

Packing-Inspired Algorithms for Periodic Scheduling Problems with Harmonic Periods

Josef Grus^{1,4}, Claire Hanen^{2,3} and Zdeněk Hanzálek⁴

¹*DCE, FEE, Czech Technical University in Prague, Czech Republic*

²*Sorbonne Université, CNRS, LIP6, F-75005 Paris, France*

³*UPL, Université Paris Nanterre, F-92000 Nanterre, France*

⁴*IID, CIIRC, Czech Technical University in Prague, Czech Republic*

Keywords: Periodic Scheduling, Harmonic Periods, Height-Divisible 2D Packing, First Fit Heuristic.

Abstract: We tackle the problem of non-preemptive periodic scheduling with a harmonic set of periods. Problems of this kind arise within domains of periodic manufacturing and maintenance and also during the design of industrial, automotive, and avionics communication protocols, where an efficient static schedule of messages is crucial for the performance of a time-triggered network. We consider the decision variant of the periodic scheduling problem on a single highly-utilized machine. We first prove a bijection between periodic scheduling and a particular 2D packing of rectangles (we name it height-divisible 2D packing). We formulate the problem using Constraint Programming and compare it with equivalent state-of-the-art Integer Linear Programming formulation, showing the former's superiority on difficult instances. Furthermore, we develop a packing-inspired first fit heuristic, which we compare with methods described in the literature. We justify our proposed methods on problem instances inspired by the communication of messages on one channel.

1 INTRODUCTION AND RELATED WORK

Periodic Scheduling Problems (*PSPs*) often arise when dealing with manufacturing and maintenance schedules. Another important application area is the configuration of communications protocols in the industrial, automotive, and avionics domains. The configuration is based upon an offline generated schedule specifying start times of so-called time-triggered messages, typically used for safety-relevant functions. It is important to have an efficient solution because the demand for communication bandwidth has been increasing significantly when messages from a camera, lidar, or radar become a part of safety-relevant functions in modern self-driving cars (Hanzálek et al., 2023).

In this paper, we study the decision problem of scheduling non-preemptive jobs on a single machine, where the periods of the jobs are harmonic; for any two distinct periods, it holds that the longer period

is divisible by the shorter one. We assume that the schedule has to be strictly periodic: the interval between two occurrences of the same job is equal to the job period.

PSPs with a harmonic set of periods and processing times are easier to solve than the more general case (Korst et al., 1996). However, the case with harmonic periods and non-harmonic processing times was proven to be strongly NP-complete (Cai and Kong, 1996). Schedules with harmonic periods were applied in many industrial settings. For example, in FlexRay standard (a communication protocol used in higher-class cars (Lukasiewicz et al., 2009; Dvoračák and Hanzálek, 2014)), we aim to find a periodic schedule for a so-called static segment where the jobs represent the messages and the machine represents the physical communication channel. Furthermore, Boeing's avionics-related problem instances were solved using methods for harmonic periods in (Eisenbrand et al., 2010b) as well.

The general *PSPs* are usually modeled using time-indexed formulations or relative order formulations, implemented using Integer Linear Programming (ILP) in (Eisenbrand et al., 2010b) or Satisfiability

^a <https://orcid.org/0000-0002-1136-370X>

^b <https://orcid.org/0000-0003-2482-5042>

^c <https://orcid.org/0000-0002-8135-1296>

ity Modulo Theory in (Minaeva and Hanzalek, 2021). In (Hladik et al., 2020), *PSP* with harmonic periods, precedence constraints, and dedicated machines was tackled using a local search first-fit heuristic.

Other methods exploit characteristics of the *PSP* with harmonic periods. In this case, the feasible schedule can be visualized using the bin tree representation, shown in (Eisenbrand et al., 2010a), or using the 2D packing approach described in this paper. In (Eisenbrand et al., 2010a), authors use the compressed bin trees to develop a first fit heuristic to solve a machine minimization problem. Later in (Eisenbrand et al., 2010b), the authors presented an alternative well-performing ILP formulation. Similar first fit heuristic and ILP formulation, used for the FlexRay bus optimization, were shown in (Lukasiewicz et al., 2009). In (Dvorak and Hanzalek, 2014), a first-fit-based heuristic was utilized to create a FlexRay schedule that could be embedded into multiple models of vehicles. In (Marouf and Sorel, 2011), authors performed a schedulability analysis of both harmonic and non-harmonic *PSPs*.

The crucial correspondence between the *PSP* with harmonic periods and the height-divisible 2D packing problem was outlined in our previous paper (Hanan and Hanzalek, 2020). This allows us to get inspiration from the algorithms developed for the packing and cutting problems. In (Fleszar and Hindi, 2002), authors presented a bin-focused heuristic. A similar method, generalized for a variable-sized bin packing, was presented in (Haouari and Serairi, 2009).

In the 2D case, the guillotine cutting and k-stage packing problems can be considered a generalization of the height-divisible 2D packing later defined in this paper. On the other hand, height-divisible 2D packing is more general than 1D packing. In (Ntene and van Vuuren, 2009), authors provide the overview of guillotine heuristics for strip packing and present their own size alternating stack heuristic, further improved in (Ortmann et al., 2010). Finally, we can specialize the ideas of general packing heuristics, such as the work of (Crainic et al., 2008) for 3D bin-packing, into our height-divisible problem environment.

With respect to the relevant research, the main contributions of our paper are the following:

- We give a new proof of the equivalence of the *PSP* with harmonic periods on a single machine and the height-divisible 2D packing problem, and we show that any solution obtained for the 2D packing problem can be transformed into a canonical form.
- We present a Constraint Programming (CP) formulation of the problem and compare it with the state-of-the-art ILP model.

- We present a packing-inspired heuristic, which enables us to solve more highly utilized problem instances in comparison with baseline methods.

2 PROBLEM DEFINITION

We consider a set of n independent jobs, J_1, \dots, J_n to be performed repeatedly on a single machine. Each job J_i is characterized by its integer processing time p_i and its integer period T_{α_i} . Periods belong to a harmonic set T_0, \dots, T_{r-1} of size r , so that $\forall k \in \{1, \dots, r-1\}, T_k = b_k T_{k-1}$ where b_k is an integer. We define the least period $w = T_0$, and we define the hyper-period \hat{H} to be the least common multiplier of periods. Due to the harmonic nature of the period set, \hat{H} is equal to the longest period T_{r-1} .

Any periodic schedule defines for each job J_i a start time $first_i \leq T_{\alpha_i}$ of its first occurrence. The j -th occurrence of J_i will start at time $first_i + (j-1)T_{\alpha_i}$, where $j \geq 1$. We consider the existence of a feasible schedule here. Let o be the starting time of the first job J_f with the least period w . We choose o as the origin. So we define s_i such that:

$$s_i = \begin{cases} first_i - o & \text{if } first_i \geq o \\ first_i + T_{\alpha_i} - o & \text{otherwise} \end{cases}$$

We call time interval $[o, o + \hat{H})$ the "observation interval". Notice that for any integer $k \geq 1$, the schedule of time interval $[o + (k-1)\hat{H}, o + k\hat{H})$ is the same as the schedule of the observation interval shifted by $(k-1)\hat{H}$ time units. In the rest of the paper, we consider without loss of generality that the first occurrence of a job J_i starts at s_i in the observation interval.

Let us now decompose s_i with respect to the least period w : $s_i = u_i + v_i \cdot w$, with $u_i < w$, and integer v_i . Also, $u_i + p_i \leq w$, since an occurrence of job J_i needs to fit between two occurrences of J_f at o and $o + w$.

