A Hybrid Path-Relinking Method for Solving a Stochastic Lot Sizing and Scheduling Problem

P. Amorim[a,*], A. M. Costa[b], B. Almada-Lobo[a]

[a] *INESC TEC, Faculdade de Engenharia, Universidade do Porto, Rua Dr. Roberto Frias, s/n, 4600-001 Porto, Portugal*

[b] *Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, Av. do Trabalhador São-Carlense, 400, CP. 668, São Carlos, SP 13560-970, Brazil*

Abstract

Path-relinking has been used to help solving deterministic problems by exploring the neighbourhood of elite solutions in an intelligent way. In this paper, we present an algorithm that combines a mixed-integer linear solver with a truncated path-relinking method in order to solve a stochastic lot sizing and scheduling problem dealing with perishable products. This supply chain planning problem may be seen as a two-stage stochastic integer problem with complete recourse and first stage integer variables. In the first stage the decision maker decides about the production sequence and the production quantities. Afterwards, the uncertain demand is unveiled and the second stage decisions concerned about inventory usage are taken. The key idea of this method is to take advantage of the possible scenario-based decomposition in an innovative way, which can be generalized to problems with a similar structure. Therefore, path-relinking is used to combine optimised solutions from different scenarios in pursuing good stochastic solutions. Computational results show a clear advantage of the proposed method in solving this stochastic problem when compared to a state-of-the-art mixed-integer solver, especially for the medium and larger instances.

Keywords: Path-Relinking, Mixed-Integer Solver, Stochastic Programming, Lot Sizing and Scheduling, Demand Uncertainty

1. Introduction

Within supply chain planning tasks, the lot sizing and scheduling problem is responsible for determining the size of each production lot and the sequence in which these

*Corresponding author

Email addresses: amorim.pedro@fe.up.pt (P. Amorim), alysson@icmc.usp.br (A. M. Costa), almada.lobo@fe.up.pt (B. Almada-Lobo)

lots are produced in a medium to short term planning horizon. Stochastic lot sizing and scheduling problems appear when parameter's uncertainty is taken into account. Lot sizing and scheduling problems are known to be NP-hard and, therefore, provably optimal solutions are difficult to obtain even for medium-sized instances. Moreover, the stochastic version of this deterministic problem, where demand is unknown in advance, turns the problems to be solved even harder. We focus on the stochastic lot sizing and scheduling problem of perishable products (S-LSP-PP) arising in the food consumer goods industries that use *direct store delivery*. The main impact of this delivery practice is an increasing responsibility of producers for all downstream processes of the supply chain until reaching the final customer. This responsibility entails a heavier burden in case the companies produce perishable products that have a limited shelf-life, inhibiting the abusive use of intermediate storage to hedge against demand uncertainty. This problem is presented in more detail in Amorim et al. (2012).

The S-LSP-PP can be classified as a two-stage stochastic integer problem with complete recourse and first stage integer variables. In this problem structure, first stage decisions are responsible for the production planning, with lot sizing and scheduling decisions, whereas second stage variables allocate the inventory and production output to the uncertain demand. This work aims to propose a hybrid solution procedure to solve the S-LSP-PP, which can be generalized to other problems with a similar model structure. This algorithm combines a truncated path-relinking method and a mixed-integer programming-based method. The truncated path-relinking starts the same path of a standard path-relinking, but stops the search in between the starting and guiding solutions. We take advantage of the inherent structure of this kind of problems in order to develop a hybrid heuristic approach. This procedure combines deterministic integer solutions coming from each demand scenario in order to approach the global optimum to the stochastic model.

