

The selection of clients for promotion campaigns by means of mathematical programming

F. Talla Nobibon¹, R. Leus¹, F. Spieksma¹

¹ Department of Decision Sciences and Information Management, K.U.Leuven, Leuven, Belgium

Abstract - This paper presents an optimization model for the selection of sets of clients that will receive an offer for one or more products during a promotion campaign. The complexity of the problem makes it very difficult to produce optimal solutions using standard optimization methods. We propose an alternative set covering formulation and develop a branch-and-price algorithm to solve it. We also describe seven heuristics to approximate an optimal solution. We perform extensive computational experiments for the two formulations as well as for the seven heuristics.

Keywords - promotion campaign, integer programming, branch-and-price algorithm, non-approximability, heuristics.

I. INTRODUCTION

Promotion campaigns are fundamental direct marketing tools for improving the economic profit of a firm, either by acquiring new customers or by generating additional revenue from existing customers [8]. The former action is called “acquisition” while the latter is “retention” [11]. In this paper we are concerned with the latter case: campaigns that generate additional revenue by offering new products to existing customers. This study is justified by at least two practical facts. Reinartz et al. [11] point out that “When firms trade off between expenditures for acquisition and those for retention, a suboptimal allocation of retention expenditures will have a greater impact on long-term customer profitability than will suboptimal acquisition expenditures”. Moreover, models and methods used for data analysis are more suited for retention [7] since more information is available. Retention boosts the customer lifetime value, which is defined by Kumar et al. [9] to be “the sum of cumulated cash flows – discounted using the weighted average cost of capital – of a customer over his or her entire lifetime with the firm”. Customer lifetime value usually serves as a metric for a ranking or segmentation of the firm’s customers [12]. During the last decades, the advances in data analysis coupled with the availability of customer data have pushed firms to develop a more customer-oriented strategy. Nowadays, such a strategy is globally accepted, but its practical implementation is far from being accomplished. This implementation delay is observed both in business-to-business and business-to-consumer settings, and is particularly pronounced in financial institutions, which often have a large number of customers with full data available but may lack sophisticated tools that efficiently take into account these advantages in decision making [5].

In literature, promotion campaign models are also frequently referred to as optimal product targeting models [7]. The

latter examine “Which products should be targeted to which customers to maximize profits, under the constraints that only a limited number can be targeted to each customer, and each product has a minimum sales target”, which is mentioned by [7] as an interesting issue for future research. A promotion campaign problem is essentially characterized by two steps, which are “data analysis” and “problem formulation and solution”. The first step, which is mainly statistical, has received increasing attention with the advances in data analysis. Recently, numerous models that carry the name “response models” have been developed and are currently being used in practice [7], [3]. Although this step is necessary for an application in financial institutions (as its outputs are used as inputs for the second step), its use can be less important for an application in other areas like advertisement, see [4].

This paper investigates the development of optimization models for promotion campaigns based on integer programming. Motivation for studying this problem comes from a case occurring at FORTIS [6], one of the leading banks in Belgium. We aim to maximize the profit subject to business constraints such as the campaign’s return on investment hurdle rate that must be met (the hurdle rate is the minimum acceptable rate of return that management will accept for the campaign), a limitation on the funding available for each product, a restriction on the maximum number of possible offers to a client, and a minimum quantity commitment (MQC) on the number of units of a product to be offered in order for that product to be part of the campaign. This constraint has been briefly mentioned by Cohen [3] as a technical issue for an application in a bank. However, he did not explicitly incorporate it into his model. Our model takes into account this constraint, making it an extension of the model used by Cohen. In our formulation, we also impose a more general version of the MQC constraint, allowing the fixed minimum quantity to depend on the product.

In this paper, we present a basic integer programming formulation for the optimization of promotion campaigns and show the non-approximability of the problem. We next present a set covering formulation and develop a branch-and-price algorithm for solving it. The size of instances that can be solved optimally using this algorithm allows its efficient use for business-to-business promotion campaigns and for sampling approaches in financial institutions [3]. We present seven heuristics which can be used for business-to-consumer promotion campaigns. These heuristics are either variants of the algorithms used in practice for application in a bank

(see [6], [7]) or developed based on the structure of the problem.