A feasible schedule has to be collision-free. We call collision between two jobs J_i, J_j a situation when two occurrences of these jobs are processed simultaneously. An example of the observation interval of the collision-free schedule is described in Section 3.2 and is shown in Figure 1a.

We denote by $B_k = \prod_{i=1}^k b_i$, $B_0 = 1$, so that:

$$T_k = B_k w. \quad (1)$$

We can now state the necessary and sufficient condition for a collision to occur in a periodic schedule.

Lemma 1. *A periodic schedule induces a collision between two jobs J_i, J_j with periods $T_{\alpha_i}, T_{\alpha_j}$ such that $\alpha_i \leq \alpha_j$ if and only if the following conditions hold:*

$$u_i < u_j + p_j \quad \text{and} \quad u_j < u_i + p_i \quad (2)$$

$$\exists \kappa \in \mathbb{N} \cap [0, \frac{B_{\alpha_j}}{B_{\alpha_i}}), \quad v_j = v_i + \kappa B_{\alpha_i} \quad (3)$$

Proof. Suppose that the occurrences k, k' of J_i, J_j collide. For this to happen, the two following conditions must hold:

$$s_i + kT_{\alpha_i} + p_i > s_j + k'T_{\alpha_j} \\ \text{and } s_j + k'T_{\alpha_j} + p_j > s_i + kT_{\alpha_i}. \quad (4)$$

We decompose $s_i = u_i + v_i \cdot w$ and $s_j = u_j + v_j \cdot w$:

$$u_i + (v_i + kB_{\alpha_i})w + p_i > u_j + (v_j + k'B_{\alpha_j})w \\ \text{and} \\ u_i + (v_i + kB_{\alpha_i})w < u_j + (v_j + k'B_{\alpha_j})w + p_j \quad (5)$$

$$\Leftrightarrow u_i + p_i - u_j > (v_j - v_i + k'B_{\alpha_j} - kB_{\alpha_i})w \\ \text{and} \\ (v_j - v_i + k'B_{\alpha_j} - kB_{\alpha_i})w > u_i - u_j - p_j \quad (6)$$

As $u_i + p_i - u_j < w$ and v_i, v_j are integers, this can occur only if $v_j - v_i + k'B_{\alpha_j} - kB_{\alpha_i} \leq 0$.

Similarly, as $u_i - u_j - p_j > -w$ the second condition can occur only if $v_j - v_i + k'B_{\alpha_j} - kB_{\alpha_i} \geq 0$.

Hence $v_j - v_i + k'B_{\alpha_j} - kB_{\alpha_i} = 0$. As B_{α_j} is a multiple of B_{α_i} , there exists an integer κ such that $v_j - v_i = \kappa B_{\alpha_i}$, which proves (3).

Now, if this condition holds, then a collision occurs if $u_j < u_i + p_i$ and $u_i < u_j + p_j$, which finally proves the lemma. \square

Utilization η_i of a job J_i is the proportion of the time the job uses: $\eta_i = \frac{p_i}{T_{\alpha_i}}$. The utilization of the instance is equal to $U = \sum_{i=1}^n \eta_i$. Obviously, U is not greater than 1 in feasible instances.

3 2D PACKING AND PERIODIC SCHEDULING

In this section, we first present a special case of a 2D packing problem and then show the packing problem associated with a periodic scheduling instance. Finally, we prove the equivalence between the two problems by establishing the correspondence between schedule times and 2D coordinates of rectangles.

3.1 Height-Divisible 2D Packing Problem

We define the Height-Divisible 2D (*HD2D*) packing problem. An instance of *HD2D* packing is defined by n rectangles. Each rectangle R_i has width ℓ_i and height h_i . These rectangles are to be packed in a 2D bin of width w and height H . We assume that there are r different heights $H_0 \geq \dots \geq H_{r-1}$, each one divides H and that they form a harmonic set: each a height H_k divides H_j if $j < k$. Note, that $H_0 = H$ and $H_{r-1} = 1$.

A feasible packing defines the bottom left corner coordinates (x_i, y_i) of the rectangles, so there are no collisions between two rectangles, and they all fit into the bin ($x_i + \ell_i \leq w, y_i + h_i \leq H$).

A packing is said to be **height divisible** if the height coordinate y_i is divisible by the height of the rectangle:

$$y_i \bmod h_i = 0.$$

So, the *HD2D* packing problem is to find a feasible height-divisible packing. We can now state the conditions of a collision between two rectangles in a height-divisible packing.

Lemma 2. *In a height-divisible packing, a collision between rectangles R_i and R_j such that $h_j \leq h_i$ occurs iff all following conditions hold:*

$$x_j < x_i + \ell_i, \quad \text{and} \quad x_i < x_j + \ell_j \quad (7)$$

$$y_i \leq y_j < y_i + h_i, \quad (8)$$

Proof. Obviously, a collision occurs if:

$$x_i - \ell_j < x_j < x_i + \ell_i \quad (9)$$

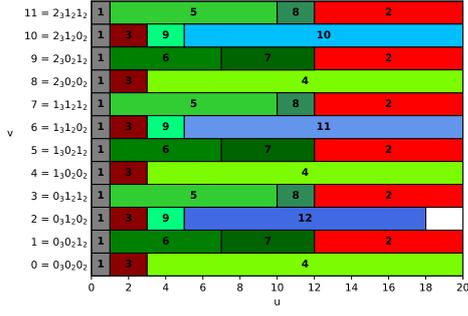
$$y_i - h_j < y_j < y_i + h_i \quad (10)$$

Assuming $h_j \leq h_i$, h_j divides h_i and also divides y_j and y_i . Consequently, the inequality $y_i - h_j < y_j$ reformulates to $y_i \leq y_j$ and the other inequalities remain the same. \square

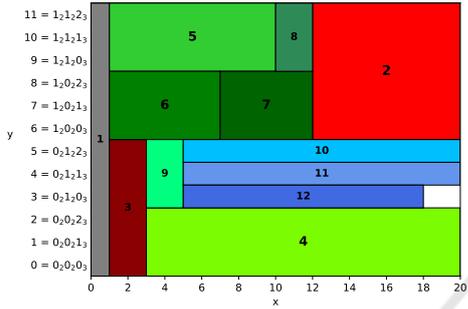
3.2 HD2D Packing Instance Associated with a PSP Instance

We can now define a 2D packing instance associated with a periodic scheduling instance with harmonic periods. Consider an instance I of the scheduling problem. We define an associated instance $\phi(I)$ of the *HD2D* packing problem.

We set $H = \frac{H}{w} = B_{r-1}$. To each job J_i we associate a rectangle R_i of width $\ell_i = p_i$ and height $h_i = \frac{H}{T_{\alpha_i}}$, to fit in a bin of width w and height H .



(a) Periodic scheduling problem.



(b) Height-divisible 2D packing problem.

 Figure 1: Solutions to *PSP* and *HD2D* packing problem, with corresponding jobs and rectangles presented using the same colors. Row indices use the mixed-radix system.

To give an intuition of the correspondence between a periodic schedule and a height-divisible packing, we first develop it on an example in Figures 1. The main idea is to arrange the periodic schedule in rows, as shown in Figure 1a. A row indexed by v represents the schedule of the time window $[vw, (v+1)w)$ of the observation interval.

