



HAL
open science

Optimal Speed Profile of a DVFS Processor under Soft Deadlines

Jonatha Anselmi, Bruno Gaujal, Louis Sébastien Rebuffi

► **To cite this version:**

Jonatha Anselmi, Bruno Gaujal, Louis Sébastien Rebuffi. Optimal Speed Profile of a DVFS Processor under Soft Deadlines. *Performance Evaluation*, 2021, 152, 10.1016/j.peva.2021.102245 . hal-03364880

HAL Id: hal-03364880

<https://hal.science/hal-03364880>

Submitted on 5 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal Speed Profile of a DVFS Processor under Soft Deadlines

Jonatha Anselmi^a, Bruno Gaujal^a, Louis Sébastien Rebuffi^{a,*}

^a*Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, Grenoble, 38000, France.*

Abstract

We consider a Dynamic Voltage and Frequency Scaling (DVFS) processor executing jobs with obsolescence deadlines: A job becomes obsolete and is removed from the system if it is not completed before its deadline. The objective is to design a dynamic speed policy for the processor that minimizes its average energy consumption plus an obsolescence cost per deadline miss. Under Poisson arrivals and exponentially distributed deadlines and job sizes, we show that this problem can be modeled as a continuous time Markov decision process (MDP) with unbounded state space and unbounded rates. While this MDP admits a continuous time optimality equation for its average cost, the standard uniformization approach is not applicable. Inspired by the scaling method introduced by Blok and Spieksma, we first define a family of truncated MDPs and we then show that the optimal speed profiles are increasing in the number of jobs in the system and are uniformly bounded by a constant. Finally, we show that these properties are inherited from the original (infinite) system. The proposed upper bound on the optimal speed profile is tight and is used to develop an extremely simple policy that accurately approximates the optimal average cost in heavy traffic conditions.

Keywords:

Energy minimization, Optimal speed, DVFS, Markov decision processes.

*Supported by the ANR under the “Investissements d’avenir” program (ANR-15-IDEX-02) and the LabEx PERSYVAL (ANR-11-LABX-0025-01)

*Corresponding author

Email addresses: jonatha.anselmi@inria.fr (Jonatha Anselmi),
bruno.gaujal@inria.fr (Bruno Gaujal),
louis-sebastien.rebuffi@univ-grenoble-alpes.fr (Louis Sébastien Rebuffi)

1. Introduction

Minimizing the energy consumption of embedded systems with real-time execution constraints is becoming more and more important. More functionalities and better performance/cost trade-offs are expected from such systems because of the increased use of real-time applications and the fact that batteries are becoming standard power supplies. Dynamically changing the speed of the processor is a common and efficient way to reduce energy consumption and remarkable gains can be obtained when considering cache-intensive and/or CPU-bound applications as the CPU energy consumption may dominate the overall energy consumption Snowdon et al. (2005). In fact, this is the reason why modern processors are equipped with Dynamic Voltage and Frequency Scaling (DVFS) technology Weiser et al. (1994).

In the deterministic case where job sizes and arrival times are known, a vast literature addressed the problem of designing both off-line and on-line algorithms to compute speed profiles that minimize the energy consumption subject to hard real-time constraints (deadlines) on job execution times; see, e.g., Yao et al. (1995); Bansal et al. (2007); Li et al. (2017) and the references therein. In a *stochastic* environment where only statistical information is available about job sizes and arrival times, it turns out that combining hard deadlines and energy minimization via DVFS-based techniques is much more difficult. In fact, forcing hard deadlines requires to be very conservative, i.e., to consider the worst cases. In spite of these difficulties, this problem has been investigated in Lorch and Smith (2001) for a single job and in Gaujal et al. (2020) for multiple jobs. The former approach constructs the optimal speed profile explicitly in “closed form” while the latter relies on the numerical solution of a discrete time Markov Decision Process (MDP) Puterman (2014). The latter approach has several drawbacks: i) it requires a discretization of both time and space, which introduces by itself an approximation on the optimal solution, ii) deadlines and job sizes need to be bounded, and iii) the size of the state space of the underlying MDP is exponential in the size of the maximal deadline. These issues makes this approach unusable in practice.

The approach followed in this paper circumvents the difficulties described above by replacing the hard real-time constraints, i.e., jobs have hard deadlines that must be satisfied, by *soft* real-time constraints, i.e., jobs may miss their deadlines, at some cost. While the hard deadline of a job must be known at the job arrival, soft deadlines allow for a different information structure: here, only the deadline distribution is known at the job arrival. In this paper,

we further assume that jobs missing their deadlines become obsolete and are dropped. Obsolescence is often found in real-time systems where the information carried by jobs may not be valid any longer after their deadline as it will be replaced by fresher input coming from other jobs. Therefore, obsolete jobs become useless and can get discarded from the queue. Dropping obsolete jobs can also model impatient customers: customers wait for service for some time (deadline) and quit (are dropped) if not served before that time.

1.1. Contribution

We investigate the problem above in a Markovian setting where jobs join the system following a Poisson process and both the deadlines and sizes of jobs are exponentially distributed. Under these assumptions, our goal is to minimize the *average cost*, i.e., the average energy spent by the processor per second plus the penalty due to jobs missing their deadlines. We formulate this problem as an MDP in continuous time where the state is the number of jobs in the system and the action is the processor speed.

Our main result, Theorem 2.2, shows the existence of an optimal speed profile that is increasing in the number of jobs in the system and upper bounded by some constant. This constant is defined in (3) as the minimizer of a function that comes out from our analysis. Surprisingly, our bound does not depend on the deadlines and arrival rates. In other words, our bound on the optimal speed does not change upon variations of these job characteristics. In addition, it yields a simple approximation for the optimal policy and several numerical tests show that such approximation is accurate in heavy-traffic conditions. Finally, the proposed approximation is used to control the proportion of jobs that leave the system because they missed their deadline in a simple manner.

Underlying the proof of our main result, there are some technical challenges that we now discuss. The proposed MDP satisfies the regularity assumptions (stability, unichain) needed to establish an optimality equation as described in Guo and Hernandez-Lerma (2009). However, this is not enough to show structural properties of the optimal policy. In fact, the classical approach to do this is to uniformize the MDP and to investigate the properties of the corresponding discrete time value iteration operator. Unfortunately, this is not possible in our case because the transition rates are unbounded. To uniformize the MDP, a typical approach consists of truncating the state space. Indeed, this is the approach we follow. However, we notice that a naive truncation will not help because the truncation barrier has a strong impact

on the structure of the optimal policy in the sense that it would not preserve any monotonicity property that it may have without truncation. This is shown in Figure 2. Instead, we use the technique proposed in Blok and Spieksma (2015), which smoothly scales down the upward rates of the truncated system as a function of the size of its state space. This technique has been successfully used in Hyon and Jean-Marie (2020); Bhulai et al. (2014) to show structural properties of controlled queueing systems. However, these works focus on discounted costs. Here, we use the same truncation technique but we apply it to the average cost. To the best of our knowledge, this has never been done before. In our specific case, the convergence to the infinite system will be guaranteed by the monotone convergence theorem.

Our paper is organized as follows. In Section 2, our model and the corresponding MDP are described in detail. We also present our main result (Theorem 2.2) as well as some hindsight on the construction of the proof. Section 3 shows how the MDP is truncated and scaled and shows the proof of the monotonicity as well as the construction of the upper bound on the optimal speed. Section 4 focuses on the convergence when the truncation point goes to infinity. Section 5 uses Theorem 2.2 to provide an approximation on the optimal policy and estimate the deadline-miss probability. Finally, Section 6 draws the conclusions of our work and addresses further research.

2. System Description, Problem Statement and Main Result

The system described here is a model for the dynamics of a real-time device composed of a single computing resource (a processor) where incoming jobs need to be executed under a constraint on the amount of time that they spend in the system.

Processor. This is a DVFS processor whose speed can continuously vary in the interval $[0, S_{\max}]$. We consider that speed changes are immediate and induce no energy cost. When the processor works at speed s , it processes s units of work per second while its power dissipation is $w(s)$ watts. The classic simple model for the dynamic power dissipation of any CMOS circuit is $w = K\alpha V^2 f$, see, e.g., Snowdon et al. (2005), where K is a constant, α measures the activity of the logical gates, V is the supply voltage and f is the clock frequency. The clock frequency of the gates is often linearly related to the voltage and therefore DVFS processors adjust both variables together. Within the model above, this means that $w(s)$ is cubic in the speed s . In

this paper, we just require that $w(s)$ is continuous, increasing and strictly convex in the speed s .

