



HAL
open science

Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks

Bruno Gaujal, Alain Girault, Stéphan Plassart

► **To cite this version:**

Bruno Gaujal, Alain Girault, Stéphan Plassart. Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks. [Research Report] RR-9101, UGA - Université Grenoble Alpes; Inria Grenoble Rhône-Alpes; Université de Grenoble. 2017, pp.1-35. hal-01615835v2

HAL Id: hal-01615835

<https://inria.hal.science/hal-01615835v2>

Submitted on 22 Nov 2017

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.



Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks

GAUJAL Bruno, GIRAULT Alain, PLASSART Stephan

**RESEARCH
REPORT**

N° 9101

September 2017

Project-Teams Polaris and Spades

ISRN INRIA/RR--9101--FR+ENG

ISSN 0249-6399



Dynamic Speed Scaling Minimizing Expected Energy Consumption for Real-Time Tasks

GAUJAL Bruno, GIRAULT Alain, PLASSART Stephan *

Project-Teams Polaris and Spades

Research Report n° 9101 — September 2017 — 37 pages

Abstract: This paper proposes a Markov Decision Process (MDP) approach to compute the optimal on-line speed scaling policy to minimize the energy consumption of a processor executing a finite or infinite set of jobs with real-time constraints. The speed policy is computed off-line but used on-line. We provide several qualitative properties of the optimal policy: monotonicity with respect to the jobs parameters, comparison with on-line deterministic algorithms. Numerical experiments show that our proposition performs well when compared with off-line optimal solutions and out-performs on-line solutions oblivious to statistical information on the jobs. Several extensions are proposed in appendix to take into account the cost of processor speed changes and of task context switches, as well as non-convex power functions to model power leakage.

Key-words: Optimization, Real-Time Systems, Markov Decision Process, Dynamic Voltage and Frequency Scaling

* This work has been partially supported by the LabEx PERSYVAL-Lab.

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Sélection en-ligne de la vitesse minimisant l'énergie dans les systèmes temps-réels

Résumé : Cet article propose d'utiliser la technique des processus de décision markovien (PDM) pour calculer la politique optimale en-ligne de choix de vitesses afin de minimiser l'énergie consommée par un processeur exécutant un ensemble de tâches avec des contraintes temps-réel. La politique de vitesse est calculée avant l'exécution du système temps réel (hors ligne), mais utilisée en ligne. Nous démontrons plusieurs propriétés qualitatives de la politique optimale : monotonie par rapport aux paramètres des tâches, comparaison avec les algorithmes déterministes. Les expérimentations montrent que notre méthode est efficace et proche de la solution optimale hors-ligne, et qu'elle est plus performante que les solutions en-ligne qui ne prennent pas en compte les informations statistiques sur les tâches futures. Nous présentons plusieurs extensions pour prendre en compte les coûts de changement de vitesse du processeur et de changement de contexte entre les tâches, ainsi que les fonctions non convexes de puissance consommée pour modéliser les fuites énergétiques.

Mots-clés : Optimisation, systèmes temps réel, processus de décision markovien, ajustement dynamique de la fréquence et du voltage

Contents

1	Introduction	4
2	Presentation of the problem	5
2.1	Jobs, Processor Speeds, and Power	5
2.2	Problem Statement	6
2.3	Description of the State Space	8
3	Markov Decision Process Solution	9
3.1	Finite Case: Dynamic Programming	10
3.1.1	Dynamic Programming Algorithm (DP)	10
3.1.2	Runtime Process: Table Look-up (TLU_{DP})	11
3.2	Infinite Case: Value Iteration	11
3.2.1	Stationary assumptions and existence of the limit	11
3.2.2	Value Iteration Algorithm (VI)	12
3.2.3	Runtime Process: Table Look-Up (TLU_{VI})	12
3.3	Schedulability Issues	12
3.4	Properties of the Optimal Policy	13
3.4.1	Comparison with Optimal Available (OA)	13
3.4.2	Monotonicity Properties	14
4	Numerical Experiments	15
4.1	Application Scenarios	15
4.2	Implementation Issues	15
4.3	Experimental Set-up, Finite Case	15
4.3.1	Comparison with the Off-line Solution	16
4.3.2	Comparisons with (OA), Sporadic Tasks	17
4.3.3	Comparisons with (OA), Periodic Tasks	19
4.4	Computation Experiments, Infinite Case	22
4.4.1	Lower bound	24
4.4.2	Comparison of TLU_{DP} and TLU_{VI}	24
5	Conclusion and Perspectives	26
A	Convexification of the Power Consumption Function	27
B	Taking into Account the Cost of Speed Changes	28

C	Taking into Account the Cost of Context Switches	30
D	Schedulability of the (OA) Policy	32
E	Size of the State Space	35

1 Introduction

Minimizing the energy consumption of embedded system is becoming more and more important. This is due to the fact that more functionalities and better performances are expected from such systems, together with a need to limit the energy consumption, mainly because batteries are becoming the standard power supplies.

The starting point of this work is the seminal paper of Yao et al. [1] followed by the paper of Bansal et al. [2], which both solve the following problem: Let $(r_i, c_i, d_i)_{i \in \mathbb{N}}$ be a set of jobs, where r_i is the release date (or *arrival time*) of job i , c_i is its *WCET* (or *workload*) *i.e.*, its execution time by the processor operating at its maximal speed, and d_i is its relative *deadline*, *i.e.*, the amount of time given to the processor to execute job i .

The problem is to choose the speed¹ of the processor as a function of time $s(t)$ such that the processor can execute all jobs before their deadlines, and such that the total energy consumption J is minimized. In our problem, J is the *dynamic* energy consumed by the processor: $J = \int_0^T j(s(t))dt$, where T is the time horizon of the problem (in the finite case) and $j(s)$ is the power consumption when the speed is s .

This problem has been solved in Yao et al. [1] in the *off-line* case, *i.e.*, when *all* jobs are known in advance, and when the power function $j(\cdot)$ is a *convex* function of the speed. But of course, the off-line case is unrealistic.

Several solutions for the *on-line* case (only the jobs released at or before time t can be used to select the speed $s(t)$) have been investigated by Bansal et al. in [2]. The authors prove that the on-line algorithm, called Optimal Available (OA), has a competitive ratio of α^α when the power dissipated by the processor working at speed s is of the form $j(s) = s^\alpha$. In CMOS circuits, the value of α is typically 3. In this case, (OA) may spend 27 times more energy than an optimal schedule in the worst case. The principle of the on-line algorithm (OA) is to choose, at each time t , the *smallest* processor speed such that all jobs released at or before time t meet their deadline, under the assumption that no more jobs will arrive after time t .

However, the assumption made by (OA) is questionable. Indeed, the speed selected by (OA) at time t will certainly need to be compensated (*i.e.*, increased) in the future due to jobs released after t , incurring an energetic inefficiency when the j function is convex. In contrast, our intuition is that the best choice is to select a speed *above* the one used by (OA) to *anticipate* on those future job arrivals.

The goal of the paper is to give a precise solution to this intuition by using statistical knowledge of the job arrivals (that could be provided by the user) to select the speed that optimizes the *expected* energy consumption. Other constructions based on statistical knowledge have been done in [3, 4] in a simpler framework, namely for one single job whose execution time is uncertain, or in [5] by using heuristic schemes. In this paper, we show that this general constrained

¹Different communities use the term “speed” or “frequency”, which are equivalent for a processor. In this paper, we use the term “speed”.

optimization problem can be modeled as an unconstrained Markov Decision Process (MDP) by choosing a proper state space that also encodes the constraints of the problem. In particular, this implies that the optimal speed at each time can be computed using a *dynamic programming* algorithm and that the optimal speed at any time t will be a deterministic function of the current state at time t .

Application in Practical Scenarios

Our approach is usable in several practical cases. The first one concerns real-time systems whose tasks are *sporadic*, with no *a priori* structure on the job release times, sizes, and deadlines. In such a case, a long observation of the jobs features can be used to estimate the statistical properties of the jobs: distribution of the inter-release times, distribution of the job sizes, and deadlines.

Another case where our approach is efficient is for real-time systems consisting of several *periodic* tasks, each one with some randomly missing jobs. The uncertainty on the missing jobs may be due, for example, to faulty sensors and/or electromagnetic interference causing transmission losses.

These two cases are explored in the experimental section where our solution is compared with current solutions. Our numerical simulations report a 5% improvement in the sporadic tasks case, and 30% to 50 % improvement in the periodic tasks case.

Outline of the Paper

In the first part of this paper (§ 2), we present our job model and the problem addressed in the paper. In a second part (§ 3), we construct a Markov decision process model of this problem. We investigate its state space and its complexity (§ 2.3), we propose an explicit *dynamic programming* algorithm to solve it when the number of jobs is finite (§ 3.1), and a *Value Iteration* algorithm [6] for the infinite case (§ 3.2). We prove that both solutions have stochastic monotonicity properties with respect to the sizes and the deadlines of the jobs (§ 3.4). Finally, we compute numerically the optimal policy in the finite and infinite horizon cases, and compare its performance with off-line policies and “myopic” policies like (OA), oblivious to the arrival of future jobs (§ 4). Moreover we present, in appendix, some generalizations such as taken into account the cost of processor speed changes, the cost of task context switches, and the non-convex power functions of model power leakage.

2 Presentation of the problem

2.1 Jobs, Processor Speeds, and Power

We consider a real-time system with one uni-core processor that executes real-time jobs, sporadic and independent. Each job i is defined by 3 integer values: (r_i, c_i, d_i) where r_i is the release time, c_i is the WCET, and d_i is the relative deadline. We assume that all jobs have a bounded size and a bounded deadline:

$$\exists C \text{ s.t. } \forall i, c_i \leq C, \quad \exists \Delta \text{ s.t. } \forall i, d_i \leq \Delta.$$

We further assume that several jobs may arrive simultaneously but that the total work brought to the processor at time t is also bounded by C .

The CPU processing speed $s(t)$ can vary in time over a finite number of rational speeds between 0 and s_{max} : $s(t) \in \mathcal{S} = \{0, s_1, \dots, s_k, s_{max}\}$, s_{max} being the maximal processor speed.

With no loss of generality, we will scale the speeds (as well as the WCETs), so that s_1, \dots, s_k, s_{max} are all integer numbers. Typically, a set of admissible speeds $\{0, 1/4, 1/2, 3/4, 1\}$ is replaced by the set $\{0, 1, 2, 3, 4\}$, while all WCETs are multiplied by 4. This is done to avoid manipulating rational numbers.

We consider that the power dissipated by the CPU working at speed $s(t)$ at time t is $j(s(t))$, so that the energy consumption of the processor from time 0 to time T is $J = \int_0^T j(s(t))dt$. Classical choices for the power consumption j are convex increasing functions of the speed (see [1, 2]), based on classical models of CMOS circuits for power dissipation, or star-shaped functions [7] to further take into account static leakage. Here, the function j is *arbitrary*. However several structural properties of the optimal speed selection will only hold when the function j is convex. In the numerical experiments (§ 4), several choices of j are used, to take into account different models of power consumption.

For the sake of simplicity, we only consider the following simple case: context switching time is null, speed changes are instantaneous, and the power consumption function $j(\cdot)$ is convex. However, preemption times, time lags for speed changes as well as non-convex energy costs can be taken into account with minimal adaptation to the current model. A detailed description of all these generalizations is available in appendix of this report.

2.2 Problem Statement

The objective is to choose at each time t the speed $s(t)$, in order to minimize the total energy consumption over the total time horizon, while satisfying all the real-time constraints. Furthermore, the choice must be made on-line, *i.e.*, it can only be based on past and current information. In other words, only the jobs released at or before time t are known.

The information (or history) $\mathcal{H}(t)$ at time t is the set of all the past and current jobs together with the past speeds selection:

$$\mathcal{H}(t) = \{(r_i, c_i, d_i) | r_i \leq t\} \cup \{s(u), u \leq t\} \quad (1)$$

Notice that in this model, unlike in [3, 4], the workload c_i and the deadline d_i are known at the release time of job i ².

