Combining heuristics and integer linear programming to compute a solution to the Robotic Process Automation (RPA) problem

S. Séguin, I. Benkalaï, H. Tremblay,
X. Lebeuf, D.-E. Perron-Chouinard

G–2021–08

March 2021

Combining heuristics and integer linear programming to compute a solution to the Robotic Process Automation (RPA) problem

Sara Séguin [a, b]

Imène Benkalaï [a, b]

Hugo Tremblay [b]

Xavier Lebeuf [c]

David-Emmanuel Perron-Chouinard [c]

[a] *GERAD, Canada, H3T 2A7*

[b] *Département d'informatique et de mathématique, Université du Québec à Chicoutimi, Canada, G7H 2B1*

[c] *Département des sciences appliquées, Université du Québec à Chicoutimi, Canada, G7H 2B1*

sara_seguin@uqac.ca
imene1_benkalai@uqac.ca
hugo_tremblay2@uqac.ca
xavier_lebeuf1@uqac.ca
david-emmanuel_perron-chouinard1@uqac.ca

Abstract :   Robotic process automation (RPA) is used in various fields of human activity in order to implement faster and more secure processes through a reduction in the risks or errors but also an increase in the productivity rates. The increase of its use and importance calls for evermore efficient solution methods for this problem. In this paper, the RPA is addressed in the context of a financial institution. The problem consists in assigning transactions to software robots, whereas each type of transaction has a different clearance day and processing time. First, four types of heuristics are used to compute an upper bound on the number of required software robots. Then, this bound is given as a parameter to an integer linear program, which is used to assign the transactions to the different robots. The quality of the solutions are assessed by an extensive experimental study on a set of 39000 instances. The results show that two heuristics outperform the others and that the LP problem solved with a timeout of 60 seconds allows to find the optimal solution for most of the instances.

Keywords: Robotic process automation, linear integer programming, bin packing, heuristics, upper bound

1   Introduction

Throughout history, new inventions have allowed mankind to achieve ever-better productivity. Nowadays, advanced software that completes complicated tasks is widely available. This allows to automate certain processes in several fields of application. This automation is often achieved by writing codes that can interact and process data directly. In various industries, well-established software relies on human intervention through a Graphical User Interface (GUI) to complete tasks. Such a software can be expensive or unsafe to replace, especially in fields where older, tested and more secure software is preferred over newer, more efficient programs that might have security flaws. Robotic Process Automation (RPA) increases productivity in such situations and is considered a highly promising approach with an increasing number of real-life applications [15]. Indeed, more and more companies rely on this technology to complement their working teams and optimize their internal processes [15]. This form of automation relies on software that mimics human mouse movements and other inputs when interacting with a GUI. These software robots are programmed to perform specific tasks faster and with more accuracy than humans [14, 16]. Furthermore, the software runs continuously and does not feel boredom or exhaustion, which are human risk factors. Many companies now use RPA successfully [2, 6, 9, 11]. A specific field in which RPA is used to automate critical operations is banking. Financial institutions must handle a large volume of transactions daily. Each transaction is associated with a type, which has specific market hours, volumes, clearance dates and processing times. In the context of RPA implementation, these transactions are to be assigned to software robots that will process them. While RPA can lead to better efficiency and accuracy, it comes with a cost. Indeed, each software robot requires a paid license to operate. Therefore, there is a significant incentive to minimize the number of software robots necessary to process a certain number of transactions. As stated by the authors of [8], even though industry is interested in the RPA problem, academic literature is lacking on this specific problem, but recently gained attention [13]. In this paper, the technical implementation of such problem is proposed.

Several recent papers deal with applying RPA on real case studies; all of them highlight a given gain in productivity achieved by the company after implementing RPA. In [1], the authors compare assigning back-office tasks to robots and to humans. Back-office tasks being often tedious and repetitive, when the are assigned to humans, the latter tend to get exhausted or bored after a given time which increases the risks of error and extends the processing times while the robots do not have these drawbacks. Indeed, the research showed a gain of up to 20% in the number of cases treated by the company with the implementation of RPA in complement of their existing teams. In [9], Lacity et al. showed that using RPA at the company Telefònica O2 both enhanced consumer experience and lowered the costs incurred. Another study by Lamberton et al. [10] showed that the number of people required to process claims of an insurance company would be up to 300% larger without the implementation of RPA.