Jobs. They form a stochastic point process, with Poisson arrivals with rate λ , i.i.d. deadlines exponentially distributed with rate δ and i.i.d. sizes exponentially distributed with rate μ . Without loss of generality, we assume that $\mu = 1$.

Dynamics. At any point in time t , the processor chooses its speed $s(t)$ and executes one of the jobs in its backlog queue. We notice that the choice of the job in execution, named *active* in the following description, is irrelevant here because of the memory-less property of the deadlines and of the sizes. Thus, at any point in time, at most one job can be active. As mentioned above, this induces an instantaneous energy cost of $w(s(t))$. Now, three events can happen in continuous time:

1. A new job may join the queue.
2. The active job is completed before its deadline. In this case, the job leaves the system.
3. One job (active or inactive) reaches its deadline. In this case, this job becomes obsolete, it is removed from the queue and an immediate cost equal to C is paid.

Cost Function. If we denote by M_T the number of missed deadlines in the time interval $[0, T]$, the objective of this paper is to study the speed profile $s(t)$ of the processor that minimizes the long-run average cost given by the missed deadlines plus power consumption, say \mathcal{J} . Specifically, this is given by

$$\mathcal{J} := \limsup_{T \rightarrow \infty} \frac{1}{T} \left(CM_T + \int_0^T w(s(t)) dt \right).$$

At this point, we claim that this problem can be modeled by a continuous time Markov decision process with a discrete state space. To see this, let us consider the system at time t under the speed profile $s(\cdot)$ and let $i(t)$ denote the number of jobs present in the system at time t . Each job is characterized by two values: (a_k, b_k) , respectively, the sojourn time of job k (how long the job has been present) and its service quantity (how many elementary operations have already been executed on this job by the processor). Also let r_t be the time elapsed since the last arrival. Then the state of the

system at time t is the tuple $(r_t, (a_1, b_1), \dots, (a_{i(t)}, b_{i(t)}))$. From time t , the cost function only depends on the expected future behavior of the system, namely, the future arrivals, future services and future drops of packets due to deadline misses. The memory-less property of the exponential distribution implies that this expected future behavior does not really depend on $(r_t, (a_1, b_1), \dots, (a_{i(t)}, b_{i(t)}))$, but only on $i(t)$, the current number of jobs. As a by-product, this implies that the choice of the optimal speed at time t only depends on $i(t)$, so that the optimal speed profile can only change when $i(t)$ changes, i.e., when an arrival, a service or a drop occur.

Markov Decision Process. We now formulate the problem of interest as an MDP. As mentioned before, the state space is \mathbb{N} and a state represents the number of jobs in the system. The action space is $[0, S_{\max}]$, i.e., the set of available speeds for the DVFS processor. Let $\sigma = (\sigma_i)_{i \in \mathbb{N}}$ denote a stationary and deterministic speed policy adopted by the processor, i.e., $\sigma_i \in [0, S_{\max}]$ is the speed used in state i . It is well known that focusing on stationary and deterministic policies can be done with no loss of optimality in our case (Guo and Hernandez-Lerma, 2009, Theorem 5.9); in other words, we may consider the broader set of *history dependent randomized* policies and then use that there exist stationary and deterministic policies that are optimal with respect to the cost defined below. For $i, j \in \mathbb{N}$, the transition rates are given by:

$$q_{i,j}(\sigma) = \begin{cases} \lambda & \text{if } j = i + 1 \text{ (arrival of a new job)} \\ \sigma_i + i\delta & \text{if } i \geq 1 \text{ and } j = i - 1 \text{ (completion of the active} \\ & \text{job or obsolescence of one job)} \\ -\lambda - i\delta - \sigma_i & \text{if } j = i \\ 0 & \text{otherwise.} \end{cases}$$

Under the speed policy $\sigma = (\sigma_i)_{i \in \mathbb{N}}$, the induced Markov chain, denoted by X^σ , is a birth and death process that resembles an $M/M/\infty$ queue but with an additional decreasing rate, which comes from the processing speed; see Figure 1.

By the ergodicity of the Markov chain X^σ under all policies σ , the long-run cost \mathcal{J} is equal to the long-run expected cost. Letting $\mathbb{E}_{i_0}^\sigma$ denote the expectation given a speed policy σ and starting state i_0 , we have

$$\mathcal{J} = J(i_0, \sigma) := \limsup_{T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{E}_{i_0}^\sigma c(X^\sigma(t), \sigma) dt.$$




Figure 1: Markov chain diagram under policy σ .

In this equation, the immediate cost function $c(\cdot, \cdot)$ is the expected cost incurred by the system at time t . It only depends on the current state and the current speed. Conditional on the state ($X^\sigma(t) = i$), the obsolescence rate is $i\delta$. Thus, the expected cost is:

$$c(i, \sigma) := Ci\delta + w(\sigma_i).$$

With a slight abuse of notation, we will use both notations $c(i, \sigma)$ or $c(i, \sigma_i)$ since c only depends on the speed used in state i and not on the whole policy.

As mentioned before, for given speed policy σ , X^σ is ergodic so that the MDP is unichain (all states are positive recurrent under all policies). This implies that the cost can be defined independently of the starting state. Moreover, for a given σ , there exists a unique stationary measure π^σ for X^σ so that we can define the cost independently of the initial state and express it as a function of the stationary measure:

$$\forall i_0, \quad J(\sigma) := J(i_0, \sigma) = \mathbb{E}^{\pi^\sigma} c(X^\sigma, \sigma) = \sum_i \pi_i^\sigma c(i, \sigma_i). \quad (1)$$

Here, \mathbb{E}^{π^σ} is the expectation with respect to the invariant measure of X^σ . Stationary policies that minimize (1) are optimal speed policies for the model. In particular, they are also optimal over all policies (history dependent and randomized) (Guo and Hernandez-Lerma, 2009, Theorem 5.9). Also, our MDP satisfies all the conditions given in (Guo and Hernandez-Lerma, 2009, Theorem 5.9) to assert the existence of an optimal stationary deterministic

policy σ^* and an optimality equation of the form

$$\begin{aligned} J^* = J(\sigma^*) &= c(i, \sigma_i^*) + \sum_j h^*(j) q_{i,j}(\sigma_i^*) \\ &= \min_{s \in [0, S_{\max}]} c(i, s) + \sum_j h^*(j) q_{i,j}(s), \quad \forall i \in \mathbb{N}, \end{aligned} \quad (2)$$

where h^* is a real function defined on \mathbb{N} , usually referred to as *bias* of the optimal policy.

Main result. The goal of this paper is to investigate structural properties on σ^* and J^* . First, let us define B as

$$B := \arg \min_{s \in \mathbb{R}^+} (w(s) + C(\lambda - s)). \quad (3)$$

This constant is well defined in $\mathbb{R}^+ \cup \{+\infty\}$ because w is strictly convex. In addition, we have the following remark.

Remark 2.1. If w is super-linear, i.e., $\lim_{s \rightarrow \infty} \frac{w(s)}{s} = \infty$, then $w(s) + C(\lambda - s)$ is also super-linear and B is finite. In practice, all models of power dissipation are super-linear in s , e.g., Chandrakasan et al. (1992). In the simple case where $w(s) = Ks^3$, we first notice that the constant K can be set to 1 without loss of generality because to compensate its effect one can adjust the missed deadline cost C accordingly. Then, in this case, we obtain $B = \sqrt{\frac{C}{3}}$.

Our main result is the following.

Theorem 2.2. *There exists a deterministic optimal policy $\sigma^* = (\sigma_i^*)_{i \in \mathbb{N}}$ that is increasing in i and upper bounded by B .*

Remark 2.3. The optimal speed policy of the processor is always bounded by a finite constant, namely $\min(B, S_{\max})$. By definition, B only depends on w (the power dissipation of the processor) and C the cost of each missed deadline. Thus, we remark that B is *independent of the job characteristics* (arrival rate, deadline and size distributions). This is both surprising and helpful in practice. Indeed, if B is finite, one can set *a priori* the maximal speed of the processor to $S_{\max} := B$. This guarantees that in all cases, no cost reduction would be possible by using a more powerful processor. Further discussion about parameter settings, in particular the link between C and the probability that jobs miss their deadline under the optimal policy, will be discussed in Section 5.