Stochastic mixed-integer problems have been used to formulate numerous planning problems and a multitude of solution approaches are available to solve them. In the review of Bianchi et al. (2008) an exhaustive coverage of exact methods and metaheuristics used to solve such problems is performed. Basically, almost every traditional metaheuristic and exact method has been adapted to solve these hard problems, such as tabu search (Costa & Silver, 1998), ant colony optimization (Gutjahr, 2003) or branch and bound (Gutjahr & Strauss, 2000). Among all these solution methods, some take advantage of the scenario sampling structure that characterizes a large portion of these problems. Our hybrid path-relinking positions itself along with these methods that use a scenario-based decomposition. A paradigmatic example of this kind of algorithms is the Progressive Hedging that was first proposed by Lokketangen & Woodruff (1996). This method considers a set of representative scenarios that grasp the stochasticity in the second-stage parameters. Each of these scenarios is solved by means of a deterministic model that captures the related sub-problem. By the end of this first step, a pool of solutions optimizing independently the scenarios is obtained. Afterwards, through an

averaging procedure of all considered solutions, a compromise is obtained for all possible uncertain outcomes. Several improvements and enhancements of the base algorithm have been suggested and the readers are referred to Watson & Woodruff (2010) for an updated overview of this method. Our proposed method distinguishes itself from the mentioned ones by handling through a different methodology, path-relinking, the combination of individual scenario solutions.

The remainder of this paper is as follows. In Section 2, a mathematical description of the S-LSP-PP is given. The formulation will be used afterwards by the hybrid path-relinking described in Section 3. In Section 4 the computational results are reported. Finally, Section 5 resumes the main findings and points out paths for future research.

2. Formal Problem Description

Consider products $i, j = 1, ..., N$ that are to be scheduled on $l = 1, ..., L$ parallel production lines over a finite planning horizon consisting of macro-periods $t = 1, ..., T$ with a given length. The changeover time and cost between products on a line is dependent upon the sequence.

A macro-period is divided into a fixed number of non-overlapping micro-periods with variable length. Since the production lines can be independently scheduled, this division is done for each line separately. Let $[S_{lt}]$ denote the set of micro-periods $s = 1, ..., S_{lt}$ belonging to macro-period t and production line l. The number of micro-periods of each macro-period defines the upper bound on the number of products to be produced on each line. The length of a micro-period is linked to the decision variable accounting for the quantity of products produced. A product lot may continue over several micro and macro-periods since setup carry-over is considered. Thus, a lot is independent of the discrete time structure of the macro-periods. Note that this lot sizing and scheduling time structure is based on the general lot sizing and scheduling structure for parallel lines (Meyr, 2002).

The demand for each product j at its fresher state in macro-period t (d_{jt}^{0v}) is stochastic and obtained through the sampling of discrete scenarios $v = 1, ..., V$. Each of these scenarios has an associated probability ϕ^v, such that $\phi^v > 0$, $\forall v$ and $\sum_v \phi^v = 1$. In Amorim et al. (2012), the authors study the influence of different consumer purchasing behaviours reflected in the demand profile over the age of the perishable products (until reaching the shelf-life, u_j). The present study focuses on a linear demand shape for customers that are rather sensitive to product freshness. With equation (1) the demand parameter for product j with age a in period t according to scenario v (given by d_{jt}^{av}), is calculated.

$$d_{jt}^{av} = d_{jt}^{0v} - \frac{0.5 d_{jt}^{0v} a}{u_j - 1}. \tag{1}$$

Consider the following indices, parameters, and decision variables.

Indices

l	parallel production lines
i, j	products
t	macro-periods
s	micro-periods
a	age (in macro-periods)
v	scenarios

Parameters

$[S_{lt}]$	set of micro-periods s within macro-period t and line l
C_{lt}	capacity (time) of production line l available in macro-period t
a_{lj}	capacity consumption (time) needed to produce one unit of product j on line l
c_{lj}	production costs of product j (per unit) on line l
p_j	price of each product j sold
$\bar{p}_j$	cost incurred for each product j spoiled
u_j	shelf-life duration of product j right after being produced (time)
m_{lj}	minimum lot size (units) of product j when produced on line l
$s_{lij}(\tau_{lij})$	sequence dependent setup cost (time) of a changeover from product i to product j on line l
y_{lj0}	equals 1, if line l is set up for product j at the beginning of the planning horizon (0 otherwise)