This paper is organized as follows. Section II describes the basic integer programming formulation for deciding on the composition of promotion campaigns and related complexity results. We present the set covering formulation in Section III and develop a branch-and-price algorithm to solve it. In Section IV, we describe seven heuristics to approximate an optimal solution. The experimental results for the two formulations (basic formulation and set covering formulation) as well as for the seven heuristics are presented in Section V followed by some conclusions in Section VI.

II. BASIC FORMULATION

The objective of a promotion campaign is to achieve a maximum profit by offering n different products to m customers while taking into account various business constraints. We incorporate the following restrictions: the return on investment hurdle rate must be met for the campaign, the budget allocated to each product is limited, an upper bound is imposed on the number of products that can be offered to each client and there is also an MQC constraint for each product. We define the parameter r_{ij} as the probability that client i accepts an offer of product j and DFV_{ij} as the return to the firm when client i responds positively to the offer of product j . The latter is termed the Delta Financial Value by FORTIS [6]. These two parameters are the basis for the computation of customer lifetime value [12]. Practically, these parameters are estimated using response models based on historical data [3], [7], [10] and are assumed to be available within the firm. We denote by p_{ij} the expected return to the firm (revenue) when client i accepts the offer of product j , so $p_{ij} = r_{ij}DFV_{ij}$. Further, there is a variable cost c_{ij} associated with the offer of product j to client i , the upper bound M_i of offers that can be made to a client i (this quantity is related to the status of the client), the minimum quantity commitment bound O_j associated with product j , the budget B_j allocated to the product j , a fixed cost f_j needed if product j is used for the campaign and finally the corporate hurdle rate R . The value of R is dependent on the firm and the riskiness of the investment. In practice, most firms use their weighted average cost of capital (WACC) as an estimation of R [2]. We define the decision variables $y_j \in \{0, 1\}$, equal to 1 if product j is used during the campaign, 0 otherwise, and $x_{ij} \in \{0, 1\}$, which is equal to 1 if product j is offered to client i and 0 otherwise. A basic formulation for the promotion campaign problem can be expressed as:

$$(M1) \quad \max \sum_{i=1}^m \sum_{j=1}^n (p_{ij} - c_{ij}) x_{ij} - \sum_{j=1}^n f_j y_j \quad (1)$$

subject to

$$\sum_{i=1}^m \sum_{j=1}^n p_{ij} x_{ij} \geq (1 + R) \left[\sum_{i=1}^m \sum_{j=1}^n c_{ij} x_{ij} + \sum_{j=1}^n f_j y_j \right] \quad (2)$$

$$\sum_{i=1}^m c_{ij} x_{ij} \leq B_j \quad j = 1, \dots, n, \quad (3)$$

$$\sum_{j=1}^n x_{ij} \leq M_i \quad i = 1, \dots, m, \quad (4)$$

$$\sum_{i=1}^m x_{ij} \leq m y_j \quad j = 1, \dots, n, \quad (5)$$

$$\sum_{i=1}^m x_{ij} \geq O_j y_j \quad j = 1, \dots, n, \quad (6)$$

$$y_j, x_{ij} \in \{0, 1\} \quad i = 1, \dots, m, j = 1, \dots, n. \quad (7)$$

The objective function (1) is the maximization of the total net benefit received from the offer of products to clients minus the fixed cost of using the products for the campaign. The first constraint (2) is the corporate hurdle rate constraint, which makes sure that the campaign's return on investment is at least R , and which was first suggested by Cohen [3] for an application in a bank. The set of constraints (3) enforces that we should not exceed the budget B_j allocated to the product j . Here, the product dependency of the budget reflects the situation in large firms where an individual business unit is responsible for the production and the sale of a product. Hence, each business unit has its own budget. The set of constraints (4) states that we cannot propose more than a certain number M_i of products to client i ; the sets of constraints (5) and (6) constitute the MQC constraint, which specifies that when a product is not part of the campaign, no clients will receive an offer, while if product j takes part in the campaign then at least $O_j > 0$ clients receive an offer, and finally the last set of constraints (7) is the integrality constraint.

