MIP-Based heuristics for the CLSP

Performance analysis and additions

Tom Dogger (302086)

July 2009

Erasmus School of Economics

Erasmus University Rotterdam

Bachelor thesis

Econometrie & Besliskunde

Logistic major

Supervisor: Dr. W. van den Heuvel

Co-supervisor: Prof.dr.ir. R. Dekker

Summary

The Capacitated Lot Sizing Problem (CLSP) is an extension to the Economic Lot Sizing Problem (ELSP), introducing production capacity restrictions to the model. As this problem has been proven to be NP-hard, heuristics have been developed to estimate solutions.

This paper focuses on MIP-based heuristics for the CLSP. Two existing MIP-heuristics (LP-and-Fix and Relax-and-Fix) are tested on the CLSP and extensions and improvements to these heuristics are introduced.

Improved versions of the Relax-And-Fix heuristic yield the best results with step and integer size up to 15 and 30 respectively, even with a less tight formulation of the CLSP.

Table of contents

Introduction

The Economic Lot Sizing Problem (ELSP) is a well studied problem, first studied by Wagner and Within (1958). It is an extension to the Economic Order Quantity (EOQ) model. The ELSP allows several products to be produced on the same machine. This means there are two decisions to be made: whether to produce in a certain period and if so, how much. This creates a Mixed Integer Problem (MIP) with the objective to minimize the total costs.

The Capacitated Lot Sizing Problem (CLSP) is an extension to the ELSP, introducing production capacity restrictions for every period. Bitran and Yanasse (1982) show that several instances of the CLSP are NP-hard. Over time, several algorithms have been developed to obtain optimal solutions for CLSP instances. Chen, Hearn and Lee (1994) describe such an algorithm based on dynamic programming.

Several costs can be identified for the CLSP. For instance, there are costs related to the setup of production in a period and to store inventory between periods. The ultimate goal of the CLSP is to minimize total costs given all capacity, demand and inventory restrictions.

This paper focuses on applying MIP-based heuristics to the CLSP. Two existing heuristics (LP-and-Fix and Relax-and-Fix as described in Pochet and Wolsey (2006)) will be tested on the CLSP. Besides that, several combinations and additions to these heuristics are developed to improve the performance of the existing heuristics.

Next, several tools will be applied to increase the performance of all heuristics. A tighter formulation of the problem, introduced by Krarup and Bilde (1977), will be applied. For this formulation they proved that, without a production capacity restriction, its LP relaxation has an optimal solution with integer production decision variables. Wagelmans et al. (1992) derived a dynamic programming algorithm based on this formulation.

Another method that will be applied to increase the performance of the heuristics is the demand transformation proposed by Bitran and Yanasse (1982). This method is used to reallocate demand to previous periods if the production capacity in the current period does not suffice. The inventory costs involved with this reallocation have been proven to be independent of the decision process and can thus be added to the total costs afterwards.

All the CLSP problem instances in this paper all have varying demand, constant production capacity and costs and constant inventory and production costs.

The main research question of this paper can be defined as follows:

Can new MIP based heuristics be developed for the CLSP?

The following question will be answered as well:

How do existing MIP based heuristics perform on the CLSP?

Mathematical formulation

CLSP formulation

The capacitated lot sizing problem can mathematically be defined as follows:

$$\min \sum_{t=1}^{T} (s_t y_t + h_t i_t + p_t x_t)$$

s.t.

$i_t = i_{t-1} + x_t - d_t$	$t = 1,...,T$	(1)
$i_0 = 0$		(2)
$x_t \leq c_t y_t$	$t = 1,...,T$	(3)
$x_t, i_t \geq 0$	$t = 1,...,T$	(4)
$y_t \in \{0,1\}$	$t = 1,...,T$	(5)

The following interpretations apply for the parameters and variables:

s_t Production setup costs in period *t*

h_t Inventory holding costs in period *t*

p_t Unit production costs in period *t*

d_t Demand in period *t*

c_t Production capacity in period *t*

y_t Decision variable whether production takes place in period *t*

i_t Decision variables for inventory level in period *t*

x_t Decision variable for number of units produced in period *t*

The objective function minimizes the production setup and inventory holding costs over all periods. Constraint (1) implies that the current inventory equals the inventory of the previous period plus the production of the current period minus the demand of the current period. Constraint (2) sets the inventory level before the first period to 0. Constraint (3) makes sure that the production does not exceed the production capacity if production takes place. Constraint (4) keeps the production and inventory level positive and constraint (5) makes the decision variable whether to produce or not binary.