In [12], an integer linear program is proposed to minimize the total costs. Periods are defined at the intersections of market operating hours for all of the transactions, leading to a reduced variable space. A first phase allows to compute the total number of robots required, whereas the second phase assigns the transactions to the robots. Preliminary results on only two test cases led to feasible results, but further experimentation's showed that transactions could not be divided between robots, therefore leading to a higher number of actual robots required. Also, the actual use of a period is not considered, which could lead to having periods where robots are only active, for example, 5% of the available time. Another approach [3] solves a network flow in the first phase and an assignment problem in the second phase. The network flow allows to find a lower bound on the number of robots. The complexity of the problem, but also the computational time to obtain a solution is significantly lower than the first proposed approach [12]. Extensive experiments also show that this method suffers from the indivisibility of the transactions between the robots.

It is also worth mentioning that since RPA is a relatively new technology, many companies still tend to have some difficulty to grasp its effects, the means to implement it in practice and also its challenges. In a attempt to fill that gap, Wewerka and Reichert present in [15] a systematic literature review and a

framework to analyze and compare publications on RPA. That paper highlights the differences between RPA and other similar technologies and details where RPA can be implemented, the tools that are used to implement it, its effects and the future improvements that should be considered in future studies.

The RPA problem presented in this paper consists in assigning financial transactions to robots. Each transaction type is constrained by market operation hours and a clearance date. It also has a volume (number) of transactions to be completed, each of which require a given processing time. The solution to the problem provides a schedule for the robots that minimizes the maximal number of active robots at a given period for different transactions combinations.

This paper is based on a previous work completed by the authors in [12] and describes an improved method to solve the RPA assignment problem. To do so, four different heuristics are proposed to compute an upper bound on the total number of required robots. In a second phase, an integer linear program uses the computed upper bound to find an optimal assignment of the transactions to the robots. Numerical results show that the upper bounds found by the heuristics are of good quality and computed very fast.

The paper is organized as follows. Section 2 describes the RPA problem and the proposed linear integer formulation. Section 3 details the proposed heuristics to compute the upper bound on the number of robots. Numerical results are presented in Section 4 and final remarks are presented in Section 5.

2   Problem description

The RPA assignment problem is stated as follows. Several transactions are to be completed using the lowest possible number of robots for each period. A period is defined by the different types of operations that can be processed between the beginning and the end of that period. In other words, periods are defined by the intersections of market operating hours for all of the transactions; if the set of transaction types that can be processed changes at a certain time, a new period is defined.

Each transaction type is characterised by its market hours, clearance date, processing time and volume. The available market hours of a type is what determines the set of periods during which it can be processed. Transactions that share the same characteristics are regrouped under a type, and the number of such operations of the same type is referred to as the volume.

The aim is to assign each transaction to a given period. The number of transactions completed in a period is constrained by the length of the period, the number of robots assigned to the period, as well as the processing time of each operation. Every robot has the time limit of the period it is assigned to at any moment. The objective of the RPA assignment problem is to find an assignment that minimizes the largest number of robots required among all periods, for any given set of operations to complete under their respective clearance dates.

In this paper, the operations are financial transactions and each robot is a software license that incurs a fixed cost. Therefore, minimizing the largest number of robots among all periods reduces license costs. In what follows, a linear integer program is first presented. It aims at improving the previous models studied in [3, 12] and providing a one-phase simpler model.

2.1   Mathematical model

This integer linear program assigns the transactions to the robots and requires an upper bound on the total number of robots. The upper bound N^{ub} is computed using heuristics and is discussed in Section 3.

Sets are defined as follows:

P = set of transaction types,

K = set of periods,

R = set of robots.

Parameters are described as follows:

v_p = volume of transactions of type $p \in P$,

t_p = process time of transaction type $p \in P$,

l_k = length of period $k \in K$,

$w_{kp} = \begin{cases} 1 \text{ if transaction type } p \in P \text{ can be processed in period } k \in K, \\ 0 \text{ otherwise,} \end{cases}$

N^{ub} = upper bound on the number of robots required (calculated with heuristic).