The instance contains 12 jobs J_1, \dots, J_{12} . J_1 is the only job with the shortest period $T_{\alpha_1} = T_0 = 20$. Jobs J_2, J_3 have period equal to 40, jobs J_4, \dots, J_9 have period equal to 80, and finally, jobs J_{10}, J_{11}, J_{12} have period equal to 240. The chart of the periodic schedule in Figure 1a therefore has width $w = T_0 = 20$, and its height is equal to $\hat{H}/w = B_3 = 12$. The 2D bin of the corresponding *HD2D* packing problem shown in Figure 1b has the same width w and height H .

Figure 1a shows the chart of the feasible periodic schedule over a single observation interval of length $T_3 = 240$. Origin 0 is set by the start time of job J_1 , and we stack each time window of length $w = 20$ on top of the previous one (see the rows in Figure 1a with their three-digit row numbers using mixed-radix system formally defined in Section 3.3). We can see that, e.g., job J_4 has a period of 80, as its subsequent occurrence is at the fourth row from its previous occurrence. Note that jobs with the same period use similar colors.

Figure 1b shows the corresponding height-divisible packing view of the same solution. The packing is obtained by interchanging the rows as follows: we swap the three-digit row numbers in Figure 1a such that the most significant digit becomes the least significant digit and vice versa. Finally, we obtain Figure 1b by ordering the rows according to their new three-digit numbers. For example, the row with J_{11} has number $1_3 1_2 0_2 = 6$ in Figure 1a and after the swap, the same row has number $0_2 1_2 1_3 = 4$ in Figure 1b.

We can see the occurrences of the original jobs have been merged, and the dimensions of the resulting rectangles correspond to those described in Section 3.2. But this is just an observation based on one example, which does not guarantee a compact shape of rectangles in *HD2D* packing. The formal proof of the equivalence of both problems is derived in Section 3.3 and uses positions of jobs and rectangles rather than three-digit numbers of rows used in this Section 3.2.

Note that there exists a bin tree representation of *PSP* with harmonic periods. It was used in (Eisenbrand et al., 2010a) and is closely related to our *HD2D* packing perspective. The bin tree represents the *PSP* with a tree, with each level of nodes corresponding to one of the periods. Examples of the bin tree and *HD2D* packing representations are shown in Figures 7, 8 in (Minaeva and Hanzalek, 2021).

3.3 Flip Operation and Equivalence Between *HD2D* Packing Problem and *PSP*

In this section, we will formally introduce a flip operation, first introduced in our paper (Hanan and Hanzalek, 2020). The flip operation allows us to provide new proof of a bijection between feasible periodic schedules and *HD2D* packing of the associated packing instance in Theorem 1.

Theorem 1. *To any periodic feasible schedule of the instance I , a feasible height-divisible 2D packing of $\phi(I)$ is associated.*

$$\forall J_i, \begin{cases} x_i = u_i \\ y_i = h_i \cdot \text{flip}(v_i, \alpha_i, \mathbf{b}) \end{cases} \quad (11)$$

Conversely, from a height-divisible 2D packing of $\phi(I)$ we can define a feasible schedule as follows:

$$\forall J_i, \begin{cases} u_i = x_i \\ v_i = \text{flip}(\frac{y_i}{h_i}, \alpha_i, Bf(\mathbf{b}, \alpha_i)) \end{cases} \quad (12)$$

We first introduce the theoretical framework used to prove Theorem 1. Let y be any natural number.

It is known that $y < H$ can be decomposed uniquely according to the mixed radix numerical system $\mathbf{b} = (b_1, \dots, b_{r-1})$ as follows:

$$y = y_1 B_0 + y_2 B_1 + y_3 B_2 + \dots + y_{r-1} B_{r-2}$$

with $\forall i > 0, y_i < b_i$. This decomposition is usually denoted with the most significant digits to the left as follows:

$$[y]_{\mathbf{b}} = [y_{r-1}]_{b_{r-1}} \dots [y_1]_{b_1} \quad (13)$$

Notice that in the usual base decomposition (for example, base 2 for a binary decomposition), the components of the base vector \mathbf{b} are all equal (for example, equal to 2 for the binary decomposition).

Let us generalize the transformation proposed by (Lukasiewicz et al., 2009) for the binary decomposition of a number by defining two operators: Let $Bf(\mathbf{b}, k)$ be the operator that flips the k first component of the base vector \mathbf{b} :

$$Bf(\mathbf{b}, k) = (b_k, b_{k-1}, \dots, b_1, b_{k+1}, \dots, b_{r-1}) \quad (14)$$

Let now $flip(y, k, \mathbf{b})$ be the number constructed by flipping the k least significant digits of the decomposition $[y]_{\mathbf{b}}$ to give a number expressed with respect to base vector $\mathbf{b}' = Bf(\mathbf{b}, k)$:

$$[flip(y, k, \mathbf{b})]_{\mathbf{b}'} = [y_{r-1}]_{b_{r-1}} \dots [y_{k+1}]_{b_{k+1}} [y_1]_{b_1} \dots [y_k]_{b_k} \quad (15)$$

Notice that the $r-1-k$ most significant digits are not modified.

Hence, we have:

$$flip(y, k, \mathbf{b}) = y_k \cdot 1 + y_{k-1} \cdot b_k + \dots + y_1 (b_2 \dots b_k) + y_{k+1} (b_1 \dots b_k) + \dots + y_{r-1} (b_1 \dots b_{r-2}) \quad (16)$$

Example: consider the base vector $\mathbf{b} = (2, 2, 3)$ used in Figure 1a. So $H = 12$. Let us consider $y = 10$. Its decomposition along \mathbf{b} is $[y]_{\mathbf{b}} = [2]_3 [1]_2 [0]_2 = 0 + 1 \cdot 2 + 2 \cdot 4$. $Bf(\mathbf{b}, 3) = (3, 2, 2)$. $flip(y, 3, \mathbf{b}) = [0]_2 [1]_2 [2]_3 = 2 + 1 \cdot 3 + 0 \cdot 6 = 5$.

This flip operation has some important properties that will be used to transform the scheduling problem into an equivalent packing problem. The following lemma expresses the reversibility of the flip and the fact that equidistant integers (with distance B_{k-1}) become consecutive integers after the flip operation.

Lemma 3. For any integer $y < H$,

$$flip(flip(y, k, \mathbf{b}), k, Bf(\mathbf{b}, k)) = y \quad (17)$$

if $[y_k]_{b_k} < b_k - 1$, then

$$flip(y + B_{k-1}, k, \mathbf{b}) = flip(y, k, \mathbf{b}) + 1 \quad (18)$$

and if $[y_k]_{b_k} > 0$,

$$flip(y - B_{k-1}, k, \mathbf{b}) = flip(y, k, \mathbf{b}) - 1 \quad (19)$$

If y is a multiple of B_k , $z < B_k$ and $k' \geq k$ then

$$flip(y + z, k', \mathbf{b}) = flip(y, k', \mathbf{b}) + flip(z, k', \mathbf{b}) \quad (20)$$

Proof. Obviously, $Bf(Bf(\mathbf{b}, k), k) = \mathbf{b}$, since a double flip of the k first digits lead to the initial digits. This proves (17). Let us consider (18). If we add B_{k-1} to y , the digit $[y_k]_{b_k}$ increases by one unit, so after the flip operation, this digit will be the least significant; hence, the corresponding number increases by 1. Similar arguments hold for (19). Finally, if y is a multiple of B_k , its k least significant digits are null with respect to base \mathbf{b} . If $z < B_k$ its digits from $k+1$ to the last one are null, so the sum $y+z$ consists in the k first digits of z followed by the next digits of y . If $k' \geq k$ $flip(y, k', \mathbf{b})$ has $k' - k$ first flipped digits from the digits k to k' of y and $flip(z, k', \mathbf{b})$ starts with $k' - k$ null digits followed by the flipped k digits of z . So, the sum is exactly the flipped digits from $y+z$. \square

We are now able to prove the equivalence between *PSP* and *HD2D* packing problem.

