

Uniprocessor Feasibility of Sporadic Tasks Remains coNP-complete Under Bounded Utilization

Pontus Ekberg and Wang Yi
Uppsala University, Sweden
Email: {pontus.ekberg | yi}@it.uu.se

Abstract—A central problem in real-time scheduling theory is to decide whether a sporadic task system with constrained deadlines is feasible on a preemptive uniprocessor. It is known that this problem is strongly coNP-complete in the general case, but also that there exists a pseudo-polynomial time solution for instances with utilization bounded from above by any constant c , where $0 < c < 1$. For a long time it has been unknown whether the bounded case also has a polynomial-time solution. We show that for any choice of the constant c , such that $0 < c < 1$, the bounded feasibility problem is (weakly) coNP-complete, and thus that no polynomial-time solution exists for it, unless $P = NP$.

I. INTRODUCTION

We let a sporadic task system be defined as a finite multiset T of *tasks*, where each task is a triple $(e, d, p) \in \mathbb{N}_+^3$, representing its worst-case execution time, relative deadline and minimum inter-arrival separation (or period), respectively.

A sporadic task *generates* potentially unbounded sequences of *jobs*. A job is an instance of the task’s workload, characterized by a release time, an execution time and an absolute deadline. A job from task (e, d, p) has execution time not larger than e time units and absolute deadline exactly d time units after its release time. Release times of two consecutive jobs from (e, d, p) are separated by at least p time units. A sporadic task system T generates any interleaving of job sequences that can be generated by each of the tasks $(e, d, p) \in T$.

For a sequence of jobs to be successfully scheduled, every job must be executed for a total duration equal to its execution time, between its release time and its absolute deadline. In the following we assume a preemptive uniprocessor, meaning that only one job can be executed at a time, but a job can be paused and resumed at a later time at no additional cost.

Definition I.1 (Feasibility). *A task system T is feasible if and only if there is some scheduling algorithm that will successfully schedule all job sequences that can be generated by T .* \square

A task system T is said to have *implicit deadlines* if $d = p$ for all $(e, d, p) \in T$, and said to have *constrained deadlines* if $d \leq p$ for all $(e, d, p) \in T$.

The *utilization* of a task system T is a measure of its asymptotic resource requirements, and is defined as

$$U(T) \stackrel{\text{def}}{=} \sum_{(e,d,p) \in T} \frac{e}{p}. \quad (1)$$

Liu and Layland [1] showed that a task system T with implicit deadlines is feasible if and only if $U(T) \leq 1$, but this is not a sufficient condition with constrained deadlines.

Dertouzos [2] showed that Earliest Deadline First (EDF) is an optimal scheduling algorithm on preemptive uniprocessors, for all sequences of independent jobs. This means that the terms “feasibility” and “EDF-schedulability” can be used interchangeably here.

Theorem I.2 ([2]). *A task system T is feasible if and only if it is EDF-schedulable.* \square

A related workload model is that of (strictly) periodic tasks. A periodic task system is a multiset of quadruples $(s, e, d, p) \in \mathbb{N} \times \mathbb{N}_+^3$. The difference to the sporadic task model is that the release times of two consecutive jobs from task (s, e, d, p) must be exactly p time units apart, and the release time of the first job is fixed at time point s . A periodic task system is *synchronous* if $s = 0$ for all tasks (s, e, d, p) , and *asynchronous* otherwise. It is known [3] that feasibility testing of sporadic tasks is equally hard as that of synchronous periodic tasks, which means that the terms “sporadic” and “synchronous periodic” can be used interchangeably for the results in this paper as well.

Theorem I.3 ([3]). *A sporadic task system T is feasible if and only if the synchronous periodic task system $\{(0, e, d, p) \mid (e, d, p) \in T\}$ is feasible.* \square

A classic result [4], [5] is that the feasibility problem for asynchronous periodic task systems with constrained deadlines is strongly coNP-complete¹, even if restricted to instances where the utilization is bounded from above by a constant c , such that $0 < c < 1$. However, Baruah et al. [5], [3] found a pseudo-polynomial time algorithm for the special case of synchronous periodic (or sporadic) task systems with such a priori bounded utilization.

Except for the existence of this pseudo-polynomial time algorithm, the complexity of the sporadic feasibility problem remained open for many years. It was listed as one of the “most important open algorithmic problems in real-time scheduling” by Baruah and Pruhs [6]. A few years ago, Eisenbrand and Rothvoß [7] showed that the general case (with unbounded utilization) is weakly coNP-complete. This was recently strengthened by Ekberg and Yi [8], who showed that the general case is strongly coNP-complete.

¹Recall that the problems that are NP- or coNP-complete in the strong sense do not permit even pseudo-polynomial time solutions, unless $P = NP$, while those that are NP- or coNP-complete in the weak (or ordinary) sense might have such solutions.

Theorem I.4 ([8]). *Deciding whether a sporadic task system with constrained deadlines is feasible on a preemptive uniprocessor is coNP-complete in the strong sense.* \square

For the case with bounded utilization, it was still unknown whether a polynomial-time solution existed, and Eisenbrand and Rothvoß [7] conjectured that it did. In this paper, however, we show that for any choice of the constant c , such that $0 < c < 1$, the feasibility problem for sporadic task systems restricted to instances with utilization at most c is weakly coNP-complete, and thus that no polynomial-time solutions exist for it, unless $P = NP$. As it is known that this problem has a pseudo-polynomial time solution, this is a complete classification. A summary of the complexity of the periodic feasibility problem is shown in Figure 1.

	General case	Utilization bounded by some constant c , such that $0 < c < 1$
Asynchronous periodic tasks	coNP-complete in the strong sense. [5]	coNP-complete in the strong sense. [5]
Synchronous periodic tasks (or sporadic)	coNP-complete in the strong sense. [8]	coNP-complete in the weak sense (from this work).

Fig. 1. Complexity classification of the feasibility problem of periodic tasks with constrained deadlines on preemptive uniprocessors.

Let FEASIBILITY denote the general case of the decision problem of whether a given sporadic task system with constrained deadlines is feasible on a preemptive uniprocessor. Let c -FEASIBILITY denote the same problem restricted to instances with utilization at most c . Our hardness proof consists of a polynomial transformation (a polynomial-time many-one reduction) from FEASIBILITY to c -FEASIBILITY for any constant c , such that $0 < c < 1$. The transformation produces some parameter values that grow exponentially in the size of the FEASIBILITY instance. Therefore, c -FEASIBILITY is only shown to be weakly coNP-hard (as expected, since a pseudo-polynomial time solution is known) despite the strong coNP-hardness of FEASIBILITY.

