Multimedia Systems
DOI 10.1007/s00530-012-0284-y

REGULAR PAPER

Allocation algorithms for personal TV advertisements

Ron Adany · Sarit Kraus · Fernando Ordonez

Abstract In this paper we consider the problem of allocating personal TV advertisements to viewers. The problem's input consists of ad requests and viewers. Each ad is associated with a length, a payment, a requested number of viewers, a requested number of allocations per viewer and a target population profile. Each viewer is associated with a profile and an estimated viewing capacity which is uncertain. The goal is to maximize the revenue obtained from the allocation of ads to viewers for multiple periods while satisfying the ad constraints. First, we present the integer programming (IP) models of the problem and several heuristics for the deterministic version of the problem where the viewers' viewing capacities are known in advance. We compare the performances of the proposed algorithms to those of the state-of-the-art IP solver. Later, we discuss the multi-period uncertain problem and, based on the best heuristic for the deterministic version, present heuristics for low and high uncertainty. Through computational experiments, we evaluate our heuristics. For the deterministic version, our best heuristic attains 98 % of the possible revenue and for the multi-period uncertain version our heuristics performances are very high, even in cases of high uncertainty, compared to the revenue obtained by the deterministic version.

R. Adany (✉) · S. Kraus
Computer Science Department, Bar-Ilan University,
Ramat-Gan 52900, Israel
e-mail: adanyr@cs.biu.ac.il

S. Kraus
e-mail: sarit@cs.biu.ac.il

F. Ordonez
Industrial Engineering, University of Chile, Santiago, Chile
e-mail: fordon@dii.uchile.cl

Keywords TV advertisements · Personalization · Allocation · Heuristics · Uncertainty

1 Introduction

Recent technological advances, such as the penetration of digital TV standards and proliferation of diverse devices to watch TV anywhere, make it possible to personalize the TV experience [6]. This move toward a higher degree of personalization in the TV medium, which had not changed for decades, has created an opportunity also for personalized advertisement.

Personal advertisements are better for the advertisers, the service suppliers and even the viewers. The advertisers spend their budget efficiently by choosing the target population with a high resolution, i.e. a single viewer, and controlling the manner and the format of exposure. The service suppliers, i.e. the media companies and the operators, can request a higher price for these advertisements and thus increase their revenues [10, 16]. Moreover, viewers will also benefit from this change since they will be able to watch ads more relevant to them.

According to Nielsen [22, 24], $67B was spent on TV advertisement in the US in 2010 and the time spent watching TV has been consistently increasing, e.g. during the first quarter of 2010, on average, each viewer watched 158 h per month. These trends and changes are the motivations for our research.

Though personalized data for a specific end-user, namely a viewer, using the current technology and infrastructure is possible [6], the issue of personal TV advertisement is still an open problem for which, to date, no adequate solution has been proposed. Personalization of ads requires making more complicated advertisement

Published online: 01 November 2012

decisions: the selection of the advertisement campaigns that will be broadcast needs to be compounded with the decisions of how these ads will be assigned to individual users. A service supplier company will aim to maximize its revenues through these binary decisions, i.e. select an ad campaign or not and assign an ad to a viewer or not. We refer to this optimization problem with binary variables as the personal TV Ads Allocation problem (and for short the Ads Allocation problem).

We note that these problems can be quite large, since TV broadcast markets involve millions of viewers that are exposed to hundreds, if not more, of different possible ad campaigns in any given period. In addition, since it is difficult to accurately estimate how much TV individual users will view in a given period, it is reasonable to assume that there is uncertainty in the personal TV Ads Allocation problem. Both the size and the presence of uncertainty make this a very challenging problem to solve.

Many studies concern the problem of selecting personal advertisements most suitable to each individual viewer, e.g. [15, 18, 20, 34], and many others focus on how to deliver them, e.g. [6, 20]. Our research supplements these studies, using their results as input with the goal of optimizing the allocation of ads. In this work, we propose algorithmic solutions and do not deal with the hardware or infrastructure problems.

Throughout this research, we assume that the infrastructure is similar to the framework of the iMedia system, which is designed for personal advertisement in the interactive TV environment [6, 30]. Based on this framework, the entire process of the personalized advertisement is as follows. In every sequence of time periods, e.g. a month, some set of advertisement requests will be given. In some centralized computing centers a subset of ads will be selected. These ads will be allocated to viewers and playlists of ads will be generated to each viewer for all the time periods. Then, advertisement contracts will be signed with the advertisers according to the allocations. The playlists will be delivered and stored in the set-top-box (STB) units with which each viewer is equipped, as is common today. During the time periods, viewers will watch TV and on commercial breaks, each STB will air ads based on the viewer's playlist. At the end of each time period, each STB will send an ads' viewing report back to the centralized computing center, detailing the actual airing of the ads from the viewer's playlist. Reallocation and modifications of the playlists will be permitted in the next periods. At the end of all the planned time periods the billing process will be activated according to the signed contracts.