A proof of Theorem 2.2 is developed in Sections 3 and 4. Before delving into the proof, we devote the remainder of this section to explain the technical difficulties underlying our problem and the general approach that we follow.

The optimality equation (2) cannot be uniformized because the rates $q_{ii}(\sigma)$ are unbounded in i . Therefore, the study of structural properties of the optimal policy must be done by constructing a sequence of truncated MDPs whose optimal policies converge to σ^* and for which we can prove monotonicity and boundedness. This approach has been proposed for the first time in Blok and Spieksma (2015) for MDPs with discounted cost by truncating the state space and scaling the rates of all the events that take the system out of the truncated space. This has been successfully applied in, e.g., Bhulai et al. (2014); Hyon and Jean-Marie (2020), to show that threshold type policies yield optimal admission control in one queue. However, all these applications consider discounted costs. To the best of our knowledge, no work has been done for the average cost. In the following, we will show that in our case, the scaling technique of Blok and Spieksma (2015) also works for the average cost and is the key ingredient to show Theorem 2.2, which gives new insights on the optimal policy. Thus, our result is another evidence of the power of this scaling approach, though our proof is quite different from the approach used in the discounted case. In fact, the classical approach is to show that the value iteration operator preserves structural properties of the cost and of the policy (typically convexity properties of the cost and level sets of the policy), so that successive iterations of the operator will also preserve the properties and converge to the optimal cost/policy. Here, we will directly consider the fixed point optimality equation and prove monotonicity of the policy by induction on the state (see Sections 3.1, 3.2, 3.3).

3. Truncated Model

As mentioned before, the original MDP cannot be uniformized because the transition rates $q_{ii}(\sigma)$ grow to infinity when i goes to infinity. To construct a discrete time model, we truncate the state space following the guidelines from Blok and Spieksma (2015) to construct a finite state MDP \mathcal{M}_N with linearly decaying arrival rates. The new state space is $\{0, \dots, N\}$, with the same continuous action space $[0, S_{\max}]$ and an average cost per second $J_N(\sigma)$ defined as

$$J_N(\sigma) := \limsup_{T \rightarrow \infty} \frac{1}{T} \int_0^T \mathbb{E}_{i_0}^\sigma c(X^\sigma(t), \sigma) dt \quad (4)$$


and with transition rates for all $i, j \leq N$ given by

$$q_{i,j}^N(\sigma) := \begin{cases} \lambda_i^N & \text{if } i < N \text{ and } j = i + 1 \\ \sigma_i + i\delta & \text{if } i > 0 \text{ and } j = i - 1 \\ -\lambda_i - i\delta - \sigma_i & \text{if } j = i \\ 0 & \text{otherwise,} \end{cases}$$

where the decaying arrival rate is $\lambda_i^N := \lambda(1 - \frac{i}{N})$. Using decaying arrival rates will be a key ingredient in this paper. To illustrate this, let us also consider a naive truncated MDP, \mathcal{M}'_N with fixed arrival rates in each state given by λ . The state space is $\{0, \dots, N\}$, with the same continuous action space $[0, S_{\max}]$ and the same cost function but with modified transition rates given by

$$q_{i,j}^{N'}(\sigma) := \begin{cases} \lambda & \text{if } i < N \text{ and } j = i + 1 \\ \sigma_i + i\delta & \text{if } i > 0 \text{ and } j = i - 1 \\ -\lambda - i\delta - \sigma_i & \text{if } j = i \\ 0 & \text{otherwise,} \end{cases}$$

The two finite MDPs \mathcal{M}_N and \mathcal{M}'_N have been solved numerically using the following parameters: $\lambda = 10, \delta = 0.14, C = 200, w(v) = v^3, N = 60$.



(a) \mathcal{M}_N with decaying arrival rates (b) \mathcal{M}'_N with no decay of the arrival rates

Figure 2: Optimal policies σ^* and $\sigma^{*'}$ for the two truncated MDPs, respectively, \mathcal{M}_N (a) and \mathcal{M}'_N (b).

The respective optimal policies σ^* and $\sigma^{*'}$ are displayed in Figures 2(a)-(b). As one can see, the two models behave very differently. At the last

state of \mathcal{M}'_N , the speed does not need to be as high as the maximal speed, as the arrival rate drops from λ to 0, so that the speed plummets near the last state. In contrast, the optimal policy for \mathcal{M}_N is increasing from 0 to some bound slightly below $\sqrt{\frac{C}{3}}$. Additional numerical tests where we let N grow (reported in Figure 3) further suggest that σ^* is increasing in the state i as well as in the level of truncation N . The bound B also appears to be rather tight when N is large in the example reported in Figure 3.




Figure 3: Two optimal speed policies in \mathcal{M}_N , for $N = 60$ and $N = 1000$ and the bound $B = \sqrt{\frac{C}{3}}$.

This makes the study of \mathcal{M}_N promising and, in the remainder, we focus on this MDP with decaying arrival rates.

Since N is fixed here, we may remove it in the notation for simplicity. As the state space is finite, we can uniformize this MDP to get a discrete time MDP. Choosing

$$U := \lambda + N\delta + S_{\max} \quad (5)$$

as uniformization constant, we get a discrete time MDP \mathcal{D}^N with transition probabilities given by

$$p_{i,j}^N(\sigma) = \begin{cases} \frac{1}{U}\lambda_i & \text{if } i < N \text{ and } j = i + 1 \\ \frac{1}{U}(\sigma_i + i\delta) & \text{if } i > 0 \text{ and } j = i - 1 \\ \frac{1}{U}\bar{U}_i & \text{if } j = i \\ 0 & \text{otherwise,} \end{cases}$$

where the complementary probability to stay in state i is $\bar{U}_i := U - \lambda_i - \delta i - \sigma_i$. Again, we focus on stationary policies σ that minimize the cost $J_N(\sigma)$. The

long-run average cost per step for the discrete time MDP is

$$g(\sigma) := \limsup_{K \rightarrow \infty} \frac{1}{K} \sum_{k=1}^K \mathbb{E}_{i_0}^{\sigma} c(X^{\sigma}(k), \sigma).$$

A classical result for MDPs (see for example (Puterman, 2014, Section 11.5.3)) says that the discrete and continuous time MDPs are related by the following relations: for any policy σ , $Ug(\sigma) = J_N(\sigma)$. This has the following consequences:

- Both MDPs have the same optimal policies.
- Optimal long-run average costs coincide up to the multiplicative uniformization constant: $Ug^* = Ug(\sigma^*) = J_N(\sigma^*) = J_N^*$.

We will first show the counterpart of Theorem 2.2 in the finite case. Let us define

$$B^N := \arg \min_{s \in \mathbb{R}^+} \left(w(s) + \frac{C(\lambda - s)}{1 + \frac{\lambda}{\delta N}} \right), \quad (6)$$

which is well defined in $\mathbb{R}^+ \cup \{\infty\}$ and unique because w is strictly convex.

We have the following properties for the optimal speed policy.

Theorem 3.1. (i) *The optimal speed policy σ^* minimizing (4) is unique.*

(ii) *The optimal speed policy is increasing in i : $\forall i \leq N, \sigma_i^* < \sigma_{i+1}^*$.*

(iii) *The optimal speed policy is upper-bounded: $\forall i \leq N, \sigma_i^* \leq B^N$.*

The proof of this result will be the object of Section 3.1 (monotonicity, item (ii)), 3.2 (upper-bound, item (iii)), 3.3 (uniqueness, item (i)). Then, the relation with the original infinite MDP will be shown in Section 4.

3.1. Proof of Theorem 3.1(ii): Monotonicity of the Optimal Speed

In this subsection, we denote by σ^* any optimal speed policy, X^* the associated Markov chain and J_N^* the optimal cost induced by σ^* . Thus, the optimal long-run average gain per step is

$$g^* := \lim_{K \rightarrow \infty} \frac{1}{K} \mathbb{E}^{\sigma^*} \sum_{k=1}^K \frac{1}{U} c(X^{\sigma^*}(k), \sigma^*). \quad (7)$$

When the state space is finite, the *bias* $H \in \mathbb{R}^{N+1}$ for the optimal speed σ^* is defined up to an additive constant by

$$H(i) := \mathbb{E}_i^{\sigma^*} \sum_{k=1}^{\infty} (c(X^{\sigma^*}(k), \sigma^*) - Ug^*), \quad \forall i \leq N. \quad (8)$$

To fix the value of the bias vector, we set $H(0) := 0$.