II. PRELIMINARIES

In the hardness proof in the next section we will make use of the following theorem due to Baruah et al. [3]

Theorem II.1 ([3]). *A sporadic task system \mathbb{T} with constrained deadlines is feasible on a preemptive uniprocessor if and only if $U(\mathbb{T}) \leq 1$ and*

$$\forall \ell \in \{0, 1, \dots, B\}, \quad \text{dbf}(\mathbb{T}, \ell) \leq \ell, \quad (2)$$

where $B = \mathcal{P}(\mathbb{T}) + \max\{d \mid (e, d, p) \in \mathbb{T}\}$, and where

$$\text{dbf}(\mathbb{T}, \ell) \stackrel{\text{def}}{=} \sum_{(e, d, p) \in \mathbb{T}} \left(\left\lfloor \frac{\ell - d}{p} \right\rfloor + 1 \right) e \quad (3)$$

is the demand bound function of \mathbb{T} in interval lengths ℓ , and

$$\mathcal{P}(\mathbb{T}) \stackrel{\text{def}}{=} \text{lcm}\{p \mid (e, d, p) \in \mathbb{T}\} \quad (4)$$

is \mathbb{T} 's hyper-period. \square

For our purposes it will be helpful to reformulate Eq. (2) in the above theorem in terms of *slack* and slightly tighten the upper bound B used for the values of ℓ .² First we prove a simple lemma that will also be useful later on.

Lemma II.2. *If \mathbb{T} is a sporadic task system with constrained deadlines and $k \in \mathbb{N}$, then*

$$\text{dbf}(\mathbb{T}, k\mathcal{P}(\mathbb{T}) + \ell) = k\mathcal{P}(\mathbb{T})U(\mathbb{T}) + \text{dbf}(\mathbb{T}, \ell).$$

Proof: From the fact that each of the periods divides the hyper-period, we get

$$\begin{aligned} \text{dbf}(\mathbb{T}, k\mathcal{P}(\mathbb{T}) + \ell) &= \sum_{(e, d, p) \in \mathbb{T}} \left(\left\lfloor \frac{k\mathcal{P}(\mathbb{T}) + \ell - d}{p} \right\rfloor + 1 \right) e \\ &= \sum_{(e, d, p) \in \mathbb{T}} \left(\frac{k\mathcal{P}(\mathbb{T})}{p} + \left\lfloor \frac{\ell - d}{p} \right\rfloor + 1 \right) e \\ &= \sum_{(e, d, p) \in \mathbb{T}} \left(\frac{ek\mathcal{P}(\mathbb{T})}{p} + \left(\left\lfloor \frac{\ell - d}{p} \right\rfloor + 1 \right) e \right) \\ &= k\mathcal{P}(\mathbb{T})U(\mathbb{T}) + \text{dbf}(\mathbb{T}, \ell). \quad \blacksquare \end{aligned}$$

As we reformulate Theorem II.1, we use this lemma to tighten the bound on the values of ℓ that we need to consider.

Corollary II.3. *A sporadic task system \mathbb{T} with constrained deadlines is feasible on a preemptive uniprocessor if and only if $U(\mathbb{T}) \leq 1$ and*

$$\forall \ell \in \{0, 1, \dots, \mathcal{P}(\mathbb{T}) - 1\}, \quad \text{slack}(\mathbb{T}, \ell) \geq 0, \quad (5)$$

where

$$\text{slack}(\mathbb{T}, \ell) \stackrel{\text{def}}{=} \ell - \text{dbf}(\mathbb{T}, \ell). \quad (6)$$

Proof: To see that the smaller range of values for ℓ considered in Eq. (5) is sufficient, note that from Lemma II.2 and $U(\mathbb{T}) \leq 1$ we have, for any $k \in \mathbb{N}$,

$$\begin{aligned} \text{dbf}(\mathbb{T}, k\mathcal{P}(\mathbb{T}) + \ell) &= k\mathcal{P}(\mathbb{T})U(\mathbb{T}) + \text{dbf}(\mathbb{T}, \ell) \\ &\leq k\mathcal{P}(\mathbb{T}) + \text{dbf}(\mathbb{T}, \ell). \end{aligned}$$

Therefore,

$$\begin{aligned} \text{slack}(\mathbb{T}, k\mathcal{P}(\mathbb{T}) + \ell) &= k\mathcal{P}(\mathbb{T}) + \ell - \text{dbf}(\mathbb{T}, k\mathcal{P}(\mathbb{T}) + \ell) \\ &\geq k\mathcal{P}(\mathbb{T}) + \ell - k\mathcal{P}(\mathbb{T}) - \text{dbf}(\mathbb{T}, \ell) \\ &= \text{slack}(\mathbb{T}, \ell) \end{aligned}$$

and so if $\text{slack}(\mathbb{T}, \ell) < 0$ holds for any $\ell \in \mathbb{N}$, then it must be the case that $\text{slack}(\mathbb{T}, \ell \bmod \mathcal{P}(\mathbb{T})) < 0$ also holds. \blacksquare

²In [3], the considered tasks could have relative deadlines larger than their periods, which is why the larger value B is used there.

III. REDUCING FEASIBILITY TO c -FEASIBILITY

In this section we describe a polynomial transformation from FEASIBILITY to c -FEASIBILITY, for any given constant c , such that $0 < c < 1$. We start by providing a high-level picture of the transformation in Section III-A, and follow with the details and proofs in Section III-B.

A. An overview

For the transformation we take an arbitrary sporadic task system \mathcal{T}_1 with constrained deadlines and compute a new task system \mathcal{T}_c , also with constrained deadlines, such that

- $U(\mathcal{T}_c) \leq c$ and
- \mathcal{T}_c is feasible if and only if \mathcal{T}_1 is feasible.

Note that if $U(\mathcal{T}_1) > 1$, then \mathcal{T}_1 is infeasible and the transformation is trivially computable by producing any infeasible \mathcal{T}_c , for example $\mathcal{T}_c = \{(1, 1, \lceil 2/c \rceil), (1, 1, \lceil 2/c \rceil)\}$. In the remainder of this paper we assume $U(\mathcal{T}_1) \leq 1$, without loss of generality.

The construction of \mathcal{T}_c follows these four high-level steps.

Step 1: Add a filler task to \mathcal{T}_1 , resulting in task system $\mathcal{T}_{\text{fill}}$, without affecting feasibility. The purpose is to ensure that $\text{slack}(\mathcal{T}_{\text{fill}}, \ell)$ exactly repeats every hyper-period.

Step 2: Scale the deadlines and periods of $\mathcal{T}_{\text{fill}}$ uniformly with a given constant factor σ , producing task system \mathcal{T}_σ . The purpose is to lower utilization while keeping a simple relation between $\text{slack}(\mathcal{T}_{\text{fill}}, \ell)$ and $\text{slack}(\mathcal{T}_\sigma, \ell)$.

Step 3: Construct a “boosting” task system \mathcal{T}_b , which is designed to fill up the extra slack created by the scaling in step 2 at certain sparsely selected points ℓ . The purpose is to effectively negate the scaling done in step 2 without increasing the utilization too much. The extra slack is filled up at points that are offset by different multiples of \mathcal{T}_σ 's hyper-period. This is possible thanks to the repetitive nature of $\text{slack}(\mathcal{T}_\sigma, \ell)$ that we ensured in step 1.

Step 4: Take \mathcal{T}_c as the union of \mathcal{T}_σ and \mathcal{T}_b . Even though $U(\mathcal{T}_c) \leq c$, we will have that there exists some $\ell \in \mathbb{N}$, such that $\text{slack}(\mathcal{T}_1, \ell) < 0$ if and only if there exists some $\ell' \in \mathbb{N}$, such that $\text{slack}(\mathcal{T}_c, \ell') < 0$.

B. The transformation

Here we describe the transformation in detail and prove a number of lemmas along the way.