The model variables are as follows:

$y_{rk} = \begin{cases} 1 \text{ if robot } r \in R \text{ is active in period } k \in K, \\ 0 \text{ otherwise,} \end{cases}$

x_{rkp} = number of transactions of type $p \in P$ processed by robot $r \in R$ in period $k \in K$,

N = number of robots.

The model minimizes the maximal number of robots required for all periods:

$$\min \quad N \tag{1}$$

$$s.t. \quad N \geq \sum_{r \in R} y_{rk}, \qquad \forall k \in K, \tag{2}$$

$$\sum_{k \in K} \sum_{r \in R} x_{rkp} = v_p, \qquad \forall p \in P, \tag{3}$$

$$\sum_{r \in R} x_{rkp} \leq w_{pk} v_p, \qquad \forall p \in P, \forall k \in K, \tag{4}$$

$$\sum_{p \in P} x_{rkp} t_p \leq l_k y_{rk}, \qquad \forall k \in K, \forall r \in R, \tag{5}$$

$$\sum_{p \in P} x_{rkp} \geq y_{rk}, \qquad \forall k \in K, \forall r \in R \tag{6}$$

$$N \leq N^{ub} \tag{7}$$

$$y_{rk} \in \mathbb{B}, \qquad \forall k \in K, \forall r \in R, \tag{8}$$

$$x_{rkp} \in \mathbb{R}^+, \qquad \forall k \in K, \forall r \in R, \forall p \in P. \tag{9}$$

Equation (2) are used to minimize the maximal number of robots N. Equation (3) are used to assign every transaction of each type. Equation (4) sets the value of x_{rkp} to 0 when w_{pk} is 0, so that no transaction is assigned outside of its market hours. Equation (5) are used to enforce the respect of the period lengths. They also ensure that x_{rkp} is 0 if y_{kr} is 0, which means that transactions can only be assigned to active robots. Finally, Equation (6) state that a robot can only be active if transactions are assigned to that robot and the bound on the number of robots is given by Equation (7). Variables' domains are given by Equations (8)–(9).

3 Heuristics to compute upper bound

It is worth mentioning that since the RPA assignment problem is NP-complete, any exact method for its resolution would have exponential complexity. The mathematical model presented in Section 2.1 requires a parameter N^{ub} that is an upper bound on the number of actual robots required and is obtained with heuristics which are detailed in this section. Solving the model where N^{ub} is a variable would require an enormous amount of time given the complexity of the problem.

In order to calculate the value of the upper bound, heuristics which are known to provide close to optimal solutions rather fast are used. The aforementioned heuristics provide an upper bound on the number of robots for the presented mathematical model. To this end, four heuristics are developed. Using the provided upper bound, solving the mathematical model allows to find an optimal assignment of the transactions to the robots.

3.1 Proof of NP-completeness

In this section, it is shown that the RPA assignment problem is NP-complete. In order to do so, a reduction from the well-known bin packing problem to the RPA assignment problem is proposed.

3.1.1 The bin packing problem

The bin packing is an optimization problem that is known to be NP-hard [7]. In this problem, p indivisible packages of sizes $s_1, s_2, ..., s_P$ must each be assigned to bins. The objective is to assign all packages using the least possible number of bins of equal capacity c. The Figure 1 presents an example of such a problem where five packages must be placed in three bins.


Figure 1: Example of a bin packing problem.

3.1.2 Proof of NP-completeness

Proof. It is clear the RPA assignment problem is in NP. Indeed, any solution can be verified in polynomial time. Now, consider a restriction of the RPA assignment problem where every transaction type has a volume of one and can be processed in a single period of length L. Then, the objective is to find the minimum number of robots in the period that can complete all transactions. Each robot can work for at most L units of time, and each transaction has an associated processing time requirement $t_1, t_2, ..., t_P$. This restriction of the RPA assignment problem is equivalent to a bin-packing problem for which each package of size $s_p : p \in P$ is a transaction with a processing time $t_p : p \in P$, and where a new bin with capacity c is equivalent to a robot active in the period if $L = c$.

A visual representation of this reduction is presented in Figure 2. By comparison with Figure 1, it is obvious that the two problems are equivalent. Bins correspond to robots and packages to different transactions. Therefore, the RPA assignment problem is NP-hard. □


Figure 2: RPA problem equivalent to the bin packing example.