FAL formulation

Krarup and Bilde (1977) suggest an alternative formulation commonly known as the facility location-based (FAL) formulation, for which they have proven that, without the capacity constraint, its LP relaxation has an optimal solution in which all *y* variables are integer. With the capacity constraint added, this is not the case, but it at least results in a tighter formulation. The FAL formulation is defined as follows:

$$\min \sum_{t=1}^{T}(s_t y_t + \sum_{i=t}^{T} c_{ti} X_{ti})$$

$$s.t.$$

$$\sum_{i=1}^{T} X_{it} = d_t \qquad \forall t \qquad (1)$$

$$X_{qt} \le d_t Y_q \qquad \forall q,t \qquad (2)$$

$$\sum_{t_1=t}^{T} X_{t,t_1} \le C_t \qquad \forall t \qquad (3)$$

$$X_{qt} \ge 0 \qquad \forall q,t \qquad (4)$$

$$y_t \in \{0,1\} \qquad \forall t \qquad (5)$$

The following interpretation applies for the parameters and variables:

s_t	Production setup costs for period *t*
c_{qt}	$h_q + h_{q+1} + \ldots + h_{t-2} + h_{t-1}$ $(q \le t)$
	Holding costs for producing for period *t* in period *q*
d_t	Demand in period *t*
C_t	Production capacity in period *t*

y_t	Decision variable whether production takes place in period *t*
X_{qt}	Decision variable deciding number of units produced for period *t* in period *q*

The objective function minimizes the production setup and inventory holding costs over all periods. Constraint (1) makes sure that demand in every period *t* is met (by production in periods up to and including *t*). Constraint (2) ensures that units only get produced when production actually takes place. Constraint (3) is the capacity constraint for every period and constraints (4) and (5) make unit count positive and production binary respectively.

Methodology

The two MIP formulations in the previous paragraph will be used to solve CLSP instances. MIP formulations and MIP-based heuristics are very useful for solving CLSP instances due to the nature of the CLSP. The CLSP has two main decisions, whether two produce in a period or not and if production takes place, how much this should be. The first decision is a binary one whilst the second decision only needs to be nonnegative. As some of the decision variables are required to be binary variables, a MIP is suited to solve this problem.

Instead of using the most common CLSP formulation, a formulation without unit production costs will used by removing the last part of the objective function. For the explanation of the existing and new heuristics, only the implementation on the CLSP formulation will be shown. The implementation on the FAL formulation follows the same reasoning.

Existing heuristics

LP-and-Fix

This first heuristic implemented is the LP-and-Fix heuristic. In this heuristic, the LP version of the initial problem is solved first by dropping integrality constraint (5) in the CLSP formulation, resulting in solution $(\hat{x}, \hat{y}, \hat{i})$. If no integer values are found in $\hat{y}$, the heuristic has failed. If any of the values in $\hat{y}$ are 0 or 1, these values are fixed and the problem is solved with integrality constraint (5) and the fixed values of *y*, resulting in the following mathematical representation ($\hat{y}_t$ indicates the optimal LP *y* variable in period *t).*

$$\min \sum_{t=1}^{T} (s_t y_t + h_t i_t)$$

$$\begin{array}{lll}
s.t. & & \\
i_t = i_{t-1} + x_t - d_t & t = 1,...,T & (1) \\
i_0 = 0 & & (2) \\
x_t \le c_t y_t & t = 1,...,T & (3) \\
x_t, i_t \ge 0 & t = 1,...,T & (4) \\
y_t \in \{0,1\} & t = 1,...,T & (5) \\
y_t = \hat{y}_t & \forall\, t : \hat{y}_t = \{0,1\} & (6)
\end{array}$$

For large problems this problem may still be hard to solve. A time limit of 120 seconds has been enforced to solve a problem, using the best feasible solution found at that moment as the solution. If no solution to this problem is found, the heuristic has failed. If a solution exists, this is the LP-and-Fix heuristic solution.

Relax-and-Fix

The second heuristic implemented is the Relax-and-Fix heuristic. This heuristic divides the problem in *R* sets $Q^1,\ldots,Q^r$ of decreasing importance, dividing the *T* periods into *R* subsets: $1,\ldots,t_1$, …, $t_{R-1}+1,\ldots,t_R$, where t_R equals *T*. After this *R* MIPs, denoted MIP^1 to MIP^r, are solved sequentially.

The size of all R sets is identical. Two sets will be forced to integrality every run. After this, the y values of the first set will be fixed and the second and the third set will be forced to integrality. For this reason, the size of one set will be called the *step size* and the amount of sets forced to integrality multiplies by the step size will be called the *integer size*.