► Step 1: Ensuring the repetitiveness of the slack function:

In the first step we simply add a filler task τ_{fill} to \mathcal{T}_1 , if needed, to create a new task system $\mathcal{T}_{\text{fill}}$ with a utilization of 1. Figure 2 illustrates this step. Let

$$\mathcal{T}_{\text{fill}} \stackrel{\text{def}}{=} \begin{cases} \mathcal{T}_1 \uplus \{\tau_{\text{fill}}\}, & \text{if } U(\mathcal{T}_1) < 1, \\ \mathcal{T}_1, & \text{otherwise,} \end{cases} \quad (7)$$

where $\tau_{\text{fill}} \stackrel{\text{def}}{=} (\mathcal{P}(\mathcal{T}_1) - \text{dbf}(\mathcal{T}_1, \mathcal{P}(\mathcal{T}_1)), \mathcal{P}(\mathcal{T}_1), \mathcal{P}(\mathcal{T}_1))$ and where \uplus is the multiset union operator.




Fig. 2. The demand bound functions of a small example task system \mathcal{T}_1 and the corresponding $\mathcal{T}_{\text{fill}}$. Note that the difference between the diagonal and $\text{dbf}(\mathcal{T}_{\text{fill}}, \ell)$, which is $\text{slack}(\mathcal{T}_{\text{fill}}, \ell)$, repeats every hyper-period.

Lemma III.1. $U(\mathcal{T}_{\text{fill}}) = 1$.

Proof: If $U(\mathcal{T}_1) = 1$, the lemma follows directly. Assume instead $U(\mathcal{T}_1) < 1$ and note that

$$\begin{aligned} U(\{\tau_{\text{fill}}\}) &= \frac{\mathcal{P}(\mathcal{T}_1) - \text{dbf}(\mathcal{T}_1, \mathcal{P}(\mathcal{T}_1))}{\mathcal{P}(\mathcal{T}_1)} \\ &= 1 - \frac{\sum_{(e,d,p) \in \mathcal{T}_1} \left(\left\lfloor \frac{\mathcal{P}(\mathcal{T}_1) - d}{p} \right\rfloor + 1 \right) e}{\mathcal{P}(\mathcal{T}_1)} \\ &= 1 - \frac{\sum_{(e,d,p) \in \mathcal{T}_1} \left(\frac{e\mathcal{P}(\mathcal{T}_1)}{p} + \left(\left\lfloor \frac{-d}{p} \right\rfloor + 1 \right) e \right)}{\mathcal{P}(\mathcal{T}_1)} \\ &\stackrel{(*)}{=} 1 - U(\mathcal{T}_1), \end{aligned}$$

where for (*) we used the fact that $\lfloor -d/p \rfloor = -1$. Clearly, $U(\mathcal{T}_{\text{fill}}) = U(\mathcal{T}_1) + U(\{\tau_{\text{fill}}\}) = 1$. ■

Now, the property of $\mathcal{T}_{\text{fill}}$ that will be useful for our transformation is captured by the next lemma.

Lemma III.2. For any $\ell \in \{0, 1, \dots, \mathcal{P}(\mathcal{T}_1) - 1\}$ and $k \in \mathbb{N}$,

$$\text{slack}(\mathcal{T}_{\text{fill}}, k\mathcal{P}(\mathcal{T}_1) + \ell) = \text{slack}(\mathcal{T}_1, \ell).$$

Proof: Note that $\mathcal{P}(\mathcal{T}_{\text{fill}}) = \mathcal{P}(\mathcal{T}_1)$, and therefore it

follows from Lemmas II.2 and III.1 that

$$\begin{aligned}
\text{slack}(\mathcal{T}_{\text{fill}}, k\mathcal{P}(\mathcal{T}_1) + \ell) & \\
&= k\mathcal{P}(\mathcal{T}_1) + \ell - \text{dbf}(\mathcal{T}_{\text{fill}}, k\mathcal{P}(\mathcal{T}_1) + \ell) \\
&= k\mathcal{P}(\mathcal{T}_1) + \ell - (k\mathcal{P}(\mathcal{T}_1)U(\mathcal{T}_{\text{fill}}) + \text{dbf}(\mathcal{T}_{\text{fill}}, \ell)) \\
&= \ell - \text{dbf}(\mathcal{T}_{\text{fill}}, \ell) \\
&\stackrel{(*)}{=} \ell - \text{dbf}(\mathcal{T}_1, \ell) \\
&= \text{slack}(\mathcal{T}_1, \ell),
\end{aligned}$$

where for (*) we have $\text{dbf}(\mathcal{T}_{\text{fill}}, \ell) = \text{dbf}(\mathcal{T}_1, \ell)$ because $\mathcal{T}_{\text{fill}}$ and \mathcal{T}_1 at most differ on one task $\tau_{\text{fill}} = (e, d, p)$, for which $d = \mathcal{P}(\mathcal{T}_1) > \ell$ and therefore $\text{dbf}(\{\tau_{\text{fill}}\}, \ell) = 0$. ■

► **Step 2: Scaling down the utilization:** In the second step we apply a scaling factor to the relative deadlines and periods of all the tasks in $\mathcal{T}_{\text{fill}}$. Figure 3 serves as an illustration. Let

$$\sigma \stackrel{\text{def}}{=} \left\lfloor \frac{2}{c} \right\rfloor \quad (8)$$

be the constant scaling factor, where c is the constant defining the set of problems c -FEASIBILITY. Then let

$$\mathcal{T}_\sigma \stackrel{\text{def}}{=} \{(e, \sigma d, \sigma p) \mid (e, d, p) \in \mathcal{T}_{\text{fill}}\} \quad (9)$$

be the scaled task system. The utilization of \mathcal{T}_σ is now bounded by half of c .

Lemma III.3. $U(\mathcal{T}_\sigma) \leq c/2$.

Proof: Directly from Lemma III.1, Eq. (8) and (9). ■

Because of the uniform way in which \mathcal{T}_σ was scaled, we have the following lemma.

Lemma III.4. For any $\ell \in \mathbb{N}$,

$$\text{slack}(\mathcal{T}_\sigma, \sigma\ell) = \text{slack}(\mathcal{T}_{\text{fill}}, \ell) + (\sigma - 1)\ell.$$

Proof: From the definition of \mathcal{T}_σ we get

$$\begin{aligned}
\text{slack}(\mathcal{T}_\sigma, \sigma\ell) &= \sigma\ell - \text{dbf}(\mathcal{T}_\sigma, \sigma\ell) \\
&= \sigma\ell - \sum_{(e,d,p) \in \mathcal{T}_\sigma} \left(\left\lfloor \frac{\sigma\ell - d}{p} \right\rfloor + 1 \right) e \\
&= \sigma\ell - \sum_{(e,d,p) \in \mathcal{T}_{\text{fill}}} \left(\left\lfloor \frac{\sigma\ell - \sigma d}{\sigma p} \right\rfloor + 1 \right) e \\
&= \sigma\ell - \sum_{(e,d,p) \in \mathcal{T}_{\text{fill}}} \left(\left\lfloor \frac{\ell - d}{p} \right\rfloor + 1 \right) e \\
&= \sigma\ell - \text{dbf}(\mathcal{T}_{\text{fill}}, \ell) \\
&= \text{slack}(\mathcal{T}_{\text{fill}}, \ell) + (\sigma - 1)\ell. \quad \blacksquare
\end{aligned}$$

