.

Gaps

In the exploration schemes presented in Section 3, we compared at each step the remainder with the value of a digit in the currently handled position. For example, in binary representation, the value of the digit 1 in the first position to the right of the radix point is $\frac{1}{2}$, while its value in the third position is $\frac{1}{2} \cdot (\frac{1}{2})^2 = \frac{1}{8}$. Analogously, we can always compare a fixed value of the digit, say $\frac{1}{2}$ for the digit 1 in binary, with the “normalized-remainder”. In the above example, instead of checking whether $\frac{1}{2} \cdot (\frac{1}{2})^2 \leq \frac{1}{8}$, we would check whether $\frac{1}{2} \leq \frac{1}{8} \cdot (2)^2$. We call this “normalized-remainder” the *gap*. (A similar notion is used in [4].)

When exploring a representation and working with gaps, we do not need to multiply, at each step, the remainder with the position-exponent of the base (as demonstrated above when comparing $\frac{1}{2}$ and $\frac{1}{8} \cdot (2)^2$). Instead, we can compute the new gap based on the current gap and the chosen digit, completely ignoring the remainder and the current position.

More formally, a left-to-right exploration in base β (with digits only to the right of the radix point), using gaps, is done as follows: The initial gap is the target number, and in every step with a gap g , the gap of the next step is $g' = \beta g - m$, where m is the chosen digit. (In case of the greedy exploration, it is the biggest m , such that $\frac{m}{\beta} \leq g$.)

There are two main advantages for exploring representations using gaps rather than remainders:

- There is no need to store the current position, which generally grows to infinity, but only the current gap. In some cases, for example when the base β is in the form of $\frac{1}{n}$ where $n \in \mathbb{N}$, there are only boundedly many possible gaps (Theorem 15).
- Once having a gap g in some position p , the suffix of the representation (the string from position p onwards) only depends on the gap g , and not on the position p . This is a central tool in analyzing the representation of a number, for example allowing to infer that the representation is eventually periodic when getting the same gap twice (Lemma 12).

Unique representation with 0's and 1's

With a nonintegral base β , a number may have several, and even infinitely many, β -representations [21, 33]. In the following, we show that if $\beta > 2$ and the representation can only use the 0 and 1 digits, then the representation, if exists, is unique.

Lemma 10. Consider a base $\beta > 2$ and a number t . If t has a β -representation with only 0's and 1's then this representation is unique.

Proof. Assume towards contradiction that a number t has two such β -representations w and w' . As $w \neq w'$, there is a finite word u and infinite words v and v' , such that $w = 0.u0v$ and $w' = 0.u1v'$. Now, observe that $w_{[\beta]} \leq 0.u01_{[\beta]}^\omega = 0.u_{[\beta]} + (\frac{1}{\beta})^{|u|+2}(\frac{\beta}{\beta-1})$, which is strictly smaller than $w'_{[\beta]} \geq 0.u10_{[\beta]}^\omega = 0.u_{[\beta]} + (\frac{1}{\beta})^{|u|+1}$, leading to a contradiction. \square

The possible scenarios of the exploration

Establishing that the representation, if exists, is unique, there are exactly three possible scenarios when running the greedy exploration with a rational base β and a target number t :

- The exploration stops when reaching a gap g , such that $g > \frac{1}{\beta-1}$. The conclusion is that there is no representation of t in base β , since even if we use, from this position further, only the 1 digit, the resulting number will be too small, as $\sum_{i=1}^{\infty} \beta^i = \frac{1}{\beta-1} < g$.
- The exploration stops when reaching a gap g , such that g already appeared as the gap in a previous step. The conclusion is that there is an eventually-periodic representation of t in base β , as formalized in Lemma 12.
- The exploration never stops, which happens in the case that t has a β -representation that is not eventually periodic.

Using the observation on the possible scenarios, we get the following result.

Theorem 11. *TDS is co-recursively-enumerable.*

Eventually-periodic representations

A key question is whether a number has an eventually-periodic representation. It is known that with every integral base (such as decimal, binary, etc.), a number n has an eventually-periodic representation iff n is rational. This no longer holds when the base is not integral. While an eventually-periodic representation implies a rational number, there are also rationals with non-eventually-periodic representations. Further, Schmidt [34] showed that if all rationals have periodic representations in some base $\beta > 1$, for a real number β , then β is an algebraic integer, and more precisely, it is a Pisot number or a Salem number. (The opposite direction is an open problem, known as ‘‘Schmidt’s conjecture’’ [22].) The only algebraic integers among the rational numbers are the integers, implying that if all rationals have periodic representations in a rational base β then β is an integer.

We show below how to decide whether a given number has an eventually-periodic representation in a rational base $\beta > 2$, when the representations can only use the 0 and 1 digits.

We first formalize the direct connection between eventually-periodic representations and a bounded set of gaps in the exploration.

Lemma 12. *Consider a rational number t and a rational base $\beta > 2$. Then, t has an eventually-periodic β -representation with only 0’s and 1’s iff there are finitely many different gaps in the greedy-exploration of t .*

Proof.

- Assume that the gaps in the greedy-exploration are eventually periodic. Then, there is a step i in the exploration in which we get a gap g , such that g was also the gap in some step $j < i$. Hence, the exploration will use in step $i + 1$ the same digit that was used in step $j + 1$, and by induction, the used digits will be eventually periodic, providing an eventually-periodic representation.
- As for the other direction, assume that t has an eventually-periodic representation, in which the repeated sequence of digits is the finite word u , starting in position p . By Lemma 10, t has a unique β -representation, implying that the greedy exploration should provide this representation. Now, let g_p be the gap of the exploration in position p and $g_{p+|u|}$ be the gap of the exploration in position $p + |u|$. By the definition of the gaps in the exploration, we have $g_p = u_{[\beta]}^\omega = g_{p+|u|}$, leading to an eventually periodic sequence of gaps.

□

We show below that by analyzing the gaps of the greedy exploration, we can decide whether there exists an eventually-periodic representation.