Proof of Theorem 1.

Proof. The proof relies on the properties of the flip operation stated in Lemma 3 and the expressions of collisions stated in Lemma 1 and 2.

Assume that there is a feasible schedule of instance I . Now, assume that a collision occurs in the associated 2D packing between rectangles R_i and R_j such that $h_i \geq h_j$, so that $T_{\alpha_i} \leq T_{\alpha_j}$.

So according to equation (7), and as $x_i = u_i, x_j = u_j$ we have $u_j < u_i + p_i$ and $u_i < u_j + p_j$. This implies that the condition (2) of Lemma 1 holds, and thus, as there is no collision in the schedule, the second condition cannot hold: $v_j - v_i$ cannot be a multiple of B_{α_i} .

Now according to Lemma 2, equation (9), the collision in the packing implies that:

$$y_i \leq y_j < y_i + h_i.$$

Assume that $y_j = y_i + \Delta$. Notice that as h_j divides h_i , and as h_j divides y_j , it should also divide Δ . Moreover:

$$\begin{aligned} y_j &= h_j \cdot flip(v_j, \alpha_j, \mathbf{b}) = h_i \cdot flip(v_i, \alpha_i, \mathbf{b}) + \Delta \\ flip(v_j, \alpha_j, \mathbf{b}) &= \frac{h_i}{h_j} flip(v_i, \alpha_i, \mathbf{b}) + \frac{\Delta}{h_j} \\ &\text{with } \frac{\Delta}{h_j} < \frac{h_i}{h_j} \end{aligned}$$

Now according to Lemma 3,

$$\begin{aligned} v_j &= flip\left(\frac{y_j}{h_j}, \alpha_j, Bf(\mathbf{b}, \alpha_j)\right) \\ &= flip\left(\frac{h_i}{h_j} flip(v_i, \alpha_i, \mathbf{b}) + \frac{\Delta}{h_j}, \alpha_j, Bf(\mathbf{b}, \alpha_j)\right) \end{aligned}$$

Observe that $\frac{h_i}{h_j} = \frac{T_{\alpha_j}}{T_{\alpha_i}} = b_{\alpha_i+1} \dots b_{\alpha_j}$. And $\frac{\Delta}{h_j} < \frac{h_i}{h_j} = b_{\alpha_i+1} \dots b_{\alpha_j}$. So applying Lemma 3 with base $Bf(\mathbf{b}, \alpha_j)$ we can say that:

$$v_j = flip\left(\frac{h_i}{h_j} flip(v_i, \alpha_i, \mathbf{b}), \alpha_j, Bf(\mathbf{b}, \alpha_j)\right) + flip\left(\frac{\Delta}{h_j}, \alpha_j, Bf(\mathbf{b}, \alpha_j)\right)$$

In the base representation $Bf(\mathbf{b}, \alpha_j)$ the $\alpha_j - \alpha_i$ least significant digits of $\frac{h_i}{h_j} flip(v_i, \alpha_i, \mathbf{b})$ equal 0. And the next α_i digits are the ones of v_i flipped. So that the α_i least significant digits of $flip\left(\frac{h_i}{h_j} flip(v_i, \alpha_i, \mathbf{b}), \alpha_j, Bf(\mathbf{b}, \alpha_j)\right)$ in the base representation \mathbf{b} are equal to the α_i least significant digits of v_i in base representation \mathbf{b} followed by digits 0. Hence $v_i = flip\left(\frac{h_i}{h_j} flip(v_i, \alpha_i, \mathbf{b}), \alpha_j, Bf(\mathbf{b}, \alpha_j)\right)$.

Moreover the number $flip\left(\frac{\Delta}{h_j}, \alpha_j, Bf(\mathbf{b}, \alpha_j)\right)$ has its α_i least significant digits equal to 0 in the base representation \mathbf{b} . So that it is a multiple of B_{α_i} :

$$v_j = v_i + \kappa B_{\alpha_i}. \quad (21)$$

Since we know that $v_j < \frac{T_{\alpha_j}}{w} = B_{\alpha_j}$, we have $\kappa < \frac{B_{\alpha_j}}{B_{\alpha_i}}$. So condition (3) of Lemma 1 holds, which contradicts our assumption that no collision occurs.

Conversely, assume that we have a feasible packing and let us assume that there is a collision in the associated schedule between jobs J_i and J_j with $T_{\alpha_i} \leq T_{\alpha_j}$. The first condition of Lemma 1 holds so that the first condition of Lemma 2 holds either. This implies that the condition (8) of Lemma 2 does not hold, and either $y_i > y_j$ or $y_j \geq y_i + h_i$.

But if a collision occurs in the schedule, we should have: $v_j = v_i + \kappa B_{\alpha_i}$.

So,

$$\begin{aligned} \frac{y_j}{h_j} &= flip(v_j, \alpha_j, \mathbf{b}) \\ &= flip(v_i + \kappa B_{\alpha_i}, \alpha_j, \mathbf{b}) \\ &= flip\left(flip\left(\frac{y_i}{h_i}, \alpha_i, Bf(\mathbf{b}, \alpha_i)\right) + \kappa B_{\alpha_i}, \alpha_j, \mathbf{b}\right) \end{aligned}$$

Notice that as $y_i < H$ so that $\frac{y_i}{h_i} < B_{\alpha_i}$ and thus when performing the operation $flip\left(\frac{y_i}{h_i}, \alpha_i, Bf(\mathbf{b}, \alpha_i)\right)$, we get a number still less than B_{α_i} . We can then apply Lemma 3, and we get:

$$\begin{aligned} &flip\left(flip\left(\frac{y_i}{h_i}, \alpha_i, Bf(\mathbf{b}, \alpha_i)\right) + \kappa B_{\alpha_i}, \alpha_j, \mathbf{b}\right) \\ &= flip\left(flip\left(\frac{y_i}{h_i}, \alpha_i, Bf(\mathbf{b}, \alpha_i)\right), \alpha_j, \mathbf{b}\right) + \kappa \\ &= \frac{B_{\alpha_j}}{B_{\alpha_i}} \cdot \frac{y_i}{h_i} + \kappa \end{aligned}$$

Now $h_i B_{\alpha_i} = h_j B_{\alpha_j} = H$, so that:

$$y_j = y_i + \kappa \cdot h_j \quad (22)$$

Now we know that $\kappa < \frac{B_{\alpha_j}}{B_{\alpha_i}} = \frac{h_i}{h_j}$. So $y_j < y_i + h_i$ and condition (8) of Lemma 2 holds, a contradiction. \square

3.4 Canonical Form of a HD2D Packing

In this section, we prove that any HD2D packing can be put in a canonical form for which the packed rectangles' heights decrease from the left to the right.

Definition 1. A height-divisible packing is said to be in canonical form if for all $i, j \in \{1, \dots, n\}$

$$h_i > h_j, \quad y_j \geq y_i, \quad y_j + h_j \leq y_i + h_i \implies x_i + \ell_i \leq x_j \quad (23)$$

Such a packing corresponds to a schedule of the observation interval where in each interval $[kw, (k+1)w)$ (i.e., the k -th row in Figure 1a), jobs are scheduled in increasing order of periods.