The following result shows that there is little hope for finding a polynomial time algorithm for solving (M1).

Proposition 1: The promotion campaign problem (M1) is strongly NP-hard, even for $O_j = 1$ for all j .

The proof of the results enunciated in this paper can be found in [14].

Moreover, (M1) is difficult to solve even approximately. We prove this non-approximability result by showing that it is NP-hard to find a non-trivial feasible solution for (M1). A *non-trivial feasible solution* for the basic formulation (M1) is a feasible solution which achieves a non-zero objective value.

Proposition 2: Finding a non-trivial feasible solution for the basic formulation (M1) is NP-hard.

The results of Proposition 1 and Proposition 2 justify the intensive use of heuristics in practice [3], [4], [7].

The basic formulation (M1) can be strengthened by using the following disaggregate version of constraints (5):

$$x_{ij} \leq y_j \quad i = 1, \dots, m, j = 1, \dots, n. \quad (8)$$

The following result allows the relaxation of the integrality constraint $y_j \in \{0, 1\}$ to $0 \leq y_j \leq 1$ for all j .

Proposition 3: The convex hull of the feasible solutions to the integer program (M1) is identical to the convex hull of the solutions that satisfy the constraints (2), (3), (4), (6), (8) and $0 \leq y_j \leq 1, x_{ij} \in \{0, 1\}$ for all i, j .

In the rest of this paper, the formulation obtained from (M1) by replacing (5) by (8) and by substituting $0 \leq y_j \leq 1$ for $y_j \in \{0, 1\}$ for all j , is denoted by (M2).

III. BRANCH-AND-PRICE

This section is devoted to the application of a branch-and-price algorithm for solving the promotion campaigns problem.

A. A set covering formulation

For every product j , let k_j be the number of distinct subsets of clients of cardinality at least O_j which can receive the offer of product j within the budget limit B_j . Explicitly, we define S_{pj} ($p = 1, \dots, k_j$) as a set of at least O_j clients satisfying $\sum_{i \in S_{pj}} c_{ij} \leq B_j$. We use the binary variable z_{pj} to indicate whether the product j is offered to the set of clients S_{pj} ($z_{pj} = 1$) or not ($z_{pj} = 0$). A *set covering formulation* of the promotion campaign problem is given by:

$$(M3) \quad \max \sum_{j=1}^n \left[\sum_{p=1}^{k_j} \left(\sum_{i \in S_{pj}} (p_{ij} - c_{ij}) - f_j \right) z_{pj} \right] \quad (9)$$

subject to

$$\sum_{j=1}^n \left[\sum_{p=1}^{k_j} \left(\sum_{i \in S_{pj}} \Delta_{ij} - (1+R)f_j \right) z_{pj} \right] \geq 0 \quad (10)$$

$$\sum_{j=1}^n \sum_{p: i \in S_{pj}} z_{pj} \leq M_i \quad i = 1, \dots, m, \quad (11)$$

$$\sum_{p=1}^{k_j} z_{pj} \leq 1 \quad j = 1, \dots, n, \quad (12)$$

$$z_{pj} \in \{0, 1\} \quad j = 1, \dots, n, p \in \{1, \dots, k_j\}, \quad (13)$$

where $\Delta_{ij} = p_{ij} - (1-R)c_{ij}$. The first constraint (10) enforces that the campaign's return on investment must be at least R , the set of constraints (11) ensures that at most M_i products are offered to client i and the set of constraints (12) states that at most one nonempty set of clients is selected for each product. We prove in [14] that (M3) is obtained by applying Dantzig-Wolfe decomposition to the LP relaxation of (M2).