The on-line energy minimization problem (\mathcal{P}) is:

*Find online speeds $s(t)$ (*i.e.* $s(t)$ can only depend on the history $\mathcal{H}(t)$) and a scheduling policy $R(t)$ to minimize $\mathbb{E} \int_0^T j(s(t))dt$ under the constraint that no job misses its deadline.*

Since all release times and job sizes are integer numbers, the information $\mathcal{H}(t)$ only changes at integer points, if we consider that the speed $s(t)$ can only change at integer points. In the following we focus on integer times $t \in \mathbb{N}$.

Let (s^*, R^*) be an optimal solution to problem (\mathcal{P}). Since the energy consumption does not depend on the schedule (preemption is assumed to be energy-free) and since the *Earliest*

²When the actual workload can be smaller than WCET, our approach still applies by modifying the state evolution Eq. (2), to take into account early termination of jobs.

Deadline First (EDF) scheduling policy is optimal for schedulability [8], then (s^*, EDF) is also an optimal solution to problem (\mathcal{P}) . In the following, we will always assume with no loss of optimality that the processor uses EDF to schedule its jobs. This implies that the only useful information to compute the optimal speed at time t , out of the whole history $\mathcal{H}(t)$, is simply the *remaining work*. The remaining work at time t is an increasing function $w_t(\cdot)$ such that $w_t(u)$ is the amount of work that remains to be done before time $t + u$. Since all available speeds, job sizes and deadlines are integer numbers, the remaining work $w_t(u)$ is an integer valued *càdlàg* staircase function.

This is illustrated in Figure 1 that shows the set of jobs released just before $t = 4$, namely $J_1 = (0, 2, 4)$, $J_2 = (1, 1, 5)$, $J_3 = (2, 2, 6)$, $J_4 = (3, 2, 4)$, $J_5 = (4, 0, 6)$, as well as the speeds chosen by the processor up to time $t = 4$: $s_0 = 1$, $s_1 = 0$, $s_2 = 2$, $s_3 = 1$. Function $A(t)$ is the amount of work that has arrived before time t . Function $D(t)$ is the amount of work that must be executed before time t . This requires a detailed explanation: the first step of $D(t)$ is the deadline of J_1 at $t = 0 + 4 = 4$; the second step is for J_2 at $t = 1 + 5 = 6$; the third step is for J_4 at $t = 3 + 4 = 7$; the fourth step is for J_3 at $t = 2 + 6 = 8$. Hence the step for J_4 occurs *before* the step for J_3 . This is because Figure 1 depicts the situation at $t = 4$. At $t = 3$ we would only have seen the step for J_3 . Finally, function $e(t)$ is the amount of work already executed by the processor at time t ; in Figure 1, the depicted function $e(t)$ has been obtained with an arbitrary policy (*i.e.*, non optimal). Finally, the remaining work function $w_t(u)$ is exactly the portion of $D(t)$ that remains “above” $e(t)$. In Fig. 1, we have depicted in red the staircase function $w_t(u)$ for $t = 4$.




Figure 1: Construction of the remaining work function $w_t(\cdot)$ at $t = 4$, for jobs $J_1 = (0, 2, 4)$, $J_2 = (1, 1, 5)$, $J_3 = (2, 2, 6)$, $J_4 = (3, 2, 4)$, $J_5 = (4, 0, 6)$, and processor speeds $s_0 = 1$, $s_1 = 0$, $s_2 = 2$, $s_3 = 1$. $A(t)$ is the amount of work that has arrived before time t . $D(t)$ is the amount of work that must be executed before time t . $e(t)$ is the amount of work already executed by the processor at time t .

Remark 1. *The on-line algorithm Optimal Available (OA) mentioned in the introduction is also based on the remaining work function: The speed of the processor at time t is the smallest slope of all linear functions above w_t . This is illustrated in Figure 1: the speed that (OA) would choose at time $t = 4$ is the slope of the green dotted line marked (OA); in the discrete speeds case (finite number of speeds), the chosen speed would be the smallest available speed just above the green dotted line.*

Since the remaining work function w_t is the only relevant information at time t , out of the whole history $\mathcal{H}(t)$, needed by the processor to choose its next speed, we call w_t the *state* of the

system at time t .

2.3 Description of the State Space

To formally describe the state space \mathcal{W} (i.e., all the possible remaining work functions) and the evolution of the state over time, we introduce several constructors.

Definition 1. We define the following operators:

– The time shift operator $\mathbb{T}f$ is the shift on the time axis of function f , defined as: $\forall t \in \mathbb{R}$, $\mathbb{T}f(t) = f(t + 1)$.

– The positive part of a function f is $f^+ = \max(f, 0)$.

– The unit step function (Heaviside function), denoted H_t , is the discontinuous step function such that $\forall u \in \mathbb{R}$:

$$H_t(u) = \begin{cases} 0 & \text{if } u < t \\ 1 & \text{if } u \geq t \end{cases}$$

Lemma 1. Let (r_n, c_n, d_n) be a job that arrives at time $t = r_n$. If the processor speed at time $t - 1$ is s_{t-1} , then at time t the remaining work function becomes:

$$w_t(\cdot) = \mathbb{T}[(w_{t-1}(\cdot) - s_{t-1})^+] + c_n H_{d_n}(\cdot) \quad (2)$$

Proof. Between $t - 1$ and t , the processor working at speed s_{t-1} executes s_{t-1} amount of work, so the remaining work decreases by s_{t-1} . The remaining work cannot be negative by definition, hence the term $(w_{t-1}(\cdot) - s_{t-1})^+$. After a time shift by one unit, a new job is released at time t , bringing c_n additional work with deadline $t + d_n$, hence the additional term $c_n H_{d_n}(\cdot)$. \square

Let us illustrate the state change over an example. In Fig. 2, a new task arrival leads to a modification of the remaining work. The red line corresponds to the previous remaining work function in r_n , and the blue line corresponds to the new remaining work function. This new state has been modified by the job arrival (r_n, c_n, d_n) . A quantity s_{n-1} of work has been executed by the processor.

The state space is finite, as stated by the following proposition. This proposition also gives the size of the state space, which will play a major role in the complexity of the dynamic programming algorithm to compute the optimal speed.

Proposition 1. If C is the maximal size of a job and Δ its maximal deadline, then the size $Q(C, \Delta)$ of the state space \mathcal{W} is:

$$Q(C, \Delta) = \sum_{y_1=0}^C \sum_{y_2=y_1}^{2C} \sum_{y_3=y_2}^{3C} \cdots \sum_{y_\Delta=y_{\Delta-1}}^{\Delta C} 1. \quad (3)$$

It can be computed in closed form³:

$$Q(C, \Delta) = \frac{1}{1 + C(\Delta + 1)} \binom{(C + 1)(\Delta + 1)}{\Delta + 1} \quad (4)$$

$$\approx \frac{e}{\sqrt{2\pi}} \frac{1}{(\Delta + 1)^{3/2}} (eC)^\Delta \quad (5)$$

³In Eq. (4), the notation $\binom{n}{k}$ is the binomial coefficient.

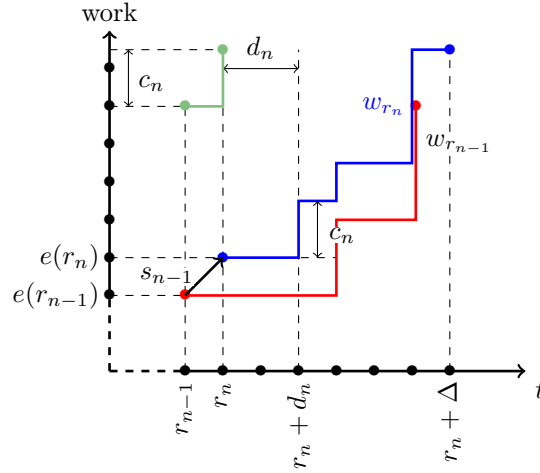


Figure 2: Task arrival at time r_n and the corresponding state change. The red line corresponds to the previous remaining work function. The blue line corresponds to the new remaining work function.

Proof. The proof is postponed to appendix E. \square

3 Markov Decision Process Solution

Since the state space \mathcal{W} is finite, one can effectively compute the optimal speeds in each possible state. In this section, we provide such algorithms to compute that optimal speed selection in two cases: when the time horizon is finite and when it is infinite. In the finite case, we minimize the *total* energy consumption, while in infinite case we minimize the *average* energy consumed per time unit.

In both cases, we compute offline the optimal *policy* σ_t^* that says which speed the processor should use at time t in all its possible states. At runtime, at each time t , the processor chooses the speed that corresponds to its current remaining work w as $s := \sigma_t^*(w)$.

The algorithms to compute the policy σ^* are based on a *Markovian evolution* of the jobs. Without loss of generality, we will assume $r_n = n$. This means that a new job arrives at each time slot, this job being of size 0 ($c_n = 0$) if no real work actually arrived at time n . Under this point of view, the job stream is given by a family of distributions $(\phi_w(\cdot, \cdot))_{w \in \mathcal{W}}$ that give the size and the deadline of a job arriving at time t : For any $0 \leq \gamma \leq C$, and for any $0 \leq \delta \leq \Delta$,

$$\phi_w(\gamma, \delta) = \Pr(c_t = \gamma, d_t = \delta | w_t = w) \quad (6)$$

Once ϕ is given, the transition matrix $P_s(w, w')$ from state w to w' when the speed chosen by the processor is s is:

$$P_s(w, w') = \begin{cases} \phi_w(\gamma, \delta) & \text{if } w' = \mathbb{T}[(w - s)^+] + \gamma H_\delta \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

3.1 Finite Case: Dynamic Programming

We suppose in this Section that the time horizon is finite and equal to T . It means that we only consider a finite number of jobs. The goal here is to minimize the processor energy consumption during the time interval $[0, T]$.

We compute the minimal total expected processor energy consumption (J^*) from 0 to T . This means that we want to solve the following problem.

If the initial state is w_0 , then

$$J^*(w_0) = \min_{\sigma} \left(\mathbb{E} \left(\sum_{t=0}^T j(\sigma_t(w_t)) \right) \right) \quad (8)$$

where σ is taken over all possible *policies* of the processor: $\sigma(w, t)$ is the speed used at time t if the state is w . The only constraint on $\sigma(w, t)$ is that it must belong to the set of admissible speeds, *i.e.*, $\sigma(w, t) \in \mathcal{S}$, and it must be large enough to execute the remaining work at the next time step:

$$\forall t, \sigma(w, t) \geq w(1) \quad (9)$$

The set of *admissible speeds* in state w is denoted $\mathcal{A}(w)$ and is therefore defined by:

$$\mathcal{A}(w) = \{s \in \mathcal{S} \text{ s.t. } s \geq w(1)\} \quad (10)$$

J^* can be computed using a backward induction. Let $J_t^*(w)$ be the minimal expected energy consumption from time t to T , if the state at time t is w ($w_t = w$). We present in the next section an algorithm to compute J^* .

3.1.1 Dynamic Programming Algorithm (DP)

We use a backward induction (Dynamic Programming) to recursively compute the expected energy consumption J^* and the optimal speed policy σ^* . We use the finite Horizon-Policy Evaluation Algorithm from [6] (p. 80). We obtain an *optimal policy*, which corresponds to the processor speed that one must apply in order to minimize the energy consumption (Algorithm 1).

Algorithm 1 Dynamic Programming Algorithm (DP) to compute the optimal speed for each state and each time.

```

 $t \leftarrow T$  % time horizon
for all  $w \in \mathcal{W}$  do  $J_t^*(w) \leftarrow 0$  end for
while  $t \geq 1$  do
  for all  $w \in \mathcal{W}$  do
     $J_{t-1}^*(w) \leftarrow \min_{s \in \mathcal{A}(w)} \left( j(s) + \sum_{w' \in \mathcal{W}} P_s(w, w') J_t^*(w') \right)$ 
     $\sigma_{t-1}^*[w] \leftarrow \arg \min_{s \in \mathcal{A}(w)} \left( j(s) + \sum_{w' \in \mathcal{W}} P_s(w, w') J_t^*(w') \right)$ 
  end for
   $t \leftarrow t - 1$  % backward computation
end while
return all tables  $\sigma_t^*[\cdot] \quad \forall t = 0 \dots T - 1$ .