Theorem 2. If there exists a feasible packing of width ℓ , then there exists a feasible canonical packing of width ℓ .

Proof. By induction on r the number of different heights of rectangles. Of course, if $r = 1$, the theorem holds. Assume this is true for any height-divisible packing with at most r different heights. Consider an instance with $r + 1$ different heights. We claim that the rectangles of maximal height $H = H_0$ can be packed at the left of the bin, and the other rectangles can be right-shifted accordingly. Let ℓ' be the sum of the lengths of the rectangles of height H_0 . Consider H_1 the second largest height. If R_j is a rectangle of height $h_j \leq H_1$, we know that its vertical coordinate is a multiple of h_j . So there exists a non-negative integer q such that $y_j = qh_j = q'H_1 + y'_j$, with $y'_j \geq 0$ and $y'_j + h_j \leq H_1$.

So we can consider $\frac{H_0}{H_1}$ so-called sub-bins of height H_1 and width $\ell - \ell'$ in which all remaining rectangles are packed. Here, rectangle R_j is packed in sub-bin q' . The set of rectangles packed into a sub-bin has at most r different heights. By induction, all the rectangles packed in a sub-bin can be packed with a canonical packing. \square

The canonical form of the packing shown in Figure 1b is presented in Figure 2. The canonical form also shows sub-bins, which are crucial for the description of heuristics. The original 2D bin of size $w \times H$ starts from the left as a single sub-bin of height H . Once all rectangles with height $H_0 = H$ are packed

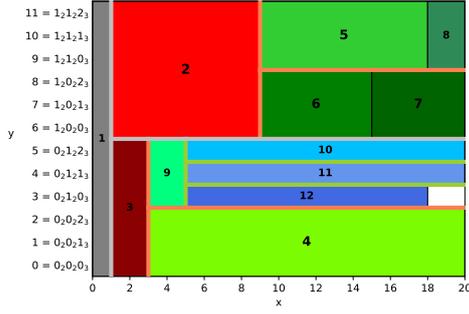


Figure 2: Canonical form of the packing shown in Figure 1b. Sub-bin division is highlighted by special T-like guillotine cuts in gray, red, and green colors, that are applicable in the canonical form only.

(grey rectangle R_1 in Figure 2), the remaining free space to the right of the parent sub-bin is divided to $\frac{H_0}{H_1}$ sub-bins of height H_1 . Once rectangles with height H_1 are packed (reddish rectangles R_2, R_3), we divide each sub-bin into $\frac{H_1}{H_2}$ sub-bins of height H_2 . This continues until sub-bins of height $H_{r-1} = 1$ are created, and the rectangles with the smallest height are packed.

In Figure 2, the lower sub-bin of height 6 contains rectangles $R_3, R_4, R_9, R_{10}, R_{11}, R_{12}$. The division of sub-bins is also highlighted by the colored T-like contours. For example, the sub-bin of height 6, where rectangle R_2 is packed, divides itself into two sub-bins, each containing two rectangles of height 3.

4 CP MODEL

We use the bin formulation, found in (Lukasiewicz et al., 2009; Eisenbrand et al., 2010b), to model and solve the single machine *PSP* with harmonic periods. We formulate the problem from the *HD2D* packing perspective, and the solution to the *PSP* is obtained using the equations from Section 3.3. The bin formulation only assigns the rectangles to suitable sub-bins, as described in the previous section. Once such an assignment is obtained, the packing solution in the canonical form is recovered.

We model the bin formulation using CP. We rely on the **pack** constraint (Shaw, 2004). The constraint **pack**($\mathbf{L}, \mathbf{a}, \mathbf{l}$) couples together list \mathbf{L} of groups' load variables, list \mathbf{a} of items' assignment variables, and list \mathbf{l} of items' sizes, so:

$$a_i \in \{0, \dots, |\mathbf{L}| - 1\} \quad \forall i \in \{0, \dots, |\mathbf{a}| - 1\} \quad (24)$$

$$L_i = \sum_{\forall j: a_j=i} l_j \quad \forall i \in \{0, \dots, |\mathbf{L}| - 1\} \quad (25)$$

Thus, each item is assigned to one of the groups, and each group's load variable contains the sum of its assigned items' sizes.

To model the *HD2D* packing problem, we use one pack constraint for each distinct height H_k . This pack constraint uses load variables $\mathbf{L}^k = (L_0^k, \dots, L_{H_0/H_k-1}^k)$ to model $\frac{H_0}{H_k}$ sub-bins of height H_k . The sub-bins are indexed from bottom to top. Lists of widths \mathbf{l}^k and assignment variables \mathbf{a}^k are associated with all rectangles with height $h_j = H_k$. With these definitions, the used pack constraints are:

$$\mathbf{pack}(\mathbf{L}^k, \mathbf{a}^k, \mathbf{l}^k) \quad \forall k \in \{0, \dots, r-1\} \quad (26)$$

Such constraint distributes the rectangles with height H_k among sub-bins with height H_k , and the load variable L_i^k is equal to the sum of widths of rectangles that are assigned to i -th sub-bin of height H_k . Then, for each row of height 1 in the packing, we sum together loads of sub-bins that intersect such row (see Figure 2). We ensure that such sum does not exceed the width of the 2D bin w :

$$\sum_{k=0}^{r-1} L_{[i/H_k]}^k \leq w \quad \forall i \in \{0, \dots, H_0 - 1\} \quad (27)$$

In the case of line 4 in Figure 2 we sum up $L_0^0 = 1$ plus $L_1^1 = 2$ plus $L_2^2 = 2$ plus $L_3^3 = 15$.

These two sets of constraints are sufficient to model the problem. Once solved, the assignment of rectangles to sub-bins is determined from values of variables \mathbf{a}^k .

5 HEURISTICS

While the exact approaches completely search the solution space, their performance diminishes with an increasing number of jobs to schedule. On the other hand, even simple heuristics are often able to find a solution for large instances (Hladik et al., 2020).

All studied heuristics process the jobs one by one, ordered first non-decreasingly by their period and then non-increasingly by their processing times. The period ordering (known as rate-monotonic in the time-triggered community) is natural, several heuristics rely on it, and it produces a solution in canonical form. The secondary ordering by processing times is used by the bin packing and scheduling approximation algorithms (e.g., in (Della Croce and Scatamacchia, 2018)).

5.1 Baseline Heuristics

5.1.1 Time-Wise First Fit

In (Hladik et al., 2020), authors originally tackled the problem with dedicated machines and precedence

constraints. Their algorithm schedules each job as soon as possible, minimizing the job’s start time in the yet-unfinished schedule. Due to the time-oriented aspect of the heuristic, we denote it **T-FF**.

5.1.2 Spatial First Fit and Best Fit

In (Eisenbrand et al., 2010a), the first fit is also used. However, the rectangles of the corresponding *HD2D* packing problem are packed to the lowest available sub-bin instead. We refer to this approach as a spatial first fit. Furthermore, we also tested best fit as a sub-bin selection policy; thus, we obtained two distinct variants of the algorithm, denoted **S-FF** and **S-BF** respectively. Note that the best-fit selection can be viewed as a single-dimension "Residual Space" merit function, used to pack items in (Crainic et al., 2008).

5.1.3 Longest Processing Time First

The last baseline approach is the well-known Longest Processing Time First list scheduling algorithm for parallel machines scheduling (Della Croce and Scatamacchia, 2018). We follow the same spatial approach of Section 5.1.2, but for the current job, its corresponding rectangle is packed into the least occupied sub-bin. We denote this method **LPT**.