First Stage Decision Variables

q_{ljs}	quantity of product j produced in micro-period s on line l
y_{ljs}	equals 1, if line l is set up for product j in micro-period s (0 otherwise)
z_{lijs}	equals 1, if a changeover from product i to product j takes place on line l at the beginning of micro-period s (0 otherwise)

Second Stage Decision Variables

w^{av}_{jt}	inventory of product j with age a available at macro-period t in scenario v
ψ^{av}_{jt}	quantity of product j with age a delivered at macro-period t in scenario v

We denote a given set $\{1, 2, ..., M\}$ as $[M]$. Further note that variables w^{av}_{jt} and ψ^{av}_{jt} are only instantiated for a certain domain to ensure that no perished product is kept in stock or used to fulfil demand. Hence, the dynamic set $[A_w] = \{a \in \mathbb{Z}^+ | a \leq \min\{u_j, t-1\}\}$ and the set $[A_\psi] = \{a \in \mathbb{Z}^+ | a \leq \min\{u_j - 1, t-1\}\}$ are used depending on the corresponding decision variable. Without loss of generality, we assume that both inventory at the beginning and at the ending of the planning horizon are null.

S-LSP-PP

$$\max \sum_{j,t,a,v} \phi^v (p_j \psi_{jt}^{av} - \bar{p}_j w_{jt}^{u_j,v}) - \sum_{l,i,j,s} s_{lij} z_{lijs} - \sum_{l,j,s} c_{lj} q_{ljs} \tag{2}$$

subject to:

$$\sum_{a} \psi_{jt}^{av} \leq d_{jt}^{0v} \quad \forall j \in [N], t \in [T], v \in [V] \tag{3}$$

$$\psi_{jt}^{av} \leq d_{jt}^{av} \quad \forall j \in [N], t \in [T], a \in [A_\psi], v \in [V] \tag{4}$$

$$w_{jt}^{av} = (w_{j,t-1}^{a-1,v} - \psi_{j,t-1}^{a-1,v}) \quad \forall j \in [N], t \in [T+1], a \in [A_w] \setminus \{0\}, v \in [V] \tag{5}$$

$$\sum_{l,s \in [S_{lt}]} q_{ljs} = w_{jt}^{0v} \quad \forall j \in [N], t \in [T], v \in [V] \tag{6}$$

$$q_{ljs} \leq \frac{C_{lt}}{a_{lj}} y_{ljs} \quad \forall l \in [L], j \in [N], t \in [T], s \in [S_{lt}] \tag{7}$$

$$\sum_{i,j,s \in [S_{lt}]} \tau_{lij} z_{lijs} + \sum_{j,s \in [S_{lt}]} a_{lj} q_{ljs} \leq C_{lt} \quad \forall l \in [L], t \in [T] \tag{8}$$

$$\sum_{j} y_{ljs} = 1 \quad \forall l \in [L], t \in [T], s \in [S_{lt}] \tag{9}$$

$$q_{ljs} \geq m_{lj}(y_{ljs} - y_{lj,s-1}) \quad \forall l \in [L], j \in [N], t \in [T], s \in [S_{lt}] \tag{10}$$

$$z_{lijs} \geq y_{li,s-1} + y_{ljs} - 1 \quad \forall l \in [L], i, j \in [N], t \in [T], s \in [S_{lt}] \tag{11}$$

$$\psi_{jt}^{av},\ w_{jt}^{av},\ q_{ljs},\ z_{lijs} \geq 0;\ y_{ljs} \in \{0, 1\} \tag{12}$$

Objective function (2) maximizes the planning horizon profit by subtracting to the revenue of the sold products the spoilage, the setup and the variable production costs. The quantity of spoiled products that reach the end of the shelf-life without being sold is obtained in a straightforward manner with the proposed formulation since it corresponds to the inventory that reaches an age of u_j (it is given by $w_{jt}^{u_j,v}$).

Equations (3) do not allow the quantity of sold products of a given age to be above the demand curve derived from the customer willingness to pay. Equations (4) force the sum of all sold products of different ages not to exceed the total demand for the product with the fresher state for each scenario. Hence, a strong assumption is made in this formulation regarding the perfect control of the retailer over the inventory available

to customers. The discussion about other inventory policies is out of the scope of this work.