Since the MDP is finite, unichain, the action space is compact and the costs and transition probabilities are continuous and bounded in the actions, (Puterman, 2014, Theorem 8.4.7) guarantees the existence of the optimality equations for the optimal cost and for the bias. Specifically, for any state $i \in \{0, \dots, N\}$,

$$g^* + H(i) = \frac{1}{U} \min_{s \in [0, S_{\max}]} \{c(i, s) + (\delta i + s)H(i-1) + \bar{U}_i H(i) + \lambda_i H(i+1)\} \quad (9)$$

with $H(-1) = H(N+1) = 0$ by convention.

For each state i , an optimal action σ_i^* is the choice of a speed minimizing the right hand side term. Notice that necessarily, $\sigma_0^* = 0$ (the speed of the processor must be 0 when there is no work to do).

Using (9), for $i \geq 1$, we can subtract $H(i-1)$ from $H(i)$ and choose $s = \sigma_{i-1}^*$ in (9) to get

$$\begin{aligned} U(H(i) - H(i-1)) &\leq \delta C + \lambda \left(1 - \frac{i}{N}\right) (H(i+1) - H(i)) \\ &+ \left(\bar{U}'_i - \frac{\lambda}{N}\right) (H(i) - H(i-1)) + (\delta(i-1) + \sigma_{i-1}^*) (H(i-1) - H(i-2)) \end{aligned} \quad (10)$$

with $\bar{U}'_i := \bar{U}_i + \sigma_i^* - \sigma_{i-1}^*$. Similarly, doing the same subtraction for $s = \sigma_i^*$, we obtain

$$\begin{aligned} U(H(i) - H(i-1)) &\geq \delta C + \lambda \left(1 - \frac{i}{N}\right) (H(i+1) - H(i)) \\ &+ \left(\bar{U}_i - \frac{\lambda}{N}\right) (H(i) - H(i-1)) + (\delta(i-1) + \sigma_i^*) (H(i-1) - H(i-2)). \end{aligned} \quad (11)$$

Combining both inequalities together we get the inequality

$$(\sigma_i^* - \sigma_{i-1}^*) (H(i) - 2H(i-1) + H(i-2)) \geq 0.$$

From this, we can deduce the following property of the model.

Proposition 3.2. *If H satisfies the following notion of discrete convexity:*

$$\forall i \geq 2, (H(i) - 2H(i-1) + H(i-2)) > 0, \quad (12)$$

then the optimal speed policy σ^ is increasing.*

Here are some additional properties of the finite MDP that will be needed for the proof of the main result.

Lemma 3.3. *The asymptotic cost per second is upper-bounded: $Ug^* \leq \frac{C\lambda}{1 + \frac{\lambda}{N\delta}}$.*

Proof. The cost Ug^* is the optimal asymptotic cost per unit of time. Therefore, we have that $Ug^* \leq Ug_0$, with Ug_0 being the asymptotic cost per unit of time when the speed is 0 for each state. Let $\pi^{(0)}$ be the asymptotic distribution in that case. It is not difficult to see that $\pi_i^{(0)} = \left(1 + \frac{\lambda}{N\delta}\right)^{-N} \binom{N}{i} \left(\frac{\lambda}{N\delta}\right)^i$ and thus we can then compute the associated asymptotic cost as follows:

$$\begin{aligned} g_0 &= \left(1 + \frac{\lambda}{N\delta}\right)^{-N} \sum_{i=0}^N \pi_i^{(0)} \delta i C = \delta C \left(1 + \frac{\lambda}{N\delta}\right)^{-N} \sum_{i=0}^N i \binom{N}{i} \left(\frac{\lambda}{N\delta}\right)^i \\ &= \delta C \left(1 + \frac{\lambda}{N\delta}\right)^{-N} \sum_{i=1}^N N \binom{N-1}{i-1} \left(\frac{\lambda}{N\delta}\right)^i \\ &= C\lambda \left(1 + \frac{\lambda}{N\delta}\right)^{-N} \left(1 + \frac{\lambda}{N\delta}\right)^{N-1} = \frac{C\lambda}{1 + \frac{\lambda}{N\delta}}, \end{aligned}$$

which concludes the proof. \square

We now want to show by backward induction on i that H satisfies (12). For $0 \leq i \leq N-1$, the exact property $\mathcal{P}(i)$ that we will show is

$$H(i) - H(i-1) < H(i+1) - H(i) < \frac{C}{1 + \frac{\lambda}{N\delta}}. \quad (13)$$

Before investigating the initialization step, we show the following preliminary inequality on the optimal cost.

Lemma 3.4. *Let $\bar{\sigma} := \mathbb{E}^{\pi^*}[\sigma_{X^*}^*]$ be the average speed under the stationary measure π^* . Then,*

$$Ug^* > w(\bar{\sigma}) + C \frac{\lambda - \bar{\sigma}}{1 + \frac{\lambda}{N\delta}}.$$




Figure 4: The transition rates of the Markov chain X^* .

Proof. Consider the optimal Markov chain X^* displayed in Figure 4, where for clarity we represented the continuous Markov chain, as its behavior is similar as the uniformized discrete time Markov chain.

Using Jensen's inequality, since w is strictly convex and the stationary measure π^* is non trivial, we get the following strict inequality:

$$\begin{aligned} Ug^* &= J_N^* = \mathbb{E}^{\pi^*}[w(\sigma_{X^*}^*)] + C\delta\mathbb{E}^{\pi^*}[X^*] \\ &> w(\mathbb{E}^{\pi^*}[\sigma_{X^*}^*]) + C\delta\mathbb{E}^{\pi^*}[X^*] \\ &= w(\bar{\sigma}) + C\delta\mathbb{E}^{\pi^*}[X^*]. \end{aligned}$$

Now, we calculate $\mathbb{E}^{\pi^*}[X^*]$. The underlying detailed balance equations are

$$\pi_{i+1}^*[\delta(i+1) + \sigma_{i+1}^*] = \pi_i^* \left(\lambda - i \frac{\lambda}{N} \right).$$

The LHS is 0 for $i = -1$, and the RHS is also 0 for $i = N$. When we sum these equations from 0 to $N - 1$, we can therefore write

$$\sum_{i=0}^N \pi_i^* (\delta i + \sigma_i^*) = \sum_{i=0}^N \pi_i^* \left(\lambda - i \frac{\lambda}{N} \right),$$

which gives $\delta\mathbb{E}^{\sigma^*}[X^*] + \bar{\sigma} = \lambda - \mathbb{E}^{\sigma^*}[X^*]\frac{\lambda}{N}$, i.e.,

$$\mathbb{E}^{\sigma^*}[X^*] = \frac{\lambda - \bar{\sigma}}{\delta + \frac{\lambda}{N}}. \quad (14)$$

Thus, we finally have

$$Ug^* > w(\bar{\sigma}) + C \frac{\lambda - \bar{\sigma}}{1 + \frac{\lambda}{N\delta}}$$

as desired. \square

Remark 3.5. The computation of the expectation that gave (14) remains true for any speed policy, and in that case $\bar{\sigma}$ becomes the average speed of that policy. This formula of the expectation will be needed later to rewrite the cost J_N^* in a different way.

For the initialization of the induction, we need the following lemma.

Lemma 3.6. $H(N) - H(N - 1) < \frac{C}{1 + \frac{\lambda}{N\delta}}$.