In MIP^1 integrality is forced on the periods in Q^1 and Q^2 and integrality is relaxed on the other variables, which leads to the following formulation:

$$\min \sum_{t=1}^{T} (s_t y_t + h_t i_t)$$

$$
\begin{array}{lll}
s.t. & & \\
i_t = i_{t-1} + x_t - d_t & t = 1,\ldots,T & (1) \\
i_0 = 0 & & (2) \\
x_t \le c_t y_t & t = 1,\ldots,T & (3) \\
x_t, i_t \ge 0 & t = 1,\ldots,T & (4) \\
y_t \in \{0,1\} & t = 1,\ldots,t_2 & (5) \\
y_t \in [0,1] & t = t_2 + 1, \ldots, T & (6)
\end{array}
$$

In MIP^r for $r = 2,\ldots,R$ the y variables in Q^{r-1}, denoted y^{r-1}, are additionally fixed to their optimal values in MIP^{r-1}, integrality is forced on variables in Q^r and Q^{r+1} and integrality is relaxed on the other variables, leading to the following formulation (y_t^{r-1} is the y value for period t found in MIP^{r-1}):

$$\min \sum_{t=1}^{T} (s_t y_t + h_t i_t)$$

$$
\begin{array}{lll}
s.t. & & \\
i_t = i_{t-1} + x_t - d_t & t = 1,\ldots,T & (1) \\
i_0 = 0 & & (2) \\
x_t \le c_t y_t & t = 1,\ldots,T & (3) \\
x_t, i_t \ge 0 & t = 1,\ldots,T & (4) \\
y_t = y_t^{r-1} & t = 1,\ldots,t_{r-1} & (5) \\
y_t \in \{0,1\} & t = t_r,\ldots,t_{r+1} & (6) \\
y_t \in [0,1] & t = t_{r+1} + 1, \ldots, T & (7)
\end{array}
$$

If MIP^r is infeasible for any r, the heuristic has failed. If a solution exists the outcome of MIP^R is the Relax-and-Fix heuristic solution.

New heuristics

LP-Relax-And-Fix

The first new heuristic is a combination of the two existing heuristics. First the LP relaxation of the problem is solved: Any y_t that has an integer value will be fixed (just like the LP-And-Fix method).

Contrary to the LP-And-Fix method, not the whole problem is solved as a MIP, but instead the Relax-And-Fix method is applied to solve the problem, with as additional constraints the y_t variables as supplied by the first part of this heuristic.

RepeatedLP

This method is much like the Relax-And-Fix method, only with a slight change in the relaxed part of the problem. In every run the relaxed part of the model is first solved by an LP. Values of y_t that are integer will be fixed and then the Relax-and-Fix MIP is solved with these integer values as additional constraint. This is done for every MIP in the Relax-and-Fix method with a relaxed part.

Improved Relax-And-Fix

This heuristic is an improvement to the Relax-And-Fix method. In the Relax-And-Fix method, if any MIP^r is infeasible, the heuristic has failed.

In this improved method if MIP^r is infeasible, the variables of Q^{r-1} that were fixed after MIP^{r-1} will now be unfixed and forced to integer again. This means the amount of integer values increases. If the resulting MIP has a solution, the variables in Q^{r-1} and Q^r will be fixed to integer and the Relax-And-Fix method continues. If the resulting MIP is still unfeasible, the heuristic has failed. This is done because if one step is infeasible it is likely to be caused by the variables fixed in the previous run. Unfixing and recalculating these variables, now taking a step size more into account, can fix this problem.

Another method that might improve the Relax-And-Fix method is to change the step and integer size. Initially a step size of 5 periods was chosen, changing this to for example 10 or 20 may have effects on the end solution. Also, initially 2 sets are fixed to integer values and after this 1 set is fixed to the calculated integer value. Initially taking more sets as integer may also have effects on the solution.

Computational results

Setup

Data generation

Test problems as presented in Helmrich (2005) are used in this paper. These problem instances are a combination of several parameters of the problem, like for example the demand pattern, the time horizon and the production capacity. All costs are taken constant over time.

The demand is generated with the following formula:

$$d_t = \mu + \sigma z_t + a \sin\left(\frac{2\pi}{b}\left(t + \frac{b}{4}\right)\right)$$

The following interpretation of parameters applies:

d_t = demand in period t
μ = mean demand
σ = standard deviation of the demand
z_t = standard normal random variable
a = amplitude of seasonal influence
b = length of the seasonal cycle in periods

This formulation was first proposed by Baker (1989) and later used by Chen, Hearn and Lee (1994).

The mean demand μ is always 200. Both a seasonal (a=125) and a non-seasonal (a=0) case are considered. Three time horizons are considered: 12, 30 and 250 periods. For seasonal demand patterns, the length of the seasonal cycle is either 12 or T periods and the standard deviation of the demand is 67. For non-seasonal patterns the standard deviation of the demand is either 67 or 237. If d_t is negative for any t, the demand is set to 0 in that period. For every setting demand is generated 5 times to eliminate randomness.

Production capacity c is taken as constant over time, with values of 250, 400, 600, 800, 1000 and 1200 respectively. Holding costs h are set to €1,- per unit per period and production setup costs are either €100,-, €900,- or €3600,- per setup. Only changing the production setup costs gives an insight in how total costs react when the ratio of holding costs versus production setup costs changes.

In total this creates 2*3*2*5*6*3 = 1080 different problem instances. All heuristics are used on every problem instance.