Lemma 13. *For a rational number $t = \frac{a}{b}$, and a rational base $\beta = \frac{p}{q} > 2$, we can decide in space polynomial in the binary representation of a, b, p and q , whether t has an eventually-periodic β -representation with only 0's and 1's. Moreover, every eventually-periodic β -representation of $\frac{c}{d}$ is of the form uv^ω with $|u| + |v| \leq d$.*

Proof. We consider the greedy exploration, and analyze the gap $g = \frac{c}{d}$ (where c and d are co-prime) at every step of the exploration. Note that having a gap $g = \frac{c}{d}$, and adding a digit m , the next gap is $g' = \frac{p}{q} \cdot \frac{c}{d} - m = \frac{pc - qmd}{qd}$. Let $\beta = \frac{p}{q}$, where p and q are co-prime.

We prove the required decidability by showing the following two claims:

1. If the gap, $g = \frac{c}{d}$, is such that c is not divided by q then the representation cannot be eventually periodic.
2. For an initial gap, $g = \frac{c}{d}$, after at most d exploration steps, the gap will either
 - exceed $\frac{1}{\beta-1}$ (implying that there is no representation, as $1_{[\beta]}^\omega = \frac{1}{\beta-1}$); or
 - be the same as a previous gap (implying, by Lemma 12, an eventually-periodic representation); or
 - will be such that c is not divided by q (implying, by the previous claim, that there is no eventually-periodic representation).

Indeed:

1. Assume that c is not divided by q . Consider the prime factorization $q = f_1^{e_1} \cdot f_2^{e_2} \cdots f_n^{e_n}$ of q . Then, by the assumption, there is some $1 \leq i \leq n$, such that c is not divided by $f_i^{e_i}$. Thus, the numerator of g' (which is $pc - qmd$) is not divided by q , as (i) pc is not divided by $f_i^{e_i}$, since p is co-prime with q and c is not divided by $f_i^{e_i}$, and (ii) qmd is divided by $f_i^{e_i}$, since q is. Hence, the exponent of f_i in the denominator of g' is bigger than its exponent in g . Therefore, by induction, the exponent of f_i in the denominators of the gaps monotonically increases, precluding the possibility of two equivalent gaps.
2. Assume that c is divided by q . Then, the numerator of g' is divided by q , meaning that $g' = \frac{c'}{d}$, for some integer c' . Thus, as long as the numerators of the gaps are divided by q , the denominator cannot grow. Hence, within d steps of exploration, we must reach a gap that satisfies one of the following properties: (i) Its numerator exceeds d , and therefore it is bigger than 1, which is bigger than $\frac{1}{\beta-1}$; or (ii) It is the same as a previous gap; or (iii) Its numerator is not divided by q .

As for the space complexity, by the above, all the gaps are in the form of $\frac{x}{d}$, where $x \leq d$. Thus, the procedure can be done using a polynomial space. □

Using Lemma 13, we get the decidability of both TDS and CTDS for eventually-periodic sequences. The following theorem only holds for a discount factor $\lambda < \frac{1}{2}$, and will be generalized in Theorem 23 to hold for an arbitrary discount factor.

Theorem 14. *[TDS (resp. CTDS) w.r.t. eventually-periodic sequences is in PSPACE.] For a discount factor $\lambda < \frac{1}{2}$, the problem, given a TDS (resp. CTDS) instance \mathcal{P} , decide whether there exists an eventually-periodic sequence that is a solution to \mathcal{P} , is in PSPACE.*

Proof. The TDS case is a corollary of Theorem 3 and Lemma 13.

As for the CTDS case, note that the proof of Lemma 13 also generates the (unique) representation, when it exists. The representation is given in the form of uv^ω with $u, v \in \{0, 1\}^*$. Hence, for checking whether there is a representation that satisfies an ω -regular expression e , one can check whether the generated representation belongs to the language of e . Although u, v can be long, they can be generated, letter by letter, in polynomial space, and the membership test can be done on-the-fly [19, 37]. □

Another corollary of Lemma 13 is the case when the discount factor λ of the form $\frac{1}{n}$, where n is a natural number.

Theorem 15. *TDS (resp. CTDS) is in PSPACE for every discount factor λ of the form $\frac{1}{n}$ where n is a natural number.*

Proof. Reducing TDS to TDS₀₁ (Theorem 3), we consider a representation of a target rational number in base $\frac{1}{\lambda} = n$, which is a natural number. As a representation of a rational number in an integral base is always eventually periodic [33], we can decide the existence of a representation using Lemma 13.

As for CTDS, we use an analogous argument to the one provided for Theorem 14. \square

Growth rate of blocks of the same digit

By Lemma 13, given a rational number, we know whether it has an eventually-periodic representation with only 0's and 1's. However, when it does not have it, we do not know whether it has a non-eventually-periodic representation, or not. We do not even know whether, for the specific base $\frac{5}{2}$, there exists a rational number with a non-eventually-periodic representation that only uses 0's and 1's.

One can ask a related question from the opposite direction – given a non-eventually-periodic representation, does it represent a rational number? For example, consider the number 0.10100001... in base $\frac{5}{2}$, where 1's appear exactly at positions 3^n . Is it a rational number?

In some cases, such as the example above, we know that the answer is negative. We show below that in a representation of a rational number, the number of consequent 0's and consequent 1's is bounded by an exponential rate.

Lemma 16. *Consider a representation of a rational number $\frac{c}{d}$ in base $\frac{p}{q} > 2$, using only 0's and 1's. Then, i) The first 1 occurs within the first $\log d$ positions, and ii) if there is a 1 in the n -th position, then the next 1 occurs at position at most $n(\log q + 1) + \log d$.*

Proof. We consider the greedy exploration, and analyze the minimal gap that might occur after using a 1. Note that having a gap $g = \frac{a}{b}$, and adding a digit m , the next gap is $g' = \frac{p}{q} \cdot \frac{a}{b} - m = \frac{pa - qmb}{qb}$. Thus, the denominator of the gap at position n is smaller than or equal to dq^n , while the nominator is, obviously, bigger than or equal to 1. Hence, the gap is bigger than or equal to $\frac{1}{dq^n}$.

Now, the maximal gap that has a representation starting with 0 is smaller than $0.1_{[\frac{p}{q}]} = \frac{q}{p}$. Since the gap is multiplied by $\frac{p}{q}$ after every 0 digit, it means that the next 1 digit must occur within i steps, such that $\frac{1}{dq^n} \cdot (\frac{p}{q})^i \geq \frac{q}{p}$. Hence, the next 1 must occur within $\log_{\frac{p}{q}}(\frac{q}{p} \cdot dq^n) < \log(\frac{q}{p} \cdot dq^n) < \log(dq^n) = n \log q + \log d$ steps. \square

We shall call a maximal subword consisting of the same letter a *block*.