```

The complexity to compute the optimal policy $\sigma_t^*(w)$ for all possible states and time steps is $\mathcal{O}(T|S|C\Delta Q(C, \Delta))$. The combinatorial explosion of the state space makes is very large when the sizes of the jobs and the maximum deadline are large. Note however that this pre-computation is done *off-line*. At runtime, the processor simply considers the current remaining work w_t at time t and uses the pre-computed speed $\sigma_t^*(w_t)$ to execute the job with the earliest deadline.

3.1.2 Runtime Process: Table Look-up (TLU_{DP})

At runtime, the processor computes the current state w and simply uses a table look-up (TLU) to obtain its optimal speed $\sigma_t^*[w]$, the speed tables having been computed offline by (DP). This is algorithm 2.

Algorithm 2 Runtime process (TLU_{DP}) used by the processor to apply the optimal speed

For Each $t = 0 \dots T - 1$
 Update w_t using Eq. (2)
 Set $s := \sigma_t^*[w_t]$
 Execute the job(s) with earliest deadline job(s) at speed s for one time unit
End

The runtime cost for TLU_{DP} is $\mathcal{O}(1)$.

3.2 Infinite Case: Value Iteration

When the time horizon is infinite, the total energy consumption is also infinite whatever the speed policy. Instead of minimizing the total energy consumption, we minimize the *average* expected energy consumption per time-unit, denoted g . We therefore find the optimal policy σ^* that minimizes g . In mathematical terms, we want to solve the following problem. Compute

$$g^* := \min_{\sigma} \mathbb{E} \left(\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T j(\sigma(w_t)) \right) \quad (11)$$

under the constraint that no job misses its deadline.

3.2.1 Stationary assumptions and existence of the limit

In the following we will make the following additional assumption on the jobs: The size and the deadline of the next job have *stationary distributions* (*i.e.*, they do not depend on time). We further assume that the probability that the next job is of size 0 (meaning that no work arrives in the next time slot) is strictly positive.

Under these two assumptions, the state space transition matrix is *unichain* (see [6] for a precise definition). Basically, this means that starting from state $w_0 = (0, \dots, 0)$, it is possible to go back to state w_0 , no matter what speed choices have been made and what jobs have occurred. This is possible because, with positive probability, jobs of size 0 can arrive for long enough a time so that all past deadlines have been met and the state goes back to w_0 .

When the state space is unichain, the limit in Eq. (11) always exists [6] and can be computed with an arbitrary precision using a *value iteration* algorithm (VI), presented in the next section.

3.2.2 Value Iteration Algorithm (VI)

This algorithm finds a *stationary* policy σ (here σ will not depend on t), which is optimal, and gives an approximation of the gain (average reward value g^*) with an arbitrary precision. In the following value iteration algorithm, the quantity u^n can be seen as the total energy up to iteration n :

Algorithm 3 Value Iteration Algorithm (VI) to compute the optimal speeds in each state and the average energy cost per time unit.

```

 $u^0 \leftarrow (0, 0, \dots, 0), u^1 \leftarrow (1, 0, \dots, 0)$ 
 $n \leftarrow 1$ 
 $\varepsilon > 0$  % stopping criterion
while  $\text{span}(u^n - u^{n-1}) \geq \varepsilon$  do
  for all  $w \in \mathcal{W}$  do
     $u^{n+1}(w) \leftarrow \min_{s \in \mathcal{A}(w)} \left\{ j(s) + \sum_{w' \in \mathcal{W}} P_s(w, w') u^n(w') \right\}$ 
  end for
   $n \leftarrow n + 1$ 
end while
Choose any  $w \in \mathcal{W}$  and let  $g^* \leftarrow u^n(w) - u^{n-1}(w)$ 
for all  $w \in \mathcal{W}$  do
   $\sigma^*[w] \in \arg \min_{s \in \mathcal{A}(w)} \left\{ j(s) + \sum_{w' \in \mathcal{W}} P_s(w, w') u^n(w') \right\}$ 
end for
return  $\sigma^*$ 

```

In Algorithm 3, the *span* of a vector is the difference between its maximal value and its minimal value: $\text{span}(x) = \max_i(x_i) - \min_i(x_i)$. A vector with a span equal to 0 has all its coordinates equal.

Algorithm 3 computes both the optimal average energy consumption per time unit (g^*) with a precision ε as well as an ε -optimal speed to be selected in each state ($\sigma^*[w]$).

The time complexity to compute the optimal policy depends exponentially on the precision $\frac{1}{\varepsilon}$. The numerical experiments show that convergence occurs reasonably fast (see Section 4).

3.2.3 Runtime Process: Table Look-Up (TLU_{VI})

As for (TLU_{DP}), at any integer time $t \in \mathbb{N}$, the processor computes its current state w using Eq.(2) and retrieves its optimal speed $s := \sigma^*[w]$ by looking-up in the table σ^* , pre-computed by (VI). This is identical to Algorithm 2).

3.3 Schedulability Issues

Depending on the maximum available speed (s_{max}) and on the jobs characteristics (C and Δ), all sets of jobs might not be schedulable. For instance, if $s_{max} < C$, then no speed policy can guarantee the schedulability: indeed, if a job of size C and relative deadline 1 arrives, then the processor cannot process it before its deadline.

The case where $s_{\max} = C$ is borderline because there exists a unique speed selection that guarantees that no job will miss its deadline. Any valid on-line algorithm must always set $s(t) = c_t$ at all times t , where by convention c_t is the WCET of the job arrived at time t if it exists, or 0 if no job arrives at time t : Choosing a speed smaller than c_t at time t will imply missing a deadline in some future time slot if a sequence of jobs of size C and deadline 1 arrive at times $t+1, t+2, \dots, t+\Delta$. These jobs will preempt the processor at each time slot $t+1, \dots, t+\Delta$ if they do not miss their deadline, and some work present at time t will not be finished at time $t+\Delta$, so one job deadline will be missed. Therefore our optimization problem \mathcal{P} is only interesting when $s_{\max} > C$.

Now let us consider the case where $s_{\max} > C$. When the (DP) algorithm is used, the optimal energy consumption in state w satisfies:

$$J_{t-1}^*(w) \leftarrow \min_{s \in \mathcal{A}(w)} \left(j(s) + \sum_{w' \in \mathcal{W}} P_s(w, w') J_t^*(w') \right) \quad (12)$$

If $\mathcal{A}(w)$ is empty (this is equivalent to the fact that, in state w , some job will miss its deadline at time t) one can set $J_{t-1}^*(w) = +\infty$. Using this convention, we can state that the optimal solution will never miss a deadline starting from a system with no pending job at time 0 is equivalent to $J_0^*(0) < \infty$.

Starting from an empty system with no pending job, we know that the speed policy $\tilde{\sigma}$ that sets $\forall t \ s_t = c_t \leq C$ is schedulable. So its energy consumption will be finite. Since the optimal solution σ^* yields a smaller energy consumption than $\tilde{\sigma}$, its energy consumption will also be finite, hence the optimal solution will also be schedulable.

To conclude, we get the following result:

Theorem 1. *Starting from an empty system, schedulability with (DP) or (VI) is guaranteed if and only if $s_{\max} \geq C$.*

We also note that schedulability is another advantage of the (DP) solution in comparison with (OA). Contrary to (DP), (OA) is not always schedulable. We present in appendix D a worst-case scenario where (OA) fails and we prove that (OA) is schedulable if and only if

$$s_{\max} \geq C \left(\frac{\Delta}{\Delta-1} + h_{\Delta-2} \right) \quad (13)$$

where h_n is the n -th harmonic number.

3.4 Properties of the Optimal Policy

In this section, we show several structural properties of the optimal policy σ^* that are true for both the finite and infinite case.

3.4.1 Comparison with Optimal Available (OA)

Optimal Available (OA) is an on-line policy that chooses the speed $s^{(OA)}(w_t)$ in state w_t to be the optimal speed at time t , should no further jobs arrive in the system. More precisely, at time t and in state w_t , the (OA) policy uses the speed

$$s^{(OA)}(w_t) = \max_u \frac{w_t(u)}{u} \quad (14)$$

We first show that, under any state $w \in \mathcal{W}$, the optimal speed $\sigma^*(w)$ is always larger than $s^{(OA)}(w)$.

Proposition 2. *In the finite or infinite case, the optimal speed policy satisfies $\sigma^*(w) \geq s^{(OA)}(w)$ for all state $w \in \mathcal{W}$, if the power consumption j is a convex function of the speed.*

Proof. The proof is based on the observation that (OA) uses the optimal speed assuming that no new job will come in the future. Should some job arrive later, then the optimal speed will have to increase. We first prove the result when the set of speeds \mathcal{S} is the whole *real* interval $[0, s_{\max}]$ (continuous speeds).

Two cases must be considered. If $s^{(OA)}(w_t) = \max_u \frac{w_t(u)}{u}$ is reached for $u = 1$ (i.e., $s^{(OA)}(w_t) = w_t(1)$), then $\sigma^*(w) \geq s^{(OA)}(w)$ by definition because the set of admissible speeds $\mathcal{A}(w_t)$ only contains speeds larger than $w_t(1)$ (see Eq. (14)).

If the maximum is reached for $u > 1$, then $\mathcal{A}(w_t)$ may enable the use of speeds below $w_t(1)$.

Between time t and $t + u$, some new job may arrive. Therefore, the optimal policy should satisfy $\sum_{i=0}^{u-1} \sigma^*(w_{t+i}) \geq w_t(u)$.

The convexity of the power function j implies⁴ that the terms of the optimal sequence $\sigma^*(w_t), \dots, \sigma^*(w_{t+u-1})$ must all be above the average value (which is larger than $w_t(u)/u = s^{(OA)}(w_t)$). In particular, for the first term, $\sigma^*(w_t) \geq s^{(OA)}(w_t)$.

Now, if the set of speeds is finite, then the optimal value of $\sigma^*(w_t)$ must be one of the two speeds in \mathcal{S} surrounding $\sigma^*(w_t)$, $s_1 < \sigma^*(w_t) \leq s_2$. If the smallest one s_1 is chosen, this implies that further choices for $\sigma^*(w_{t+i})$ will have to be larger or equal to s_2 , to compensate for the work surplus resulting from choosing a speed below $\sigma^*(w_t)$. This implies that it is never sub-optimal to choose s_2 in the first place (by convexity of the j function).

This trajectorial argument is true almost surely, so that the inequality $\sigma^*(w_t) \leq s^{(OA)}(w_t)$ will also hold with the *expected* energy over both a finite or infinite time horizon. \square

3.4.2 Monotonicity Properties

Let us consider two sets of jobs \mathcal{T}_1 and \mathcal{T}_2 for which we want to apply our speed scaling procedure. We wonder which of the two sets uses more energy than the other when optimal speed scaling is used for both.

Of course, since jobs have random features, we cannot compare them directly, but instead we can compare their distributions. We assume in the following that the sizes and deadlines of the jobs in \mathcal{T}_1 (resp. \mathcal{MT}_2) follow a distribution ϕ_1 (resp. ϕ_2) independent of the current state w and of the time t .

Definition 2. *Let us define a stochastic order (denoted \leq_{st}) between the two sets of jobs \mathcal{T}_1 and \mathcal{T}_2 as follows. $\mathcal{T}_2 \leq_{st} \mathcal{T}_1$ if the respective distributions ϕ_1 and ϕ_2 are comparable: For any job (c_1, d_1) with distribution ϕ_1 and any job (c_2, d_2) with distribution ϕ_2 , we have:*

$$\forall \gamma, \delta, \quad \Pr(c_2 \geq \gamma, d_2 \leq \delta) \leq \Pr(c_1 \geq \gamma, d_1 \leq \delta) \quad (15)$$

Proposition 3. *If $\mathcal{T}_2 \leq_{st} \mathcal{T}_1$, then:*

1. *over a finite time horizon T , the total energy consumption satisfies $J^{(2)} \leq J^{(1)}$ (computed with Eq. (8));*

⁴Actually, we use the fact that the sum $\sum_{i=0}^{u-1} j(s)$ is Schur-convex when j is convex (see [9]).