► **Step 3: Generating a boosting task system:** The task system \mathcal{T}_σ created in the last step clearly does not preserve feasibility with regards to $\mathcal{T}_{\text{fill}}$ because of the extra slack introduced by the scaling of parameters. In this step we will craft a task system \mathcal{T}_b , designed to effectively negate this extra slack. The main challenge is in doing so while using only polynomially many tasks and keeping the utilization low enough. This step is more involved than the previous steps and

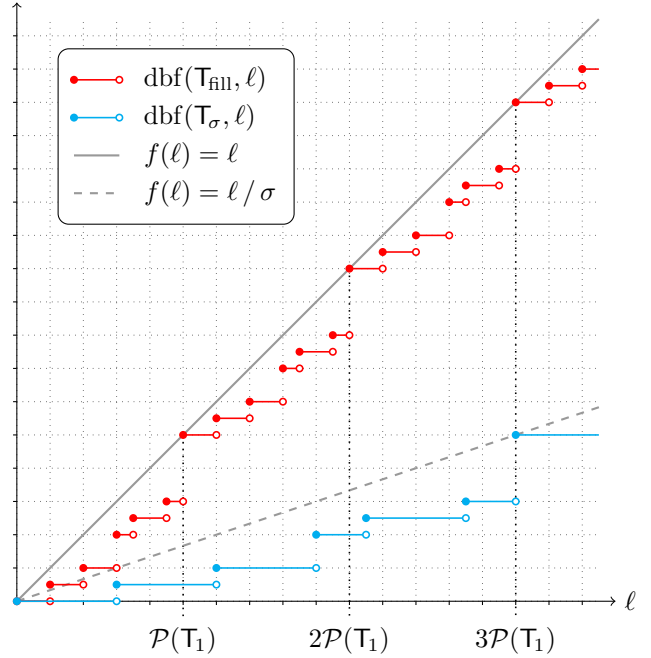


Fig. 3. The scaled task system \mathcal{T}_σ has essentially the same pattern as $\mathcal{T}_{\text{fill}}$, but with a lower utilization. Here we have $\sigma = 3$.

also a bit harder to visualize, but Figure 4 illustrates some of the key concepts. The boosting task system will contain β different tasks that are referred to using their indices,

$$\mathcal{T}_b \stackrel{\text{def}}{=} \{\tau_0, \dots, \tau_{\beta-1}\}, \quad (10)$$

where

$$\beta \stackrel{\text{def}}{=} \lceil \log_2(\mathcal{P}(\mathcal{T}_1)) \rceil. \quad (11)$$

For each boosting task τ_i , let $\tau_i \stackrel{\text{def}}{=} (e_i, d_i, p_i)$, where

$$d_i \stackrel{\text{def}}{=} (\sigma\mathcal{P}(\mathcal{T}_1) + 2)^i \sigma, \quad (12)$$

$$p_i \stackrel{\text{def}}{=} (\sigma\mathcal{P}(\mathcal{T}_1) + 2)^{i+1} \sigma, \quad (13)$$

$$e_i \stackrel{\text{def}}{=} \left(\frac{\sigma - 1}{\sigma} - \sum_{j=0}^{i-1} \frac{e_j}{p_j} \right) d_i. \quad (14)$$

The first thing we need to show is that the boosting task system, as defined above, consists of valid constrained-deadline sporadic tasks. The next lemma establishes this and some slightly stronger constraints on the parameters.

Lemma III.5. For all $\tau_i \in \mathcal{T}_b$, we have

$$(e_i, d_i, p_i) \in \mathbb{N}_+^3 \text{ and } e_i < d_i < p_i.$$

Proof: It follows directly from Eq. (12) and (13) that for any $\tau_i \in \mathcal{T}_b$, we have $d_i, p_i \in \mathbb{N}_+$ and $d_i < p_i$. What remains to be shown is that

$$e_i \in \mathbb{N}_+ \text{ and } e_i < d_i.$$

We use strong induction over i to prove this statement for all $\tau_i \in \mathbb{T}_b$. As there is nothing to show for $i \geq \beta$, we assume $i < \beta$ in the following.

For the base case, note that

$$\begin{aligned} e_0 &= \left(\frac{\sigma - 1}{\sigma} \right) d_0 \\ &= \left(\frac{\sigma - 1}{\sigma} \right) \sigma \\ &= \sigma - 1. \end{aligned}$$

From Eq. (8) it is clear that $\sigma \geq 3$, and therefore $\sigma - 1 \in \mathbb{N}_+$ and $e_0 = \sigma - 1 < \sigma = d_0$.

We split the induction step into three parts, showing that $e_i \in \mathbb{Z}$, that $e_i > 0$ and that $e_i < d_i$, respectively.

Part 1 ($e_i \in \mathbb{Z}$): First, we rewrite e_i in a form without fractions.

$$\begin{aligned} e_i &= \left(\frac{\sigma - 1}{\sigma} - \sum_{j=0}^{i-1} \frac{e_j}{p_j} \right) d_i \\ &= (\sigma - 1)(\sigma \mathcal{P}(\mathbb{T}_1) + 2)^i - \sum_{j=0}^{i-1} e_j \frac{(\sigma \mathcal{P}(\mathbb{T}_1) + 2)^i}{(\sigma \mathcal{P}(\mathbb{T}_1) + 2)^{j+1}} \\ &= (\sigma - 1)(\sigma \mathcal{P}(\mathbb{T}_1) + 2)^i - \sum_{j=0}^{i-1} e_j (\sigma \mathcal{P}(\mathbb{T}_1) + 2)^{i-j-1} \end{aligned}$$

Now, from the induction hypothesis we have $e_j \in \mathbb{N}_+$ for $j < i$, and therefore the rewritten form of e_i consists of only additions, subtractions and (repeated) multiplications of integers. As \mathbb{Z} are closed under these operations, it follows that $e_i \in \mathbb{Z}$.

Part 2 ($e_i > 0$): We have

$$\begin{aligned} e_i &= \left(\frac{\sigma - 1}{\sigma} - \sum_{j=0}^{i-1} \frac{e_j}{p_j} \right) d_i \\ &\stackrel{(*)}{\geq} \left(\frac{\sigma - 1}{\sigma} - \sum_{j=0}^{i-1} \frac{d_j}{p_j} \right) d_i \\ &= \left(\frac{\sigma - 1}{\sigma} - \sum_{j=0}^{i-1} \frac{1}{(\sigma \mathcal{P}(\mathbb{T}_1) + 2)} \right) d_i \\ &= \left(\frac{\sigma - 1}{\sigma} - \frac{i}{(\sigma \mathcal{P}(\mathbb{T}_1) + 2)} \right) d_i \\ &\stackrel{(**)}{>} \left(\frac{\sigma - 1}{\sigma} - \frac{\mathcal{P}(\mathbb{T}_1)}{(\sigma \mathcal{P}(\mathbb{T}_1) + 2)} \right) d_i \\ &> \left(\frac{\sigma - 1}{\sigma} - \frac{1}{\sigma} \right) d_i \\ &\geq \left(\frac{1}{\sigma} \right) d_i \\ &> 0, \end{aligned}$$

where for $(*)$ we used $e_j < d_j$ from the induction hypothesis and for $(**)$ we used $i < \beta = \lceil \log_2(\mathcal{P}(\mathbb{T}_1)) \rceil < \mathcal{P}(\mathbb{T}_1)$.