Theorem 17. *Let $\beta \in \mathbb{Q}$ and $w \in \{0, 1\}^\omega$. The number $(0.w)_{[\beta]}$ is rational only if for every n , the block of w starting at the n th position is linearly bounded in n .*

Proof. Let $\beta = \frac{p}{q}$ and $w \in \{0, 1\}^\omega$. Assume that $(0.w)_{[\frac{p}{q}]} = \frac{c}{d}$, where $c, d \in \mathbb{N}$. Lemma 16 implies that for every n , the length of a block of 0's starting at the n th position is linearly bounded.

As for blocks of 1's, consider the value $t' = (0.1^\omega)_{[\frac{p}{q}]} - (0.w)_{[\frac{p}{q}]} = \frac{1}{\beta-1} - \frac{c}{d}$. Note that t' is a rational, and its representation is obtained from w by swapping all 0's with 1's and vice versa. Thus, by applying Lemma 16 on t' , for every n , the length of a block of 1's starting at the n th position of w is linearly bounded. \square

5 Generalized TDS – Analysis and Results

We extend below the results that were shown for TDS and CTDS to the generalized versions GTDS and CGTDS, in which there are arbitrarily many weights. The main difficulty here is that there is no analogous to the combination of Theorem 8 and Lemma 10, stating that a relevant representation is either guaranteed, or else guaranteed to be unique, if it exists. Indeed, allowing enough digits, a number might not have, or may have many, and even infinitely many, different representations in a nonintegral

base [14, 21]. We overcome the problem by using boundedly-branching trees, rather than unique words, and applying König's Lemma, as well as a few other observations.

We begin with establishing a normal form of GTDS, in which the weights are natural numbers and the least weight is 0. Such a normal form generalizes TDS₀₁.

Theorem 18. *GTDS polynomially reduces to GTDS with weights from \mathbb{N} , where the least weight is 0.*

Proof. Consider an instance \mathcal{P} of GTDS with a discount factor λ , a target t , and weights $a_1, \dots, a_k \in \mathbb{Q}$. Assume that $a_1 < \dots < a_k$. Let M be the least common denominator of the weights. We define a GTDS \mathcal{P}' with a discount factor λ , a target $t \cdot M$, and weights $a_1 \cdot M, \dots, a_k \cdot M$. Note that all the weights in \mathcal{P}' are natural numbers. Now, observe that \mathcal{P} has a solution iff \mathcal{P}' has a solution, since for every sequence w , we have $\sum_{i=0}^{|w|} M \cdot w(i) \lambda^i = M \cdot \sum_{i=0}^{|w|} w(i) \lambda^i$. Finally, we define a GTDS instance \mathcal{P}'' with a discount factor λ , a target $(t \cdot M) - \frac{a_1 \lambda}{1-\lambda}$, and weights $0, a_1 \cdot M, \dots, a_k \cdot M$. \square

In the remaining part of this section we will assume that all instances of GTDS and CGTDS are in normal form. This allows us to consider the problem in the setting of β -expansions.

Proposition 19. *A GTDS with a discount factor λ , a target value t , and weights $a_1 = 0 < a_2 < \dots < a_k \in \mathbb{N}$ has a solution iff t has a representation in base $\frac{1}{\lambda}$, using only the digits a_1, \dots, a_k .*

With TDS, Theorem 8 shows that all relevant target numbers have a solution when the discount factor is equal to or bigger than half. This obviously also holds for GTDS. Further, having more weights, there are additional cases that guarantee a solution, as characterized below.

Theorem 20. *Consider a GTDS instance \mathcal{P} with a discount factor λ and weights $a_1 = 0 < a_2 < \dots < a_k \in \mathbb{N}$. Then \mathcal{P} has a solution for every target $t \in [(0.0^\omega)_{[\frac{1}{\lambda}]}, (0.a_k^\omega)_{[\frac{1}{\lambda}}]]$ iff for every $i \in \{1, \dots, k-1\}$ we have $a_{i+1} - a_i \leq (0.a_k^\omega)_{[\frac{1}{\lambda}]}$.*

Proof.

\Rightarrow : Suppose there exists $i \in \{1, \dots, k-1\}$, such that $a_{i+1} - a_i > (0.a_k^\omega)_{[\frac{1}{\lambda}]}$. Then, $(0.a_i a_k^\omega)_{[\frac{1}{\lambda}]} < (0.a_{i+1})_{[\frac{1}{\lambda}]}$. Accordingly, all the numbers that are bigger than $(0.a_i a_k^\omega)_{[\frac{1}{\lambda}]}$ and smaller than $(0.a_{i+1})_{[\frac{1}{\lambda}]}$ have no representation in base $\frac{1}{\lambda}$ with the digits $0, a_2, \dots, a_k$.

\Leftarrow : Consider an execution of the greedy exploration with gaps on the target number t . Note that having a gap g and adding a digit m , the next gap is $g' = \frac{1}{\lambda} \cdot g - m$. It suffices to show that the following invariant holds, stating that the gaps are always not too small and not too big:

(*) if a gap $g \in [0, (0.a_k^\omega)_{[\frac{1}{\lambda}}]]$, and $g' = \frac{1}{\lambda} g - a_i$, where a_i is the maximal among a_1, \dots, a_k for which $g' \geq 0$, then $g' \in [0, (0.a_k^\omega)_{[\frac{1}{\lambda}}]]$.

To prove (*) we consider two cases. If $i < k$, then by the maximality of i , $\frac{1}{\lambda} g - a_{i+1} < 0$. It follows that $g' < a_{i+1} - a_i < (0.a_k^\omega)_{[\frac{1}{\lambda}]}$. Therefore, $g' \in [0, (0.a_k^\omega)_{[\frac{1}{\lambda}}]]$. Otherwise, if $i = k$, $g \leq (0.a_k^\omega)_{[\frac{1}{\lambda}]}$ implies $g' = \frac{1}{\lambda} g - a_k \leq (0.a_k^\omega)_{[\frac{1}{\lambda}]}$. \square

For TDS, Lemma 10 shows that when the discount factor does not guarantee a solution it guarantees that the solution, if exists, is unique. This no longer holds for GTDS. Yet, as in the case of TDS (Theorem 11), GTDS is also co-recursively-enumerable.