2. in the infinite time horizon case, the average energy consumption per time unit satisfies $g^{(2)} \leq g^{(1)}$ (computed with Eq. (11)).

Proof. Case 1: The definition of $\mathcal{T}_2 \leq_{st} \mathcal{T}_1$ implies that we can couple the set of jobs \mathcal{T}_1 with the set of jobs \mathcal{T}_2 , such that at each time $t \leq T$, $J_t^{(1)} = (t, c_t^1, d_t^1)$ and $J_t^{(2)} = (t, c_t^2, d_t^2)$ with $c_t^2 \leq c_t^1$ and $d_t^2 \geq d_t^1$ (see [10]). In turn, this implies that the optimal sequence of speeds selected for \mathcal{T}_1 is admissible for \mathcal{T}_2 , hence the optimal sequence for \mathcal{T}_2 should have a better performance. Since this is true for any set of jobs generated using ϕ_1 , it is also true in expectation, hence $J^{(2)} \leq J^{(1)}$.

Case 2: We just use the fact that the optimal sequence for \mathcal{T}_2 is better than the optimal sequence for \mathcal{T}_1 over any finite horizon T . Letting T go to infinity shows that the average energy cost per time unit will also be better for \mathcal{T}_2 . \square

4 Numerical Experiments

4.1 Application Scenarios

In this part, we come back to the possible scenarios described in the introduction. The numerical experiments are divided in 2 cases: For § 4.3.1 and § 4.3.2, we consider a situation with a single sporadic task whose jobs are identical. The inter-arrival times of the jobs are estimated to follow a geometric distribution⁵, while their sizes and deadlines remain constant.

The second set of experiments deals with another type of real-time systems made of several periodic tasks. Each task is characterized by its offset, period, size, and deadline. In addition, jobs may be lost (for example due to sensor failures and/or perturbations on transmission links). These losses are modeled by a loss probability.

All the experiments reported below are based on these two scenarios.

4.2 Implementation Issues

The state space \mathcal{W} has a rather complex structure and is very large. Therefore, the data structure used in the implementations of Algorithms 1 and 3 must be very efficient to sweep the state space as well as to address each particular state when state changes occur. This is done by using a hashing table to retrieve states according to a multi-dimensional *key* that represents the state (the vector $(w(1), w(2) - w(1), \dots, w(\Delta) - \sum_{k=1}^{\Delta-1} w(k))$) and a recursive procedure based on Eq. (3) to traverse the state space.

The implementation of Algorithms 1 and 3 has been done in R to take advantage of the possibility to manipulate linear algebraic operation easily, and in C when the state space was too large to be efficiently handled in R.

4.3 Experimental Set-up, Finite Case

Our experiments are done in two steps: Firstly we compute the optimal speeds for each possible state. To compute this, we use Algorithm 1 (DP) or 3 (VI), and we store in the σ^* table the optimal speed for each possible state of the system.

⁵Such an estimation can be obtained through the measurement of many traces of the real-time system.

Secondly, to compare different speed policies, we simulate a sequence of jobs (produced by our real-time tasks, see § 4.1) over which we use our (TLU_{DP}) solution or other solutions (*e.g.*, off-line or (OA)) and we compute the corresponding energy consumption.

In a nutshell, the experiments show that our solution performs very well in practice, almost as well as the optimal *off-line* solution (§ 4.3.1). As for the comparison with (OA), in most of our experiments, (TLU_{DP}) outperforms (OA) by 5% on average in the sporadic case when job inter-arrival times are iid⁶ (§ 4.3.2). In the periodic case, where jobs are more predictable, the gap with (OA) grows to about 50% (§ 4.3.3).

4.3.1 Comparison with the Off-line Solution

To evaluate our on-line algorithm, we compare it with the off-line solution computed on a simulated set of jobs, generated using the distribution described below. We draw the aggregated work done by the processor (the respective speeds are the slopes) in two cases: the optimal off-line solution that only uses speeds in the finite set \mathcal{S} , and the (TLU_{DP}) solution. Figure 3 displays the result of our simulations in the following case.

Example 1. *One sporadic task τ_1 with jobs of size $c_1 = 2$ and deadline $d_1 = 5$. The inter-arrival times are random, modeled by a geometric distribution of parameter $p_1 = 0.6$.*

As for the processor, we consider that it can only use 4 speeds $\mathcal{S} = \{0, 1, 2, 5\}$ and that its energy consumption per time unit is given by the function $j(s) = s^3$. Here, the maximal speed is large enough so that schedulability is not an issue: $5 = s_{max} > C = 2$ (§3.3). (Note that (OA) is not schedulable).

The result over one typical simulation run is displayed in Figure 3.

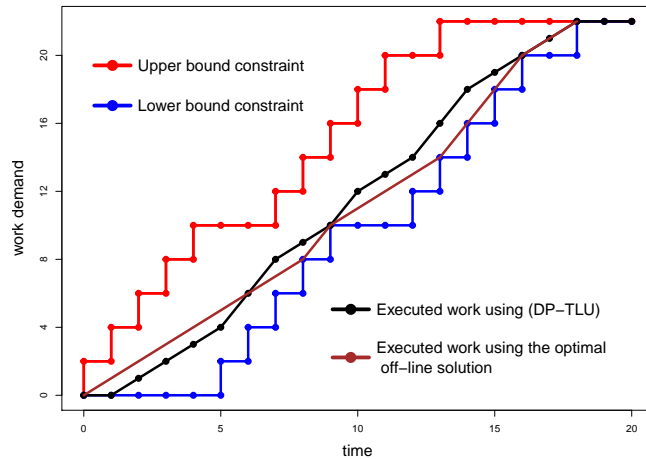


Figure 3: Comparison of the executed work of off-line and (TLU_{DP}) solution on one simulation of Example 1. Red curve: workload arrived between 0 and T ; Blue curve: workload deadlines from 0 to T ; Brown curve: work executed using the optimal off-line speeds; Black curve: work executed using the speed selection computed by (TLU_{DP}).

⁶Independent and identically distributed random variables.

The associate energy consumption curve is:

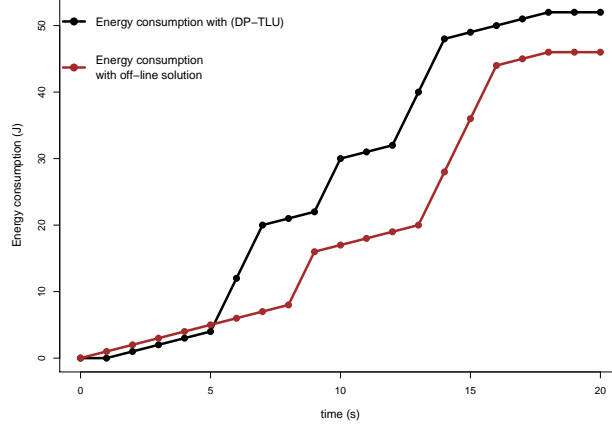


Figure 4: Comparison of the energy consumption of the off-line solution (in red) and the (TLU_{DP}) solution (in black).

As expected, in Figure 3 (TLU_{DP}) consumes more energy than the off-line case. However, the off-line and the (TLU_{DP}) solution only differ in few points: speed 0 is used once by (TLU_{DP}) but is never used by the off-line solution. Speed 2 is used 5 times by (TLU_{DP}) and only 4 times in the off-line case. So the energy consumption gap between the two is $2^3 + 0^3 - 1^3 - 1^3 = 6 J$. The total energy consumption under the off-line solution is 46 J, while the total energy consumption under the (TLU_{DP}) solution is 52 J, so the gap is less than 12 % of the total energy consumption.

4.3.2 Comparisons with (OA), Sporadic Tasks

Recall that, under (OA), the processor speed at time t in state w_t is set to $\max_u \frac{w_t(u)}{u}$. However, when the number of speeds is finite, (OA) must be adapted: the speed $s^{(OA)}(w_t)$ is thus set to the smallest speed in \mathcal{S} larger than $\max_u \frac{w_t(u)}{u}$.

As a consequence, to compare both solutions, the number of possible speeds must be large enough to get a chance to see a difference between the two. Here, we have simulated two sporadic tasks with the following parameters.

Example 2. Two sporadic tasks τ_1 and τ_2 . τ_1 has job of size $c_1 = 3$, deadline $d_1 = 3$, and arrival rate $p_1 = 0.6$, while τ_2 has jobs of size $c_2 = 6$, deadline $d_2 = 3$, and an arrival rate $p_2 = 0.2$. The processor can use 4 processor speeds $\mathcal{S} = \{0, 1, 2, 3, 4\}$ and its energy consumption per time unit follows the function $j(s) = s^3$.

We ran an exhaustive experiment made of 10,000 simulations of sequences of jobs over which we computed the relative energy gain of (TLU_{DP}) over (OA) in percentage. The gain percentage of (TLU_{DP}) was in the range [5.17, 5.39] with a 95% confidence interval and an average value of 5.28%.

Even if this gain is not very high, one should keep in mind that it comes for free once the (DP) solution has been computed. Indeed, using (TLU_{DP}) online takes a constant time to select

the speed (table look-up) while using (OA) online takes $\mathcal{O}(\Delta)$ to compute $\max_u \frac{w_t(u)}{u}$.

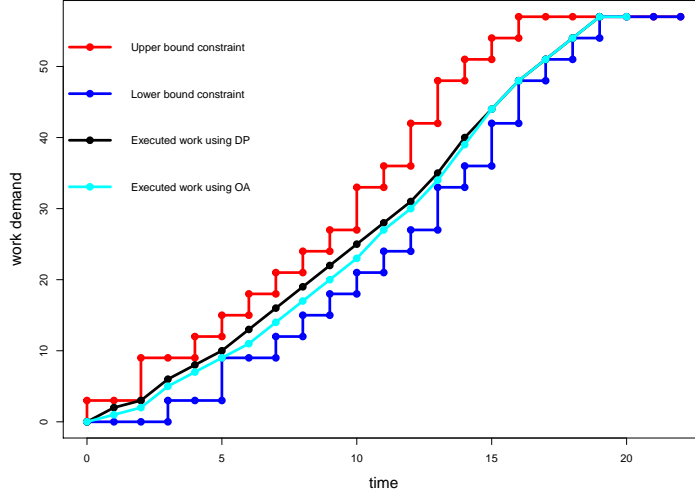


Figure 5: Comparison of the executed work between (OA) and (TLU_{DP}) solutions, with fixed deadlines $d_n = 3$, size $c_n \in \{0, 2, 4\}$ with respective probabilities (0.2, 0.6, 0.2), and processor speeds in $\{0, 1, 2, 3, 4\}$.

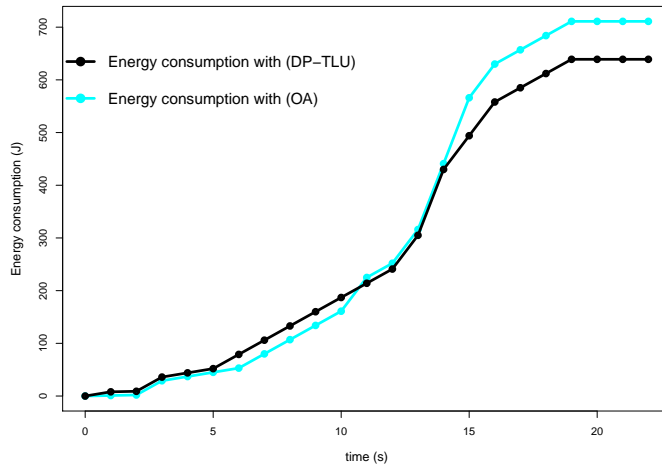


Figure 6: Comparison of the energy consumption between (OA) and (TLU_{DP}) solutions, with fixed deadlines $d_n = 3$, size $c_n \in \{0, 2, 4\}$ with respective probabilities (0.2, 0.6, 0.2), and processor speeds in $\{0, 1, 2, 3, 4\}$.