Optimal solutions

All test cases can be solved to optimality. Chen, Hearn and Lee (1994) describe an algorithm to generate the optimal solution of a test case. Optimal solutions to all the test cases are known and are used to test the performance of all heuristics.

Implementation

All heuristics and test cases are programmed in Java with Eclipse after which they are solved with CPLEX 10.1.

For implementation ease, the formulation of the CLSP was slightly altered when implemented in Java. The inventory i in period 0 has to be 0, to get all the indexes equal in Java, x_0 and y_0 were also created and forced to 0, which leads to the following CLSP formulation:

$$\min \sum_{t=0}^{T} (s_t y_t + h_t i_t)$$

$$s.t.$$

$$i_t = i_{t-1} + x_t - d_t \qquad t = 0,\dots,T \qquad (1)$$

$$x_t \le c_t y_t \qquad t = 0,\dots,T \qquad (2)$$

$$x_t, i_t \ge 0 \qquad t = 0,\dots,T \qquad (3)$$

$$y_t \in \{0,1\} \qquad t = 0,\dots,T \qquad (4)$$

$$i_0 = x_0 = y_0 = 0 \qquad (5)$$

With d_0, s_0, h_0 and c_0 set to 0 as well. This reformulation should not result in any changes in the solutions. The same reasoning applies to the FAL formulation.

Demand smoothing

Bitran and Yanasse (1982) provide a way (including proof) to get a solution for every problem instance that has a solution. This is done by transforming the demand before a heuristic is executed. The transformation starts at the last period. If in period T the demand d_T exceeds the production capacity c_T, the part of the demand that exceeds the production capacity should be produced earlier. This should be done in period T-1, which makes $d_{T-1} = d_{T-1} + (d_T - c_T)$. Obviously holding costs h have to be added to the total costs because supply has to be stored for (at least) 1 period. The total costs thus increase with $h(d_T - c_T)$. Excess demand in period T-1 should be produced in period T-2 et cetera.

If after this transformation $d_0 > 0$ the problem instance is infeasible because not all demand can be produced in time. If d_0 remains zero, an optimal solution for the problem exists. Bitran and Yanasse (1982) provide proof that the additional holding costs should be added to the total costs afterwards and do not need to be included in the decision making process as they are independent of this process.

CLSP results

General

All heuristics have been applied on the 1080 test cases. 4 different versions of the Improved Relax-And-Fixed method were used, all with different step sizes and integer counts, as shown in table 1.

Method name	Step size	Integer size
Relax-And-Fix	5	10
Improved	5	10
Improved 5-15	5	15
Improved 10-20	10	20
Improved 15-30	15	30

Table 1: Different variants of the Improved Relax-And-Fix method.

Table 2 shows an overview of the results of the heuristics. 159 of the test problems were infeasible, leaving 921 test problems to be solved.

Demand Smoothing	Heuristic	Feasible	Optimality	Not solved	Gap	Average gap
No	LP-And-Fix	921	450	369	102	0.41%
	Relax-And-Fix	921	610	0	311	0.98%
	LP-Relax-And-Fix	921	610	0	311	0.98%
	RepeatedLP	921	552	114	255	1.13%
	Improved	921	610	0	311	0.98%
	Improved 5-15	921	711	0	210	0.55%
	Improved 10-20	921	739	0	182	0.36%
	Improved 15-30	921	791	0	130	0.25%
Yes	LP-And-Fix	921	444	369	108	0.40%
	Relax-And-Fix	921	610	0	311	0.98%
	LP-Relax-And-Fix	921	610	0	311	0.98%
	RepeatedLP	921	610	0	311	0.98%
	Improved	921	610	0	311	0.98%
	Improved 5-15	921	711	0	210	0.55%
	Improved 10-20	921	739	0	182	0.36%
	Improved 15-30	921	791	0	130	0.25%

Table 2: Heuristic results

For table 2 and the following tables it holds that the *Feasible* column gives the amount of problem instances that are feasible, the *Optimal* column gives the amount of problem instances solved to optimality, the *Not solved* column gives the amount of problem instances that are feasible but not solved by the heuristic, the *Gap* column gives the amount of problem instances that have a solution which is not optimal and finally the *Average gap* column gives the average gap of the problem instances with a non-optimal solution.

The LP-And-Fix heuristic leaves some problems unsolved. This is due to the fact that if in the first step of the heuristic (the LP) no integer values are found, the heuristic has failed by definition and will not solve the MIP.

The Relax-And-Fix, LP-Relax-And-Fix and Improved heuristics lead to the same results with demand smoothing on and off. RepeatedLP also gives exactly the same results with demand smoothing on.

The RepeatedLP heuristic leaves some problems unsolved when demand smoothing is turned off. This is probably due to the LP's which are performed every step. The LP solved before every step of the Relax-And-Fix procedure is independent of the part of the problem solved (fixed to a value) so far. This may lead to infeasibility because the past is not taken into account. The production capacity of the first period of the LP may be insufficient to produce demand in that period, which makes the LP infeasible.