Constraints (5) establish the inventory balance constraints that are modelled here in a propagation form that updates the age of the inventories throughout the planning horizon for each demand scenario and discounts the products sold. Equations (6) establish that the production done in a certain macro-period on all lines for a given product is equal to the available stock of that same product with age 0 (maximum freshness). These constraints link the production planning with the demand fulfilment requirements and, thus, the first and the second stage decision variables of the stochastic model. Constraints (7) ensure that in order to produce a certain product, the necessary setup is performed. Moreover, at each moment only one product may be produced on a certain line (9) and each product lot is subject to a minimum lot size (10). Constraints (8) limit the use of the capacity with setups and production in each macro-period and, finally, constraints (11) are responsible for tracking the changeover between products. Note that the integrality condition of variables z_{lijs} is not necessary.

3. The Hybrid Path-Relinking Method

Path-relinking is a method that explores the neighbourhood of elite solutions in a systematic way. This method was proposed by Glover & Laguna (1993) and readers are referred to Resende et al. (2010) to a deeper understanding of this method and its variants, such as forward, backward and mixed path-relinking. The key reasoning behind path-relinking is that good solutions should share a similar structure among them. Thus, in the basic version of path-relinking, two solutions are chosen at each time and the solution elements of one of them are changed through a path that finalizes when the starting solution converts into the target one.

In this paper, we reframe the idea of path-relinking to be applied in stochastic problems represented by different probability scenarios. As mixed-integer linear solvers do not take advantage of the structure of this kind of problems, we tackle heuristically this issue by using a hybridization of path-relinking with a solver that independently explores the deterministic scenarios. Our reasoning is that an optimal solution should potentially lie somewhere in between different scenario solutions. Hence, our first step is to optimize individually the deterministic problem associated with each scenario. Afterwards, the most promising solutions are combined via path-relinking and the intermediate solutions are evaluated by the stochastic objective function. Figure 1 shows the main outline of the method where the first step is performed within a branch-and-bound scheme.

Throughout the algorithm execution two different types of calls to the mixed-integer solver are executed: $solveSCE(v)$ and $solveEVA(\vec{y}_v)$. In $solveSCE(v)$ a scenario v is explored by the solver without regarding the stochastic nature of the problem. The output is the vector $\vec{y}_v$ of integer decision variables obtained from the best solution found. Routine $solveEVA(\vec{y}_v)$ solves the linear problem of the stochastic model derived from fixing