Proof. For any σ , the optimality equation (9) at $i = N$ gives

$$U(g^* + H(N)) \leq w(s) + CN\delta + (\delta N + s)H(N - 1) + (U - \delta N - s)H(N),$$

which holds if and only if

$$Ug^* + (\delta N + s)H(N) \leq w(s) + CN\delta + (\delta N + s)H(N - 1),$$

which holds if and only if

$$H(N) - H(N - 1) \leq \frac{w(s) + CN\delta - Ug^*}{\delta N + s}, \quad (15)$$

with equality if the chosen σ is optimal. We therefore want to show the following inequality for some s :

$$h(s) := \frac{w(s) + CN\delta - Ug^*}{\delta N + s} < \frac{C}{1 + \frac{\lambda}{N\delta}}$$

To show this, we use the inequality (3.4) on Ug^* and choose $s = \bar{\sigma}$. This gives

$$h(\bar{\sigma}) < C \frac{N\delta - \frac{\lambda - \bar{\sigma}}{1 + \frac{\lambda}{N\delta}}}{\delta N + \bar{\sigma}} < C \frac{N\delta + \bar{\sigma}}{(\delta N + \bar{\sigma})(1 + \frac{\lambda}{N\delta})} < \frac{C}{1 + \frac{\lambda}{N\delta}}.$$

This concludes the proof by (15). \square

Therefore, for the initialization step we write the inequality (11) for $i = N$ to obtain

$$\begin{aligned} U(H(N) - H(N - 1)) &\geq \delta C + \left(U - \delta N - s - \frac{\lambda}{N} \right) (H(N) - H(N - 1)) \\ &\quad + (\delta(N - 1) + s)(H(N - 1) - H(N - 2)), \end{aligned}$$

which implies

$$\begin{aligned} (\delta(N-1) + s)(H(N) - H(N-1)) &\geq \delta C - \left(\frac{\lambda}{N} + \delta\right)(H(N) - H(N-1)) \\ &\quad + (\delta(N-1) + s)(H(N-1) - H(N-2)), \end{aligned}$$

which implies (using Lemma 3.6)

$$(\delta(N-1) + s)(H(N) - H(N-1)) > (\delta(N-1) + s)(H(N-1) - H(N-2)),$$

which as desired gives $H(N) - H(N-1) > H(N-1) - H(N-2)$. This new inequality and Lemma 3.6 imply $\mathcal{P}(N-1)$.

Now, let us assume that the property $\mathcal{P}(i)$ is true for some $i \leq N-1$. We then have, writing (11) with $s = \sigma_i^*$ and using the first inequality of $\mathcal{P}(i)$ (13),

$$\begin{aligned} U(H(i) - H(i-1)) &\geq \delta C + \left(U - \delta i - s - \frac{\lambda}{N}\right)(H(i) - H(i-1)) \\ &\quad + (\delta(i-1) + s)(H(i-1) - H(i-2)) \end{aligned}$$

which implies

$$\left(\delta i + s + \frac{\lambda}{N}\right)(H(i) - H(i-1)) \geq \delta C + (\delta(i-1) + s)(H(i-1) - H(i-2)),$$

which implies

$$(\delta(i-1) + s)(H(i) - H(i-1)) > (\delta(i-1) + s)(H(i-1) - H(i-2)).$$

This inequality comes from the second inequality of $\mathcal{P}(i)$ (13), and from its first inequality we finally obtain both inequalities of $\mathcal{P}(i-1)$.

Thus, $H(i) - H(i-1) > H(i-1) - H(i-2)$ and the backward induction is complete. Since H is strictly convex, with Proposition 3.2 we deduce that σ^* is increasing.

3.2. Proof of Theorem 3.1(iii): Upper Bound on the Optimal Speed

By definition, we recall that B^N is the unique minimum of u , see (6). We also have

$$J_N(\sigma) = \sum_{i \in S} \pi_i u(\sigma_i),$$

where for all $a \geq 0$,

$$u(a) := w(a) + C \frac{\lambda - a}{1 + \frac{\lambda}{N\delta}}.$$

Proposition 3.7. *If σ^* is an optimal increasing speed policy, then it is upper-bounded by B^N , which means $\sigma_i^* \leq B^N$ for all $i \leq N$.*

Proof. We show this statement by contradiction. Let σ^* be an optimal increasing speed policy. As σ^* is increasing, assume that $\sigma_N^* > B^N$. Let π be the associated stationary measure. As $\sigma_N^* > B^N \geq 0 = \sigma_0^*$, we can define $i_0 > 0$ as the smallest state such that $\sigma_{i_0}^* > B^N$. Now, we can define the following policy

$$\tilde{\sigma}_i = \begin{cases} \sigma_i^* & \text{if } i < i_0 \\ B^N & \text{if } i \geq i_0 \end{cases}$$

and let $\tilde{\pi}$ denote its associated stationary measure. By definition, we notice that the minimum of the function u is reached at B^N . Moreover, for all $i < i_0$, we get from the local balance equations that $\pi_i > \tilde{\pi}_i$. Indeed, when $i \geq i_0 - 1$:

$$\frac{\pi_{i+1}}{\pi_i} = \frac{\lambda \left(1 - \frac{i}{N}\right)}{(i+1)\delta + \sigma_{i+1}^*} < \frac{\tilde{\pi}_{i+1}}{\tilde{\pi}_i},$$

with equality only when $i < i_0 - 1$. It remains to show that the cost associated to this new speed is lower than the original one. We have

$$\begin{aligned} J(\sigma^*) - J(\tilde{\sigma}) &= \sum_{i=0}^N \pi_i u(\sigma_i^*) - \sum_{i=0}^N \tilde{\pi}_i u(\tilde{\sigma}_i) \\ &= \sum_{i=0}^{i_0-1} (\pi_i - \tilde{\pi}_i) u(\sigma_i^*) + \sum_{i=i_0}^N \pi_i u(\sigma_i^*) - \sum_{i=i_0}^N \tilde{\pi}_i u(B^N) \\ &> \sum_{i=0}^{i_0-1} (\pi_i - \tilde{\pi}_i) u(B^N) + \sum_{i=i_0}^N \pi_i u(B^N) - \sum_{i=i_0}^N \tilde{\pi}_i u(B^N) > 0. \end{aligned}$$

This contradicts the optimality of σ^* and concludes the proof. \square

3.3. Proof of Theorem 3.1(i): Uniqueness of the Optimal Policy

The following lemma gives a different expression for σ_i^* using the notion of generalized inverse of the derivative w' of w . First, since w is strictly convex, w' is continuous and well defined everywhere but on a countable subset $D \subseteq [0, S_{max}]$. Moreover, w' is strictly increasing and diverges, so that we can correctly define the inverse

$$w'^{-1} : y \mapsto \inf\{s \in \mathbb{R}, w'(s) \geq y\},$$

and w'^{-1} is increasing.

Lemma 3.8. *If σ^* is an optimal speed policy with bias function H , we can relate the variation of the bias to the speed at a given state:*

$$\sigma_i^* = w'^{-1}(\Delta_i), \quad \forall i \geq 1$$

with $\Delta_i := H(i) - H(i-1)$ for $i \in [1, N]$.

Proof. Let σ^* be an optimal speed policy with bias function H . For $i \geq 1$, the optimality equation (9) can be written as

$$J_N^* = \min_s \{w(s) + C\delta - (\delta i + s)\Delta_i + \lambda_i \Delta_{i+1}\}, \quad (16)$$

so that for any speed s :

$$w(\sigma_i^*) - \sigma_i^* \Delta_i \leq w(s) - s \Delta_i.$$

This yields

$$\begin{cases} \frac{w(\sigma_i^*) - w(s)}{\sigma_i^* - s} \leq \Delta_i & \text{for } s \leq \sigma_i^* \\ \frac{w(\sigma_i^*) - w(s)}{\sigma_i^* - s} \geq \Delta_i & \text{for } s \geq \sigma_i^*. \end{cases}$$

Let $\bar{D} = [0, S_{\max}] \setminus D$. These inequalities on Δ_i give: $w'(\sigma_i^{*-}) \leq \Delta_i \leq w'(\sigma_i^{*+})$, so that there are two possibilities. Either i) $\Delta_i \in w'(\bar{D})$ and $w'(\sigma_i^*) = \Delta_i$, or ii) $\Delta_i \notin w'(\bar{D})$, so that $\sigma_i^* = w'^{-1}(\Delta_i) \in D$. In both cases, we have $\sigma_i^* = w'^{-1}(\Delta_i)$. \square

Proposition 3.9. *The optimal speed policy is unique, and therefore Blackwell optimal.*

Proof. Let $\sigma, \tilde{\sigma}$ be two optimal speed policies, H, \tilde{H} their respective biases and $\Delta, \tilde{\Delta}$ the respective variations of the biases. We will show by induction that the variation of the speed and biases are equal.

We already have that $\sigma(0) = \tilde{\sigma}(0) = 0$. The optimality equation (9) for $i = 0$ then gives

$$\Delta_1 = \tilde{\Delta}_1 = \frac{Ug^*}{\lambda} = \frac{J_N^*}{\lambda}.$$

Using Lemma 3.8, we then have $\sigma_1 = w'^{-1}(\Delta_1) = w'^{-1}(\tilde{\Delta}_1) = \tilde{\sigma}_1$.