Theorem 21. *GTDS is co-recursively-enumerable.*

Proof. Assume that a GTDS instance \mathcal{P} has no solution and consider a tree consisting of all runs of the nondeterministic exploration. Infinite paths in that tree correspond to solutions to \mathcal{P} , therefore each failing path is finite. Since the exploration tree has a finite degree and all of its paths are failing, König's Lemma implies that it is finite. Therefore, exhausting all the possibilities of the nondeterministic exploration on \mathcal{P} , we are guaranteed to stop with the result that \mathcal{P} has no solution. \square

We continue with solving GTDS w.r.t. eventually-periodic sequences. The following lemma generalizes Lemma 12.

Lemma 22. *Consider a rational number t , a rational base $\beta > 1$, and digits $a_1, \dots, a_k \in \mathbb{N}$, for $k \in \mathbb{N}$. Then, t has an eventually-periodic β -representation iff there are finitely many different gaps in some run of the corresponding nondeterministic exploration of t .*

Proof. Analogous to the proof of Lemma 12. \square

We are now in place to provide the result on eventually-periodic sequences.

Theorem 23. *GTDS (resp. CGTDS) w.r.t. eventually-periodic sequences is in PSPACE. Moreover, every eventually-periodic sequence of integer weights that forms a solution to an integer target $t = \frac{c}{d}$ is of the form uv^ω with $|u| + |v| \leq d$.*

Proof. Consider a GTDS instance with a discount factor λ , a target $t = \frac{c}{d}$, and weights $a_1 < a_2 < \dots < a_k \in \mathbb{N}$.

The decision procedure is analogous to the proof of Lemma 13, with the following differences: i) Rather than considering the presentation generated by the greedy exploration, we consider some representation generated by the nondeterministic exploration; ii) The bounds on the gaps are: the lower bound $\frac{\lambda a_1}{1-\lambda}$ (can be negative) and the upper bound $\frac{\lambda a_k}{1-\lambda}$; and iii) A decision is nondeterministically reached within $d \cdot a_k$ steps.

Hence, the procedure runs in nondeterministic polynomial space in the binary representations of $\lambda, c, d, a_1, \dots, a_k$. By Savitch's Theorem, we get the PSPACE complexity.

As for the generalization to CGTDS, the above procedure nondeterministically generates, on-the-fly, a representation in the form of uv^ω with $u, v \in \{a_1, \dots, a_k\}^*$. For checking whether it belongs to the language of a given regular expression e , we can check, on-the-fly, whether it is accepted by a nondeterministic Büchi automaton that is equal to e [19, 37]. \square

As a corollary of Lemma 22, we get the decidability for the case that the discount factor is of the form $\frac{1}{n}$ for $n \in \mathbb{N}$.

Theorem 24. *GTDS (resp. CGTDS) is in PSPACE for every discount factor λ of the form $\frac{1}{n}$ where n is a natural number.*

Proof. Consider an instance of GTDS with a discount factor $\lambda = \frac{1}{n}$ and a target number $t = \frac{c}{d}$. Observe that a nondeterministic exploration can only visit gaps whose denominator is d . Indeed, if a gap $g = \frac{a}{d}$ then the next gap, after using a digit m , is $g' = \frac{an}{d} - m$, also having the denominator d .

Therefore, analogously to the arguments given in the proof of Theorem 23, we get a bound on the gap numerators, implying decidability in PSPACE of both the GTDS and the CGTDS instances. \square

We conclude the section with the desired result on the decidability of the generalized target discounted-sum problem over finite words.

Theorem 25. *GTDS^F (resp. CGTDS^F) is in PSPACE. Moreover, every finite sequence of integer weights that forms a solution to an integer target $t = \frac{c}{d}$ is of length at most d .*

Proof. By Theorem 6, an instance \mathcal{P} of CGTDS^F can be polynomially reduced to the question of whether an instance \mathcal{P}' of CGTDS has an eventually periodic solution. By Theorem 23, the latter question can be resolved in PSPACE, and a solution, if exists, is of length at most d . \square

6 Results on Discounted-Sum Automata

In this section, we establish the connection between TDS and discounted-sum automata, and use our results about TDS for solving some of the latter's open problems. In particular, we solve the exact-value problem for nondeterministic automata over finite words and the universality and inclusion problems for functional automata.

We start with the definitions of discounted-sum automata and their related problems.

Discounted-sum automata

A *discounted-sum automaton* (DSA) is a tuple $\mathcal{A} = \langle \Sigma, Q, q_{in}, Q_F, \delta, \gamma, \lambda \rangle$ over a finite alphabet Σ , with a finite set of states Q , an initial state $q_{in} \in Q$, a set of accepting states $Q_F \subseteq Q$, a transition function $\delta \subseteq Q \times \Sigma \times Q$, a weight function $\gamma : \delta \rightarrow \mathbb{Q}$, and a rational discount factor $0 < \lambda < 1$.

A run of an automaton on a word $w = \sigma_1 \sigma_2 \dots$ is a sequence of states and letters, $q_0, \sigma_1, q_1, \sigma_2, q_2, \dots$, such that $q_0 = q_{in}$ and for every i , $(q_i, \sigma_{i+1}, q_{i+1}) \in \delta$. The length of a run r , denoted $|r|$, is n for a finite

run $r = q_0, \sigma_1, q_1, \dots, \sigma_n, q_n$, and ∞ for an infinite run. A finite run of an automaton \mathcal{A} is *accepting* if the last state of r belongs to Q_F . In the infinite case, we assume that every run is *accepting* (and Q_F is irrelevant).

The value of a run r is $\gamma(r) = \sum_{i=0}^{|r|-1} \lambda^i \cdot \gamma(q_i, \sigma_{i+1}, q_{i+1})$. The value of a word w (finite or infinite) is $\mathcal{A}(w) = \inf\{\gamma(r) \mid r \text{ is an accepting run of } \mathcal{A} \text{ on } w\}$.

A DSA \mathcal{A} over finite words is said to be *functional* if for every word w , all accepting runs of \mathcal{A} on w have the same value [16].