3.2   Heuristic algorithms

Four Heuristic methods are tested. Each are inspired by the *First fit* [5] and *First Fit Decreasing* [4] algorithms used to solve the bin packing problem.

1. **First Fit (FF) Algorithm. (Algorithm 1)**
   This algorithm is very similar to the FF used in the offline bin packing problem found in the literature. Transaction types are considered in the order they are given, and periods are considered in chronological order. The entire volume of a certain transaction type is assigned before moving on to another type. Each transaction is assigned to the first possible period that is in the set of valid periods and has enough time to complete it. If there is no period in which a particular transaction can be executed, a new robot is added to all periods. However, a new robot for any period is only used if all previously added robots in all other possible periods are full.

2. **FF with Decreasing Transaction Processing Time Algorithm. (Algorithm 2)**
   The only difference with this algorithm and the FF is that the different transaction types are ordered in decreasing order of processing time before being assigned to a period.

3. **FF with Decreasing Period Length Algorithm. (Algorithm 3)**
   This algorithm is also similar to the FF. However, the periods are ordered in decreasing order of length. This means that, for any given transaction, the algorithm tries to assign it to the longest period in which it can be executed, before checking the other ones.

4. **FF with Decreasing Transaction Processing Time and Decreasing Period Length Algorithm. (Algorithm 4)**
   This algorithm is essentially a mix of both previously described algorithms. It is similar to the first one described, except both the transactions and the periods are sorted in decreasing order.

4   Experimental results

This section details how the synthetic test cases are generated and describes the testing environment for the different algorithms. Numerical results are then presented, followed by a discussion.

4.1   Synthetic test cases generation

The RPA assignment problem in this study is applied in the context of a financial institution. A bank in North America provided a real test case with 12 transaction types. In order to further assess the quality of our methods, synthetic test cases are generated based on the distribution of data of the real test case. For each synthetic transaction type generated, the market operating hours and the clearance date were selected randomly from those of the set of transaction types of the real case. As for the processing time and volume, bounds were randomly selected from the real case, then random values were generated for both attributes, while being limited by the upper and lower bounds previously selected.

Synthetic cases with between 12 and 50 transaction types are generated in order to observe the impact of the instance sizes on both the solution quality and the running time of methods. In order to obtain meaningful results, 1000 instances are generated for each number of different transaction types ranging from 12 to 50. A total of 39 000 instances are used to test the proposed methodology.

All tests are conducted on an "E2-Standard" Google Cloud Platform virtual machine. It consists of 4 virtual CPUs, 16 GB of main memory and an Intel Haswell platform. The operating system used is Debian 10 (Buster). The C programming language is used to implement the heuristics, wihthout parallelism or multhithreading. It is worth mentioning that computing time would be reduced if parallelism, multithreading or another sorting algorithm was used. For this paper, the qsort function is used. For the integer linear program, Python programming language and Xpress solver are used.

4.2   Numerical results and analysis

As this study is based on previous work, the number of robots and calculation times are of most importance, given that the assignment of the transactions to the robots has been assessed before. Therefore, numerical results focus on the these elements.

First, the heuristics are used to determine an upper bound on the number of robots that is then used as a parameter in the optimization model. Each heuristic is run until it finds a solution and the linear integer program is first limited to a 60 seconds timeout in order to find a solution quickly. For the instances that did not reach optimality after 60 seconds, the timeout is set to 15 minutes and the instances are re-solved.