Looking at the Improved heuristics, the amount of optimal solutions increases when the step size and integer count increase, and the problems with a gap as well as the average gap decrease. This will be discussed in detail in the following paragraphs.

Periods

Table 3 shows the heuristic results split to periods.

		T = 12					T = 30					T = 250				
Demand Smoothing	Heuristic	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap
No	LP-And-Fix	300	117	183	0	-	303	168	135	0	-	318	165	51	102	0.41%
	Relax-And-Fix	300	271	0	29	2.50%	303	226	0	77	1.08%	318	113	0	205	0.73%
	LP-Relax-And-Fix	300	271	0	29	2.50%	303	226	0	77	1.08%	318	113	0	205	0.73%
	RepeatedLP	300	264	9	27	2.66%	303	201	30	72	1.13%	318	87	42	156	0.87%
	Improved	300	271	0	29	2.50%	303	226	0	77	1.08%	318	113	0	205	0.73%
	Improved 5-15	300	300	0	0	-	303	261	0	42	0.99%	318	150	0	168	0.44%
	Improved 10-20	300	300	0	0	-	303	279	0	24	0.53%	318	160	0	158	0.34%
	Improved 15-30	300	300	0	0	-	303	303	0	0	-	318	188	0	130	0.25%
Yes	LP-And-Fix	300	117	183	0	-	303	168	135	0	-	318	165	51	102	0.41%
	Relax-And-Fix	300	271	0	29	2.50%	303	226	0	77	1.08%	318	113	0	205	0.73%
	LP-Relax-And-Fix	300	271	0	29	2.50%	303	226	0	77	1.08%	318	113	0	205	0.73%
	RepeatedLP	300	271	0	29	2.50%	303	226	0	77	1.08%	318	113	0	205	0.73%
	Improved	300	271	0	29	2.50%	303	226	0	77	1.08%	318	113	0	205	0.73%
	Improved 5-15	300	300	0	0	-	303	261	0	42	0.99%	318	150	0	168	0.44%
	Improved 10-20	300	300	0	0	-	303	279	0	24	0.53%	318	160	0	158	0.34%
	Improved 15-30	300	300	0	0	-	303	303	0	0	-	318	188	0	130	0.25%

Table 3: Heuristic results split to periods

LP-And-Fix leaves the most problems unsolved when 12 periods are used. As explained, the heuristic fails by definition if no integer values are found in the LP. This means that in smaller problems less integer values occur in the LP, which is as expected. Bigger problems result in more integer values.

Relax-And-Fix, LP-Relax-And-Fix and Improved yield the same result. Improved is expected to yield the same result as no problem instance with Relax-And-Fix is not solved, thus there is no need to increase the integer count by adding the last solved period. The fact that LP-Relax-And-Fix yields the same result as Relax-And-Fix has two reasons. First of all, the heuristic does not fail if no integers are found in the LP, instead it will just do a regular Relax-And-Fix procedure. However, as the results are exactly the same for every problem, this also indicates that if the LP does find integer values, they are the same values as the Relax-And-Fix method would find.

Small problems (12 periods) are all solved to optimality with an Improved heuristic with an integer count of 15 or higher. This is as expected as in every step of the heuristic a complete integer answer is required. All values are required to be integer but not all values are fixed right away. The same hold for the Improved 15-30 heuristic on medium sized problems of 30 periods.

Increasing the step and integer size in Improved heuristics always leads to a higher amount of optimal solutions, together with a lower amount of solutions with a gap and a lower average gap. This may be caused by the costs of setting up production in a period. This will be discussed thoroughly in the next paragraph.

With demand smoothing enabled, the RepeatedLP heuristic yields the same results as the Relax-And-Fix variants with the same step and integer size. This means that if integer values are found in the LP, these are the same values Relax-And-Fix would find on its own.

Production setup costs

Table 4 shows the heuristic results split to production setup costs.

		F = 100					F = 900					F = 3600				
Demand Smoothing	Heuristic	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap
No	LP-And-Fix	307	184	123	0	-	307	164	123	20	0.18%	307	102	123	82	0.47%
	Relax-And-Fix	307	307	0	0	-	307	170	0	137	0.57%	307	133	0	174	1.30%
	LP-Relax-And-Fix	307	307	0	0	-	307	170	0	137	0.57%	307	133	0	174	1.30%
	RepeatedLP	307	269	38	0	-	307	157	38	112	0.66%	307	126	38	143	1.50%
	Improved	307	307	0	0	-	307	170	0	137	0.57%	307	133	0	174	1.30%
	Improved 5-15	307	307	0	0	-	307	263	0	71	0.13%	307	168	0	139	0.76%
	Improved 10-20	307	307	0	0	-	307	251	0	56	0.08%	307	181	0	126	0.49%
	Improved 15-30	307	307	0	0	-	307	279	0	28	0.04%	307	205	0	102	0.31%
Yes	LP-And-Fix	307	184	123	0	-	307	164	123	20	0.18%	307	102	123	82	0.47%
	Relax-And-Fix	307	307	0	0	-	307	170	0	137	0.57%	307	133	0	174	1.30%
	LP-Relax-And-Fix	307	307	0	0	-	307	170	0	137	0.57%	307	133	0	174	1.30%
	RepeatedLP	307	307	0	0	-	307	170	0	137	0.57%	307	133	0	174	1.30%
	Improved	307	307	0	0	-	307	170	0	137	0.57%	307	133	0	174	1.30%
	Improved 5-15	307	307	0	0	-	307	263	0	71	0.13%	307	168	0	139	0.76%
	Improved 10-20	307	307	0	0	-	307	251	0	56	0.08%	307	181	0	126	0.49%
	Improved 15-30	307	307	0	0	-	307	279	0	28	0.04%	307	205	0	102	0.31%