Part 3 ($e_i < d_i$): Lastly, note that

$$\begin{aligned} e_i &= \left(\frac{\sigma - 1}{\sigma} - \sum_{j=0}^{i-1} \frac{e_j}{p_j} \right) d_i \\ &< \left(1 - \sum_{j=0}^{i-1} \frac{e_j}{p_j} \right) d_i \\ &\stackrel{(*)}{\leq} d_i, \end{aligned}$$

where for $(*)$ we used $e_j \in \mathbb{N}_+$ for $j < i$ from the induction hypothesis.

Taken together, the three parts imply $e_i \in \mathbb{N}_+$ and $e_i < d_i$, which concludes the proof. \blacksquare

Having shown that \mathbb{T}_b is a valid sporadic task system, we can show that it has the properties we seek. We start by showing that the utilization of \mathbb{T}_b is low, namely bounded by half of c .

Lemma III.6. $U(\mathbb{T}_b) < c/2$.

Proof: Consider any $\tau_i \in \mathbb{T}_b$ and note that

$$\begin{aligned} U(\{\tau_i\}) &= e_i / p_i \\ &< d_i / p_i \\ &< 1 / (\sigma \mathcal{P}(\mathbb{T}_1)). \end{aligned}$$

Now, because $\mathcal{P}(\mathbb{T}_1) \in \mathbb{N}_+$, we have $\lceil \log_2(\mathcal{P}(\mathbb{T}_1)) \rceil < \mathcal{P}(\mathbb{T}_1)$ and therefore

$$\begin{aligned} U(\mathbb{T}_b) &< \beta / (\sigma \mathcal{P}(\mathbb{T}_1)) \\ &< 1 / \sigma \\ &\leq c/2. \end{aligned} \quad \blacksquare$$

Now we study the demand of \mathbb{T}_b to see that it provides the right amount of it to selectively fill up the extra slack that was created for \mathbb{T}_σ by the scaling in step 2. With the next lemma we ensure that the demand of \mathbb{T}_b never fills up more than this extra slack.

Lemma III.7. For all $\ell \in \mathbb{N}$,

$$\text{dbf}(\mathbb{T}_b, \ell) \leq \frac{\sigma - 1}{\sigma} \ell.$$

Proof: First, let \mathbb{T}_b^m denote the subset consisting of the smallest $\min(m, \beta)$ tasks in \mathbb{T}_b ,

$$\mathbb{T}_b^m \stackrel{\text{def}}{=} \{\tau_i \in \mathbb{T}_b \mid i < m\}.$$

Then we use induction over m to prove that

$$\forall \ell \in \mathbb{N}, \text{dbf}(\mathbb{T}_b^m, \ell) \leq \frac{\sigma - 1}{\sigma} \ell,$$

for all $m \in \mathbb{N}_+$. Because $\text{dbf}(\mathbb{T}_b^m, \ell) = \text{dbf}(\mathbb{T}_b, \ell)$ for $m \geq \beta$, the lemma will follow.

The base case holds as we have $\mathbb{T}_b^1 = \{\tau_0\}$ and

$$\begin{aligned} \text{dbf}(\{\tau_0\}, \ell) &= \left(\left\lfloor \frac{\ell - d_0}{p_0} \right\rfloor + 1 \right) e_0 \\ &\leq \left(\left\lfloor \frac{\ell - d_0}{d_0} \right\rfloor + 1 \right) e_0 \\ &= \left\lfloor \frac{\ell}{d_0} \right\rfloor e_0 \\ &\leq \frac{e_0}{d_0} \ell \\ &= \frac{\sigma - 1}{\sigma} \ell. \end{aligned}$$

For the induction step we need to show for all $m \in \mathbb{N}_+$ that

$$\begin{aligned} \forall \ell \in \mathbb{N}, \text{dbf}(\mathbb{T}_b^m, \ell) &\leq \frac{\sigma - 1}{\sigma} \ell \implies \\ \forall \ell \in \mathbb{N}, \text{dbf}(\mathbb{T}_b^{m+1}, \ell) &\leq \frac{\sigma - 1}{\sigma} \ell. \end{aligned}$$

The induction step trivially holds for $m \geq \beta$. Assume instead $m < \beta$ and note that $\mathbb{T}_b^{m+1} \setminus \mathbb{T}_b^m = \{\tau_m\}$. Take any $\ell \in \mathbb{N}$ and write $\ell = kd_m + \ell'$ for some $k \in \mathbb{N}$ and $\ell' < d_m$. Then,

$$\begin{aligned} \text{dbf}(\{\tau_m\}, \ell) &= \text{dbf}(\{\tau_m\}, kd_m + \ell') \\ &= \left(\left\lfloor \frac{kd_m + \ell' - d_m}{p_m} \right\rfloor + 1 \right) e_m \\ &\stackrel{(*)}{\leq} \left(\left\lfloor \frac{kd_m}{p_m} \right\rfloor + 1 \right) e_m \\ &\stackrel{(**)}{\leq} ke_m \\ &= k \left(\frac{\sigma - 1}{\sigma} - \sum_{j=0}^{m-1} \frac{e_j}{p_j} \right) d_m \\ &= kd_m \frac{\sigma - 1}{\sigma} - kd_m U(\mathbb{T}_b^m), \end{aligned}$$

where for (*) we used $\ell' < d_m$ and for (**) we used $d_m < p_m$. Now, from Eq. (12) and (13) it is clear that $d_m = p_{m-1} = \mathcal{P}(\mathbb{T}_b^m)$. Using Lemma II.2 and the induction hypothesis we then get

$$\begin{aligned} \text{dbf}(\mathbb{T}_b^m, \ell) &= \text{dbf}(\mathbb{T}_b^m, kd_m + \ell') \\ &= kd_m U(\mathbb{T}_b^m) + \text{dbf}(\mathbb{T}_b^m, \ell') \\ &\leq kd_m U(\mathbb{T}_b^m) + \frac{\sigma - 1}{\sigma} \ell'. \end{aligned}$$

Putting the above together, we get

$$\begin{aligned} \text{dbf}(\mathbb{T}_b^{m+1}, \ell) &= \text{dbf}(\mathbb{T}_b^m, \ell) + \text{dbf}(\{\tau_m\}, \ell) \\ &\leq \frac{\sigma - 1}{\sigma} (kd_m + \ell') \\ &= \frac{\sigma - 1}{\sigma} \ell. \end{aligned}$$

To conclude, we have shown that for all $m \in \mathbb{N}_+$,

$$\forall \ell \in \mathbb{N}, \text{dbf}(\mathbb{T}_b^m, \ell) \leq \frac{\sigma - 1}{\sigma} \ell.$$

In particular, it holds for $m = \beta$ and the lemma follows. \blacksquare

We now consider the points at which \mathbb{T}_b provides exactly the right amount of demand to negate the scaling done for \mathbb{T}_σ in step 2. We call these the *boost points* of \mathbb{T}_b and denote them $\mathcal{B}(\mathbb{T}_b)$.

$$\mathcal{B}(\mathbb{T}_b) \stackrel{\text{def}}{=} \left\{ \ell \in \mathbb{N} \mid \text{dbf}(\mathbb{T}_b, \ell) = \frac{\sigma - 1}{\sigma} \ell \right\} \quad (15)$$

The following lemma lets us identify boost points.