Figure 5 shows a comparison between (OA) and (TLU_{DP}). The total work executed by

(TLU_{DP}) solution is always above the total work executed by (OA), as stated in Proposition 2. Moreover, we notice in Figure 5 that the consumed energy is more important at the beginning with (TLU_{DP}) than with (OA), because we anticipate the work to do in the future. The processor executes more work so it consumes more energy with (TLU_{DP}) before time $t = 11$; but after this time, it's the opposite, the energy consumed by (TLU_{DP}) is lower than the energy consumed by (OA). Over the whole period, (TLU_{DP}) outperforms (OA).

Indeed, on this example, the total energy consumption for (OA) is 711 J while that for (TLU_{DP}) is 639 J . (TLU_{DP}) outperforms (OA) policy by a margin of around 10 %. Even if this gain is not very high, one should keep in mind that, again, it comes for free once the (TLU_{DP}) solution has been computed.

4.3.3 Comparisons with (OA), Periodic Tasks

We now consider several examples consisting of two or more periodic tasks.

Example 3. Two periodic tasks τ_1 and τ_2 . For task τ_1 , the period is 2, the offset is 0, the job size is $c_1 = 2$, and the deadline is $d_1 = 2$. The loss probability is $p_1 = 0.2$. For task τ_2 , the period is 2, the offset is 1, the job size $c_2 = 4$, and the deadline $d_2 = 1$. The loss probability is $p_2 = 0.25$.

Figure 7 depicts a typical simulation over 20 time steps.




Figure 7: Executed work using (OA) (light blue curve) and (TLU_{DP}) (black curve) with the two task models of Example 3.

The energy consumption is as follow:

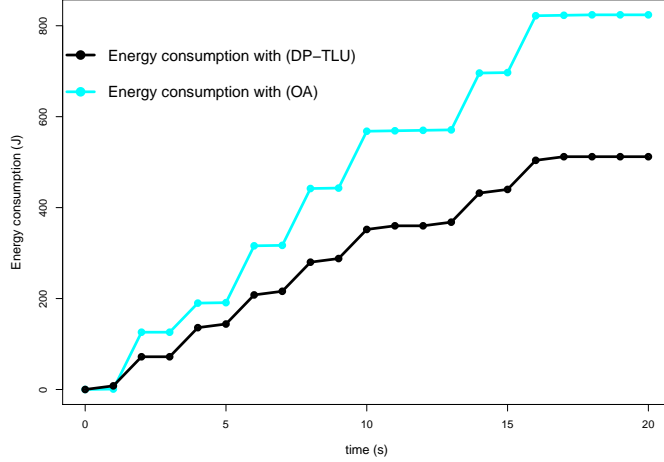


Figure 8: Energy consumption using (OA) (light blue curve) and (TLU_{DP}) (black curve) under the two task models of Example 3.

The total energy consumption over the 20 units of time displayed in Figure 7 is 513 J for (TLU_{DP}), and 825 J for (OA), so more than 60 % higher. Here (TLU_{DP}) has a clear advantage because the job stream is more predictable than in the independent case.

Example 4. Four periodic tasks $\tau_1, \tau_2, \tau_3, \tau_4$ with the same period equal to 4 and respective offsets 0, 1, 2, 3. The loss probability is the same for all tasks $p = 0.2$. The other parameters are as follows: $(c_1 = 1, d_1 = 3)$, $(c_2 = 4, d_2 = 2)$, $(c_3 = 4, d_3 = 1)$, and $(c_4 = 2, d_4 = 2)$.

With Example 4, the energy consumed by (TLU_{DP}) is on average 30% lower than the energy consumed by (OA). We performed 10.000 simulations over 40 time steps: the average gain is 29.30% with the following confidence interval at 95%: [29.10, 29.50]. In figure 9, we report one example run only over the first 20 time step. Notice that between time steps 2 and 4, the blue curve (OA) is below the black curve (DP), meaning that (OA) first uses a speed below that of (DP), but thanks to the convexity property of the energy consumption, the energy consumed by (DP) is lower.

Example 5. Seven periodic tasks τ_1 to τ_7 . Task τ_1 has period 4, offset 0, and $(c_1 = 4, d_1 = 2, p_1 = 0.2)$. All the other tasks τ_2, \dots, τ_7 have period 8, loss probability $p = 0.2$, respective offsets 1, 2, 3, 5, 6, 7 (4 being for the second job of τ_1), and respective parameters $(c_2 = 1, d_2 = 2)$, $(c_3 = 4, d_3 = 1)$, $(c_4 = 2, d_4 = 2)$ $(c_5 = 4, d_5 = 3)$ $(c_6 = 2, d_6 = 1)$ $(c_7 = 1, d_7 = 2)$.

With Example 5, the energy consumed by (TLU_{DP}) is on average 30% lower than the energy consumed by (OA). We performed 10.000 simulations over 80 time steps, the average gain is 29.98% with the following confidence interval at 95%: [29.87, 30.10]. In Figure 9, we report one example run only over the first 20 time step.

The other simulation parameters for Examples 3 to 5 are $T = 20$, $\mathcal{S} = \{0, 1, 2, 3, 4, 5\}$ and $j(s) = s^3$.




Figure 9: Executed work using (OA) (light blue curve) and (TLU_{DP}) (black curve) under the four task models of Example 4.




Figure 10: Executed work using (OA) (light blue curve) and (TLU_{DP}) (black curve) under the seven task models of Example 5.

Table 1 summarizes these results.

example	gain against (OA)	95% confidence interval
Ex. 3 (2 tasks)	49.27%	[49.04,49.50]
Ex. 4 (4 tasks)	29.30%	[29.10,29.50]
Ex. 5 (7 tasks)	29.98%	[29.87,30.10]

Table 1: Comparisons between (OA) and (TLU_{DP}).

In all these cases, (TLU_{DP}) outperforms (OA) by a larger margin than with sporadic tasks. Indeed, the job sequence is more predictable, so the statistical knowledge over which (TLU_{DP}) is based is more useful here than in the sporadic case.

We notice some gain differences between the different examples. A partial explanation comes from the following considerations. If there is no error (*i.e.*, no missing jobs), the gain of (TLU_{DP}) versus (OA) can be computed exactly, as we can notice on Figure 11. It is 75% for Ex. 3 (2 tasks), 28% for Ex. 4 (4 tasks) and 52% for Ex. 5 (7 tasks). These computations explain partly the gain differences obtained in the experimental results in Table 1 with 20% losses. Furthermore, in all examples, the system has an hyper-period (respectively 2, 4, 8). With the 7 tasks, errors accumulate over 8 time units; With 2 tasks, the error rate is larger (25 % instead of 20 % for one of the two tasks). Therefore, the job sequence in one hyper-period is less predictable in both examples than with 4 tasks. This is why the gain with 7 tasks and with 2 tasks is not as well predicted by the computation made without errors, as it is with 4 tasks.

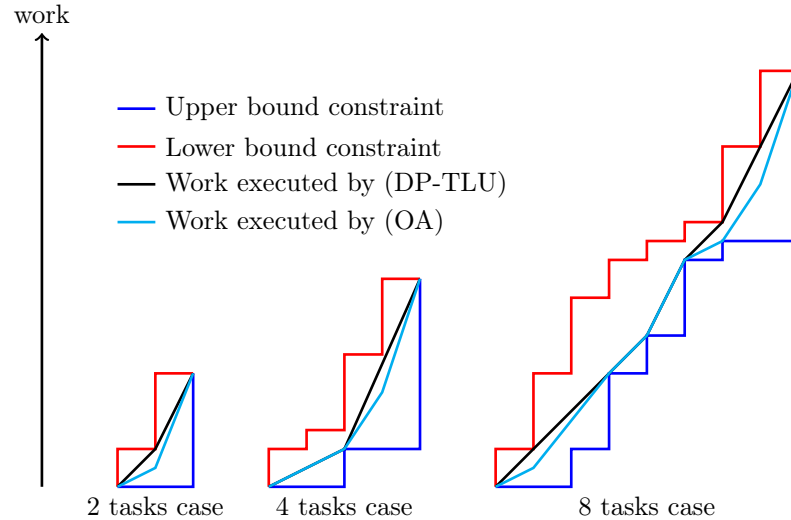


Figure 11: Executed work using (OA) (light blue curve) and (TLU_{DP}) (black curve) on Examples 4, 5 and 6 over one period in the case where no jobs are missing.

4.4 Computation Experiments, Infinite Case

In this part, we run algorithm (VI) (Algorithm 3) to compute the optimal speed to be used at each time step over an infinite horizon. We fix the stopping criterion in Algorithm 3 to

$\varepsilon = 1.0 * 10^{-5}$, so our computation of the average energy consumption is precise by at least 5 digits. We ran the program in the following cases:

Example 6. One sporadic task with jobs of size 2, deadline 3, and arrival rate p that varies from 0 to 1.

Example 7. One sporadic task with jobs of size 2, deadline 5, and arrival rate p that varies from 0 to 1.

In both examples, the available processor speeds are in $\mathcal{S} = \{0, 1, 2\}$ and the energy consumption function is $j(s) = s^2$. The only difference between Examples 6 and 7 are the deadlines.

The results of our computations are displayed in Figure 12. The three curves depict respectively the average energy consumption per time unit as a function of the arrival rate p (which varies from 0 to 1) for Examples 6 and 7, together with a theoretical lower bound.

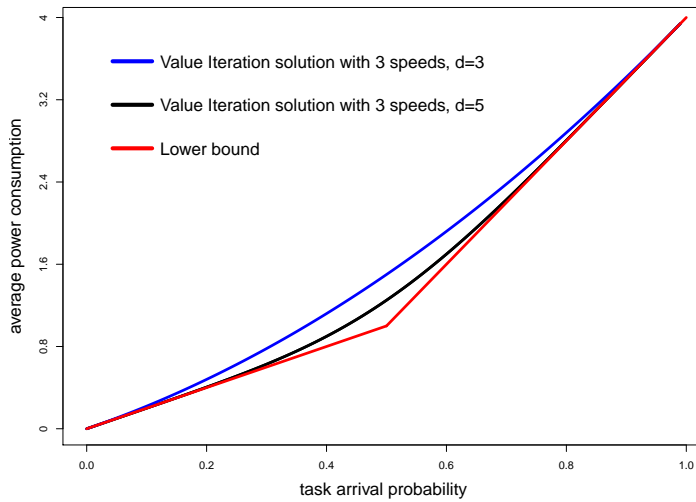


Figure 12: Average energy consumption per time unit for (TLU_{VI}) and lower bound (red curve), with deadlines equal to 3 (black curve, Example 6) and 5 (blue curve, Example 7).

The different curves in Figure 12 have the following meaning:

- The black and blue curves depict the Value Iteration solution with three processor speeds $\mathcal{S} = \{0, 1, 2\}$. These curves display g^* , as computed by Algorithm 3 (VI), as a function of p , the probability that a job of size $c = 2$ and deadline $d = 3$ (black curve) or deadline $d = 5$ (blue curve) arrives in the next instant.
- The red curve is a lower bound on g^* , oblivious of the jobs distribution and deadlines, only based on the average amount of work arriving at each time slot.

As it was expected according to Proposition 3, the higher the arrival rate, the higher the average energy consumption (both curves are increasing).

Proposition 3 also implies that larger deadlines improve the energy consumption. This is in accordance with the fact that the black curve (deadline 5) is below the blue curve (deadline 3).

What is more surprising here is how well our solution behaves when the deadline is 5. Its performance is almost indistinguishable from the theoretical lower bound (valid for all deadlines) over a large range of the rate p . More precisely, the gap between our solution with deadline equal to 5 and the lower bound is less than 10^{-3} for a rate p between 0 and 0.20, and between 0.80 and 1.

4.4.1 Lower bound

The lower bound has been obtained by solving the optimization problem without taking into account the distribution of the jobs features nor the constraint on the deadlines. Without constraints, and since the power is a convex function of the speed (here $j(s) = s^2$), the best choice is to keep the speed constant. The ideal constant speed needed to execute the jobs over a finite interval $[0, T]$ is $A(T)/T$, where $A(T)$ is the workload arrived before T . When T goes to infinity, the quantity $A(T)/T$ converges to $2p$ by the strong law of large numbers. Therefore, the optimal constant speed is $s^\infty = 2p$.