Assume that for some $i \geq 1$, $\sigma_i = \tilde{\sigma}_i$ and $\Delta_i = \tilde{\Delta}_i$. Writing (16) for both optimal speeds, and using the assumption of the induction, we have $\Delta_{i+1} = \tilde{\Delta}_{i+1}$. Then by using Lemma 3.8 again, $\sigma_{i+1} = w'^{-1}(\Delta_{i+1}) = w'^{-1}(\tilde{\Delta}_{i+1}) = \tilde{\sigma}_{i+1}$.

The induction is complete, so that the optimal speed policy is unique. \square

This completes the last part of the proof of Theorem 3.1.

4. Convergence of the Truncated MDPs

To show the convergence of the sequence of the truncated MDPs to the infinite one as N goes to infinity, we first show monotonicity properties in N . These guarantee the existence of the limit and allow us to invoke the monotone convergence theorem to show that this limit satisfies the optimality equation of the infinite MDP.

4.1. The Optimal Speed is Increasing in the Size of the State Space

The next proposition states that the optimal cost is increasing in the number of states.

Proposition 4.1. *Let $N \geq 1$. Then, $J_N^* \leq J_{N+1}^*$.*

Proof. Let $N \geq 1$ and let σ be the optimal speed policy when the state space is $\{0, \dots, N+1\}$. Let $\tilde{\sigma}$ be a speed policy for the MDP with state space $\{0, \dots, N\}$, defined as the truncation of σ : for $i \in \leq i$, $\tilde{\sigma}_i = \sigma_i$. Recall that $X^{\tilde{\sigma}}$ is the continuous time Markov chain with speed policy $\tilde{\sigma}$ on the reduced state space $\{0, \dots, N\}$, and therefore cannot be compared directly with X^σ , which is defined on $\{0, \dots, N+1\}$. Thus, let \tilde{X} be the following discrete time Markov chain, with state space $\{0, \dots, N+1\}$ and transition probabilities given by, for $i \leq N+1$,

$$\tilde{p}_{i,j} = \begin{cases} \frac{1}{U^{(N+1)}} \lambda_i^{(N)} & \text{if } 1 \leq j = i+1 \leq N \\ \frac{1}{U^{(N+1)}} (\sigma_i + i\delta) & \text{if } 0 \leq j = i-1 \leq N-1 \\ \frac{1}{U^{(N+1)}} \bar{U}_i^{(N+1)} & \text{if } 0 \leq j = i \leq N \\ \frac{1}{U^{(N+1)}} (U - (N+1)\delta - \sigma_{N+1}) & \text{if } j = i = N+1 \\ 0 & \text{otherwise,} \end{cases}$$

with the notations: $U^{(N+1)} := 2\lambda + (N+1)\delta + \sigma_{max}$, $\lambda_i^{(N)} := \lambda \left(1 - \frac{i}{N}\right)$ and $\bar{U}_i^{(N+1)} := U - \lambda_i^{(N)} - \delta i - \sigma_i$.

With a slight abuse of notation, we denote by X^σ and $X^{\tilde{\sigma}}$ the uniformized Markov chains with the same uniformization constant $U^{(N+1)}$, so that we will be able to compare both Markov chains defined with the same time step. Moreover, notice that the Markov chain \tilde{X} is not irreducible: the use of the last state is only to extend the chain to a larger number of states while keeping the behavior of \tilde{X} similar to the one of $X^{\tilde{\sigma}}$.

In order to effectively compare \tilde{X} and X^σ , we will now define a coupling (\tilde{Y}, Y) such that \tilde{Y} and Y have the same distributions as \tilde{X} and X^σ respectively. For each time step $k \in \mathbb{N}$, let $A(k)$ be a uniformly distributed random variable on $[0, U^{(N+1)}]$. Then,

$$\tilde{Y}(k+1) = \begin{cases} \tilde{Y}(k) + 1 & \text{if } A(k) \in [0, \lambda_{\tilde{Y}(k)}^{(N)}] \text{ and } \tilde{Y} \leq N \\ \tilde{Y}(k) & \text{if } A(k) \in [\lambda_{\tilde{Y}(k)}^{(N)}, U^{(N+1)} - \tilde{Y}(k)\delta - \sigma_{\tilde{Y}(k)}] \\ & \text{and } \tilde{Y}(k) \leq N \\ \tilde{Y}(k) - 1 & \text{if } A(k) \in [U^{(N+1)} - \tilde{Y}(k)\delta - \sigma_{\tilde{Y}(k)}, U^{(N+1)}] \\ & \text{and } \tilde{Y}(k) \leq N \\ \tilde{Y}(k) & \text{if } A(k) \in [0, U^{(N+1)} - \tilde{Y}(k)\delta - \sigma_{\tilde{Y}(k)}] \\ & \text{and } \tilde{Y}(k) = N + 1 \\ \tilde{Y}(k) - 1 & \text{if } A(k) \in [U^{(N+1)} - \tilde{Y}(k)\delta - \sigma_{\tilde{Y}(k)}, U^{(N+1)}] \\ & \text{and } \tilde{Y}(k) = N + 1, \end{cases}$$

and similarly:

$$Y(k+1) = \begin{cases} Y(k) + 1 & \text{if } A(k) \in [0, \lambda_{Y(k)}^{(N+1)}] \\ Y(k) & \text{if } A(k) \in [\lambda_{Y(k)}^{(N+1)}, U^{(N+1)} - Y(k)\delta - \sigma_{Y(k)}] \\ Y(k) - 1 & \text{if } A(k) \in [U^{(N+1)} - Y(k)\delta - \sigma_{Y(k)}, U^{(N+1)}]. \end{cases}$$

By construction, if $\tilde{Y}(0) = Y(0)$, then we show by induction that for all n , $\tilde{Y}(k) \leq Y(k)$.

We now check for all possible cases. For all $i \leq N + 1$, $\lambda_i^{(N)} < \lambda_i^{(N+1)}$, therefore if $\tilde{Y}(k) = Y(k)$, then $\tilde{Y}(k+1) \leq Y(k+1)$. If $\tilde{Y}(k) = Y(k) - 1$ with $\tilde{Y}(k+1) = \tilde{Y}(k) + 1$, then $Y(k+1) \geq Y(k)$ by definition of U , as $U - Y(k)\delta - \sigma_{Y(k)} \geq \lambda \geq \lambda_{\tilde{Y}(k)}^{(N)}$, and therefore $Y(k+1) \geq \tilde{Y}(k+1)$. In the remaining cases, $\tilde{Y}(k+1) \leq Y(k+1)$.

Hence, for all k , $\tilde{Y}(k) \leq Y(k)$, so that $\tilde{X} \leq_{st} X^\sigma$; here, \leq_{st} denotes the stochastic order Shaked and Shanthikumar (1994). As $i \mapsto c(i, \sigma_i)$ is increasing, this implies:

$$\mathbb{E}^\sigma c(\tilde{X}, \sigma) \leq \mathbb{E}^\sigma c(X^\sigma, \sigma) = J_{N+1}^*,$$

and we have

$$J_N(\tilde{\sigma}) = \mathbb{E}^{\tilde{\sigma}} c(X^{\tilde{\sigma}}, \tilde{\sigma}) = \mathbb{E}^\sigma c(\tilde{X}, \sigma).$$

Therefore,

$$J_N^* \leq J_N(\tilde{\sigma}) \leq J_{N+1}(\sigma) = J_{N+1}^*$$

as desired. \square

The next proposition states that the optimal speed policy is increasing in the size of the state space.

Proposition 4.2. *Let $N \geq 1$ and let $\sigma^{(N)}$ be the unique optimal speed policy for the N -th MDP. Then, $\sigma_i^{(N)} \leq \sigma_i^{(N+1)}$ and $\Delta_i^{(N)} \leq \Delta_i^{(N+1)}$ for all $i \leq N$.*

Proof. We use the expression of the optimal speed from Lemma 3.8 and the bias function to show by induction on the states $i \geq 1$ that:

$$\Delta_i^{(N)} \leq \Delta_i^{(N+1)} \text{ and } \sigma_i^{(N)} \leq \sigma_i^{(N+1)}, \quad (17)$$

where $\Delta_i^{(N)} = H(i) - H(i-1)$ and H is the bias function for the MDP with state space $[0, N]$.