Table 4: Heuristic results split to costs

LP-and-Fix performs almost equally on all possible production setup costs. The amount of unsolved problems is equal amongst all costs, which indicates that production setup costs have no influence on the *y* values in the LP. The only difference that occurs between the setup costs is that as setup costs increase, more problems are not solved to optimality but are instead solved with a (small) gap (on average not bigger than 0.5%). This could be caused by the fact that when with different setup costs the same amount of production setups is found, costs are always higher with higher setup costs, and consequently if the solution is not optimal (too many setups take place), the optimality gap is bigger with higher setup costs.

The Relax-And-Fix based heuristics perform optimally with low production setup costs (that is, because holding costs are set to €1,-, a relatively small difference between holding and setup costs). When setup costs increase, the amount of problems solved with a gap and the average gap both increase. This may be explained as follows:

When production setup costs increase, the amount of production periods should be small to minimize costs. When in a Relax-And-Fix instance production needs to take place, the integer size limits the view of the problem. It may be optimal to produce for the first period after the integer part, but as the *y* variable of that period is relaxed, this decision is not fixed and because of this, it may be cheaper in the sub problem to produce a period earlier (which is not optimal for the entire solution). When this is fixed, production capacity may no longer suffice to produce in the extra period, which results in extra production taking place. When setup costs are high, the gap to optimality can thus increase more.

Increasing the integer size reduces the average gap. As the integer periods are bigger there is less likelihood that problems as described above occur, which leads to a reduction in average gap.

Production capacity

Table 5 shows the heuristic results split to production capacity.

		C=250					C=400					C = 600					C = 800					C = 1000					C = 1200				
Demand Smoothing	Heuristic	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap	Feasible	Optimal	Not solved	Gap	Average gap
No	LP-And-Fix	42	37	0	5	0.12%	162	87	45	30	0.35%	177	74	81	22	0.36%	180	84	81	15	0.58%	180	83	81	16	0.50%	180	85	81	14	0.45%
	Relax-And-Fix	42	25	0	17	0.21%	162	110	0	52	0.46%	177	111	0	66	0.77%	180	118	0	62	1.28%	180	120	0	60	1.34%	180	126	0	54	1.25%
	LP-Relax-And-Fix	42	25	0	17	0.21%	162	110	0	52	0.46%	177	111	0	66	0.77%	180	118	0	62	1.28%	180	120	0	60	1.34%	180	126	0	54	1.25%
	RepeatedLP	42	3	39	0	-	162	84	54	24	0.65%	177	104	18	55	0.85%	180	115	3	62	1.28%	180	120	0	60	1.34	180	126	0	54	1.25%
	Improved	42	25	0	17	0.21%	162	110	0	52	0.46%	177	111	0	66	0.77%	180	118	0	62	1.28%	180	120	0	60	1.34%	180	126	0	54	1.25%
	Improved 5-15	42	31	0	11	0.12%	162	119	0	43	0.22%	177	135	0	42	0.37%	180	141	0	39	0.57%	180	144	0	36	0.63	180	141	0	39	1.12%
	Improved 10-20	42	32	0	10	0.10%	162	124	0	38	0.17%	177	139	0	38	0.24%	180	145	0	38	0.24%	180	150	0	30	0.54%	180	149	0	31	0.56%
	Improved 15-30	42	34	0	8	0.06%	162	136	0	26	0.20%	177	151	0	26	0.20%	180	155	0	25	0.29%	180	158	0	22	0.32%	180	157	0	23	0.34%
Yes	LP-And-Fix	42	37	0	5	0.12%	162	87	45	30	0.35%	177	74	81	22	0.36%	180	84	81	15	0.58%	180	83	81	16	0.50%	180	85	81	14	0.45%
	Relax-And-Fix	42	25	0	17	0.21%	162	110	0	52	0.46%	177	111	0	66	0.77%	180	118	0	62	1.28%	180	120	0	60	1.34%	180	126	0	54	1.25%
	LP-Relax-And-Fix	42	25	0	17	0.21%	162	110	0	52	0.46%	177	111	0	66	0.77%	180	118	0	62	1.28%	180	120	0	60	1.34%	180	126	0	54	1.25%
	RepeatedLP	42	25	0	17	0.21%	162	110	0	52	0.46%	177	111	0	66	0.77%	180	118	0	62	1.28%	180	120	0	60	1.34%	180	126	0	54	1.25%
	Improved	42	25	0	17	0.21%	162	110	0	52	0.46%	177	111	0	66	0.77%	180	118	0	62	1.28%	180	120	0	60	1.34%	180	126	0	54	1.25%
	Improved 5-15	42	31	0	11	0.12%	162	119	0	43	0.22%	177	135	0	42	0.37%	180	141	0	39	0.57%	180	144	0	36	0.63	180	141	0	39	1.12%
	Improved 10-20	42	32	0	10	0.10%	162	124	0	38	0.17%	177	139	0	38	0.24%	180	145	0	38	0.24%	180	150	0	30	0.54%	180	149	0	31	0.56%
	Improved 15-30	42	34	0	8	0.06%	162	136	0	26	0.20%	177	151	0	26	0.20%	180	155	0	25	0.29%	180	158	0	22	0.32%	180	157	0	23	0.34%