Lemma III.8. *If $\dot{\mathbb{T}} \subseteq \mathbb{T}_b$, then*

$$\sum_{\tau_i \in \dot{\mathbb{T}}} d_i \in \mathcal{B}(\mathbb{T}_b).$$

Proof: Consider the relation $(2^{\mathbb{T}_b}, \subset)$, i.e., the power set of \mathbb{T}_b ordered by strict inclusion, and note that it is well-founded. We prove the lemma using well-founded induction on this relation. To do so we must show that the following implication holds for any $\dot{\mathbb{T}}$, such that $\dot{\mathbb{T}} \subseteq \mathbb{T}_b$.

$$\left(\forall \ddot{\mathbb{T}}, \text{ s.t. } \ddot{\mathbb{T}} \subset \dot{\mathbb{T}}, \sum_{\tau_i \in \ddot{\mathbb{T}}} d_i \in \mathcal{B}(\mathbb{T}_b) \right) \implies \sum_{\tau_i \in \dot{\mathbb{T}}} d_i \in \mathcal{B}(\mathbb{T}_b)$$

We split the proof into two cases.

Case 1 ($\dot{\mathbb{T}} = \emptyset$): Clearly, $\sum_{\tau_i \in \emptyset} d_i = 0$ and from Eq. (15) we have $0 \in \mathcal{B}(\mathbb{T}_b)$.

Case 2 ($\dot{\mathbb{T}} \neq \emptyset$): Now, $\dot{\mathbb{T}}$ must contain a ‘‘largest’’ task τ_α , where

$$\alpha \stackrel{\text{def}}{=} \max\{i \mid \tau_i \in \dot{\mathbb{T}}\}$$

is the largest index of the tasks in $\dot{\mathbb{T}}$. Let

$$\begin{aligned} \ddot{\mathbb{T}} &\stackrel{\text{def}}{=} \dot{\mathbb{T}} \setminus \{\tau_\alpha\}, \\ \ddot{\ell} &\stackrel{\text{def}}{=} \sum_{\tau_i \in \ddot{\mathbb{T}}} d_i, \\ \dot{\ell} &\stackrel{\text{def}}{=} d_\alpha + \ddot{\ell}. \end{aligned}$$

Since $\ddot{\mathbb{T}} \subset \dot{\mathbb{T}}$, we know that $\ddot{\ell} \in \mathcal{B}(\mathbb{T}_b)$ holds from the induction hypothesis. To conclude the proof we must show that $\dot{\ell} \in \mathcal{B}(\mathbb{T}_b)$ holds as well.

Consider now the sets of boosting tasks with indices smaller than α and at least as large as α , respectively.

$$\begin{aligned} \mathbb{T}_{\alpha\downarrow} &\stackrel{\text{def}}{=} \{\tau_i \in \mathbb{T}_b \mid i < \alpha\} \\ \mathbb{T}_{\alpha\uparrow} &\stackrel{\text{def}}{=} \{\tau_i \in \mathbb{T}_b \mid i \geq \alpha\} \end{aligned}$$

Because $\dot{\mathbb{T}} \subseteq \mathbb{T}_{\alpha\downarrow}$ we have

$$\begin{aligned} \dot{\ell} &\leq \sum_{\tau_i \in \mathbb{T}_{\alpha\downarrow}} d_i \\ &= (\sigma \mathcal{P}(\mathbb{T}_1) + 2)^0 \sigma + \dots + (\sigma \mathcal{P}(\mathbb{T}_1) + 2)^{\alpha-1} \sigma \\ &\stackrel{(*)}{=} \sigma \left(\frac{(\sigma \mathcal{P}(\mathbb{T}_1) + 2)^\alpha - 1}{\sigma \mathcal{P}(\mathbb{T}_1) + 2 - 1} \right) \\ &< (\sigma \mathcal{P}(\mathbb{T}_1) + 2)^\alpha \sigma \\ &= d_\alpha, \end{aligned}$$

where for (*) we used the closed form for the sum of a finite geometric series. It follows that for all $\tau_i \in \mathbb{T}_{\alpha\uparrow}$ we have $d_i \geq d_\alpha > \ell$ and therefore

$$\text{dbf}(\mathbb{T}_{\alpha\uparrow}, \ddot{\ell}) = 0.$$

Using this and the induction hypothesis we get

$$\begin{aligned} \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \ddot{\ell}) &= \text{dbf}(\mathbb{T}_b, \ddot{\ell}) - \text{dbf}(\mathbb{T}_{\alpha\uparrow}, \ddot{\ell}) \\ &= \text{dbf}(\mathbb{T}_b, \ddot{\ell}) \\ &= \frac{\sigma - 1}{\sigma} \ddot{\ell}. \end{aligned}$$

Now, from Eq. (12) and (13) we know that p_i divides d_α for all $\tau_i \in \mathbb{T}_{\alpha\downarrow}$. Therefore,

$$\begin{aligned} \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) &= \text{dbf}(\mathbb{T}_{\alpha\downarrow}, d_\alpha + \dot{\ell}) \\ &= \sum_{\tau_i \in \mathbb{T}_{\alpha\downarrow}} \left(\left\lfloor \frac{d_\alpha + \dot{\ell} - d_i}{p_i} \right\rfloor + 1 \right) e_i \\ &= \sum_{\tau_i \in \mathbb{T}_{\alpha\downarrow}} \left(\frac{d_\alpha}{p_i} + \left\lfloor \frac{\dot{\ell} - d_i}{p_i} \right\rfloor + 1 \right) e_i \\ &= \sum_{\tau_i \in \mathbb{T}_{\alpha\downarrow}} \frac{e_i}{p_i} d_\alpha + \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) \\ &= \sum_{\tau_i \in \mathbb{T}_{\alpha\downarrow}} \frac{e_i}{p_i} d_\alpha + \frac{\sigma - 1}{\sigma} \dot{\ell} \end{aligned}$$

Finally,

$$\begin{aligned} \text{dbf}(\mathbb{T}_b, \dot{\ell}) &= \text{dbf}(\mathbb{T}_{\alpha\uparrow}, \dot{\ell}) + \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) \\ &\geq \text{dbf}(\{\tau_\alpha\}, \dot{\ell}) + \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) \\ &\geq \text{dbf}(\{\tau_\alpha\}, d_\alpha) + \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) \\ &= e_\alpha + \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) \\ &= \frac{\sigma - 1}{\sigma} d_\alpha - \sum_{j=0}^{\alpha-1} \frac{e_j}{p_j} d_\alpha + \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) \\ &= \frac{\sigma - 1}{\sigma} d_\alpha - \sum_{\tau_i \in \mathbb{T}_{\alpha\downarrow}} \frac{e_i}{p_i} d_\alpha + \text{dbf}(\mathbb{T}_{\alpha\downarrow}, \dot{\ell}) \\ &= \frac{\sigma - 1}{\sigma} (d_\alpha + \dot{\ell}) \\ &= \frac{\sigma - 1}{\sigma} \dot{\ell} \end{aligned}$$

The inequalities above must be strict equalities based on Lemma III.7, and we conclude that $\dot{\ell} \in \mathcal{B}(\mathbb{T}_b)$. ■

Now that we have a lemma for identifying boost points, we can show that $\mathcal{B}(\mathbb{T}_b)$ cover enough relevant points in \mathbb{N} .