There are diverse devices to watch TV anywhere anytime. The most important and promising one for the coming years, in addition to the traditional TV device, seems to be the mobile device, i.e. Mobile TV. First, there seems to be a delay in the current TV medium moving toward personalization, probably due to infrastructure issues. The TV's current infrastructure is based on a broadcast distribution system, which makes the delivery of personalized content problematic [30]. Delivery of personalized content to a specific end-user with the current mobile technologies and infrastructure as well as personalized content via Mobile TV services, e.g. [8, 29], is possible. Second, the mobile phone, as the most personal medium people have [27], is perfect for personalized advertising. Third, according to some media researches [14, 25] Mobile TV is predicted to be the largest ad revenue generator in the upcoming years with expected revenues of billions of dollars. The same selected framework, i.e. iMedia, is relevant for Mobile TV with some modifications. The viewer' ads playlist will not be stored locally but will be delivered to the viewers via the network on demand, i.e. on commercial breaks.

The rest of this paper is organized as follows. In Sect. 2 we provide an overview of personal advertisement in the TV medium versus the Internet medium, in addition to related work on similar deterministic and uncertain allocation problems. We formally describe the certain, i.e. deterministic, problem and the uncertain problem with the IP models in Sect. 3. Our proposed heuristic solutions are presented in Sect. 4. The description of the evaluation method and the results of our experiments are presented in Sect. 5. Finally, we summarize and present our conclusions in Sect. 6.

2 Related work

In this section, we present the differences between one of the most known mediums in which personalization is applied, i.e. the Internet, and the TV medium, which is the medium under discussion in this paper. Later, we present some related work concerning allocation and packing with and without uncertainty.

2.1 Personal ads: TV versus internet

Though personal advertisement is already the most common method applied in the Internet medium (e.g. Google AdWords [13]), it is still not the main method used in the Mobile or TV mediums. One of the reasons for the delay of its application in the TV medium is probably the current infrastructure, which is based on a broadcast distribution system making the delivery of personalized content problematic [30]. But, in addition to the technical issues, there are several key points that distinguish personalized ads for the Internet from those for the TV. In the following paragraphs we will present some of these key points.

The method of exposure is different. On the Internet, the viewer receives his required data (e.g. the required web page, the search results) and concurrently is exposed to the ads which appear beside it. However, on TV the ads are in the forefront, i.e. the only content on the screen. Thus, if a viewer chooses to ignore the Internet ads, he can do so and continue to consume the content, whereas a TV viewer cannot since he is inevitably exposed to the ads (unless he actively chooses to turn off the TV or to zap). This difference is the motivation to the multi-viewing constraint described later.

The pricing method is different, mainly due to the methods of exposure. For instance, certain pricing methods on the Internet are based on PPC (pay per click), such that if the viewer clicks on the ad, a payment is made. In contrast, on TV the pricing method is based on GRP (gross rating point), which is a measure of the advertising impact calculated as factors of the target population size and exposure frequency. Furthermore, the production costs for creating the ads are different. Whereby TV ad production costs may reach millions of dollars, Internet ad production costs are much cheaper or may incur no cost, e.g. AdWords. This fact is the motivation to the global-viewing constraint described later.

There are also differences in the ads allocation constraints. While the main constraints in the Internet medium are relevancy (e.g. keywords) and budget, in the personal TV ads there are additional constraints, such as multi-, global- and conflict-viewing. In this paper we focus on multi- and global-viewing constraints, which represent the number of times each viewer needs to view the ad and the total number of viewers who need to view the ad. The motivation for such constraints, in addition to the exposure method, pricing method and production costs, stems from basic advertisement theories. Common theories on concepts such as ads recall or efficiency of ads are based on the behavior of an individual, i.e. a single viewer. Our solutions enable advertisers to implement their advertisement campaigns taking into account the advertising theories and thus lead to better results for the campaigns.

2.2 Allocation and packing problems

The Ads Allocation problem is closely related to several basic problems in Computer Science such as the General Assignment Problem (GAP) and the Multiple Knapsack Problem (MKP). In this section, we briefly summarize the definitions of such problems, the state-of-the-art results and the similarities and differences from the Ads Allocation problem.

In GAP [17] there are m machines (viewers) and n jobs (ads). Each machine has a given capacity and each job has a required capacity and a cost incurred for being assigned to a machine. The objective is to minimize the costs of assigning the jobs, where each job needs to be assigned to exactly one machine. In [31] a 2-Approximation algorithm was presented. GAP is similar to the Ads Allocation problem with respect to the allocation task but it is different due to the job properties, the objective and the constraints. In GAP the jobs' size and cost properties depend on the machines to which they are assigned while in the Ads Allocation problem they are independent. The objective is to minimize the cost of assigning all jobs in contrast to maximizing revenues by assigning a subset of ads. In addition, the constraints are different. In GAP all jobs can be assigned to all of the machines, whereas in the Ads Allocation problem assignment restrictions exist. Moreover, in GAP each job must be assigned once, while in the Ads Allocation problem, if an ad is assigned it must be assigned multiple times, i.e. to several viewers.