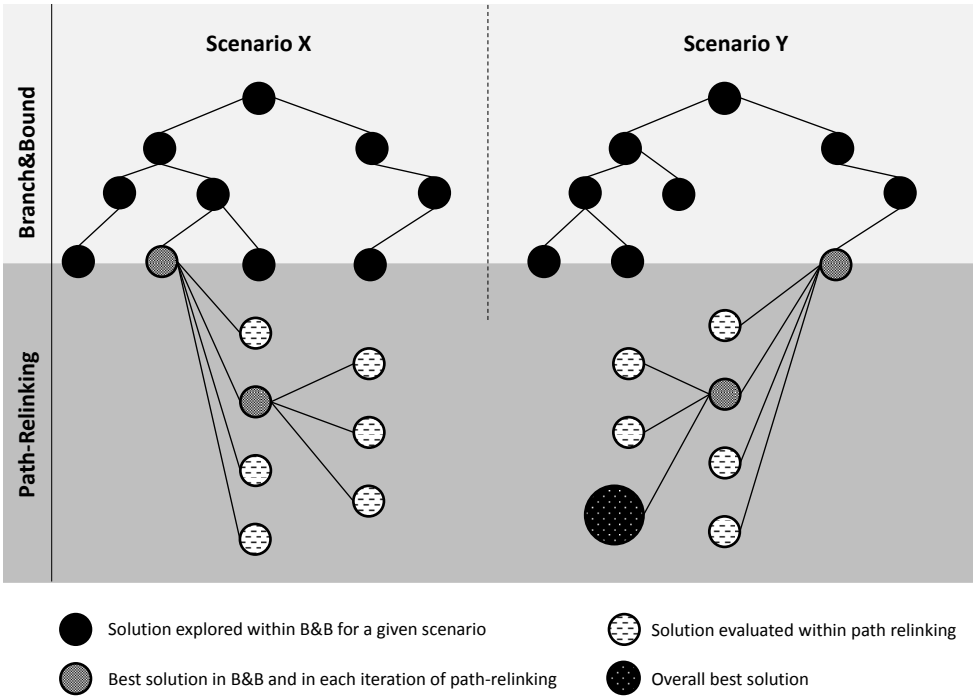


Figure 1: Outline of the hybrid path-relinking.

the first-stage integer variables with the values of $\vec{y}_v$. In the end it returns the stochastic objective function value.

Algorithm 1 describes the pseudo-code for the hybrid path-relinking where these calls are used. First, a feasible solution is found for every deterministic demand scenario within a certain time limit ($solveSCE(v)$). These deterministic solutions are then evaluated in the stochastic scenario ($solveEVA(\vec{y}_v)$). The two best solutions found within all scenarios are then explored within a truncated path-relinking strategy connecting the best solution to the second best. Basically, at each iteration an element of the starting solution is changed into one of the target solution until reaching a degree of resemblance set by the $Stop$ criterion. To calculate this criterion another parameter is used ($DimParam$) that varies between 0 and 1. Hence, the cardinality of the set containing the different positions in which the two considered integer solutions are different is calculated. The position i of the vector of integer decision variables $\vec{y}_v$ is denoted by $y_{i,v}$. This cardinality is then multiplied by $DimParam$ limiting the proportion of resemblance between the two solutions. Therefore, if $DimParam$ has the value of 1, than the truncated path-relinking converts to the standard path-relinking.

Algorithm 1: Pseudo-code for the hybrid path-relinking

for $v = 1$ *to* V **do**
  $\vec{y}_v := solveSCE(v)$;
$max := \arg\max_{v\in[V]}\{solveEVA(\vec{y}_v)\}$;
$max2 := \arg\max_{v\in[V]\setminus\{max\}}\{solveEVA(\vec{y}_v)\}$;
$\vec{\delta} = \{i = 1, ..., |\vec{y}_v| : y_{i,max} \neq y_{i,max2}\}$;
$Stop = |\vec{\delta}| \times DimParam$;
while $|\vec{\delta}| > Stop$ **do**
  for $i \in \vec{\delta}$ **do**
    $\vec{y}_{max} = \{..., y_{i-1,max}, y_{i,max2}, y_{i+1,max}, ...\}$;
    $Obj := solveEVA(\vec{y}_{max})$;
    if $Obj > ObjMax$ **then**
      $ObjMax = Obj$;
      $i_{max} = i$;
  $\vec{\delta} = \vec{\delta} \setminus \{i_{max}\}$;
  if $ObjMax > ObjBest$ **then**
    $ObjBest = ObjMax$;
Output: ObjBest

4. Computational Study

In this section the performance of the hybrid path-relinking method is assessed through computational experiments. To test the proposed method we run a C++ implementation of the algorithm with the mixed-integer programming solver CPLEX 12.2 on a PC with 4 GB of RAM and an Intel i5 with 2.67 GHz running on Windows 7.