Lemma III.9. *For all $\ell \in \{0, 1, \dots, \mathcal{P}(\mathbb{T}_1) - 1\}$, there exists some $k \in \mathbb{N}$, such that*

$$\sigma(k\mathcal{P}(\mathbb{T}_1) + \ell) \in \mathcal{B}(\mathbb{T}_b).$$

Proof: Take any $\ell \in \{0, 1, \dots, \mathcal{P}(\mathbb{T}_1) - 1\}$ and note that ℓ can be represented in binary using $\beta = \lceil \log_2(\mathcal{P}(\mathbb{T}_1)) \rceil$ bits.

Let

$$\mathbf{b}_{\beta-1} \cdots \mathbf{b}_1 \mathbf{b}_0$$

be this binary representation of ℓ , where \mathbf{b}_0 is the least significant bit. Consider the set of boosting tasks with indices matching the indices of the positive bits in the binary representation of ℓ

$$\dot{\mathbb{T}} \stackrel{\text{def}}{=} \{\tau_i \in \mathbb{T}_b \mid \mathbf{b}_i = 1\}.$$

For any $\tau_i \in \dot{\mathbb{T}}$ we can rewrite

$$\begin{aligned} d_i &= (\sigma\mathcal{P}(\mathbb{T}_1) + 2)^i \sigma \\ &\stackrel{(*)}{=} \left(\sum_{m=0}^i \binom{i}{m} \sigma^m \mathcal{P}(\mathbb{T}_1)^m 2^{i-m} \right) \sigma \\ &= \left(\sum_{m=1}^i \binom{i}{m} \sigma^m \mathcal{P}(\mathbb{T}_1)^m 2^{i-m} + 2^i \right) \sigma \\ &= (k_i \mathcal{P}(\mathbb{T}_1) + 2^i) \sigma, \end{aligned}$$

where

$$k_i = \sum_{m=1}^i \binom{i}{m} \sigma^m \mathcal{P}(\mathbb{T}_1)^{m-1} 2^{i-m}$$

and where for (*) we used the binomial theorem. Let

$$k \stackrel{\text{def}}{=} \sum_{\tau_i \in \dot{\mathbb{T}}} k_i$$

and note that $k \in \mathbb{N}$. We then have

$$\begin{aligned} \sum_{\tau_i \in \dot{\mathbb{T}}} d_i &= \sum_{\tau_i \in \dot{\mathbb{T}}} (k_i \mathcal{P}(\mathbb{T}_1) + 2^i) \sigma \\ &= \sigma \left(\sum_{\tau_i \in \dot{\mathbb{T}}} k_i \mathcal{P}(\mathbb{T}_1) + \sum_{\tau_i \in \dot{\mathbb{T}}} 2^i \right) \\ &= \sigma \left(k \mathcal{P}(\mathbb{T}_1) + \sum \{2^i \mid \mathbf{b}_i = 1\} \right) \\ &= \sigma(k\mathcal{P}(\mathbb{T}_1) + \ell). \end{aligned}$$

To conclude, we note that because $\dot{\mathbb{T}} \subseteq \mathbb{T}_b$, Lemma III.8 gives us

$$\sum_{\tau_i \in \dot{\mathbb{T}}} d_i = \sigma(k\mathcal{P}(\mathbb{T}_1) + \ell) \in \mathcal{B}(\mathbb{T}_b)$$

and the lemma follows. ■

► **Step 4: Assembling the finished task system:** In the last step we put everything together. Let

$$\mathbb{T}_c \stackrel{\text{def}}{=} \mathbb{T}_\sigma \uplus \mathbb{T}_b, \quad (16)$$

and note that \mathbb{T}_c is an instance of c -FEASIBILITY.

Lemma III.10. $U(\mathbb{T}_c) < c$.

Proof: Directly from Lemmas III.3 and III.6. ■

We can now show that the transformation correctly preserves the feasibility of the original task system. We start by showing that \mathbb{T}_c is feasible if \mathbb{T}_1 is.

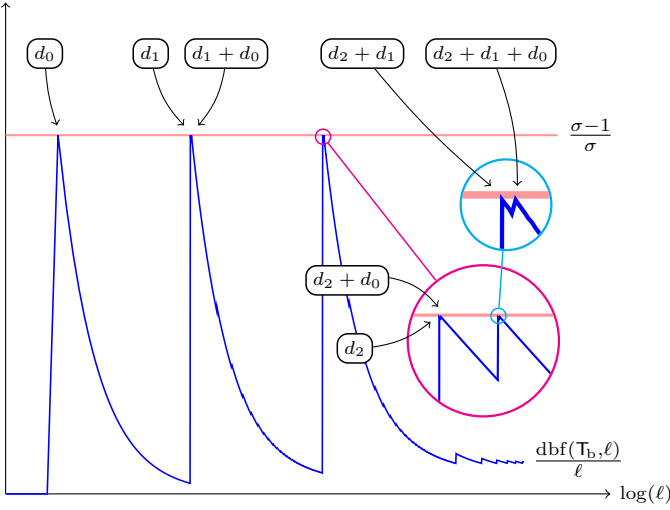


Fig. 4. A plot of $\text{dbf}(\mathbb{T}_b, \ell) / \ell$ for a small example \mathbb{T}_b with $\beta = 3$. Note the logarithmic scale on the horizontal axis. The boost points are those points at which this function touches the line $(\sigma - 1) / \sigma$. Every task $\tau_i \in \mathbb{T}_b$ adds a new “peak” to the plot, touching the line $(\sigma - 1) / \sigma$ at $\ell = d_i$. Immediately after, all previous peaks for tasks τ_j , with $j < i$, will repeat, as can be seen in the zoomed in portions of the plot. The number of boost points therefore grows exponentially in β . The labels show the value of ℓ at the peaks. Note also how the function converges towards $U(\mathbb{T}_b)$ to the right.