Let z be any feasible solution to the LP relaxation of (M3) and let $x_{ij}^* = \sum_{p: i \in S_{pj}} z_{pj}$, $y_j^* = \sum_{p=1}^{k_j} z_{pj}$, then (x^*, y^*) is a feasible solution to the LP relaxation of (M2) which achieves the same objective value as z . Furthermore, we have the following result, which is useful for the branching strategies.

Proposition 4: Given a feasible solution z to the LP relaxation of (M3), if, for a given product j , z_{pj} is fractional, then there must be an i such that $x_{ij}^* = \sum_{p: i \in S_{pj}} z_{pj}$ is fractional.

Proposition 4 implies that the bound provided by the LP relaxation of the set covering formulation is at least as strong as that obtained by the LP relaxation of the basic formulation.

B. The pricing problem

We consider the LP relaxation (LPM3) of (M3) obtained by replacing constraints (13) by

$$z_{pj} \geq 0 \quad j = 1, \dots, n, p \in \{1, \dots, k_j\}. \quad (14)$$

The formulation (LPM3) has an exponential number of variables; thus instead of solving (LPM3), we consider a restricted problem that includes only a subset of columns. Additional columns can be generated by solving for each product j the following question called the *pricing problem*:

$$\exists S_p \text{ such that } \sum_{i \in S_p} w_i + v + f(1 - (1+R)d) < 0? \quad (15)$$

where $w_i = p_i(d-1) + c_i(1 - (1+R)d) + u_i$ for all i . The variable d is the dual variable corresponding to (10), u_i corresponding to the set of constraints (11), v_j corresponding to (12).

A solution to the pricing problem (15) can be obtained by solving a variant of the *k-item knapsack problem*. We propose a pseudopolynomial time dynamic programming algorithm and a linear time 2-approximation algorithm for solving the pricing problem. For more details, see [14].

Using the column generation procedure outlined above, we can solve (LPM3) in reasonable time. There is no guarantee, however, that the solution found will be integral; if this is not the case we will proceed with a branch-and-bound algorithm.

C. Branch-and-bound

Proposition 4 allows the use of a hybrid branching policy [13], [1]. We will then fix a single variable (variable dichotomy) [15], [13].

Fixing x_{ij} to zero leads to $\sum_{p: i \in S_{pj}} z_{pj} = 0$ and fixing it to one leads to $\sum_{p: i \in S_{pj}} z_{pj} = 1$ in the set covering formulation. Hence, at the node u of the tree, let $H(u)$ be the subset of products j for which there exists a non-empty set of clients R_j^u who must receive an offer of product j . Similarly, let $L(u)$ be the set of products j for which there exists a non-empty set of clients N_j^u who cannot receive an offer of product j . The LP problem (LP $_u$) to be solved at the node u is the combination of (LPM3) and the following two constraints.

$$\sum_{p: R_j^u \subseteq S_{pj}} z_{pj} = 1 \quad j \in H(u), \quad (16)$$

$$\sum_{p: N_j^u \cap S_{pj} \neq \emptyset} z_{pj} = 0 \quad j \in L(u). \quad (17)$$

The branching scheme resulting from the variable dichotomy branching strategy is compatible with the pricing problem for it does not render the pricing problem more difficult.

An upper bound at node u is provided by the optimal solution to the master problem (LP $_u$) or estimated by dualizing some constraints; details on these computations are available in [14].