Decision problems

Given DSAs \mathcal{A} and \mathcal{B} and a value $t \in \mathbb{Q}$,

- the *exact-value* problem asks whether there exists a word w such that $\mathcal{A}(w) = t$,
- the *<-universality* (resp. *≤-universality*) problem asks whether for every word w we have $\mathcal{A}(w) < t$ (resp., $\mathcal{A}(w) \leq t$).
- the *<-inclusion* (resp. *≤-inclusion*) problem asks whether for every word w we have $\mathcal{A}(w) < \mathcal{B}(w)$ (resp., $\mathcal{A}(w) \leq \mathcal{B}(w)$).

Next, we establish the connection between the target discounted-sum problem and the above decision problems.

Results for finite words

Our techniques for resolving the target discounted-sum problem over finite words directly relate to the exact-value problem:

Theorem 26. *The exact-value problem for discounted-sum automata over finite words is decidable in PSPACE.*

Proof. Consider an exact-value problem over finite words with a target value t and a DSA \mathcal{A} with a discount factor λ and weights a_1, \dots, a_k . Let M be the least common denominator of a_1, \dots, a_k . We define a CGTDS^F \mathcal{P} with a discount factor λ , a target $t' := t \cdot M$, and weights $a_1 \cdot M, \dots, a_k \cdot M$. Note that for every finite word w , if $\mathcal{A}(w) = t$ via a run ρ of \mathcal{A} on w then the sequence of weights in ρ is a solution to \mathcal{P} .

Note also that all the weights in \mathcal{P} are integers. Let $t' = \frac{c}{d}$, where c and d are co-prime. By Theorem 25, a finite solution to \mathcal{P} is of length at most d . Hence, one can check in space polynomial in d whether there exists a word w of length at most d , such that $\mathcal{A}(w) = t$. \square

In the case of functional automata, it is shown in [16] that the non-strict versions of the inclusion and universality problems are decidable in PTIME. They leave the strict versions of these problems as an open question. Our result about the exact-value problem provides an immediate solution to these open problems.

Theorem 27. *The inclusion and universality problems of functional discounted-sum automata are decidable in PSPACE.*

Proof. The \leq -inclusion and \leq -universality problems are shown in [16] to be in PTIME. A solution to the $<$ -universality problem then directly follows from Theorem 26: Given a functional DSA \mathcal{A} and a threshold value t , for every finite word w $\mathcal{A}(w) < t$ iff for every word w we have $\mathcal{A}(w) \leq t$, and there is no word w such that $\mathcal{A}(w) = t$.

As for the $<$ -inclusion problem, given functional DSAs \mathcal{A} and \mathcal{B} , one can construct in PSPACE a functional DSA \mathcal{C} , such that for every word w we have $\mathcal{C}(w) = \mathcal{A}(w) - \mathcal{B}(w)$ [16]. Then, the $<$ -inclusion problem for \mathcal{A} and \mathcal{B} reduces to the $<$ -universality problem for \mathcal{C} and the threshold 0. \square

Results for infinite words

In the case of infinite words, TDS reduces to the $<$ -universality problem, which in turn reduces to the $<$ -inclusion problem. As for the other direction, there is a partial implication: given a DSA \mathcal{A} and a threshold t , one can define a corresponding GTDS \mathcal{P} , such that an answer that \mathcal{P} has no solution would provide a decision procedure for the universality question with respect to \mathcal{A} and t . We show these connections below.

Lemma 28. *For every instance \mathcal{P} of TDS_{01} with a discount factor λ and a target value t , one can compute in polynomial time a discounted-sum automaton \mathcal{A} , such that \mathcal{P} has a solution iff \mathcal{A} is not universal.*

Proof. First, we claim that the TDS_{01} instance \mathcal{P} is equivalent to the TDS instance \mathcal{P}' with $a = -t(1-\lambda)$, $\beta = 1 - t(1-\lambda)$ and $t' = 0$. Indeed, subtracting $t(1-\lambda)$ from every element in a discounted-sum sequence $w \in \{0,1\}^\omega$, provides a discounted-sum sequence $w' \in \{-t(1-\lambda), 1-t(1-\lambda)\}^\omega$, such that $w' = w - t(1-\lambda) \sum_{i=0}^{\infty} \lambda^i = w - \frac{t(1-\lambda)}{1-\lambda} = w - t$.

Consider the discounted-sum automaton \mathcal{A} of Figure 1 over the alphabet $\{a, b\}$. The automaton \mathcal{A} has exactly two runs over a word w – a run r_1 that is solely in q_1 and a run r_2 that is solely in q_2 . The value that \mathcal{A} assigns to w , denoted $\mathcal{A}(w)$, is the minimum between the value of r_1 on w , denoted $r_1(w)$ and the value of r_2 on w , denoted $r_2(w)$. Note that, by the automaton weights to a and b , we have for every word w that $r_1(w) = -r_2(w)$. Hence, for every word w , $\mathcal{A}(w) = 0$ iff $r_1(w) = 0$.

Now, the infinite universality problem of \mathcal{A} asks whether for all infinite words w , we have $\mathcal{A}(w) < 0$. Thus, the answer to the universality problem is “no” iff there is a sequence $w \in \{a, b\}^\omega$ such that $r_1(w) = 0$, which is true iff there answer to the given TDS_{01} problem \mathcal{P} is “true”. \square




Figure 1: The discounted-sum automaton, with a discount factor λ , whose $<$ -universality problem is equivalent to a TDS_{01} instance with a discount factor λ and a target value t .

Theorem 29. *If TDS is undecidable then so are the universality and inclusion problems of discounted-sum automata over infinite words.*

The next theorem demonstrates the close connection between the target discounted-sum problem and the core difficulty in solving the universality problem.

Theorem 30. *Consider an instance \mathcal{P} of GTDS with a discount factor λ , weights a_1, \dots, a_k , and a target value $t \in \mathbb{Q}$. If \mathcal{P} has no solution, then the $<$ -universality and \leq -universality problems with the threshold t are decidable over the class of finite-words (resp., infinite-words) DSA with a discount factor λ and weights from $\{a_1, \dots, a_k\}$.*

Proof. Let \mathcal{P}' be the GTDS problem over natural weights and a target t' that is equivalent to \mathcal{P} (Theorem 18). Consider the tree consisting of all runs of the explorations of t' in base $\frac{1}{\lambda}$, as described in Sections 4 and 5. Since \mathcal{P}' is known to have no solution, each path of the tree is finite. As the degree of every node in the tree is also finite, we get by König’s Lemma that the tree is finite, having some height H .