The different heuristic methods provided good quality results in general. Figure 3 shows the median deviation percentages of the heuristics compared to the linear program (LP), more precisely the difference, in percentage, between the number of robots found by the heuristics and the optimal number of robots in the LP. For the sake of fairness, only the instances for which optimal solutions were obtained by the LP were considered. It can be seen that the median deviation percentages does not exceed 20% which provides a rather good upper bound that allows to speed up the resolution of the LP. The best heuristic is Algorithm 4, which is FF with decreasing transaction processing time and period length. The worst heuristic is the FF algorithm and results show that Algorithm 4 greatly improves the bound.
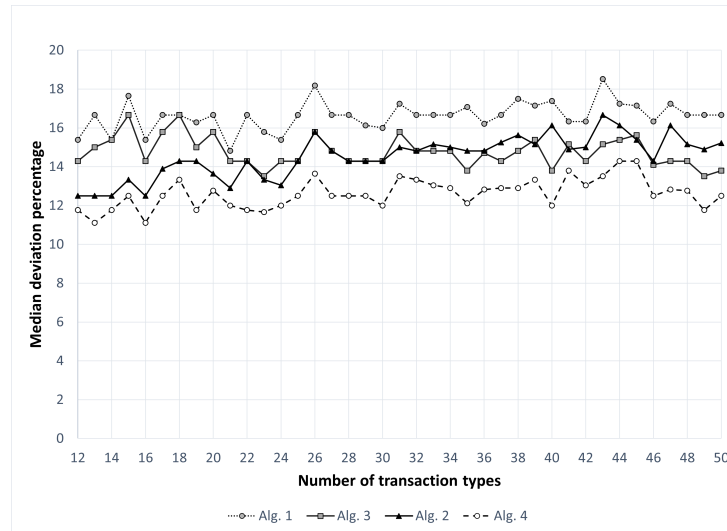

Figure 3: The median deviation percentages for all the solution methods.

Furthermore, the heuristics are, in some cases, able to find the optimal solution to several instances. Figure 4 presents the success rates of the different methods, *i.e.* the percentage of instances for which optimal solutions are found by each method. Even if this rate decreases with the increase in the number of transaction types, heuristics are still able to provide good quality upper bounds that become good starting points for the resolution of the LP. On all instances, the best algorithm is Algorithm 4 as it is the only one that is non-dominated.

It is worth mentioning that for a small percentage of instances, the LP was not able to find optimal solutions and in rarer cases, was not able to find any solution, even after 15 minutes. The percentages are presented in Figures 5 and 6. In the case of large instances, allowing a 15 minutes timeout reduces the percentage of instance with no optimal solution from 14 % to about 9 %. It is also worth mentioning that the LP almost always finds a solution after 15 minutes, and that less than 0.4 % of instances are left without a solution.
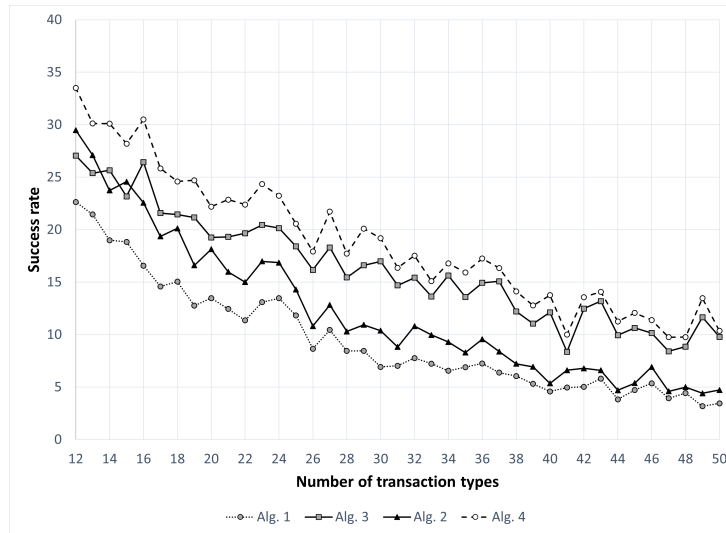
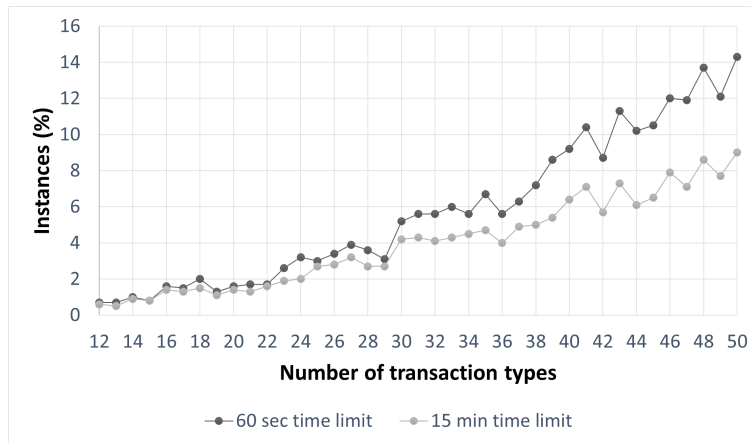**Figure 4: The success rate of the different solution methods.**

Figure 5: Percentage of instances with non-optimal solutions found by the LP solver. The dark line is after 60 seconds and the pale line after 15 minutes.