IV. HEURISTICS

We present seven heuristics for the promotion campaign problem. The first heuristic (Heuristic 1) is a variant of the algorithm developed by FORTIS [6]. It uses the average cost $\frac{1}{m} \sum_{i=1}^m c_{ij}$ and revenue $\frac{1}{m} \sum_{i=1}^m p_{ij}$ for each product and ignores the selection of products. This heuristic solves a simplified integer program and uses the obtained solution to

offer products to clients. The second heuristic (Heuristic 2) is an improvement of Heuristic 1 as it chooses the products to be offered during the campaign. The third heuristic (Heuristic 3) iteratively solves a number of successive Exact k -item knapsack problems. At each stage, it selects the product with the highest positive profit such that the hurdle rate constraint is not violated. Then, this product is offered to the set of selected clients for that product, the problem is updated and the procedure is repeated until no more product can be offered to clients. The fourth heuristic (Heuristic 4) is inspired by the Next-Product-To-Buy model proposed by Knott et al. [7] for an application in a retail bank. For each product, the clients are sorted in decreasing order of their probability to buy that product. Then the product is offered to clients who can still receive an offer until no more budget is left for that product. The fifth heuristic (Heuristic 5) is a depth-first heuristic based on the branch-and-price approach. It aims to find a feasible and high-quality solution as quickly as possible by performing a partial traversal of the nodes corresponding to a particular branching decision within our branch-and-price approach. The sixth heuristic (Heuristic 6) is a truncated call to the MIP-solver CPLEX: the solver is used to solve (M2) and is interrupted when a time limit of one hour is reached. The last heuristic (Heuristic 7) is based on the LP relaxation of (M2). It either outputs the solution of the LP if it is integer or rounds it up to an integer. The output of these heuristics is always a feasible solution as they output the trivial solution when the solution they obtain is not feasible. More details about these heuristics can be found in [14].

V. COMPUTATIONAL EXPERIMENTS

All algorithms have been coded in C using Visual Studio C++ 2005 and run on a Dell Optiplex GX620 personal computer with Pentium R processor with 2.8 GHz clock speed and 1.49 GB RAM, equipped with Windows XP. CPLEX 10.2 was used for solving the linear programs.

We used random instances for the experiments. The instances are generated in such a way as to reflect real-life situations. We have generated six groups, each with 54 instances. Each instance in group S1 has 100 clients, S2: 200 clients, S3: 300 clients, M1: 1 000 clients, M2: 2 000 clients and L: 10 000 clients. For each group, we have three different numbers of products; these are 5, 10 and 15 products. Moreover, each client can receive either few products (s) or a large number of products (l). All instances are available online¹.

We have compared two implementations of the column generation procedure for solving the LP relaxation of the set covering formulation for the group S1. One implementation adds a single column per product to the master problem and one can add up to five different columns per product. We come to the conclusion that the overall time of the former implementation is smaller than that of the latter.

We have compared the LP relaxation of the basic formulation and of the set covering formulation for the groups S1 and

TABLE I
BASIC FORMULATION AND BRANCH-AND-PRICE ALGORITHMS

	n		CPLEX		depth first		best first	
			Opt	Time	Opt	Time	Opt	Time
S1	5	s	9	57.39	9	7365.78	9	2031.13
		l	9	3.91	9	38.75	9	72.85
	10	s	9	27.90	9	1141.08	9	3035.86
		l	9	6.05	9	20172.64	9	8587.53
	15	s	9	22.57	8	39845.53	8	19521.5
		l	9	3.22	9	31117.1	9	18014.18
S2	5	s	9	253.47	9	36725.30	7	36382.61
		l	8	138.18	9	12125.45	9	15699.19
	10	s	8	2246.54	9	36138.39	8	36696.35
		l	9	20.40	9	12010.04	9	18334.06
	15	s	7	1811.47	9	12545.51	7	36209.22
		l	9	9.47	9	5742.13	9	18183.22

S2. The results obtained confirm the theoretical result obtained in Section III-A that the set covering formulation is at least as strong as the the basic formulation. The LP relaxation of the set covering formulation provides solution with LP-gap usually less than 1%, however, this quality comes with a relatively high computation time.

We have compared three different strategies of traversing the branching tree of the branch-and-price algorithm. These are breadth first, depth first, and best first. We observe that the breadth-first branch-and-price algorithm is dominated either by the depth-first branch-and-price algorithm or by the best-first branch-and-price algorithm.