5.2 Rectangle-Guided First Fit

We observed that the baseline heuristics often cannot schedule infrequent long jobs, whose corresponding rectangles are wide and low. These rectangles are packed later in the process when the packing is mostly filled. Following this observation, we designed our heuristic, called rectangle-guided first fit, using the *HD2D* packing perspective. The algorithm consists of two phases. In the first phase, so-called dummy rectangles are created, and in the second phase, the rectangles (including dummy rectangles) are packed using a modified first fit heuristic.

Partly inspired by "wide" and "narrow" rectangles of (Ntene and van Vuuren, 2009; Ortmann et al., 2010), the dummy rectangles are created for each distinct height H_k of the problem. For $H_{r-1} = 1$, there are no dummy rectangles. For each subsequent height H_k , the dummy rectangles are constructed based on both real and dummy rectangles of height H_{k+1} . As the newly created rectangles are taller, they will be processed earlier by the heuristic and will reserve space for the smaller ones. However, their only purpose is to provide a look-ahead; once all the rectangles with height H_k are packed, we remove all the dummy rectangles with height H_k from the sub-bins and continue with packing the rectangles of height H_{k+1} .

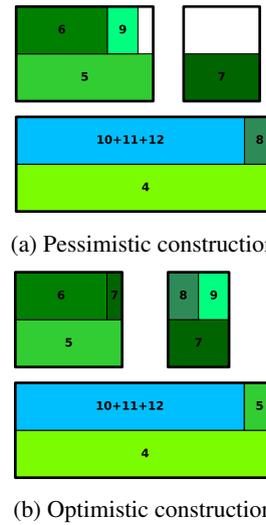


Figure 3: Bag view of dummy rectangles for height $H_1 = 6$ constructed for instance from Section 3.2.

5.2.1 On Construction of Dummy Rectangles

Dummy rectangle of height H_k and width l essentially consists of H_k/H_{k+1} bins of height H_{k+1} , each with width l . Our aim is to create dummy rectangles by filling these bins with the smaller rectangles of height H_{k+1} , so the new rectangles roughly capture how the smaller rectangles could be packed. We create the rectangles heuristically using 1D bin packing. To avoid confusion with the 2D packing, we call the one-dimensional bins **bags**. We propose two approaches for constructing dummy rectangles.

The pessimistic approach sorts the rectangles with height H_{k+1} by their width. We start with an empty list of bags. If a rectangle R_i with width l_i fits in one of the existing bags, it is packed according to the best fit policy. If not, a new dummy rectangle of height H_k is created, with width l_i . Simultaneously, H_k/H_{k+1} bags of size l_i are created, and R_i fully fills one of them. The process repeats until all the rectangles are packed in the bags. Since many bags may contain holes, not all constructed dummy rectangles are completely filled with the smaller ones. This is shown in Figure 3a. Three dummy rectangles of height 6 were created from rectangles of height 3, but two bags associated with two narrower rectangles were not fully filled. Note that the blue rectangle is the dummy rectangle created for $H_2 = 3$. We refer to our algorithm using pessimistic construction as **RG-FF-PES**.

The optimistic variant fills the bags in a "preemptive manner" and always keeps at most a single bag not full. We start with an empty list of bags. Until we exhaust all the rectangles with height H_{k+1} , we select the widest R_i with width l_i . If it fits any (i.e., the single not full) bag, we pack it to it. If the bag has

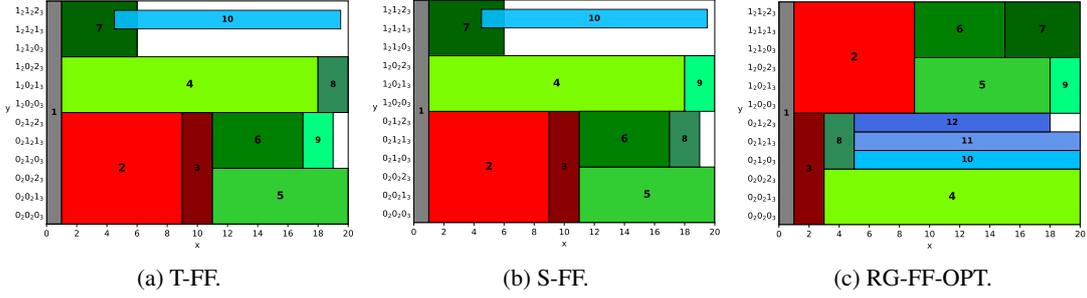


Figure 4: Unfinished (a), (b) and finished (c) packings created by different methods on instance from Section 3.2.

vacant size $0 < A < l_i$, we split R_i into two rectangles with widths A and $l_i - A$. The first rectangle is packed into the bag, and the second is pushed back among the remaining rectangles. If all the bags are full, we create a new dummy rectangle of height H_k and width l_i , push a new bag with size $l_i \cdot H_k / H_{k+1}$ to the list of bags, and pack R_i into it; this new bag is a concatenation of H_k / H_{k+1} bags of width l_i , that make up the rectangle. We denote this variant of the algorithm as **RG-FF-OPT**.

Note that only the last bag might contain holes. In Figure 3b, all the bags were actually fully filled. Also, notice that rectangles R_5, R_7 were split into two parts. Due to the splitting and the usage of a single rather than $H_1 / H_2 = 2$ bags, we cannot simply partition the packed dummy rectangle into a feasible packing of rectangles that filled the associated bag; in a pessimistic case, it is possible (see Figure 3a). However, since the dummy rectangles only provide a look-ahead, this "preemptiveness" does not cause issues with the final packing.

5.2.2 First Fit Packing

Once the dummy rectangles are constructed, our algorithm follows the first fit approach of algorithm **S-FF**. All the rectangles, both real and dummy, are sorted following the policy of Section 5; first, non-increasingly by height, then non-increasingly by width.

Our algorithm differs from **S-FF** in two ways. Once the last rectangle of height H_k is packed, we remove packed dummy rectangles from all the sub-bins, recalculate the sub-bins' utilizations, and then continue packing the rectangles of height H_{k+1} . The second difference is the packing policy. If the current rectangle R_i fits any sub-bin, we pack it according to the first fit policy. If not, we do not immediately stop the algorithm; the sub-bins may become under-utilized once the dummy rectangles are removed. Thus, if R_i is a dummy rectangle, we pack it to the least utilized sub-bin, temporarily exceeding the sub-bin's size. If R_i is a real rectangle, we find

the sub-bins that would fit R_i after removal of dummy rectangles. If such sub-bins exist, we pack R_i to the least utilized one; otherwise, the heuristic fails to find a feasible solution. If all real rectangles are packed, the algorithm finishes successfully, and we transform the packing into the schedule.

Internally, we use the compressed bin tree representation of (Eisenbrand et al., 2010a) so the complexity of the first fit packing of a single rectangle is $O(n)$ given n rectangles to pack (i.e., jobs to schedule). Since for each height H_k we need to process $O(n)$ real and dummy rectangles, the overall complexity of our solution is $O(r \cdot n^2)$.

5.3 Comparison of Methods

In Figures 4, we compare **T-FF**, **S-FF**, and **RG-FF-OPT** methods on example instance from Section 3.2, using the *HD2D* packing view. Neither **T-FF** nor **S-FF** methods were successful, as they were not able to reserve space for blue rectangle R_{10} . Notice that **T-FF** and **S-FF** put rectangle R_8 at different sub-bins. While both methods process R_8 at the same moment, **T-FF** selects the time-wise earliest start time for the associated job, while **S-FF** selects the lowest sub-bin. Finally, **RG-FF-OPT** used dummy rectangles to reserve enough space for R_{10} by putting rectangles R_2, R_3 into different sub-bins.