Table 5: Heuristic results split to capacity.

The amount of LP-And-Fix problems unsolved increases when the capacity increases. This indicates that with less tight capacity restrictions the LP finds fewer integer y variables.

For Relax-And-Fix based heuristics it holds that the amount of optimal solutions found increases when capacity increases. For every capacity it also holds that increasing the integer size increases the amount of optimal solutions found and almost everywhere the gap (if found) decreases as well.

The RepeatedLP heuristic fails sometimes when demand smoothing is turned off. This only occurs with tight capacity restrictions. This can again be explained by the fact that the LP solved before every Relax-and-Fix iteration is independent of that iteration. It may set some y variables to 0 (which can be feasible/optimal for that part of the problem), but when the entire problem is solved, this may lead to infeasibility. This can occur when demand in a period is well above the production capacity of that period. As demand is not smoothed in advance, infeasibilities are more likely to take place.

Additional results

FAL formulation

The FAL formulation has been implemented on some test problems as well. The advantage of this method is, as described above, that without the capacity constraint the LP would give the optimal integer solution. With the capacity restriction this does not hold, but the formulation is tighter than the initial formulation, which results in more integer values.

The disadvantage of this method, however, is the calculation time. The difference in problem size is displayed in table 6.

Characteristic	Formulation	
	CLSP	**FAL**
Objective variables	753	63003
Objective nonzeros*	500	31375
Linear constraints	503	63003
- Nonzeros*	1253	188003
Average integer values found in LP**	9.50	75.04
Time to solve LP***	0.12 sec	7.97 sec
Time to solve Relax-And-Fix problem***	2.06 sec	283.22 sec

* Variables that still exist after presolve.

** Measured over all feasible instances, 921 in total.

*** Problem instance: Periods: 250, Seasonal: No, Capacity: 250, Setup costs: 900, Std: 67, Instance: 4.

Table 6: Comparison between CLSP and FAL formulations.

The FAL formulation is clearly more constrained than the CLSP formulation for the same problem size. This results in a 66 times longer calculation time for an LP and even 138 times longer calculation time for the Relax-And-Fix heuristic. The advantage of this method clearly is the average amount of integer values found in the LPs of all problem instances.

Several methods have been applied to reduce the calculation time of the FAL formulation. For instance, originally c_{ti}, a 250 by 250 matrix, was re-calculated before every Relax-And-Fix iteration. Calculating this matrix in advance resulted in some performance improvement, but not significant.

Another method used to decrease the calculation time is to decrease the scope of the problem. In a problem instance of 250 periods there would always be decided how much should be produced in period 1 to supply period 250. A way to reduce this scope is to set a maximum amount of periods to schedule ahead.

When choosing a fixed number a trade-off needs to be made between calculation time and result. When choosing a low fixed number, calculation time decreases but the result will most likely be further away from the optimal solution as only a limited amount of periods can be produced ahead (while this can be optimal). When choosing a high fixed number the result will suffer less, but the calculation time increases again.

This fixed number method has been implemented. In the beginning of a Relax-And-Fix instance choosing a low value makes a big difference in calculation time, but at the end of the instance a large problem size is needed again. This shows that even with a small scope, calculation time of the FAL formulation will never be close to the calculation time of the CLSP formulation.

Demand smoothing

Demand smoothing is a useful tool when future periods are not taken into account. By smoothing the demand all problem instances that have a solution can be solved because in every period the amount of units produced will not exceed the production capacity.

This tool only appeared to be useful for the RepeatedLP heuristic. This is because, as explained before, the LP solved before every Relax-And-Fix iteration is independent of the following iteration.