We have compared the basic formulation (M2) solved using CPLEX and the set covering formulation solved with the branch-and-price algorithms (depth first and best first) in Table I. *Opt* reports the number of instances solved optimally (out of 9) and *Time* gives the average CPU time in seconds. We observe that using the MIP-solver of CPLEX 10.2 for solving (M2) (enhanced with the default cutting planes) is much faster than the branch-and-price approaches. Although the pricing problem is solved fast and LP-gap is smaller at the root node, the number of nodes in the branching tree can be excessive.

We have compared the heuristics for the groups S3, M1, M2 and L. For S3, neither Heuristic 2 (and hence Heuristic 1) nor Heuristic 4 can find any non-trivial solution. Heuristic 3 gives a feasible solution with strictly positive objective value for all the test instances. Moreover, the Gap reported for each instance is less than 10% and the computation time is at most 10s. The solutions provided by Heuristic 5 are of good quality with Gap less than 9%. However, unlike Heuristic 3, Heuristic 5 requires much more time. Heuristic 6 provides feasible solutions with strictly positive objective value for each instance in S3. Although it takes more time than Heuristic 3, the gap is the smallest compared to other heuristics able to solve all the instances. The last heuristic (Heuristic 7) has a competitive computation time (less than 4s) but provides very few feasible solutions with strictly positive objective value (less than 23%).

The results of the application of the best heuristics (Heuristic 3, 5, 6 and 7) to groups M1, M2 and L are reported in Table II. *Gap* shows the Gap with respect to the LP relaxation, *Feas* reports the percentage of non trivial solutions and *Time*

¹www.econ.kuleuven.be/public/NDBAC96/promotion_campaigns.htm

TABLE II
COMPARISON OF HEURISTICS FOR MEDIUM AND LARGE SIZE INSTANCES

n	M1			M2			L			
	5	10	15	5	10	15	5	10	15	
Heuristic 3	Gap	11.20	13.23	14.46	13.10	13.15	12.31	34.04	33.52	37.80
	Time	146.60	394.54	304.32	2125.62	2231.14	1195.45	3573.92	3143.41	3353.81
	Feas	100	100	100	100	100	100	100	100	100
Heuristic 5	Gap	8.72	13.04	14.43	12.40	12.84	12.01	13.60	24.54	34.26
	Time	2441.74	3318.04	3274.59	1978.80	3565.28	3384.92	3730.05	3711.15	3634.33
	Feas	100	100	100	100	100	100	100	100	100
Heuristic 6	Gap	4.44	2.97	23.11	2.77	17.74	23.20	14.19	13.11	—
	Time	1845.20	2414.50	3018.14	2048.79	3132.60	3523.35	3667.33	3763.12	3803.61
	Feas	100	100	77.78	100	83.33	77.78	61.11	5.56	0.00
Heuristic 7	Gap	0.00	0.00	0.00	0.00	0.39	0.27	0.00	2.75	51.10
	Time	8.63	28.30	64.26	17.63	124.36	298.37	558.66	2974.50	3558.28
	Feas	11.11	11.11	16.67	11.11	22.22	16.67	22.22	5.56	11.11

the CPU time in seconds.

Heuristic 3 behaves very well for large instances: all the instances are solved with Gap less than 15% for group M1 and group M2, and less than 38% for the instances with 10 000 clients. This good quality of the solutions is coupled with an increased running time. Heuristic 5 finds non trivial feasible solutions to all instances, with a smaller Gap than Heuristic 3. Compared to Heuristic 3, the running time is higher, however. The solutions proposed by Heuristic 6 are usually close to the optimal solution (judging from the gap), but for large instances, this heuristic is quite limited in terms of the number of non trivial feasible solutions it can find. We observe that for instances in the set L, Heuristic 6 provides at most 62% of feasible solutions when the campaign involves 5 products and at most 6% when there are 10 products. When 15 products are involved, this heuristic fails to find any non-trivial solution. The last heuristic, Heuristic 7, provides good solutions from time to time, but is very limited by the number of feasible solutions it is able to find (at most 23%).