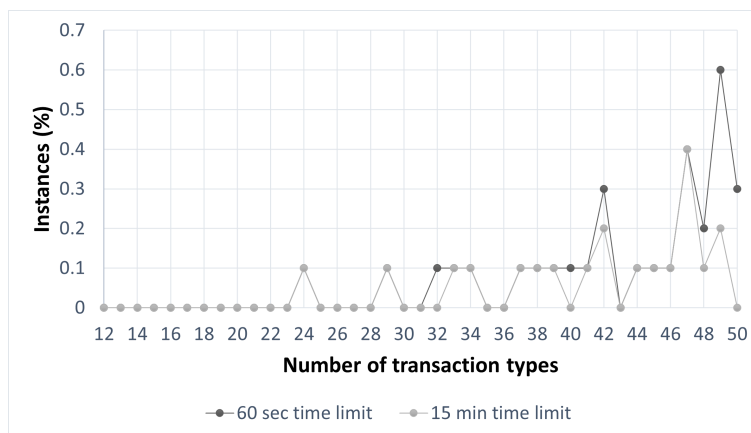
Figure 6: Percentage of instances with no solutions found by the LP solver after 60 seconds (dark line) and 15 minutes (pale line).

In the cases where non-optimal solutions are found, their gap does not exceed 7% as shown in Figure 7. For the cases where no solutions are found by the LP, the heuristics are a good alternative since they provide good quality solutions in a very short time.
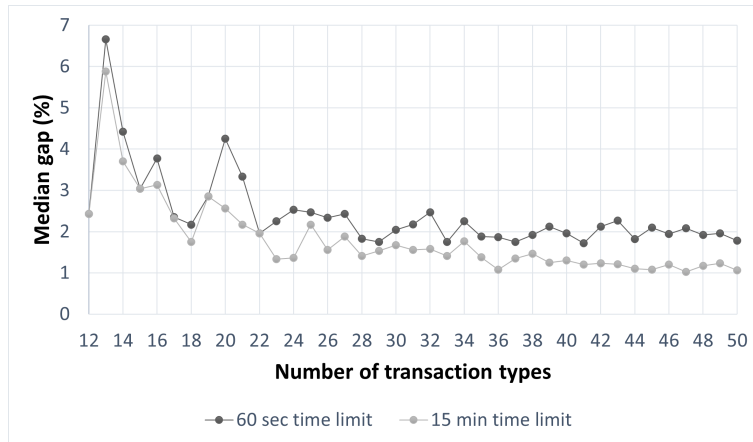
Figure 7: Median optimality gaps for the LP solver after 60 seconds (dark line) and 15 minutes (pale line).

For the LP, the running time is almost instant for 12 transaction types and 2.5 s for 50 instances. Of course, the time grows exponentially with the number of transactions. Concerning the running times for the heuristics, they can be viewed on Figure 8. Algorithm 1 and Algorithm 2 are at least two times faster than Algorithm 3 and Algorithm 4. The heuristics are three to six orders of magnitude faster than the LP.
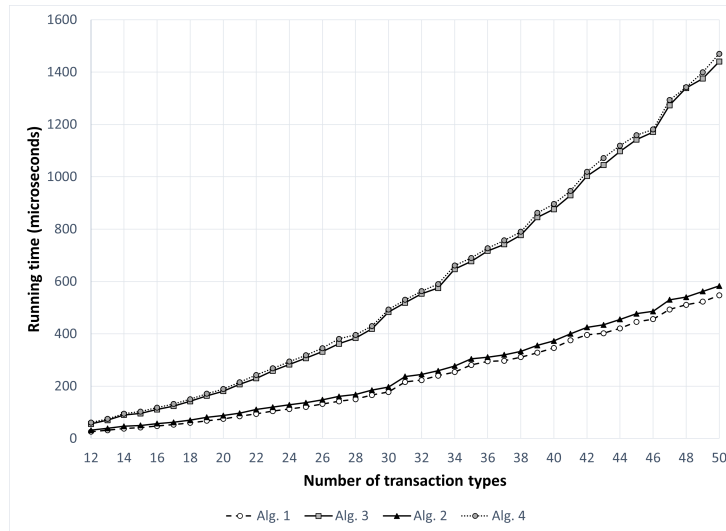

Figure 8: The median running times of heuristics in microseconds.