Other Relax-And-Fix based heuristics do not benefit at all from demand smoothing – the results are identical when it is turned off. Demand smoothing is useful when solving a problem without looking at future periods, but Relax-And-Fix heuristics take the future into account in a relaxed way. Production variables are relaxed, but for example inventory restriction (1) still holds for every period as only y is relaxed.

Step and integer size performance

Table 7 shows the calculation time difference between the different Improved heuristics.

Method name	Time (sec)*
Improved	2.06
Improved 5-15	2.82
Improved 10-20	1.99
Improved 15-30	1.81
Improved 20-40**	3.00

* Problem instance: Periods: 250, Seasonal: No, Capacity: 250, Setup costs: 900, Std: 67, Instance: 4.

** Not applied on all test cases, only for speed reference.

Table 7: Calculation time difference of the Improved Relax-And-Fix methods.

Solution time increases if only the integer size increases, but it decreases when the step and integer size are increased simultaneously up to 15 and 30 respectively. After this, the time gained by running less iterations is eliminated by the extra time required to solve the individual steps.

Conclusion

Existing MIP based heuristics and new extensions to these heuristics have been applied to the Capacitated Lot Sizing problem with promising results.

The LP-And-Fix heuristics leaves some problems unsolved due to the stopping criterion of the heuristic, but performs well on almost all test instances. Downside to this method is that the calculation time is high. This paper limits the calculation time to 120 which is utilized fully for most large problem instances.

The Relax-And-Fix heuristic and some extensions to it (LP-Relax-And-Fix and Improved) perform rather well. For large problem instances almost half of the problems are solved to optimality and the optimality gap of the remaining problems is on average below 1%. High production setup costs result in somewhat bigger optimality gaps, but still the average optimality gap with the highest tested setup costs is only 1.30%.

The RepeatedLP heuristic shows no improvement to the Relax-And-Fix heuristic, giving the exact same results with demand smoothing turned on. RepeatedLP even gets slightly worse than Relax-And-Fix when demand smoothing is turned off, leaving some problem instances unsolved.

Changing integer and step size for Relax-And-Fix based heuristics does improve the results overall and for the same period length, capacity or setup costs. Up to a step size of 15 and an integer size of 30 solution time decreases, beyond that the integer part is too big to be solved efficiently and it will thus take longer to solve a problem instance.

The FAL formulation does result in more integer values in the LP solutions, but the calculation time increases dramatically compared to the CLSP formulation, making the CLSP formulation the preferable formulation to apply the heuristics on.

Demand smoothing is not useful for the Relax-And-Fix based heuristics as it does not improve their performance. For the RepeatedLP heuristic it makes all feasible solutions solvable, but as the Improved heuristics perform better that RepeatedLP, demand smoothing is not necessary.

The final conclusion of this paper is therefore that the existing Relax-And-Fix heuristic already performs very well on the CLSP. The Improved 15-30 heuristic has the best overall performance, thus an integer size up to 30 and a step size up to 15 are beneficial for the outcome. Additional tools to improve the outcomes like the FAL formulation and demand smoothing are not needed to obtain good solutions.

This paper only uses fixed setup and holding costs and uses a fixed capacity for every period. As a suggestion for further research I would advise to try the MIP-based heuristics described in this paper on more general assumptions, like a (periodically) fluctuating capacity and more flexible holding and setup costs. Adding (variable) unit production costs may also be a valuable extension.

References

Baker, K.R. (1989). Lot-sizing procedures and a standard data set: a reconciliation of the literature. *Journal of manufacturing and operations management* volume 2 number 3. pp. 199-221.

Bitran, G.R. and Yanasse, H.H. (1982). Computational Complexity of the Capacitated Lot Size Problem. *Management Science* volume 28 number 10. pp. 1174-1186.

Chen, H.-D., Hearn, D.W. and Lee, C.-Y. (1994). A dynamic programming algorithm for dynamic lot size models with piecewise linear costs. *Journal of Global Optimization* volume 4 number 4. pp. 397-413.

Helmrich, M.R. (2005). Heuristieken voor het capacitated lot-sizing problem. *Bachelor thesis* Erasmus University Rotterdam. pp. 16-18.

Krarup, J. and Bilde, O. (1977). Plant location, set covering and economic lot size: An O(mn)-algorithm for structured problems. *International Series of Numerical Mathematics* volume 36. pp. 155–186.

Pochet, Y. and Wolsey, L.A. (2006). Primal Heuristics – Finding Feasible Solutions. *Mixed Integer Programming Algorithms*. ISBN 978-0-387-29959-4. pp. 107-113.

Wagelmans, A., Hoensel, van S. and Kolen, A. (1992). Economic Lot Sizing: An O(n log n) Algorithm That Runs in Linear Time in the Wagner- Whitin Case. *Operations Research* volume 40 supplement 1. pp. S145-S156.

Wagner, H.M. and Whitin, T.M. (1958). Dynamic Version of the Economic Lot Size Model. *Management Science* number 5. pp. 89-96.