Now, if we consider the fact that only 3 processor speeds, namely $\{0, 1, 2\}$, are available, then the ideal constant speed $s^\infty = 2p$ cannot be used. In this case the computation of the lower bound is based on the following construction.

If $0 \leq p \leq \frac{1}{2}$, the the ideal constant processor speed $s^\infty = 2p$ belongs to the interval $[0, 1]$. In that case, only speeds $\{0, 1\}$ will be used. To obtain an average speed equal to $2p$, the processor must use speed 1 during a fraction $2p$ of the time and speed 0 the rest of the time. The corresponding average energy per time unit has therefore the following form:

$$g^\infty = 2p \times 1^2 + (1 - 2p) \times 0^2 = 2p \quad (16)$$

If $p \geq \frac{1}{2}$, then $s^\infty = 2p$ belongs to the interval $[1, 2]$. In that case, the processor only uses speeds 1 or 2. To get an average speed of $2p$, it must use the speed 2 during a fraction $2p - 1$ of the time and speed 1 the rest of the time. The corresponding average energy per time unit in this case is:

$$g^\infty = (2p - 1) \times 2^2 + (2 - 2p) \times 1^2 = 6p - 2 \quad (17)$$

As a result, the lower bound on g is given by: $g^\infty = 2p$ if $p \leq 1/2$ and $g^\infty = 6p - 2$ if $p \geq 1/2$. This is the red curve in Figure 12.

4.4.2 Comparison of TLU_{DP} and TLU_{VI}

A comparison between the two algorithms TLU_{DP} and TLU_{VI} over different time horizons T was done to study the impact of this parameter. The gain in energy of (TLU_{DP}) vs (TLU_{VI}) , represented in blue in Figure 13, is computed as follow:

$$\frac{\text{Energy}_{VI} - \text{Energy}_{DP}}{\text{Energy}_{DP}} \quad (18)$$

This fraction computes the relative difference between the infinite horizon case algorithm and the the finite horizon case algorithm. Besides, the cost of (OA) vs (TLU_{DP}) , also represented in Figure 13, in a dashed red curve, is defined as follows:

$$\frac{\text{Energy}_{OA} - \text{Energy}_{DP}}{\text{Energy}_{DP}} \quad (19)$$

Computations were done on Example 2 with 10000 simulations, and are summarized in Table 2.

T	10	15	20	25	30	40
(VI) vs (DP)	4.3%	1.7%	0.95%	0.62%	0.45%	0.29%
(OA) vs (DP)	4.8%	5.3%	5.3%	5.2%	5.0%	4.9%
T	100	150	200	250	1000	
(VI) vs (DP)	0.099%	0.064%	0.046%	0.031%	$6.19 \cdot 10^{-5}\%$	
(OA) vs (DP)	4.6%	4.5%	4.3%	4.3%	4.2%	

Table 2: Influence of the time horizon T on (TLU_{DP}) in comparison with (TLU_{VI}) .

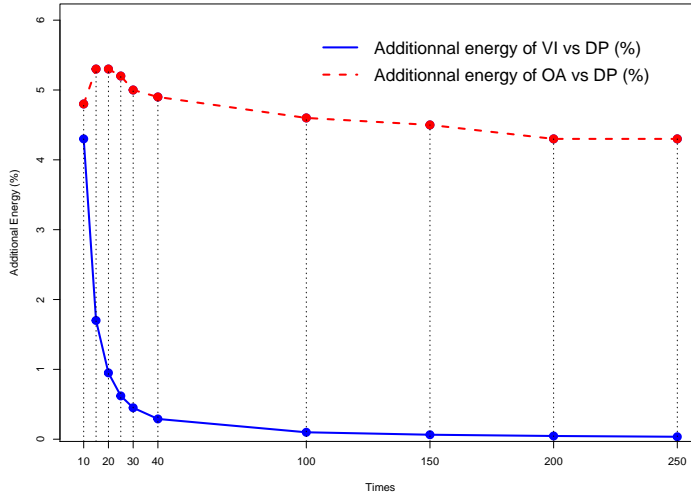


Figure 13: Influence of the time horizon T on TLU_{DP} by comparison with TLU_{VI} and (OA) on Example 2.

One can notice in Table 2 and also on the blue curve in Figure 13 that, as soon as the time horizon is greater than 20 time units, the difference of energy between TLU_{DP} and TLU_{VI} is smaller than 1%, and is negligible in comparison with the difference of energy between TLU_{DP} and (OA) . So using (VI) instead of (DP) is a good approximation even over rather short time horizons, because the results are almost as good and computing the optimal processor speeds is faster for (VI) than for (DP) . This result is rather intuitive because the only important difference between TLU_{DP} and TLU_{VI} concerns the last steps. Indeed, during these last steps, (VI) behaves as if jobs continue to arrive in the future (after $T - \Delta$), whereas (DP) considers that there is no job arrival after $T - \Delta$. (DP) can therefore adapt the chosen speeds in the last steps, whereas (VI) cannot. With these informations, the energy consumption of (DP) on the last steps is, in average, better than that of (VI) .

As for (VI) , (OA) doesnot take into account finite time horizon (except on the last Δ steps). This is why the red curve is also decreasing with the time horizon.

Remark: The red dashed curve shows that the difference of energy between (OA) and (TLU_{DP}) is almost constant, whatever the horizon time T we consider. The horizon time has no impact on the gain. Data in Table 13 confirm the results obtained in Example 2 before, because whatever the considered time horizon, the gain of (TLU_{DP}) in comparison with (OA) ranges between 4% and 5.5%.

5 Conclusion and Perspectives

In this paper, we showed how to select online speeds to execute real-time jobs while minimizing the energy consumption by taking into account statistic information on job features. This information may be collected by using past experiments or simulations as well as deductions from the structure of job sources. Our solution provides performances that are close to the optimal off-line solutions on average, and outperforms classical on-line solutions in cases where the job features have distributions with large variances.

While the goal of this study is to propose a better processor speed policy, several points are still open and will be the topic of future investigations.

The first one is about the scheduling model: In this paper we assume that jobs are executed under the *Earliest Deadline First* policy, but this is not always the case in practice. What would be the consequence of using another scheduling policy?

The second one is about the time and space complexity of our algorithms that are exponential in the deadlines of the jobs. Although our algorithms (DP) and (VI) are used off-line and can be run on powerful computers, our approach remains limited to a small range of parameters. One potential solution is to simplify the state space and to aim for a sub-optimal solution (but with proven guarantees), using approximate dynamic programming.

Finally, the statistical information gathered on the job features is crucial. When this information is not accurate or even not available, Markov Decision approaches are not possible and one should use reinforcement learning techniques (such as Q-learning [11]) to construct a statistical model of the jobs on-line and select speeds accordingly, which will converge to optimal speeds over time.

Appendix

In the first three sections of this appendix, we develop several extensions to make our model more realistic. To achieve this, we assume that the processor can change speeds at any time. This assumption is not very strong because there is no technical reason to change processor speed only at task arrival. These generalizations are the following:

1. Convexification of the power consumption function: Any non-convex power function can be advantageously replaced by its convex hull.
2. Taking into account the time penalty required to change the processor speed: This time penalty can be replaced by an additional cost on the energy consumption.
3. Taking into account the context switching time between one task to another: This switching time can also be included in the cost function.

Finally, we present in the last section, appendix D, a worst-case scenario where (OA) fails to schedule the given jobs.

A Convexification of the Power Consumption Function

Our general approach does not make any assumption on the power function $j(\cdot)$. Our algorithms (DP) and (VI) will compute the optimal speed selection for any function $j(\cdot)$. However structural properties (comparison with (OA), monotonicity) need convexity assumptions. It is therefore desirable to convexify the power function.

Let us consider a processor, whose speeds belong to the set $\mathcal{S} = \{s_0, s_1, s_2, s_{max}\}$ and the power function of the processor $j(\cdot) : \mathcal{S} \rightarrow \mathbb{R}$, represented by the black curve in Figure 14. The function $g(\cdot)$ is the convex hull of $j(\cdot)$ (the red curve in Figure 14).

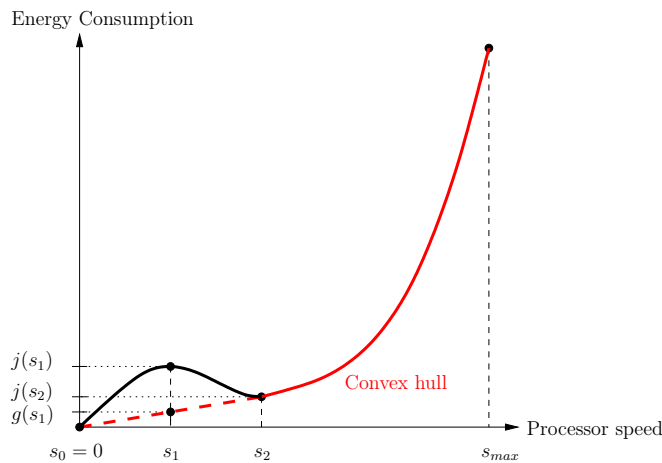


Figure 14: Convexification of the power consumption function.

If the energy function is not convex, some speeds are not relevant, because using these speeds is more expensive in term of energy than using a combination of more relevant speeds. Figure 14

depicts a non-convex power function in black, and its convex hull in red. It is better to choose speeds s_0 and s_2 (actually a combination of s_0 and s_2), rather than speed s_1 . In fact, all points of the power function curve, that are above the convex hull, should never be taken into consideration. It is always better to only select the speeds whose power consumption belongs to the convex hull of the power function. Indeed if $g(s_1) < j(s_1)$ (see Figure 14), instead of selecting speed s_1 during any time interval $[t, t + 1)$, the processor can select speed s_2 during a fraction of time α_2 and then speed s_0 during a fraction of time α_0 , such that $\alpha_0 s_0 + \alpha_2 s_2 = s_1$. The total quantity of work executed during the time interval $[t, t + 1)$ will be the same as with s_1 , but the energy consumption will be $g(s_1) = \alpha_0 j(s_0) + \alpha_2 j(s_2) < j(s_1)$ thanks to the convexity of function j . This approach uses the V_{dd} hopping technique.

So we can always consider that the power function is convex. This is very useful in practice. Indeed, the actual power consumption of a CMOS circuit working at speed s is non-convex function of the form $j(s) = Cs^\alpha + L(s)$, where the constant C depends on the activation of the logical gates, α is between 2 and 3, and $L(s)$ is the leakage, with $L(0) = 0$ and $L(s) \neq 0$ if $s > 0$. In this case, convexification removes the small values of s from the set of useful speeds.

Remark: This idea to replace one speed with a combination of two speeds (*i.e.*, V_{dd} hopping) can also be used to simulate any speed between 0 and s_{max} . Indeed if a speed doesn't exist in the set \mathcal{S} , a solution is to simulate it by combining two neighboring speeds. For our problem, this technique let the processor have more speeds to choose from so that the optimal speed computed by the (DP) algorithm will use less energy with V_{dd} hopping than without it.

B Taking into Account the Cost of Speed Changes

In our initial model, we have assumed that the time needed by the processor to change speeds is null. However, in all synchronous CMOS circuits, changing speeds takes some time; this delay is needed to resynchronize the logical gates. We denote by δ the time needed by the processor to change speeds. During this time, the circuit logical functions are altered so no computation can take place.

With time lags for speed changes, the executed work by the processor has *two* slope changes, at times t_1 and t_2 , with $t_2 - t_1 = \delta$ (see the red solid line in Figure 15). The problem is that, since in general $\delta \notin \mathbb{N}$, we cannot have both $t_1 \in \mathbb{N}$ and $t_2 \in \mathbb{N}$. As a consequence, one of the states w_{t_1} or w_{t_2} will not be integer valued. This is forbidden by our MDP approach.