Now, recall that each run of the exploration stops in one of four cases: i) reaching the target value; ii) repeating a previous gap, which guarantees an eventually-periodic solution; iii) reaching a gap that is too big to be recovered, implying that every continuation of this prefix will be below the threshold; and iv) exceeding the threshold, implying that every continuation of this prefix will be above the threshold.

As \mathcal{P}' is known to have no solution, all the exploration runs are guaranteed to stop for either reason (iii) or (iv) above. In both cases, all the (finite and infinite) continuations this finite sequence of weights is known to provide a discounted-sum that is either above or below the threshold.

Hence, the finite and infinite universality problems can be decided by exploring the above finite tree in parallel to running the DSA \mathcal{A} over input words of length up to H . For each such word, we can tell whether all of its continuations are above or below the threshold t . \square

7 Connections to Other Areas

In this section, we show the connection between TDS and various areas in mathematics and computer science. In particular, we show that if TDS is undecidable then so is the reachability problem of one-dimensional piecewise affine maps, and from the other direction, if TDS is decidable then so is the membership problem of a middle- k th Cantor set.

7.1 Piecewise Affine Maps

We shall show that TDS can be reduced to the reachability problem of one-dimensional piecewise affine maps.

A *Piecewise affine map* (PAM) of dimension d is a function $f : \mathbb{R}^d \rightarrow \mathbb{R}^d$, where the domain is partitioned into disjoint sections, and from each section there is a different affine map. That is, each affine map is $f(x) = a_i x + b_i$, where a_i and b_i are vectors of dimension d , specific to the i^{th} section, and a_i is non-zero.

The *reachability problem* asks, given a PAM f , and points s and t , whether there exists $n \in \mathbb{N}$, such that $f^n(s) = t$. The problem is undecidable for 2, or more, dimensions [25, 29, 30]. It is open for one dimension, even when only having two sections [2, 3, 26]. (For a single section, meaning when the mapping is affine but not “piecewise”, the reachability problem is known as the “orbit problem”, and it is decidable for all dimensions over the rationals [24, 35].)

Next, we establish the connection between TDS and the above reachability problem.

Lemma 31. *For every instance \mathcal{P} of TDS_{01} with a discount factor $\lambda < \frac{1}{2}$ and a target value t , one can compute a one-dimensional PAM f , such that \mathcal{P} has a solution iff 1 is not reachable from t via f .*

Proof. Let $t = \frac{a}{b}$ and $\lambda = \frac{p}{q}$. We may assume that $t \leq \frac{\lambda}{1-\lambda}$, as otherwise \mathcal{P} has no solution. We define the following one-dimensional PAM instance:

$$f(x) = \begin{cases} \frac{x}{\lambda} & x < \lambda \\ \frac{x-\lambda}{\lambda} & \lambda \leq x \leq \frac{\lambda}{1-\lambda} \\ b \cdot q \cdot x & \frac{\lambda}{1-\lambda} < x < 1 \\ x & x = 1 \\ p \cdot x & 1 < x < 2 \\ x - 1 & x \geq 2 \end{cases}$$

Consider the sequence of gaps in exploring the representation of t in base $\frac{1}{\lambda}$, using only 0’s and 1’s. Note that iterations of f exactly match that sequence if it is infinite and extend that sequence of gaps if the exploration is finite.

- The initial gap t is the starting point in iterating f .
- If the current gap is smaller than $0.10_{[\frac{1}{\lambda}]}$, which equals λ , it is multiplied by $\frac{1}{\lambda}$ for getting the new gap (as a 0 is chosen for the current position in the representation).
- If the current gap is equal to or bigger than $0.10_{[\frac{1}{\lambda}]}$ and smaller than or equal to $0.1_{[\frac{1}{\lambda}]}$, which equals $\frac{\lambda}{1-\lambda}$, it is multiplied by $\frac{1}{\lambda}$ and 1 is subtracted to get the new gap $\frac{x-\lambda}{\lambda}$ (as 1 is at the current position in the representation).
- If the current gap is bigger than $0.1_{[\frac{1}{\lambda}]}$, which equals $\frac{\lambda}{1-\lambda}$, then the exploration stops, as there is no relevant representation. Consider the first i such that $f^i(t) > \frac{\lambda}{1-\lambda}$. Observe that $f^i(t)$

belongs to the interval $(\frac{\lambda}{1-\lambda}, 1)$. Also, $f^i(t) = f^i(\frac{a}{b})$ and it is of the form $\frac{K}{b \cdot p^i}$, where $K \in \mathbb{N}$. Then, $f^{i+1}(t) = \frac{q \cdot K}{p^i} > \frac{b \cdot q \cdot \lambda}{1-\lambda} = \frac{b \cdot q \cdot p}{q-p} > 1$.

- We can show, by induction on k , that for every $t' = \frac{c}{p^k} \geq 1$ some iteration of f on t' reaches 1. Indeed, if t' is a natural number, then $f^{t'-1}(t') = t' - (t' - 1) = 1$. Otherwise, for $n = \lfloor t' \rfloor - 1$ we have $f^n(t') = t' - n$ belongs to $(1, 2)$ and $f^{n+1}(t')$ is of the form $\frac{c'}{p^{k-1}} \geq 1$.

Thus, 1 is reachable from t by iterating f iff at some point in exploring t in base $\frac{1}{\lambda}$ we have to use a digit different from 0 or 1. \square

Theorem 32. *If TDS is undecidable then so is the reachability problem of one-dimensional piecewise affine maps.*

7.2 Cantor Sets

We shall show that the membership problem of a middle- k th Cantor set can be reduced to TDS.

The *Cantor set* contains the numbers between 0 and 1 that are not removed by iteratively removing the middle third: at the first step, the numbers in $(\frac{1}{3}, \frac{2}{3})$ are removed; then, the middle third of both the upper and lower parts are removed, and so on.

With base 3, a number between 0 and 1 has a representation with only 0's and 2's if and only if it is in the Cantor set. It is easy to check if a rational number is in the Cantor set, since if so, it must have an eventually-periodic representation. A variation of the Cantor set, where at each step only the $\frac{1}{5}$ upper and lower parts remain, is very analogous.