Figure 9 presents the median differences in the number of robots between the heuristics and the LP for the instances where an optimal solution is found. As expected, this difference increases with the increase in the number of transaction types. Algorithm 4 is the best algorithm on large instances.

As previously mentioned, the solver for the LP was first run with a 60 seconds timeout, if no optimal solution is found, the timeout is set to 15 minutes. In most cases, the number of robots is improved by only one unit which represents less than 0.5% improvement in average. This leads to the conclusion that 60 seconds of timeout is enough since the heuristics provide a good quality upper bound which, combined with the LP, allows to obtain quite good solutions in a very short time.

Figure 10 shows the median optimality gap differences between solutions found after 60 seconds and after 15 minutes. The majority of instances sees a reduction of 0.5% in the optimality gap. This shows that running the LP solver with a 60 second timeout is sufficient. Also note that the heuristic used to compute the bound is the Algorithm 4 for this comparison.
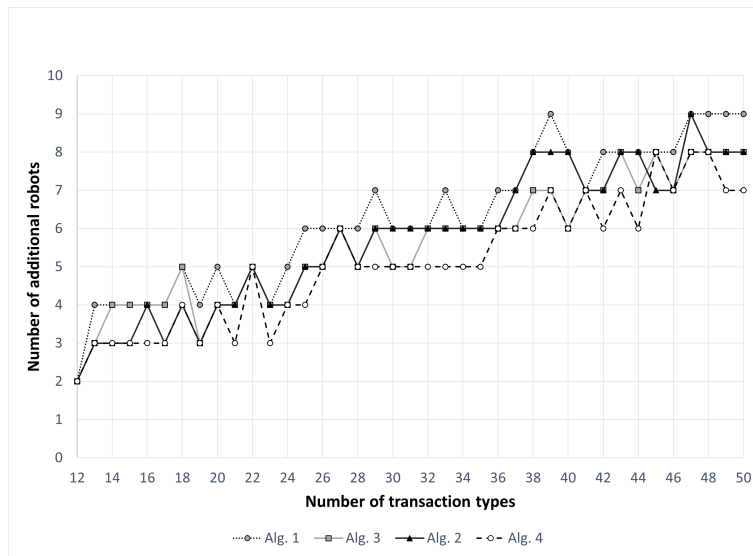
Figure 9: The median difference of the number of robots with respect to the optimal solution.
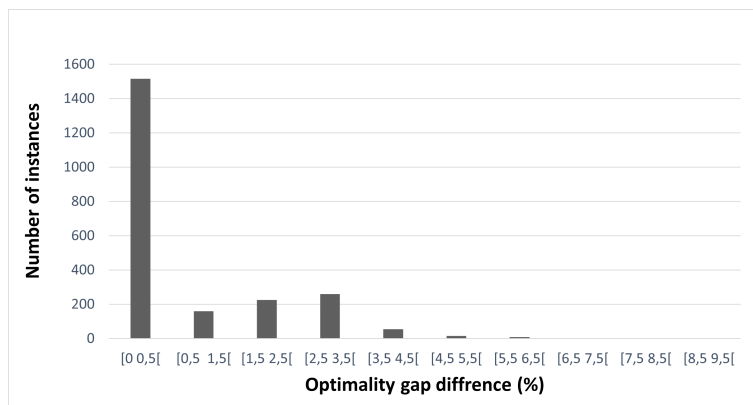

Figure 10: The improvement in the optimality gap between the 60 seconds and the 15 minutes runs of the LP solver.

4.2.1 Analysis

Results show that heuristics allow to obtain a good upper bound on the number of robots for the linear integer program. Since the heuristics run in microseconds, it is worthwhile to compute a bound before conducting the optimization. As for the LP, setting a 60 second timeout is sufficient, given that 5.75% of the instances do not lead to an optimal solution in the given timeout. Of the total 39 000 instances, this represents about 2243 instances. For most of the instances without an optimal solution in 60 seconds, the optimality gap is around 2% for the LP. The instances with no feasible solutions after both timeouts are rather rare, representing only 0.6 % of the total number of instances. Also, it is worth mentioning that using a 15 minute timeout leads to a reduction in the number of non-feasible instances as shown in Figure 6.