The solution we propose involves replacing the actual behavior of the processor (represented by the red solid line in Figure 15) by a simulated behavior, equivalent in terms of the amount of work performed (represented by the blue dashed line in Figure 15). This simulated behavior exhibits a *single* speed change and is such that the total amount of work done by the processor is identical in both cases at all integer times (*i.e.* at $t_3 - 1$, t_3 , and $t_3 + 1$ in Figure 15). The advantage is that, since there is only one state change, it can be chosen to occur at an integer time. In other words, we *choose* t_1 such that $t_3 \in \mathbb{N}$.

One issue remains, due to the fact that the consumed energy will not be identical with the real behavior and the simulated behavior; it will actually be higher for the real behavior for convexity reason. This additional energy cost of the real processor behavior must therefore be added to the energy cost of the equivalent simulated behavior.

Finally, in order to trigger the speed change at time t_1 , the processor needs to be “clairvoyant”,

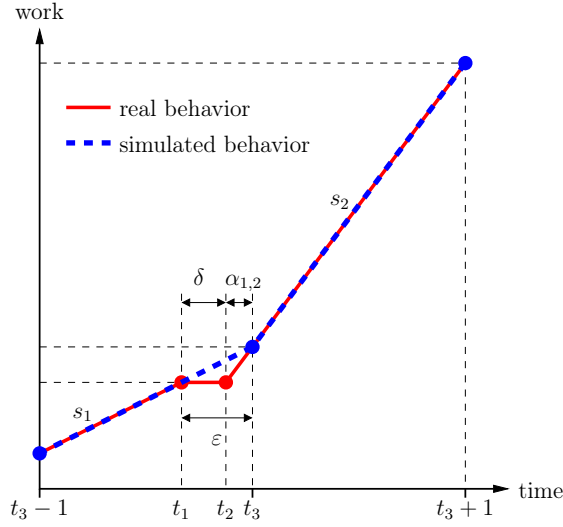


Figure 15: Transformation of the time lag into an energy additional cost (case $s_1 < s_2$). The red line represents the real behavior of the processor, which includes a δ time lag. The blue dashed line represents an equivalent behavior in terms of executed work, with no time lag.

i.e., it needs to know in advance (before t_1) the characteristics of the job arriving at time t_3 . This will allow the processor to compute the new speed s_2 and the length ε of the required interval to make sure that the work done by the processor at t_3 in the two cases (real and simulated) is identical.

The value of ε , $\alpha_{1,2}$, and the additional energy cost $h(s_1, s_2)$ of this speed change are computed as follows. In the case $s_2 > s_1$ (as in Figure 15), we have:

$$\begin{aligned} s_1 \varepsilon = s_2 \alpha_{1,2} = s_2 (\varepsilon - \delta) &\Leftrightarrow (s_2 - s_1) \varepsilon = \delta s_2 \\ &\Leftrightarrow \varepsilon = \delta + \alpha_{1,2} = \frac{\delta s_2}{s_2 - s_1} \end{aligned} \quad (20)$$

We further assume that, during the time lag δ , the energy is consumed by the processor as if the speed were s_1 . The additional energy cost incurred in the real behavior (red curve) compared with the simulated behavior (blue curve), denoted $h(s_1, s_2)$, is therefore:

$$h(s_1, s_2) = \alpha_{1,2}(j(s_2) - j(s_1))$$

Using the value of $\alpha_{1,2}$ from Eq. (20) yields:

$$h(s_1, s_2) = \delta s_1 \left(\frac{j(s_2) - j(s_1)}{s_2 - s_1} \right) \quad (21)$$

When $s_1 > s_2$, the additional cost becomes:

$$h(s_1, s_2) = \delta s_2 \left(\frac{j(s_1) - j(s_2)}{s_1 - s_2} \right) \quad (22)$$

This additional energy due to speed changes will be taken in consideration in our model in the cost function, by modifying the state space \mathcal{W} and adding the current speed to the state

at $t - 1$. Therefore the new state will have the following form:

$$(w_t, s_{t-1})$$

And the new main step of the (DP) Algorithm 1 becomes:

$$J_{t-1}^*(w, s) \leftarrow \min_{s' \in \mathcal{A}(w)} \left(j(s) + h(s, s') + \sum_{w' \in \mathcal{W}} P_s(w, w') J_t^*(w', s') \right) \quad (23)$$

with $h(s, s') = 0$ when $s = s'$ and otherwise given by Eq. (22) if $s' < s$ and Eq. (21) if $s < s'$. The rest of the analysis is unchanged.

C Taking into Account the Cost of Context Switches

We now consider that switching from the execution of one job to another one takes some non-negligible time, denoted γ , whether this corresponds to the beginning of a new job or the resuming of a job after preemption. This is essentially the time needed to upload/download the content of the execution stack. During this context switch, no useful work is being executed. In Figure 16, we consider 2 jobs with the following characteristics: $\mathcal{J}_1(r_1 = 0, c_1 = 3, d_1 = 7, 5)$ and $\mathcal{J}_2(r_2 = 3, c_2 = 2, d_2 = 3)$.




Figure 16: Impact of a context switch on the execution time.

We notice that, each time the processor starts to execute a new job, it takes some time for switching from one job context to another (context switch). This switching time is denoted by γ (see the barred area in Figure 16). We assume that each context switch takes γ , whether this corresponds to the beginning of a new job or the resume of a job after preemption.

Figure 17 illustrates the fact that, during one time step, several context switches can occur. In this example, during the time interval $[t, t + 1)$, the processor completes two jobs, \mathcal{J}_1 and \mathcal{J}_2 , and starts the execution of \mathcal{J}_3 (see the red curve). This involves two context switches, both of which occur during one time unit. This leads to a total delay of 2γ . As in Appendix B, we transform this time delay into an energy cost: In one time unit, the evolution of the executed work under speed s , with K context switches (see the green curve), is the same as the evolution of the executed work under speed $s_1 = s(1 - K\gamma)$, with no switching delay (red curve).

The state space of the system must be modified to be able to compute K , the number of context switches in each time interval. We must keep in memory the sizes of the jobs instead of only the total remaining work. Indeed, with the current state space \mathcal{W} , we do not know the number of actual different jobs composing a given amount $w(i), i \in \{0, \dots, \Delta\}$, so we cannot know the number of context switches. We denote by $\overline{\mathcal{W}}$ the new state space, where the modified state at time t , $\overline{w}(t) \in \overline{\mathcal{W}}$, has the following form:

$$\overline{w}(t) = \left[(\rho_1^1, \dots, \rho_1^{k_1}), \dots, \underbrace{(\rho_i^1, \dots, \rho_i^{k_i})}_{\substack{\text{remaining work quantity of jobs} \\ \text{with absolute deadline } t+i}}, \dots, (\rho_\Delta^1, \dots, \rho_\Delta^{k_\Delta}) \right] \quad (24)$$




Figure 17: Compensation of the impact of the context switches on the executed work by using a higher speed. The released jobs are $\mathcal{J}_i = (r_i, c_i, d_i)$, for $1 \leq i \leq 5$, with $r_5 + d_5 = \Delta$, $r_4 + d_4 = r_3 + d_3 = r_2 + d_2 = \Delta - 1$, and $r_1 + d_1 = \Delta - 2$.

where k_i is the number of jobs whose relative deadline is i time units away (hence their absolute deadline is $t + i$), and ρ_i^j is the work quantity of the j -st such such job.

Let (r_n, c_n, d_n) be a new job that arrives at time $t = r_n$. If the processor speed at time $t - 1$ is s_{t-1} , then at time t the next state $\bar{w}(t)$ becomes:

$$\begin{aligned} \bar{w}(t) = & \left[\left((\rho_2^1(t-1) - \ell(1,1))^+, \dots, (\rho_2^{k_1}(t-1) - \ell(1, k-1))^+ \right), \dots, \right. \\ & \left((\rho_{d_n}^1(t-1) - \ell(d_n,1))^+, \dots, (\rho_{d_n}^{k_{d_n}}(t-1) - \ell(d_n, k_{d_n}))^+, c_n \right), \dots, \\ & \left((\rho_\Delta^1(t-1) - \ell(\Delta,1))^+, \dots, (\rho_\Delta^{k_\Delta}(t-1) - \ell(\Delta, k_\Delta))^+ \right), \\ & \left. \left((\rho_\Delta^1(t-1) - \ell(\Delta,1))^+, \dots, (\rho_\Delta^{k_\Delta}(t-1) - \ell(\Delta, k_\Delta))^+ \right) \right] \end{aligned}$$

where $\ell(d, k) = \left(s_{t-1} - \sum_{i=1}^{d-1} \sum_{j=1}^{k_i} \rho_i^j - \sum_{j=1}^k \rho_d^j \right)^+$.

The idea of the state change is such as we set all the ρ_i^j to 0 when $s < \sum_{i,j} \rho_i^j$, except for the job, that will be executed, possibly only partially, where some residual work remains.

We further assume that the energy consumption during a context switch is the same as when some work is executed. The new main step of the (DP) Algorithm 1 now has the following form:

$$J_{t-1}^*(\bar{w}) \leftarrow \min_{s \in \mathcal{A}(\bar{w})} \left(j \left(\frac{s}{1 - K_s \gamma} \right) + \sum_{\bar{w}' \in \bar{W}} P_s(\bar{w}, \bar{w}') J_t^*(\bar{w}') \right) \quad (25)$$

Note that the speed $\frac{s}{1 - K_s \gamma}$ may not be directly available, but thanks to the remark made in Appendix A, we can easily simulate this speed with the neighboring available speeds.

We denote by K_s the number of job executed if we use the speed s :

$$K_s = \sum_{i=1}^{\alpha-1} k_\alpha + \alpha$$

where α is such that $\min_{\alpha, k_\beta} \left(s < \sum_{i,j}^{\alpha, k_\beta} \rho_i^j \right)$. The rest of the analysis is unchanged.

As a final remark, if processor sharing is enabled, which is often the case nowadays, the state space can be simplified: One only needs to keep in memory the number of jobs, which are executed in a specific $w(i)$, instead of all their sizes, so a state \bar{w} is of the following form:

$$\bar{w} = [(w(1), k_1), (w(2), k_2), \dots, (w(\Delta), k_\Delta)] \quad (26)$$

with k_i the number of jobs, whose have the same relative deadline of value i . We have also to modify the change state function accordingly. Let (r_n, c_n, d_n) be a job that arrives at time $t = r_n$. If the processor speed at time $t - 1$ is s_{t-1} , then at time t the next \bar{w} becomes:

$$w_t(\cdot) = \mathbb{T}[(w_{t-1}(\cdot) - s_{t-1})^+] + c_n H_{d_n}(\cdot)$$

and the number k that corresponds to the deadline d_n of the new task is increased by one:

$$k_{d_n} = k_{d_n} + 1$$

In the processor sharing case, the penalty cost due to the time of processor sharing is proportional to the number of executed jobs. If we denote by μ the total switching time between two jobs over one time unit, the Bellman equation is the same as Eq. (25) but replacing K_s by:

$$K'_s = \sum_i \frac{k_i}{s} (s - \sum_{j=1}^{i-1} w_j) + k_1 \frac{w_1}{s}$$

D Schedulability of the (OA) Policy

We present in this appendix a worst-case scenario where the (OA) policy fails to schedule the sequence of jobs it is given, and we show that (OA) is schedulable under a condition on the maximal speed s_{max} of the processor.