6 EXPERIMENTS

The algorithms were implemented in Python 3.10. Experiments were performed on an Intel Xeon Silver 4110@2.10 GHz, using IBM ILOG CP Optimizer v22.1.1 and Gurobi v10.0.2.

6.1 Experimental Data

We evaluated the methods on several sets of synthetically generated problem instances. Due to the success of the method by (Hladik et al., 2020), we only gen-

Table 1: Number of instances solved by each method on instance sets S_1, S_2, S_3 .

set	# instances	heuristic						exact	
		LPT	T-FF	S-FF	S-BF	RG-FF-PES	RG-FF-OPT	ILP	CP
S_1	3518	2213	3233	3248	3263	3378	3378	3518	3518
S_2	800	122	558	562	563	681	694	800	800
S_3	297	8	14	15	15	24	29	289	289

Table 2: Characteristics of the instance sets.

set	# instances	average jobs	median η_i [%]	median p_i/T_0 [%]
S_1	3518	84	0.098	0.40
S_2	800	57	1.563	6.25
S_3	297	47	1.719	9.10
D_2^6	200	84	0.278	3.50
D_3^6	200	405	0.028	3.25
D_5^3	200	3719	0.001	3.00
D_{20}^3	200	1473	0.019	6.88

erated feasible instances with utilization U equal to 100 %.

The first sets of instances were generated using the scheme proposed in (Hladik et al., 2020). Starting with a selected period T_0 , a single job with period and processing time equal to T_0 is generated. Then, for a given number of iterations, one of the jobs with period $T_{\alpha_i} = T_k$ is selected and is either split into two new jobs of the same period but shorter processing times or is divided into b_{k+1} jobs with the same processing time and period equal to T_{k+1} .

Set S_1 contains original instances from (Hladik et al., 2020). The original generation scheme often created instances with many long-period short-processing times jobs that are relatively easy to solve. We modified the generator scheme to account for this and produced sets S_2 , and S_3 .

Then, we generated 4 sets of difficult instances using *HD2D* packing perspective (in canonical form). We ensured that the majority of jobs had a processing time greater than 13. These short-jobs-free instances seemed to be the most difficult to solve; their sets were labeled by the letter D . Four sets with 200 instances were created. Each instance set D_T^r used period set $\{800 \cdot T^k \mid k \in 0, \dots, r-1\}$.

Characteristics of the instance sets are shown in Table 2. We show the number of generated instances and the average number of jobs in the instance. Median η_i is the median utilization across all the jobs of the instance set. Similarly, median p_i/T_0 is the median value of the processing time of the job, normalized by the minimum period T_0 of the associated instance. We can see that the simpler instances contain only around 80 jobs. However, thanks to their diverse characteristics, we use them to thoroughly compare our heuristics.

6.2 Success Rates of the Methods

First, we evaluated how many instances each method can successfully solve. CP and ILP solvers were lim-

Table 3: Number of instances solved by exact methods on difficult instance sets.

set	# instances	ILP	CP
D_2^6	200	95	152
D_3^6	200	1	36
D_5^3	200	0	0
D_{20}^3	200	0	0

 Table 4: Number of instances solved by RG-FF-OPT and combinations of heuristics M_i on instance sets S_1, S_2, S_3 .

set	# instances	RG-FF-OPT	M_1	M_2	M_3	M_4
S_1	3518	3378	3423	3439	3437	3459
S_2	800	694	709	726	725	742
S_3	297	29	30	30	30	33

ited to 3 minutes of computation time per instance.

We found that the heuristics could not find any feasible solutions in the case of the difficult D_T^r instances. Thus, we only used these instances to compare CP and ILP. As is reported in Table 1 both solvers can solve smaller instances S_1, S_2, S_3 . However, CP is more successful on the mentioned difficult instances shown in Table 3, given the limited computation time of 3 minutes. We suspect that the reason is an efficient formulation that relies on the powerful pack constraint. However, for difficult sets D_5^3, D_{20}^3 , neither method found any solution. We use these sets later in Section 6.3 to highlight the shortcomings of the CP in comparison with heuristics.

The results of heuristics on sets S_1, S_2, S_3 are also reported in Table 1. We can see that on all these sets, the packing-inspired methods RG-FF-PES and RG-FF-OPT defeated the baseline heuristics, with the optimistic variant RG-FF-OPT being the best. This is most pronounced for set S_2 . Our revised generation scheme made the instances harder to solve; this was highlighted by set S_3 , where heuristics solved only a few instances. Finally, except for LPT, the results of baseline heuristics do not differ significantly.

We were also interested in the possible orthogonality of the studied heuristics; if each type of the heuristic would solve a distinct subset of the instances, we could run them as a portfolio and achieve a higher overall success rate. For each pair of heuristics, we determined the number of instances of the instance set solved by at least one heuristic from the pair, and we present several interesting combinations. M_1 is combination of RG-FF-OPT and RG-FF-PES, M_2 is combination of RG-FF-OPT and S-BF, and M_3

Table 5: Results of the utilization experiment on instance sets S_1, S_2, S_3 .

set	# instances	heuristic												exact			
		LPT		T-FF		S-FF		S-BF		RG-FF-PES		RG-FF-OPT		CP10		CP60	
		# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F
S_1	3518	3452	0.962	3512	0.994	3512	0.995	3512	0.995	3518	0.998	3518	0.998	3518	0.999	3518	0.999
S_2	800	788	0.923	798	0.988	798	0.989	798	0.989	800	0.997	800	0.996	800	1.000	800	1.000
S_3	297	277	0.915	296	0.964	296	0.965	296	0.966	297	0.978	297	0.977	297	1.000	297	1.000

is combination of RG-FF-OPT and T-FF. Finally, M_A corresponds to running all the heuristics together.

From the results reported in Table 4, we can see that the combination of the two proposed methods does not offer much improvement. On the other hand, when we use our heuristic with the S-BF (M_2), the results are significantly better, increasing the number of successes by 30 in the case of set S_2 . However, set S_3 remains problematic for heuristics to solve, even when we would run all the heuristics together (M_A).

6.3 Utilization Experiment

While the results reported in Section 6.2 show the power of the CP, we also studied how each method tackles the instances with utilization $U < 1.0$. While the heuristics have trouble solving instances with $U = 1.0$, they might be successful when U is slightly lower. We study this in the following experiment.

For a given method and instance, we evaluate whether the method successfully solves the instance. If not, we remove the job with the smallest value of utilization η_i and repeat the process. When the success eventually occurs, we report the final utilization U_F of the instance. Thus, if the method solved each instance immediately or just after a few iterations, it would report all the values U_F close to 1. Note that the jobs are removed in a deterministic way, as outlined; there might exist a situation when removing a single longer job would make the instance easier than removing several shorter jobs consecutively. We remove the jobs until the utilization of the instance reaches $U < 0.7$ when the failure is reported. The experiment was performed for each instance set using all the heuristics, as well as the CP solver with a time limit: CP10 was limited to 10 seconds of computation, while CP60 was limited to 60 seconds.