Algorithm 3 and Algorithm 4 are the best since they offer the best bounds when compared to the optimal solution. In an operational perspective, these heuristics should be used before solving the optimization model. Indeed, even if they are slower than Algorithms 1 and 2, they are still very fast as they always run well under one second, even for the largest instances.

5 Conclusion

This paper presents a methodology to solve the RPA problem. First, four heuristics are compared to determine the best algorithm in terms of quality of the solution and second, this bound is used as a parameter in a linear integer optimization problem. A solution to the RPA problem is proposed to assign financial transactions to software robots and 39 000 instances are generated to test the methodology. The results show that the heuristics provide a good bound, allowing to solve the LP problem with a timeout of 60 seconds. The results are conducted on static test cases, but in a dynamic mindset, where transactions arrive in real-time, this methodology could easily be adapted to find a good solution in a short amount of time.

References

[1] S. Aguirre and A. Rodriguez. Automation of a business process using robotic process automation (RPA). In Applied Computer Sciences in Engineering. Springer International Publishing, 2017.

[2] A. Asatiani and E. Penttinen. Turning robotic process automation into commercial success – case opuscapita. Journal of Information Technology Teaching Cases, 6(2):67–74, 2016.

[3] I. Benkalaï, S. Séguin, H. Tremblay, and G. Glangine. Computing a lower bound for the solution of a robotic process automation (rpa) problem using network flows. In 7^{th} International Conference on Control, Decision and Information Technologies, Prague, Czech Republic, June 2020.

[4] G. Dósa. The tight bound of first fit decreasing bin-packing algorithm is ffd (i)$\leq$ 11/9opt (i)+ 6/9. In International Symposium on Combinatorics, Algorithms, Probabilistic and Experimental Methodologies, pages 1–11. Springer, 2007.

[5] G. Dósa and J. Sgall. First fit bin packing: A tight analysis. In 30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013). Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.

[6] J.C. Figueroa-García, E.R. López-Santana, J.L. Villa-Ramírez, and R. Ferro-Escobar. Automation of a business process using robotic process automation (RPA): A case study. In Applied Computer Sciences in Engineering. Cham. Springer International Publishing., 2017.

[7] M.R. Garey and D.S. Johnson. Computers and Intractability : A guide to the theory of NP-completeness. W.H. Freeman and Company, 1979.

[8] P. Hofmann, C. Samp, and N. Urbach. Robotic process automation. Electronic Markets, 30(1):99–106, 2020.

[9] M. Lacity, L. Willcocks, and A. Craig. Robotic process automation at telefonica o2. Technical report, London, US: University of Missouri-St-Louis - The London School of Economics and Political Science, 2015.

[10] C. Lamberton, D. Brigo, and D. Hoy. Impact of robotics, rpa and ai on the insurance industry: Challenges and opportunities. Journal of Financial Perspectives, 4(1):8–20, 2017.

[11] S. Parcells. The power of finance automation. Strategic Finance, 98(6):40–45, 2016.

[12] S. Séguin and I. Benkalaï. Robotic process automation (RPA) using an integer linear programming formulation. Cybernetics and Systems, pages 1–13, 2020.

[13] R. Syed, S. Suriadi, M. Adams, W.Bandara, S. J.J. Leemans, C. Ouyang, A. H.M. ter Hofstede, I. van de Weerd, M. Thandar Wynn, and H. A. Reijers. Robotic process automation: Contemporary themes and challenges. Computers in Industry, 115:103162, 2020.

[14] W.M.P. van der Aalst, M. Bichler, and A. Heinzl. Robotic process automation. Business & Information Systems Engineering, 60(4):269–272, 2018.

[15] J. Wewerka and M. Reichert. Robotic process automation – a systematic literature review and assessment framework. Technical report, Cornell University, 2020.

[16] L. Willcocks, M. Lacity, and A. Craig. The it function and robotic process automation. Technical report, London, US: University of Missouri-St-Louis - The London School of Economics and Political Science, 2015.