Formally, let s_{max} be the maximal processor speed, then we have:

$$\text{(OA) is schedulable} \iff \forall t, s_{max} \geq w_t^{(OA)}(1) \iff s_{max} \geq \sup_t (w_t^{(OA)}(1))$$

Let us first compute an upper bound \bar{w}_1 on the state $\sup_t (w_t^{(OA)}(1))$. By looking backward in time, at time $t > \Delta$, $w_t^{(OA)}(1)$ must satisfy:

$$w_t^{(OA)}(1) = w_{t-\Delta+1}(\Delta) + \sum_{i=t-\Delta+2}^t c_i \mathbf{1}_{\{d_i \leq t+1\}} - \sum_{i=t-\Delta+1}^{t-1} s_i^{(OA)} \quad (27)$$

However, we know by definition that $w_{t-\Delta+1}(\Delta) \leq C\Delta$ if no deadlines are missed before t , and that $\forall i, c_i \leq C$. It follows that $w_t^{(OA)}(1)$ satisfies the following inequality:

$$w_t^{(OA)}(1) \leq C\Delta + C(\Delta - 1) - \sum_{i=t-\Delta+1}^{t-1} s_i^{(OA)} \quad (28)$$

By definition of (OA), $s_i^{(OA)}$ decreases with the deadlines of the pending jobs, so $s_i^{(OA)}$ is lower bounded when the deadlines of the jobs are as late as possible (but not after $t + 1$). This is

the case when all the jobs arriving between $t - \Delta + 1$ and t have their absolute deadline at $t + 1$. Moreover, replacing the state $w_{t-\Delta+1}$ by a lower bound ($\forall i \in [t - \Delta + 1, t]$, $w'_{t-\Delta+1}(i) = 0$ and $w'_{t-\Delta+1}(t + 1) = C\Delta$) also decreases the speeds to their lower bounds, denoted s_i^* . By doing so, we obtain an upper bound \bar{w}_1 on $w_t^{(OA)}(1)$, given in Eq (29):

$$\bar{w}_1 = C\Delta + C(\Delta - 1) - \sum_{i=t-\Delta+1}^{t-1} s_i^* \quad (29)$$

This construction of an upper bound on $w_t^{(OA)}(1)$ is illustrated in the case where $C = 2$ and $\Delta = 6$ in Figure 18.




Figure 18: Construction of the bound \bar{w}_1 of $w_t^{(OA)}(1)$ for $C = 2$ and $\Delta = 6$. The black bold curve is the lower bound on the state $w_{t-\Delta+1}$. The blue bold curve is the upper bound \bar{w}_1 on the state. The arrows represent the work executed by the processor at each time slot, at speeds s_i^* .

The configuration displayed in Figure 18 is not a reachable state under (OA). However, one can construct a state that is asymptotically reachable and that makes (OA) use the same speeds s_i^* . This is illustrated in Figure 19. We consider the following case: First, jobs of type (C, Δ) arrive at each time slot during n time slots. In this case, the state under (OA) approaches the black staircase state as n grows. Then, jobs of type $(C, \Delta - i + n)$ arrive $\forall i \in [n + 1, n + \Delta - 1]$ time slots (jobs arrival are represented alternatively in blue and red in Figure 19). In that case, let us compute $w_{n+\Delta-1}(1)$. The remaining work of the black state in Figure 19 is the same as in the worst case in Figure 18, with a value of $C\Delta$. Furthermore, the speeds used in the state of Figure 19 are the same as in the upper bound displayed in Figure 18. So $w_{n+\Delta-1}(1)$ approaches \bar{w}_1 as n goes to infinity. This shows that \bar{w}_1 is tight: $\sup_t(w_t^{(OA)}(1)) = \bar{w}_1$.

Let us compute \bar{w}_1 using the above construction, displayed in Figure 19 in the case $C = 1$ and $\Delta = 5$. The speed under (OA) at time 0 is $s^{(OA)}(0) = C$. At time 1, it is $s^{(OA)}(1) = \frac{C\Delta}{\Delta-1}$. More generally, for $1 \leq k < \Delta$, we have:

$$(\Delta - k)s^{(OA)}(k) = (\Delta + k - 1)C - \sum_{j=1}^{k-1} s^{(OA)}(j) \quad (30)$$

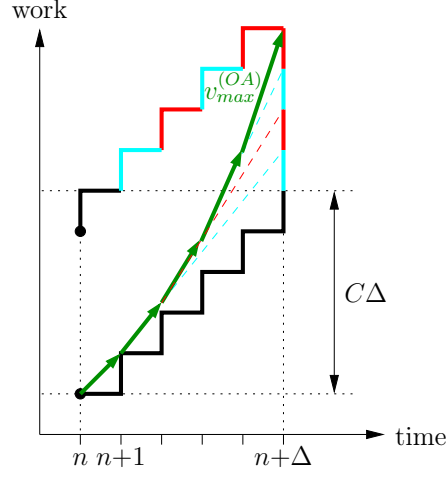


Figure 19: Asymptotic worst case state, for $C = 1$ and $\Delta = 5$. The staircase black curve represents the state reached asymptotically while the colored parts (blue and red segments represent one job at each time slot with same WCET C and different deadline) lead to the maximal $w_{n+\Delta-1}(1)$. The green arrows represent quantities of work executed by the processor under (OA).

$$(\Delta - k - 1)s^{(OA)}(k + 1) = (\Delta + k)C - \sum_{j=1}^k s^{(OA)}(j) \quad (31)$$

By subtracting Eq. 30 and 31, we obtain:

$$\begin{aligned} (\Delta - k - 1) \left(s^{(OA)}(k + 1) - s^{(OA)}(k) \right) &= C \\ \Leftrightarrow s^{(OA)}(k + 1) &= s^{(OA)}(k) + \frac{C}{(\Delta - k - 1)} \end{aligned}$$

By using a telescoping sum, we obtain $\forall k \geq 1$:

$$\begin{aligned} s^{(OA)}(k + 1) &= s^{(OA)}(1) + C \sum_{j=2}^{k+1} \frac{1}{\Delta - i} \\ &= s^{(OA)}(1) + C(h_{\Delta-2} - h_{\Delta-k-2}) \end{aligned}$$

Therefore $\bar{w}_1 = s^{(OA)}(\Delta - 1)$ has the following value:

$$\begin{aligned}
 \bar{w}_1 &= 2C\Delta - C - \sum_{j=0}^{\Delta-2} s^{(OA)}(j+1) \\
 &= 2C\Delta - C - Cs^{(OA)}(1) - \sum_{j=1}^{\Delta-2} (s^{(OA)}(1) + C(h_{\Delta-2} - h_{\Delta-j-1})) \\
 &= C(\Delta - 1) - C \sum_{j=1}^{\Delta-2} (h_{\Delta-2} - h_{\Delta-j-1}) \\
 &= C \left(\frac{\Delta}{\Delta-1} + h_{\Delta-2} \right),
 \end{aligned}$$

using $s^{(OA)}(1) = C \frac{\Delta}{\Delta-1}$ and h_n , the n -th harmonic number: $h_n = \sum_{i=1}^n 1/i$ and $h_0 = 0$.

To conclude, the (OA) policy is schedulable if and only if

$$s_{max} \geq C \left(\frac{\Delta}{\Delta-1} + h_{\Delta-2} \right) \approx C \log \Delta, \quad (32)$$

when Δ is large. Therefore (OA) is not always schedulable.

E Size of the State Space

This appendix is dedicated to the enumeration of the total number of states.

Let $w(\cdot)$ be a valid state of the system, at any time t . Since all parameters are integer numbers and the maximum deadline of a task is Δ , the maximal look-ahead at any time is Δ , hence $w(\cdot)$ is characterized by its first Δ integer values (that are non-decreasing by definition): $w(1) \leq \dots \leq w(\Delta)$.

Let us define the step sizes of w , starting from the end: $x_1 = w(\Delta) - w(\Delta - 1)$, and more generally, $x_j = w(\Delta - j + 1) - w(\Delta - j)$, for all $j = 1, \dots, \Delta$, the released work being of maximal size C at any time, $x_1 \leq C$ because x_1 must be bounded by the amount of work that was released at step t . Similarly, $x_1 + x_2$ must be bounded by the amount of work that was released at steps t and $t - 1$, namely $2C$, and so on and so forth up to $x_1 + x_2 + \dots + x_\Delta \leq \Delta C$. This is the only condition for a function w to be a possible state when deadlines and sizes are arbitrary integers bounded by Δ and C respectively.

Therefore $(x_1, x_2, \dots, x_\Delta)$ satisfy the following conditions:

$$\left\{ \begin{array}{l} x_1 \leq C \\ x_1 + x_2 \leq 2C \\ x_1 + x_2 + x_3 \leq 3C \\ \vdots \\ x_1 + x_2 + \dots + x_\Delta \leq \Delta C \end{array} \right.$$

By defining the partial sums $y_j = x_1 + \dots + x_j$, the number of states satisfies:

$$Q(C, \Delta) = \sum_{y_1=0}^C \sum_{y_2=y_1}^{2C} \sum_{y_3=y_2}^{3C} \dots \sum_{y_\Delta=y_{\Delta-1}}^{\Delta C} 1.$$

This multiple sum can be seen as a generalized Catalan number. Indeed, notice that a state characterized by its steps (x_1, \dots, x_Δ) is in bijection with a path on the integer grid from $(0, 0)$ to $(\Delta + 1, C(\Delta + 1))$, which remains *below* the diagonal of slope C (see Fig. 20).




Figure 20: A valid state seen as a path below the diagonal $(0, 0) - (6, 12)$, for $C = 2$ and $\Delta = 5$.

Counting the number of such paths has been done in [12] and corresponds to the generalized Catalan numbers, $C_n^k = \frac{1}{nk+1} \binom{nk+1}{n}$, with $k = C + 1$ and $n = \Delta + 1$. Therefore, $Q(C, \Delta) = \frac{1}{1+C(\Delta+1)} \binom{(C+1)(\Delta+1)}{\Delta+1}$.

Using the Stirling formula, we finally get $Q(C, \Delta) \approx \frac{e}{\sqrt{2\pi}} \frac{1}{(\Delta+1)^{3/2}} (eC)^\Delta$.

References

- [1] F. Yao, A. Demers, and S. Shenker, "A scheduling model for reduced CPU energy," in *Proceedings of IEEE Annual Foundations of Computer Science*, pp. 374–382, 1995.
- [2] N. Bansal, T. Kimbrel, and K. Pruhs, "Speed scaling to manage energy and temperature," *Journal of the ACM*, vol. 54, no. 1, 2007.
- [3] F. Gruian, "On energy reduction in hard real-time systems containing tasks with stochastic execution times," in *IEEE Workshop on Power Management for Real-Time and Embedded Systems*, pp. 11–16, 2001.
- [4] J. Lorch and A. Smith, "Improving dynamic voltage scaling algorithms with pace," in *ACM SIGMETRICS 2001 Conference*, pp. 50–61, 2001.
- [5] P. Pillai and K. G. Shin, "Real-time dynamic voltage scaling for low-power embedded operating systems," *SIGOPS Oper. Syst. Rev.*, vol. 35, pp. 89–102, Oct. 2001.
- [6] M. L. Puterman, *Markov Decision Process : Discrete Stochastic Dynamic Programming*. Wiley, wiley series in probability and statistics ed., February 2005.
- [7] B. Gaujal, N. Navet, and C. Walsh, "Shortest Path Algorithms for Real-Time Scheduling of FIFO tasks with Minimal Energy Use," *ACM Transactions on Embedded Computing Systems (TECS)*, vol. 4, no. 4, 2005.
- [8] W. Horn, "Some simple scheduling algorithms," *Naval Research Logistics*, vol. 21, no. 1, pp. 177–185, 1974.
- [9] A. Marshall and I. Olkin, *Inequalities: Theory of Majorization and Its Applications*, vol. 143 of *Mathematics in Science and Engineering*. Academic Press, 1979.
- [10] A. Müller and D. Stoyan, *Comparison Methods for Stochastic Models and Risks*. No. ISBN: 978-0-471-49446-1 in Wiley Series in Probability and Statistics, Wiley, 2002.
- [11] D. Bertsekas and J. Tsitsiklis, *Neuro-dynamic programming*. Belmont, Mass., USA: Athena Scientific, 1996.
- [12] P. Hilton and J. Pedersen, "Catalan numbers, their generalization, and their uses," *The Mathematical Intelligencer*, vol. 13, pp. 64–75, Mar 1991.
- [13] H. Yun and J. Kim, "On energy-optimal voltage scheduling for fixed priority hard real-time systems," *ACM Trans. Embed. Comput. Syst.*, vol. 2, no. 3, pp. 393–430, 2003.
- [14] M. Li and F. F. Yao, "An efficient algorithm for computing optimal discrete voltage schedules," *SIAM J. Comput.*, vol. 35, pp. 658–671, 2005.



**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399