However, removing, for example, the middle $\frac{1}{5}$, which generates a set that is usually termed the "middle-fifth Cantor set", is something different – at each step the remained parts should be multiplied by $\frac{5}{2}$, which makes it analogous to a representation in base $\frac{5}{2}$. In general, for every integer $k > 2$, the set of numbers between 0 and 1 that are not removed by iteratively removing the middle k th is termed the *middle- k th Cantor set* [12, 13, 18].

We show below that a middle- k th Cantor set corresponds to the set of numbers that have a representation in base $\frac{2k}{k-1}$ with only 0's and 1's.

Lemma 33. *Consider an integer $k > 2$ and a number $t \in [0, 1]$. Then t belongs to the middle- k th Cantor set iff TDS₀₁ with a discount factor $\lambda = \frac{k-1}{2k}$ and a target value $t \frac{k-1}{k+1}$ has a solution.*

Proof. Consider the set S of numbers that have a $\frac{2k}{k-1}$ -representation with only 0's and 1's. By the uniqueness of the representation (Lemma 10), the set S can be achieved by the limit of the following iterative procedure:

1. We start with the set S_1 of all numbers between 0.0^ω and $0.1^\omega_{[\frac{2k}{k-1}]}$, which is $[0, \frac{k-1}{k+1}]$.
 2. We generate the set S_2 by removing from S_1 all the numbers that cannot be represented with only 0's and 1's according to the first digit, namely the numbers that are smaller than $0.10^\omega_{[\frac{2k}{k-1}]}$, which equals $\frac{k-1}{2k}$, and bigger than $0.01^\omega_{[\frac{2k}{k-1}]}$, which equals $\frac{k-1}{2k} \cdot \frac{k-1}{k+1}$. Note that S_2 has two separate segments – $S_2^0 = [0.00^\omega_{[\frac{2k}{k-1}]}, 0.01^\omega_{[\frac{2k}{k-1}]}]$ and $S_2^1 = [0.10^\omega_{[\frac{2k}{k-1}]}, 0.11^\omega_{[\frac{2k}{k-1}]}]$.
 3. We generate the set S_3 by removing from S_2 all the numbers that cannot be represented with only 0's and 1's according to the second digit. That is, we remove from S_2^0 the numbers that are smaller than $0.010^\omega_{[\frac{2k}{k-1}]}$ and bigger than $0.001^\omega_{[\frac{2k}{k-1}]}$, and from S_2^1 the numbers that are smaller than $0.110^\omega_{[\frac{2k}{k-1}]}$ and bigger than $0.101^\omega_{[\frac{2k}{k-1}]}$.
- i) In the i -th iteration, we generate S_i by removing from S_{i-1} all the numbers that cannot be represented with only 0's and 1's at the i -th position. In consequence, S_i consists of numbers that have a representation which up to i th position consists of 0's and 1's.

Then, $S = \bigcap_{i=1}^{\infty} S_i$.

We claim that every removed segment is exactly the middle k th of the segment from which it is removed. We show it for the case of generating S_2 from S_1 , while all other cases are analogous, as they are generated in the exact same way, with just a shift to the right of the representation. Recall that the set S_1 is the segment $[0, \frac{k-1}{k+1}]$, and S_2 is generated from it by removing the segment $(\frac{k-1}{2k} \cdot \frac{k-1}{k+1}, \frac{k-1}{2k})$.

We first show that the size of the removed segment is $\frac{1}{k}$ of the size of S_1 . Indeed:

$$\frac{\frac{k-1}{2k} - \frac{k-1}{2k} \cdot \frac{k-1}{k+1}}{\frac{k-1}{k+1}} = \frac{\frac{k-1}{2k} \cdot \frac{2}{k+1}}{\frac{k-1}{k+1}} = \frac{2(k-1)}{2k(k-1)} = \frac{1}{k}$$

Next, we show that the removed segment is in the middle, meaning that the size of the lower segment of S_2 is $\frac{(k-1)/2}{k}$ of the size of S_1 . Indeed: $\frac{\frac{k-1}{2k} \cdot \frac{k-1}{k+1}}{\frac{k-1}{k+1}} = \frac{k-1}{2k}$.

Now, we showed that the numbers that have a $\frac{2k}{k-1}$ -representation with only 0's and 1's are exactly the numbers in the middle- k th Cantor set of the segment $[0, \frac{k-1}{k+1}]$. Hence, by the multiplicative nature of the middle- k th removal procedure, a number t is in the middle- k th Cantor set of the segment $[0, 1]$ iff $t \frac{k-1}{k+1}$ has a $\frac{2k}{k-1}$ -representation with only 0's and 1's. \square

Theorem 34. *If TDS is decidable then so is the membership problem in the middle- k th Cantor set.*

Remark 35. *One may wonder why a representation in base $\frac{5}{2}$, for example, that only uses the 0 and 1 digits is not similar to the standard Cantor set, with the only difference of removing the upper third rather than the middle third. This follows the intuition that, at the n th step, we remove the numbers whose $\frac{5}{2}$ -representation has a 2 in the n th position. The problem is that it will also remove numbers that do have a representation with only 0's and 1's, as with a nonintegral base, the representation need not be unique.*

8 Conclusions

The target discounted-sum problem, which is identified and analyzed for the first time in this paper, turns out to be related to several open problems in math and CS, among them are problems of β -expansions, discounted-sum automata and games, piecewise affine maps, and generalizations of the Cantor set. We established a partial solution to the target discounted-sum problem, resolving its restrictions to finite and eventually-periodic sequences, as well as to various specific discount factors, among which are the cases that $\lambda = \frac{1}{n}$, for every $n \in \mathbb{N}$. We generalized our solutions to an extended version of the target discounted-sum problem, in which there may be arbitrarily many weights and an ω -regular constraint on the allowed sequences. Using these generalized solutions, we solved some open problems on discounted-sum automata over finite words.

Acknowledgement. This research was supported in part by the European Research Council (ERC) under grant 267989 (QUAREM) and by the Austrian Science Fund (FWF) under grants S11402-N23 (RiSE) and Z211-N23 (Wittgenstein Award).