4.1. Data Generation

A total of 18 instances were systematically generated, following a methodology similar to the one proposed by Haase & Kimms (2000); therefore, L was set to 1. For all products $a_{lj} = 1$, $c_{lj} = 0.5$, $p_j = 2$ and $\bar{p}_j = 2$. Moreover, the machine is set up for product 1 at the beginning of the planning horizon. The number of products N is 5, 10 and 15. The number of macro-periods T is 5, 10 and 20. The number of micro-periods within a macro-period ($|S_{lt}|$) is set at the value of N allowing for all products to be produced in each macro-period with minimum lot-sizes (m_{lj}) of 1 unit. For the setup times between products (τ_{lij}) the interval [2,10] was used for the 15 products (except for the case where $i = j$, where the setup is 0). Shelf-lives (u_j) were generated for all 15 products for each possible planning period length choosing randomly from the interval [1, T]. The setup cost s_{lij} for a changeover from product i to product j on line l is computed as: $s_{lij} = 50\tau_{lij}$. Each element of the initial demand matrix (d_{jt}^{0v}) for the average scenario with 15 products (rows) and 20 macro-periods (columns) was randomly generated on the interval [40, 60]. Afterwards, using equations (1) the complete demand for all ages is created (d_{jt}^{av}). For the possible demand realizations 5 different scenarios were generated by multiplying the average demand elements by 0.6, 0.8, 1.2 and 1.4. We assume that every scenario has the same probability of occurrence and we test two cases: i) all 5 scenarios and ii) only the intermediate scenario and the less extreme ones (multiplying the average demand by 0.8 and 1.2). Finally, the capacity per macro-period Cap_{lt} is determined according to:

$$C_{lt} = \frac{\sum_j d_{jt}^{0v}}{U}, \quad \forall l, t, v,$$

where the capacity utilization U is equal to 0.8. It is important to notice that the utilization of capacity is a rough estimative, as setup times do not influence the computation of C_{lt}. In summary there are:

$$|5, 10, 15| \times |5, 10, 20| \times |3, 5| = 18 \text{ instances.}$$

4.2. Parameters Tuning

The proposed hybrid path-relinking method only requires a few parameters to be set. The time limit for solving each scenario ($solveSCE$) depends on the number of products handled in each instance. Hence for instances with 5, 10 and 15 products, $solveSCE$ runs for 30, 60 and 180 seconds, respectively. Routine $solveEVA$ runs the evaluation of each solution until optimality. For the parameter $DimParam$, which is responsible for controlling how truncated the path-relinking is,the value of 0.7 is adopted.

Table 1: Results of the computational study

Instance	**Path Relinking**		**Cplex**			
NxTxV	**Solution**	**Gap**	**Solution**	**Gap**	**Running Time**	**UB Diff.**
5x5x5	1933.36	14%	1918.05	15%	178	0.80%
5x10x5	4496.41	9%	4519.17	8%	206	-0.50%
5x20x5	8346.66	13%	8465.46	11%	266	-1.40%
					216.67	**-0.37%**
10x5x5	3963.86	21%	3761.99	27%	384	5.37%
10x10x5	8620.64	15%	8185.85	21%	670	5.31%
10x20x5	15879.03	20%	14926.09	28%	2059	6.38%
					1037.67	**5.69%**
15x5x5	5771.56	26%	4174.62	75%	1365	38.25%
15x10x5	11870.01	26%	10645.59	40%	2196	11.50%
15x20x5	20878.68	37%	18670.28	53%	5255	11.83%
					2938.67	**20.53%**
5x5x3	1059.82	20%	1070.22	19%	112	-0.97%
5x10x3	2550.74	14%	2622.72	11%	146	-2.74%
5x20x3	4774.37	21%	4797.97	20%	205	-0.49%
					154.33	**-1.40%**
10x5x3	2097.5	38%	1996.46	45%	262	5.06%
10x10x3	4808.97	24%	4255.96	40%	449	12.99%
10x20x3	8845.34	32%	8293.97	41%	1706	6.65%
					805.67	**8.23%**
15x5x3	3039.13	46%	2504.94	77%	754	21.33%
15x10x3	6409.11	40%	5444.13	65%	851	17.73%
15x20x3	11647.11	50%	9444.59	85%	6211	23.32%
					2605.33	**20.79%**

4.3. Results

To assess the performance of our algorithm, we compared it to the performance of CPLEX. The hybrid path-relinking ran with the specified parameters for every instance and the execution times of each solution were fed back to CPLEX, in order to obtain the best solution found in the same execution time. The final results for the 18 instances are presented in Table 1. For each instance, the best solution (**Solution**), the integrality gap (**Gap**) that is calculated with the lower bound given by CPLEX and the execution time (**Running Time**) are presented for each solution method (**Path Relinking** and **CPLEX**). Furthermore, the relative difference between the solutions of both methods is computed (**UB Diff.**). The calculation of this indicator is obtained by subtracting the CPLEX solution value to the Path Relinking and dividing this difference by the Path Relinking solution value. Finally, averages for the instances clusters are presented in bold.