We first have that $\sigma_0^{(N)} = \sigma_0^{(N+1)} = 0$ and that

$$\Delta_1^{(N)} = \frac{J_N^*}{\lambda} \leq \frac{J_{N+1}^*}{\lambda} = \Delta_1^{(N+1)},$$

where the inequality comes from Proposition 4.1. Let us now assume that for some $i \geq 1$, $\Delta_i^{(N)} \leq \Delta_i^{(N+1)}$ and $\sigma_i^{(N)} \leq \sigma_i^{(N+1)}$. $\Delta_i^{(N)} \leq \Delta_i^{(N+1)}$, we directly have that $\sigma_i^{(N)} \leq \sigma_i^{(N+1)}$.

To show the first inequality, we write the optimality equation (16) for N and $N+1$ with their respective optimal speed, and subtract one from the other to get:

$$J_{N+1}^* - J_N^* = -(i\delta + \sigma_i^{(N+1)})\Delta_i^{(N+1)} + (i\delta + \sigma_i^{(N)})\Delta_i^{(N)} + \lambda_i(\Delta_{i+1}^{(N+1)} - \Delta_{i+1}^{(N)}).$$

Using Proposition 4.1 and the induction assumption, we get

$$\lambda_i(\Delta_{i+1}^{(N+1)} - \Delta_{i+1}^{(N)}) \geq 0.$$

Then using Lemma 3.8, as $\Delta_{i+1}^{(N)} \leq \Delta_{i+1}^{(N+1)}$ and w'^{-1} is increasing,

$$\sigma_{i+1}^{(N+1)} = w'^{-1}(\Delta_{i+1}^{(N+1)}) \geq w'^{-1}(\Delta_{i+1}^{(N)}) = \sigma_{i+1}^{(N)}.$$

This completes the induction so that the optimal speed policy is increasing. \square

4.2. Convergence Results and Proof of Theorem 2.2

For the truncated MDP in discrete time, let g^* , $\sigma^{(N)}$ and $H^{(N)}$ be the optimal average cost, the optimal policy and its bias. They satisfy the optimality equation.

$$g^* + H^{(N)}(i) = \frac{1}{U} \left(c(i, \sigma_i^{(N)}) + \sum_j H^{(N)}(j) p_{i,j}^N(\sigma_i^{(N)}) \right) \quad \forall i \leq N.$$

This implies that for the truncated model in continuous time, $J_N = Ug^*$, $\sigma^{(N)}$ and $H^{(N)}$ also satisfy the optimality equation

$$J_N = c(i, \sigma_i^{(N)}) + \sum_j H^{(N)}(j) q_{i,j}^N(\sigma_i^{(N)}) \quad \forall i \leq N.$$

Furthermore, for all N , $J_N \leq C\lambda$ by Lemma 3.3, $\sigma^{(N)} \leq B^N \leq B$ by Theorem 3.1 and since the function w' is increasing, $H^{(N)}(i) \leq iw'(B)$ for all $i \leq N$ by Lemma 3.8.

Now, by the monotonicity of J_N (Proposition 4.1) and the monotonicity of $\sigma^{(N)}$ and $H^{(N)}$ (Proposition 4.2), they all converge to finite non-negative limits when N goes to infinity, denoted respectively by J_∞ , $\sigma^{(\infty)}$ and $H^{(\infty)}$.

As for the rates, $q_{i,j}^N(s)$ also converges monotonically to $q_{i,j}(s)$ and is continuous in s . Finally, the immediate cost c is continuous in σ . The monotone convergence theorem implies that these limits satisfy an optimality equation,

$$J_\infty = c(i, \sigma_i^{(\infty)}) + \sum_j H^{(\infty)}(j) q_{i,j}(\sigma_i^{(\infty)}) \quad \forall i \in \mathbb{N}.$$

This shows that these limits are respectively the optimal average cost (J_∞), an optimal policy ($\sigma^{(\infty)}$) and its bias ($H^{(\infty)}$) for the original MDP. This completes the proof of Theorem 2.2.

5. Cost and Deadline-Miss Probability Approximations

Our main result, Theorem 2.2, may suggest to consider the simple policy σ^B defined by $\sigma_i^B = B \mathbb{I}_{\{i>0\}}$ for all $i \geq 0$, where B is defined in (3) and \mathbb{I}_A is the indicator function of the event A . Thus, σ^B uses constant speed B whenever the system is busy. In this section, we numerically show that the dynamics induced by σ^B are “close” to the ones induced by σ^* in the sense that the optimal average cost $J(\sigma^*)$ is very well approximated by $J(\sigma^B)$.

Then, we investigate mathematical properties of the Markov chain induced by σ^B and we give an upper bound on the stationary probability of missing deadlines as a function of the model parameters. In particular, by varying the cost parameter C , this bound can be used to keep such probabilities below a desired threshold.

In the following, we let π^B denote the stationary probability of the Markov chain induced by policy σ^B . Using the detailed balance equations, we obtain

$$\pi_i^B = \pi_0^B \frac{\lambda^i}{\prod_{j=1}^i (\delta j + B)}, \quad i \geq 1$$

and, using $\sum_{i \geq 0} \pi_i^B = 1$,

$$\pi_0^B = \left(\sum_{i \geq 0} \frac{\lambda^i}{\prod_{j=1}^i (\delta j + B)} \right)^{-1} = \frac{\delta}{B} \frac{\left(\frac{\lambda}{\delta}\right)^{\frac{B}{\delta}}}{e^{\frac{\lambda}{\delta}} \gamma\left(\frac{B}{\delta}, \frac{\lambda}{\delta}\right)} \quad (18)$$

where $\gamma(\cdot, \cdot)$ is the lower incomplete gamma function.

We will refer to the following proposition, which can be proven by summing the detailed balance equations of the underlying Markov chain as done in the proof of Lemma 3.4; for this reason, we omit the proof.

Proposition 5.1. *Let X be distributed as π^B . Then, $\mathbb{E}[X] = \frac{1}{\delta}(\lambda - B(1 - \pi_0^B))$.*

5.1. Approximation of the Average Cost

By definition, the average cost induced by σ^B , $J(\sigma^B)$, is an upper bound on the optimal average cost, $J(\sigma^*)$. In particular,

$$\begin{aligned} J(\sigma^*) \leq J(\sigma^B) &= \sum_{i \geq 0} c(i, \sigma^B) \pi_i^B = C\delta \sum_{i \geq 1} i \pi_i^B + \sum_{i \geq 1} w(\sigma_i^B) \pi_i^B \\ &= C(\lambda - B(1 - \pi_0^B)) + w(B)(1 - \pi_0^B). \end{aligned}$$

where the last equality follows by Lemma 5.1.

5.2. Deadline-Miss Probabilities

Let us consider the probability p_M that a job misses its deadline under the stationary regime of policy σ^B . This is defined by

$$p_M := \sum_{i \geq 1} \frac{\pi_i^B}{1 - \pi_0^B} \frac{\delta i}{\delta i + B}. \quad (20)$$

Our objective here is to control p_M by fine-tuning the model parameter C . In other words, we want to design C such that p_M remains below a given threshold. Though the structure of π_0^B in (18) implies that the exact relation between C and p_M is not trivial, this problem can be clearly solved numerically. Nonetheless, we aim at developing an upper bound on p_M allowing for a simple analytic evaluation.

The following proposition provides a first upper bound on p_M .