Lemma III.11. *If \mathbb{T}_1 is feasible, then \mathbb{T}_c is also feasible.*

Proof: We prove the lemma by contradiction. Assume for this purpose the negation of the lemma’s claim—that \mathbb{T}_1 is feasible, but \mathbb{T}_c is not. By Corollary II.3, there must then exist some $\ell \in \mathbb{N}$ such that $\text{slack}(\mathbb{T}_c, \ell) < 0$. Since all $(e, d, p) \in \mathbb{T}_c$ have d and p that are multiples of σ , it is evident that there must also exist some $\ell' \in \mathbb{N}$ such that

$$\text{slack}(\mathbb{T}_c, \sigma \ell') < 0.$$

Note that

$$\begin{aligned} \text{slack}(\mathbb{T}_c, \sigma \ell') &= \sigma \ell' - \text{dbf}(\mathbb{T}_c, \sigma \ell') \\ &= \sigma \ell' - \text{dbf}(\mathbb{T}_\sigma, \sigma \ell') - \text{dbf}(\mathbb{T}_b, \sigma \ell') \\ &= \text{slack}(\mathbb{T}_\sigma, \sigma \ell') - \text{dbf}(\mathbb{T}_b, \sigma \ell'). \end{aligned} \quad (17)$$

From Lemma III.4 we have

$$\text{slack}(\mathbb{T}_\sigma, \sigma \ell') = \text{slack}(\mathbb{T}_{\text{fill}}, \ell') + (\sigma - 1)\ell' \quad (18)$$

and from Lemma III.7 we have

$$\text{dbf}(\mathbb{T}_b, \sigma \ell') \leq (\sigma - 1)\ell'. \quad (19)$$

By combining Eq. (17), (18) and (19) and using Lemma III.2, we conclude that

$$\begin{aligned} \text{slack}(\mathbb{T}_c, \sigma \ell') &\geq \text{slack}(\mathbb{T}_{\text{fill}}, \ell') \\ &= \text{slack}(\mathbb{T}_1, \ell' \bmod \mathcal{P}(\mathbb{T}_1)). \end{aligned}$$

By assumption, however, \mathbb{T}_1 is feasible and therefore

$$\text{slack}(\mathbb{T}_c, \sigma \ell') \geq \text{slack}(\mathbb{T}_1, \ell' \bmod \mathcal{P}(\mathbb{T}_1)) \geq 0.$$

The lemma follows from this contradiction. ■

We now show the other direction.

Lemma III.12. *If \mathbb{T}_1 is infeasible, then \mathbb{T}_c is also infeasible.*

Proof: Assume that \mathbb{T}_1 is infeasible, then by Corollary II.3 there exists an $\ell \in \{0, 1, \dots, \mathcal{P}(\mathbb{T}_1) - 1\}$ such that

$$\text{slack}(\mathbb{T}_1, \ell) < 0.$$

By Lemma III.9, there also exists some $k \in \mathbb{N}$ such that

$$\sigma(k\mathcal{P}(\mathbb{T}_1) + \ell) \in \mathcal{B}(\mathbb{T}_b).$$

By the definition of $\mathcal{B}(\mathbb{T}_b)$, we therefore have

$$\text{dbf}(\mathbb{T}_b, \sigma(k\mathcal{P}(\mathbb{T}_1) + \ell)) = (\sigma - 1)(k\mathcal{P}(\mathbb{T}_1) + \ell). \quad (20)$$

Note also that from Lemma III.4 we have

$$\begin{aligned} \text{slack}(\mathbb{T}_\sigma, \sigma(k\mathcal{P}(\mathbb{T}_1) + \ell)) \\ = \text{slack}(\mathbb{T}_{\text{fill}}, k\mathcal{P}(\mathbb{T}_1) + \ell) + (\sigma - 1)(k\mathcal{P}(\mathbb{T}_1) + \ell). \end{aligned} \quad (21)$$

Using Eq. (20) and (21) and Lemma III.2 we conclude that

$$\begin{aligned} \text{slack}(\mathbb{T}_c, \sigma(k\mathcal{P}(\mathbb{T}_1) + \ell)) \\ = \text{slack}(\mathbb{T}_\sigma, \sigma(k\mathcal{P}(\mathbb{T}_1) + \ell)) - \text{dbf}(\mathbb{T}_b, \sigma(k\mathcal{P}(\mathbb{T}_1) + \ell)) \\ = \text{slack}(\mathbb{T}_{\text{fill}}, k\mathcal{P}(\mathbb{T}_1) + \ell) \\ = \text{slack}(\mathbb{T}_1, \ell) \\ < 0, \end{aligned}$$

and hence that \mathbb{T}_c is infeasible by Corollary II.3. ■

Finally, our main Theorem follows.

Theorem III.13. *The c -FEASIBILITY problem is (weakly) coNP-complete for any constant c such that $0 < c < 1$.*

Proof: First, note that we have described a transformation from FEASIBILITY to c -FEASIBILITY, which produces correct outputs according to Lemmas III.10, III.11 and III.12. Secondly, note that the constructed task system \mathbb{T}_c has only polynomially many tasks,

$$|\mathbb{T}_1| + 1 + \lceil \log_2(\mathcal{P}(\mathbb{T}_1)) \rceil$$

at most. Also note that the largest parameter in \mathbb{T}_c is $p_{\beta-1}$, which has a value of

$$(\sigma \mathcal{P}(\mathbb{T}_1) + 2)^{\lceil \log_2(\mathcal{P}(\mathbb{T}_1)) \rceil} \sigma,$$

which requires only polynomial space. The transformation is then trivially computable in polynomial time.

From Theorem I.4 we know that FEASIBILITY is coNP-complete. The polynomial transformation therefore proves the coNP-hardness of c -FEASIBILITY. Because c -FEASIBILITY is a special case of FEASIBILITY, we know that c -FEASIBILITY is also in coNP. The coNP-completeness of c -FEASIBILITY follows. ■

IV. CONCLUSIONS

The sporadic task model is one of the basic formalisms for specifying real-time workload. The problem of deciding whether a given sporadic task system is schedulable on a

preemptive uniprocessor is of fundamental importance for real-time scheduling theory. The special case that is limited to task systems with utilization bounded by some constant is widely encountered in the literature, thanks to the existence of an algorithm that solves that problem in pseudo-polynomial time. A long-standing open question has been whether this special case also has a polynomial-time solution. We have shown that for any reasonable choice of the constant, the problem is coNP-complete and therefore cannot be solved in polynomial time, assuming $P \neq NP$.

ACKNOWLEDGMENT

We would like to thank the anonymous reviewers for their detailed and helpful comments.

REFERENCES

- [1] C. L. Liu and J. W. Layland, "Scheduling algorithms for multiprogramming in a hard-real-time environment," *Journal of the ACM*, vol. 20, no. 1, pp. 46–61, 1973.
- [2] M. L. Dertouzos, "Control robotics: The procedural control of physical processes," in *Proceedings of the IFIP congress*, vol. 74, 1974, pp. 807–813.
- [3] S. Baruah, A. Mok, and L. Rosier, "Preemptively scheduling hard-real-time sporadic tasks on one processor," in *Proceedings of the Real-Time Systems Symposium (RTSS)*, 1990, pp. 182–190.
- [4] J. Y.-T. Leung and M. Merrill, "A note on preemptive scheduling of periodic, real-time tasks," *Information Processing Letters*, vol. 11, no. 3, pp. 115–118, 1980.
- [5] S. K. Baruah, L. E. Rosier, and R. R. Howell, "Algorithms and complexity concerning the preemptive scheduling of periodic, real-time tasks on one processor," *Real-Time Systems*, vol. 2, no. 4, pp. 301–324, 1990.
- [6] S. Baruah and K. Pruhs, "Open problems in real-time scheduling," *Journal of Scheduling*, vol. 13, no. 6, pp. 577–582, 2010.
- [7] F. Eisenbrand and T. Rothvoß, "EDF-schedulability of synchronous periodic task systems is coNP-hard," in *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010, pp. 1029–1034.
- [8] P. Ekberg and W. Yi, "Uniprocessor feasibility of sporadic tasks with constrained deadlines is strongly coNP-complete," in *Proceedings of the Euromicro Conference on Real-Time Systems (ECRTS)*, 2015, pp. 281 – 286.