To conclude this comparison of heuristics, although Heuristic 6 and 7 provide good solutions, they are strongly limited by the number of non trivial solutions they output for large instances. Therefore, we advice the use of either Heuristic 3 or Heuristic 5 with a slight preference for the latter.

VI. CONCLUSIONS

This text explicitly models the promotion campaign problem taking into account both business constraints and customer preferences and specificities. Our work shows that the problem is strongly NP-hard and that it is unlikely that a constant factor approximation algorithm can be proposed for solving this problem. We have also presented a set covering formulation for the promotion campaign problem in which each product is associated with a subset of clients (which can be empty) in the optimal solution and developed a branch-and-price algorithm for solving it. We have shown that this last formulation is stronger than the basic formulation. Experimental results confirm that these two formulations are limited by the size of instances that they can efficiently solve, which makes their application more suited for business-to-business applications.

To extend the application to a business-to-consumer environment with considerably more customers, we have presented

seven heuristics. Some of these heuristics are currently used in practice (Heuristic 1, 2, 4), while others are new (Heuristic 3, 5, 6, 7). Based on extensive experimental results, we provide a comparison and comments on the efficiency and quality of the results obtained using the different formulations and the heuristics. These results clearly show a trade-off between computation time and solution quality and suggest the use of optimal algorithms for small and medium size instances, while heuristics are preferable for large size instances (Heuristic 5) and when time is an important factor (Heuristic 3).

REFERENCES

- [1] C. Barnhart, E. Johnson, and M. Savelsbergh. Branch-and-price: column generation for solving huge integer programs. *Operations Research*, 46:316–329, 1998.
- [2] R.C. Bruner, K.M. Eades, R.S. Harris, and R.C. Higgins. Best practices in estimating the cost of capital: survey and synthesis. *Financial Practice and Education*, 8:13–28, 1998.
- [3] M. Cohen. Exploiting response models - optimizing cross-sell and up-sell opportunities in banking. *Information Systems*, 29:327–341, 2004.
- [4] B. De Reyck and Z. Degraeve. Broadcast scheduling for mobile advertising. *Operations Research*, 51:509–517, 2003.
- [5] F.R. Dwyer. Customer lifetime valuation to support marketing decision making. *Journal of Direct Marketing*, 11:6–13, 1997.
- [6] E. Hellinckx. Customer relationship management: De optimalisatie van de planning van campagnes. Master’s thesis, KULeuven, 2004 (in Dutch).
- [7] A. Knott, A. Hayes, and S.A. Neslin. Next-product-to-buy models for cross-selling applications. *Journal of Interactive Marketing*, 16:59–75, 2002.
- [8] P. Kotler and G. Armstrong. *Principles of Marketing*. Pearson Prentice Hall, 2006.
- [9] V. Kumar, G. Ramani, and T. Bohling. Customer lifetime value approaches and best practice applications. *Journal of Interactive Marketing*, 18:60–72, 2004.
- [10] A. Prinzie and D. Van Den Poel. Predicting home-appliance acquisition sequences: Markov/Markov for discrimination and survival analysis for modeling sequential information in NPTB models. *Decision Support Systems*, 44:28–45, 2007.
- [11] W. Reinartz, J.S. Thomas, and V. Kumar. Balancing acquisition and retention resources to maximize customer profitability. *Journal of Marketing*, 69:63–79, 2005.
- [12] L. Ryals. Making customer relationship management work: The measurement and the profitable management of customer relationships. *Journal of Marketing*, 69:252–261, 2005.
- [13] M. Savelsbergh. A branch-and-price algorithm for the generalized assignment problem. *Operations Research*, 45:831–841, 1997.
- [14] F. Talla Nobibon, R. Leus, and F.C.R. Spieksma. Models for the optimization of promotion campaigns: exact and heuristic algorithms. Research report KBI_0814, Katholieke Universiteit Leuven, 2008.
- [15] L.A. Wolsey. *Integer Programming*. Wiley, 1998.