The number of successes (i.e., number of instances for which success occurred before $U_F < 0.7$) and an average value of U_F of successful attempts on instance sets S_1, S_2, S_3 are shown in Table 5. All the heuristics, with the exception of the LPT, are able to solve most of the instances with just a slight reduction of their overall utilization. This matches the results reported in (Hladik et al., 2020). On instance sets S_2, S_3 (see Figure 5), which were generated with our modified scheme, the proposed heuristics RG-FF-PES and RG-FF-OPT report average value of U_F 1 % higher

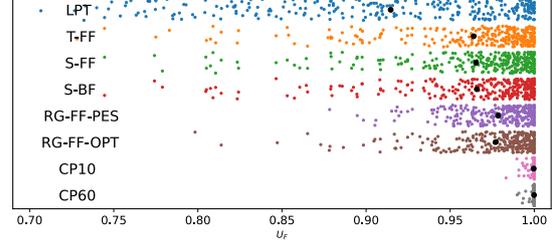


Figure 5: Utilization experiment on instance set S_3 . Colored dots mark final utilizations U_F for each instance. Black dots mark the average value per method.

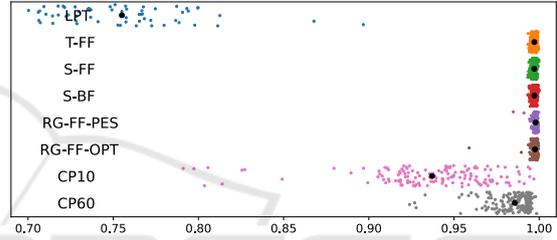


Figure 6: Utilization experiment on instance set D_5^6 . Colored dots mark final utilizations U_F for each instance. Black dots mark the average value per method.

than achieved by their baseline counterparts (with the theoretical maximum being 3 %). These results again suggest that the proposed methods perform better on the instances with infrequent long jobs.

Results on difficult sets D_T^r are reported in Table 6. Interestingly, most heuristics are able to solve them once the instances' utilization slightly drops. Regarding the performance of the CP solver, interesting results are reported for the largest instances of D_5^6, D_{20}^3 . There, the performance of CP10 and CP60 diminishes in contrast with heuristics (whose runtime is at most 0.2 seconds), as is also shown in Figure 6. Thus, as the instance size grows and with a limited computation time, it may be better to use heuristics to create a slightly less utilized schedule.

7 CONCLUSION

In this paper, we studied the *PSP* with a harmonic set of periods and non-preemptive jobs on a single machine. We formally proved that the studied problem is equivalent to the *HD2D* packing problem and that each feasible schedule and its associated packing can

Table 6: Results of the utilization experiment on difficult instance sets.

set	# instances	heuristic												exact			
		LPT		T-FF		S-FF		S-BF		RG-FF-PES		RG-FF-OPT		CP10		CP60	
		# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F	# succ	avg U_F
D_2^5	200	188	0.873	200	0.995	200	0.994	200	0.995	200	0.994	200	0.992	200	1.000	200	1.000
D_3^5	200	139	0.805	200	0.996	200	0.996	200	0.996	200	0.997	200	0.996	200	0.999	200	1.000
D_5^5	200	61	0.755	200	0.997	200	0.997	200	0.997	200	0.998	200	0.998	144	0.937	200	0.986
D_{30}^5	200	0	-	200	0.997	200	0.997	200	0.997	200	0.997	200	0.996	200	0.991	200	0.997

be transformed into the canonical form. We used the CP with a bin formulation and showed its effectiveness compared to the ILP. Then, we proposed heuristics based on the *HD2D* packing perspective, suitable for scheduling instances with infrequent long jobs.

We have shown the advantages of our method in comparison with several baseline heuristics on sets of synthetically generated instances. Furthermore, we performed an experiment where we incrementally lowered the utilization of instances. The experiments allowed us to compare the heuristics and observe the CP’s shortcomings when limited computation time is in place. The experiments have shown that the short-jobs-free instances with 100 % utilization are the most difficult to solve. However, rate-monotonic-based heuristics work very well whenever the utilization slightly drops. The empty space of 99% utilization instance behaves as infrequent, very short jobs, which can fill the fragmented space left at the end of the constructive heuristics run.

We aim to use the presented results and heuristics as a basis of evolutionary algorithms to tackle more complex problems with dedicated machines, precedence relations, and a criterion to optimize.

ACKNOWLEDGEMENTS

This work was supported by the Grant Agency of the Czech Technical University in Prague, grant No. SGS22/167/OHK3/3T/13.

This work was co-funded by the European Union under the project ROBOPROX - Robotics and advanced industrial production (reg. no. CZ.02.01.01/00/22_008/0004590).

REFERENCES

Cai, Y. and Kong, M. (1996). Nonpreemptive scheduling of periodic tasks in uni- and multiprocessor systems. *Algorithmica*, 15(6):572–599.

Crainic, T. G., Perboli, G., and Tadei, R. (2008). Extreme point-based heuristics for three-dimensional bin packing. *INFORMS J. Comput.*, 20:368–384.

Della Croce, F. and Scatamacchia, R. (2018). Longest processing time rule for identical parallel machines revisited. *Journal of Scheduling*, 23/2:163–176.

Dvorak, J. and Hanzalek, Z. (2014). Multi-variant time constrained FlexRay static segment scheduling. *2014 10th IEEE WFCS*, pages 1–8.

Eisenbrand, F. et al. (2010a). Scheduling periodic tasks in a hard real-time environment. *Automata, Languages and Programming*, 6198:299–311.

Eisenbrand, F. et al. (2010b). Solving an avionics real-time scheduling problem by advanced IP-methods. *European Symposium on Algorithms*, pages 11–22.

Fleszar, K. and Hindi, K. S. (2002). New heuristics for one-dimensional bin-packing. *Comput. Oper. Res.*, 29(7).

Hanen, C. and Hanzalek, Z. (2020). Periodic scheduling and packing problems. *CoRR*, abs/2011.01898.

Hanzalek, Z., Zahora, J., and Sojka, M. (2023). Cone slalom with automated sports car – trajectory planning algorithm. *IEEE Trans. on Vehicular Technology*.

Haouari, M. and Serairi, M. (2009). Heuristics for the variable sized bin-packing problem. *Comput. Oper. Res.*, 36(10).

Hladik, R., Minaeva, A., and Hanzalek, Z. (2020). On the complexity of a periodic scheduling problem with precedence relations. In *Proc. of COCOA 2020*.

Korst, J., Aarts, E., and Lenstra, J. K. (1996). Scheduling periodic tasks. *INFORMS J. Comput.*, 8(4):428–435.

Lukasiewicz, M. et al. (2009). FlexRay schedule optimization of the static segment. In *Proc. of IEEE/ACM CODES+ISSS 2009*, pages 363–372.

Marouf, M. and Sorel, Y. (2011). Scheduling non-preemptive hard real-time tasks with strict periods. In *ETFA2011*, pages 1–8.

Minaeva, A. and Hanzalek, Z. (2021). Survey on periodic scheduling for time-triggered hard real-time systems. *ACM Comput. Surv.*, 54(1).

Ntene, N. and van Vuuren, J. H. (2009). A survey and comparison of guillotine heuristics for the 2D oriented off-line strip packing problem. *Discrete Optimization*, 6(2):174–188.

Ortmann, F. G., Ntene, N., and van Vuuren, J. H. (2010). New and improved level heuristics for the rectangular strip packing and variable-sized bin packing problems. *EJOR*, 203(2):306–315.

Shaw, P. (2004). A constraint for bin packing. In *International Conference on Principles and Practice of Constraint Programming*.