Results point out that the proposed method outperforms CPLEX for all the instances with 10 and 15 products. Whereas, for instances with 5 products, for the same execution

time, CPLEX is able to find for the majority of the instances better integer solutions. It seems also that the advantage of the proposed method increases with the size of the instances. In fact, for instance with 15 products, the hybrid path-relinking is able to improve on average by more than 20% the solution obtained by CPLEX.

5. Conclusions and Future Work

The main contribution of this work lies on the exploration of hybrid methods to solve complex stochastic problems, such as the stochastic lot sizing and scheduling of perishable goods. Specifically, we propose a novel hybridization of a truncated path-relinking with a mixed-integer solver that takes advantages of the special structure of this problem. From a reasonable upper bound for each demand scenario separately, a path-relinking method works on the most promising solutions from the stochastic point of view. Since we are dealing with a two-stage stochastic model with only linear second stage variables, each iteration evaluation of the path-relinking is rather fast. Results show the increased advantage of this method over CPLEX alone.

Future work is to be performed in three directions. Improve the decision on which starting and guiding solutions to use, reduce the need of the MIP solver in order to improve scalability and, finally, assess the performance of other path-relinking strategies embedded in this solution scheme.

Acknowledgements

The first author appreciate the support of the FCT Project PTDC/EGE-GES/104443/2008 and the FCT Grant SFRH/BD/68808/2010.

References

Amorim, P., Costa, A., & Almada-Lobo, B. (2012). Influence of consumer purchasing behaviour on the production planning of perishable food. *Submited to ORSpectrum*, (pp. –).

Bianchi, L., Dorigo, M., Gambardella, L. M., & Gutjahr, W. J. (2008). A survey on metaheuristics for stochastic combinatorial optimization. *Natural Computing*, *8*, 239–287.

Costa, D., & Silver, E. (1998). Tabu search when noise is present: An illustration in the context of cause and effect analysis. *Journal of Heuristics*, *23*, 5–23.

Glover, F., & Laguna, M. (1993). Tabu search. In C. Reeves (Ed.), *Modern Heuristic Techniques for Combinatorial Problems* (pp. 70–141). Blackwell Scientific Publishing.

Gutjahr, W. (2003). A converging ACO algorithm for stochastic combinatorial optimization. *Stochastic algorithms: Foundations and applications*, (pp. 10–25).

Gutjahr, W., & Strauss, C. (2000). A stochastic branch-and-bound approach to activity crashing in project management. *INFORMS Journal on Computing*, (pp. 1–22).

Haase, K., & Kimms, A. (2000). Lot sizing and scheduling with sequence-dependent setup costs and times and efficient rescheduling opportunities. *International Journal of Production Economics*, *66*, 159–169.

Lokketangen, A., & Woodruff, D. L. (1996). Progressive hedging and tabu search applied to mixed integer (0, 1) multistage stochastic programming. *Journal of Heuristics*, *128*, 111–128.

Meyr, H. (2002). Simultaneous lotsizing and scheduling on parallel machines. *European Journal of Operational Research*, *139*, 277–292.

Resende, M. G., Ribeiro, C. C., Glover, F., & Martin, R. (2010). Scatter search and path-relinking: Fundamentals, advances, and applications. In M. Gendreau, & J.-Y. Potvin (Eds.), *Handbook of Metaheuristics* (pp. 87–107). Springer US volume 146 of *International Series in Operations Research & Management Science*.

Watson, J.-P., & Woodruff, D. L. (2010). Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems. *Computational Management Science*, *8*, 355–370.