References

- [1] D. Andersson. An improved algorithm for discounted payoff games. In *Proc. of ESSLLI Student Session*, pages 91–98, 2006.
- [2] E. Asarin, O. Maler, and A. Pnueli. Reachability analysis of dynamical systems having piecewise-constant derivatives. *Theoretical Computer Science*, 138(1):35 – 65, 1995.
- [3] E. Asarin and G. Schneider. Widening the boundary between decidable and undecidable hybrid systems. In *proceedings of CONCUR*, volume 2421 of *Lecture Notes in Computer Science*, pages 193–208, 2002.
- [4] U. Boker and T. A. Henzinger. Determinizing discounted-sum automata. In *proceedings of CSL*, pages 82–96, 2011.

- [5] U. Boker and T. A. Henzinger. Exact and approximate determinization of discounted-sum automata. *Logical Methods in Computer Science*, 10(1), 2014.
- [6] V. Bruyère, N. Meunier, and J. R. Raskin. Secure equilibria in weighted games. In *proceedings of LICS*, 2014.
- [7] K. Chatterjee, L. Doyen, and T. A. Henzinger. Quantitative languages. In *proceedings of CSL*, LNCS 5213, pages 385–400, 2008.
- [8] K. Chatterjee, V. Forejt, and D. Wojtczak. Multi-objective discounted reward verification in graphs and mdps. In *proceedings of LPAR*, volume 8312 of *LNCS*, pages 228–242, 2013.
- [9] J. Conway. Unpredictable iterations. *Number Theory Conference*, pages 49–52, 1972.
- [10] L. de Alfaro, M. Faella, T. A. Henzinger, R. Majumdar, and M. Stoelinga. Model checking discounted temporal properties. *Theoretical Computer Science*, 345(1):139–170, 2005.
- [11] L. de Alfaro, T. A. Henzinger, and R. Majumdar. Discounting the future in systems theory. In *Proc. of ICALP*, volume 2719 of *LNCS*, pages 1022–1037, 2003.
- [12] M. Dindoš. Generalized cantor sets and sets of sums of convergent alternating series. *Journal of Applied Analysis*, 7(1):131–150, 2001.
- [13] V. Ya. Eiderman. Capacities of generalized cantor sets. In *Selected Topics in Complex Analysis*. Birkhäuser, 2005.
- [14] P. Erdős, I. Joó, and V. Komornik. On the number of q-expansions. *Ann. Univ. Sci. Budap. Rolando Eötvös, Sect. Math.*, 37:109–118, 1994.
- [15] G. Everest, A. van der Poorten, I. Shparlinski, and T. Ward. *SRecurrence Sequences*. American Mathematical Society, 2003.
- [16] E. Filiot, R. Gentilini, and J. R. Raskin. Quantitative languages defined by functional automata. In *CONCUR 2012, CONCUR’12*, pages 132–146, 2012.
- [17] E. Filiot, R. Gentilini, and J. R. Raskin. Finite-valued weighted automata. In *FSTTCS 2014*, pages 133–145, 2014.
- [18] D. K. Ganguli and S. Ray. A note on generalised Cantor set $C_\lambda(0 < \lambda < 1)$. *Soochow Journal of Mathematics*, 19(2):119–127, 1993.
- [19] R. Gerth, D. Peled, M. Y. Vardi, and P. Wolper. Simple on-the-fly automatic verification of linear temporal logic. In *In Protocol Specification Testing and Verification*, pages 3–18. Chapman & Hall, 1995.
- [20] H. Gimbert and W. Zielonka. Limits of multi-discounted markov decision processes. In *Proc. of LICS*, pages 89–98, 2007.
- [21] P. Glendinning and N. Sidorov. Unique representations of real numbers in non-integer bases. *Mathematical Research Letters*, 8(4):535–543, 2001.
- [22] K.G. Hare. Beta-expansions of pisot and salem numbers. In *Waterloo Workshop in Computer Algebra*, 2006.
- [23] P. Hunter and J. R. Raskin. Quantitative Games with Interval Objectives. In *FSTTCS 2014*, volume 29, pages 365–377, 2014.
- [24] R. Kannan and R. J. Lipton. Polynomial-time algorithm for the orbit problem. *J. ACM*, 33(4):808–821, 1986.
- [25] Pascal Koiran, Michel Cosnard, and Max H. Garzon. Computability with low-dimensional dynamical systems. *Theor. Comput. Sci.*, 132(2):113–128, 1994.

- [26] O. Kurganskyy, I. Potapov, and F. Sancho-Caparrini. Reachability problems in low-dimensional iterative maps. *Int. J. Found. Comput. Sci.*, 19(4):935–951, 2008.
- [27] E. Lehtonen. Two undecidable variants of collatzs problems. *Theoretical Computer Science*, 407(1):596–600, 2008.
- [28] K. Mahler. An unsolved problem on the powers of $\frac{3}{2}$. *The journal of the Australian mathematical society*, 8(2):313–321, 1968.
- [29] C. Moore. Unpredictability and undecidability in dynamical systems. *Phys. Rev. Lett.*, 64(20):2354–2357, 1990.
- [30] C. Moore. Generalized shifts: Unpredictability and undecidability in dynamical systems. *Nonlinearity*, 4(2):199–230, 1991.
- [31] J. Ouaknine and J. Worrell. Positivity problems for low-order linear recurrence sequences. In *Proceedings of SODA*, pages 366–379, 2014.
- [32] W. Parry. On the β -expansions of real numbers. *Acta Mathematica Academiae Scientiarum Hungarica*, 11(3-4):401–416, 1960.
- [33] A. Rényi. Representations for real numbers and their ergodic properties. *Acta Mathematica Academiae Scientiarum Hungarica*, 8(3-4):477–493, 1957.
- [34] K. Schmidt. On periodic expansions of pisot numbers and salem numbers. *Bulletin of the London Mathematical Society*, 12(4):269–278, 1980.
- [35] H. S. Shank. The rational case of a matrix problem of harrison. *Discrete Mathematics*, 28(2):207 – 212, 1979.
- [36] N. Sidorov. Personal communication, 2013.
- [37] M. Y. Vardi. An automata-theoretic approach to linear temporal logic. In *Logics for Concurrency*, volume 1043 of *LNCS*, pages 238–266. Springer-Verlag, 1996.