Proposition 5.2.

$$p_M \leq \frac{1}{1 - \pi_0^B} \left(1 - \frac{\delta \pi_0^B}{\delta + B} \right) - \frac{B}{\lambda}. \quad (21)$$

Proof. This proposition is proven by the following inequalities

$$\begin{aligned} p_M &\leq \sum_{i \geq 1} \frac{\pi_i^B}{1 - \pi_0^B} \frac{\delta(i+1)}{\delta(i+1) + B} \\ &= \frac{\delta}{\lambda(1 - \pi_0^B)} \sum_{i \geq 1} (i+1) \frac{\lambda^{i+1}}{\prod_{j=1}^{i+1} (\delta j + B)} \pi_0^B \\ &= \frac{1}{1 - \pi_0^B} \frac{\delta}{\lambda} \sum_{i \geq 1} (i+1) \pi_{i+1}^B \\ &= \frac{1}{1 - \pi_0^B} \left(1 - \frac{B}{\lambda} (1 - \pi_0^B) \right) - \frac{\delta}{\delta + B} \frac{\pi_0^B}{1 - \pi_0^B} \\ &= \frac{1}{1 - \pi_0^B} - \frac{B}{\lambda} - \frac{\delta}{\delta + B} \frac{\pi_0^B}{1 - \pi_0^B}. \end{aligned} \quad (22)$$

In (22), we have used Proposition 5.1. □

By coupling the underlying Markov chain under σ^B with an auxiliary $M/M/\infty$ queue with arrival rate λ and service rate $\delta + B$, we notice that the state of the former is stochastically dominated by the latter. Therefore, $\pi_0^B \leq e^{-\frac{\lambda}{\delta+B}}$. Using that the mapping $x \mapsto \frac{1}{1-x} \left(1 - \frac{\delta x}{\delta+B} \right)$ is increasing over $[0, 1)$ and the previous inequality on π_0^B in (21), we obtain

$$p_M \leq \overline{p_M} = \frac{1}{1 - e^{-\frac{\lambda}{\delta+B}}} \left(1 - \frac{\delta e^{-\frac{\lambda}{\delta+B}}}{\delta + B} \right) - \frac{B}{\lambda}. \quad (23)$$

We notice that $\overline{p_M} = 1$ if $B = 0$ and that $\overline{p_M} \rightarrow 0$ as $B \rightarrow \infty$.

Now, the bound $\overline{p_M}$ can also be used to adjust the cost of missed deadlines C so that the proportion of jobs that miss their deadline will stay below some acceptable level α .

Since Equation (23) cannot be inverted in closed form, a first order Taylor expansion gives $\overline{p_M} \approx \frac{\delta}{\delta+B}$ when B goes to infinity. Using the value $B = \sqrt{C/3}$ for the classical power dissipation $w(s) = s^3$ (see Remark 2.3), we get $C \approx 3\delta^2 \left(\frac{1-\overline{p_M}}{\overline{p_M}}\right)^2$. Therefore, $C = 3\delta^2 \left(\frac{1-\alpha}{\alpha}\right)^2$ is the cost of missed deadlines that keeps the deadline-miss probability below α .

5.3. Accuracy Assessment

By means of numerical calculations, we now evaluate the accuracy of the bound in (19). For this purpose, we consider the two following set-ups. In the first one, we focus on one acceptable value for the probability to miss a deadline under policy σ^B and we consider that such an acceptable level is $p_M = 0.1$. We let the deadline rate δ and the arrival rate λ vary in $[0.1, 1]$ and $[0.1, 50]$, respectively. Each value of the couple (λ, δ) induces a unique value for B through (20) and for C through (3). Then, for each value of (λ, δ) , we compute the percentage relative error

$$\mathcal{E} := \frac{J(\sigma^B) - J(\sigma^*)}{J(\sigma^*)} 100\% \quad (24)$$

where $J(\sigma^B)$ and $J(\sigma^*)$ are computed numerically by truncating the state space to some N large enough so that increasing N does not change the average cost by more than 10^{-4} ; in all cases, $N \leq 250$.

Figure 5 depicts the level sets for the values of \mathcal{E} . We observe that \mathcal{E} decreases as the arrival rate increases and that it is small under a wide set of parameters. In particular, in heavy traffic conditions and with “large” deadlines the percentage relative error can be smaller than 3%. Thus, we conclude that our approximation is accurate within these conditions. On the other hand, in light load conditions and with “short” deadlines, \mathcal{E} can be above 20% and in this case the optimal speed profile is more complex.

In the second set of experiments, we consider a fixed cost C instead of a fixed miss probability. Specifically, we let $C = 300$, which implies that the upper bound on the speed is $B = \sqrt{C/3} = 10$. For each pair (λ, δ) , we compute the relative error \mathcal{E} for policy σ^B compared with the optimal policy, σ^* . The corresponding results are displayed in Figure 6. Unlike for the




Figure 5: Level sets of the percentage relative error \mathcal{E} with a fixed miss probability $p_M = 0.1$. The darker (lighter) the zone, the higher (lower) the error.




Figure 6: Level sets of the percentage relative error \mathcal{E} with a fixed cost per deadline, $C = 300$.

case with fixed probability, the error decreases sharply as the arrival rate λ increases. In particular, the error is very small in heavy traffic ($\lambda \geq B = 10$).

6. Conclusion and Perspectives

In a stochastic environment, it is well known that the problem of determining the load-dependent speed profile that a DVFS processor should follow to minimize the mean energy consumption under hard deadlines is difficult. In this paper, the basic idea is to follow a Lagrangian approach where hard deadlines are replaced by *soft* deadlines, meaning that jobs are allowed to miss their deadline, though each missed deadline comes with a penalty that can be fine-tuned to keep the proportion of jobs missing their deadline as small as desired. The resulting advantage of this approach stands in the ability of formulating the problem above as a Markov decision process for which we can establish constructive structural results (Theorem 2.2). Beyond the existence of monotone optimal policies, we have found that the optimal speed is bounded from above by some constant that does not depend on the deadline and the arrival rates, which may be quite surprising, and that such constant let us define an extremely simple policy whose average cost is close to the optimal one.

There are some open questions that we leave as future research. First, we have assumed that the available processing speeds vary continuously on a compact set. On the other hand, it may be convenient to consider the case where only a finite number of speeds is available. Second, it is interesting to investigate whether our results are insensitive to the job size and/or deadline distributions. This may be justified by the fact that the proposed queueing system is somewhat similar, when the speeds are bounded, to an $M/M/\infty$ queue. Furthermore, if service times were to follow a phase-type distribution, we could use again an MDP formulation to model the problem but the analysis would require more work as the state space would be much larger. Finally, we wonder whether the optimal policy could be “learned” and, in this respect, how the proposed upper bound could speed up the learning process.

References

- Nikhil Bansal, Tracy Kimbrel, and Kirk Pruhs. 2007. Speed Scaling to Manage Energy and Temperature. *J. ACM* 54, 1 (2007).
- S. Bhulai, A. C. Brooms, and F. M. Spieksma. 2014. On structural properties of the value function for an unbounded jump Markov process with an application to a processor sharing retrial queue. *Queueing Systems* 76, 4 (2014), 425–446.
- H. Blok and F. M. Spieksma. 2015. Countable state Markov decision processes with unbounded jump rates and discounted cost: optimality equation and approximations. *Advances in Applied Probability* 47, 4 (2015), 1088 – 1107.
- A. P. Chandrakasan, S. Sheng, and R. W. Brodersen. 1992. Low-power CMOS digital design. *IEEE Journal of Solid-State Circuits* 27, 4 (1992), 473–484.
- Bruno Gaujal, Alain Girault, and Stéphan Plassart. 2020. Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks. *Journal of Scheduling* (July 2020), 1–25. <https://doi.org/10.1007/s10951-020-00660-9>
- Xianping Guo and Onesimo Hernandez-Lerma. 2009. *Continuous-time Markov decision processes. Theory and applications*. Vol. 62.
- Emmanuel Hyon and Alain Jean-Marie. 2020. Optimal control of admission in service in a queue with impatience and setup costs. *Performance Evaluation* 144 (2020).
- Minming Li, Frances F. Yao, and Hao Yuan. 2017. An $O(n^2)$ Algorithm for Computing Optimal Continuous Voltage Schedules. In *Annual Conference on Theory and Applications of Models of Computation, TAMC'17 (LNCS, Vol. 10185)*. Bern, Switzerland, 389–400.
- Jacob R. Lorch and Alan Jay Smith. 2001. Improving Dynamic Voltage Scaling Algorithms with PACE. In *ACM SIGMETRICS 2001 Conference*. 50–61.
- Martin L Puterman. 2014. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons.

- Moshe Shaked and J George Shanthikumar. 1994. *Stochastic orders and their applications*. Academic Pr.
- David C. Snowdon, Sergio Ruocco, and Gernot Heiser. 2005. Power Management and Dynamic Voltage Scaling: Myths and Facts. In *Proceedings of the 2005 Workshop on Power Aware Real-time Computing*. New Jersey, USA.
- Mark Weiser, Brent Welch, Alan Demers, and Scott Shenker. 1994. Scheduling for Reduced CPU Energy. In *Proceedings of the 1st USENIX Conference on Operating Systems Design and Implementation* (Monterey, California) (*OSDI '94*). USENIX Association, USA, 2es.
- F. Yao, A. Demers, and S. Shenker. 1995. A scheduling model for reduced CPU energy. In *Proceedings of IEEE Annual Foundations of Computer Science*. 374–382.