In the Multiple Knapsack Problem (MKP) there are assignment profits instead of costs, where the capacity and the profit from jobs are independent of the machines to which they are assigned and the objective is to maximize the assignment profits. There is also an option to choose which jobs to assign and which not. In [7] multiple results on MKP were presented. Using the result for GAP from [31], a 2-Approximation algorithm was provided. In addition, the problem was proven to be a special case of GAP which is not APX-hard (nor FPTAS) and a polynomial-time approximation scheme (PTAS) was presented. The MKP problem is more similar to the Ads Allocation compared to GAP. The objective and the job properties are the same, though there are assignment restrictions and multi-assignment constraints in the Ads Allocation problem.

The Multiple Knapsack Problem with Assignment Restrictions (MKAR) is an extension of MKP which was discussed in [9, 26]. In addition to the constraints in MKP there are assignment restrictions such that each job can be assigned to a subset of machines. In [9] this problem was considered with the objective of maximizing the total assigned capacity. A 2-Approximation algorithm was proposed using heuristics and linear programming rounding methods. The MKAR problem is different from the Ads Allocation problem since it maximizes weights and not revenues and each item should be assigned to no more than one machine. Another version of the MKAR problem was discussed in [26] where the objective was to maximize the total revenue. A 2-Approximation algorithm was proposed using linear programming and rounding methods. This version of the MKAR is different from the Ads Allocation problem since, again, there are no multi-assignment constraints.

The generalized multi-assignment problem is another generalization of GAP with multi-assignment constraints.

The job's capacity and the cost are dependent on the machine to which they are assigned, the objective is to minimize the costs, and all the jobs must be assigned. The only differences from GAP are the multiple assignment constraints of each job. This problem was discussed in [28] where Lagrangian dual-based branch-and-bound algorithms were proposed. The goal in [28] was to find an exact solution, i.e. the optimal solution. The generalized multi-assignment problem is similar to the Ads Allocation problem in reference to the multiple assignment constraints but it is different with respect to other aspects: it lacks assignment restrictions, the jobs' capacity and the profit are dependent on the machines to which they are assigned and all the jobs must be assigned. Moreover, since the goal in [28] was to search for the optimal solution, and because it is an NP-hard problem it can be used only for small instances as demonstrated in [28] under the computational results section.

2.3 Allocation under uncertainty

In the following paragraphs, we will briefly describe the three major techniques that can be used to address optimization problems with uncertainty [4] in the context of our problem. The techniques are:

- chance constrained optimization,
- stochastic optimization, and
- robust optimization.

In chance constrained optimization the distribution of the uncertain data is assumed to be known. This information is used to express the constraints such that they have to be satisfied with a given probability. In order to use chance constrained optimization for the Ads Allocation problem we would have to know the probability distribution of the viewers' viewing capacity, which in itself is a very challenging problem. Thus we did not use this technique.

Stochastic optimization, in turn, typically uses discrete scenarios to describe the uncertainty, essentially repeating the problem in each scenario. This approach usually leads to massive optimization problems and the main challenge becomes a computational-algorithmic issue. However, given that the deterministic version of the problem is already NP-hard, we opted to take a different approach.

In robust optimization there is a value range for each uncertain data and the worst value is considered, i.e. the lower or the upper bound of the range depending on the case. This technique is simple to implement by changing each uncertain data to the worst value, i.e. changing the problem instance into the worst case instance. Though solving the worst case instance guarantees that the solution will be feasible in any case, the achieved objective value may be low. In this paper, we propose several algorithms that are inspired by this technique. These algorithms require information only about the range variation of the viewing capacity and lead to a problem that is comparable in size to the original deterministic version.

Robust optimization for linear programming (LP) problems under uncertainty was discussed in [2, 3, 19]. Ben-Tal and Nemirovski [2, 3] showed that LP problems can be very sensitive to uncertainty. Even with small uncertainty of 0.1 % of the constraints, the "nominal" LP solution can become infeasible. They described a new LP modeling methodology, replacing the original LP model with a "Robust Counterpart" that takes into account the uncertainty. Experiments were conducted on a large set of LP problems and the loss never exceeded 1 % of the optimal solution. Lin et al. [19] extended Ben-Tal and Nemirovski's work to mixed-integer linear optimization problems (MILP). We also followed some robust principles while adapting them to our problem and its specific constraints.

Toktas et al. [33] considered GAP with uncertain resource capacities and penalty for violations. The resource capacities were given with a distribution and samples of the problem were generated and solved with different deterministic heuristics. Their problem is similar to the Ads Allocation problem where the uncertainty is the viewers' viewing capacities but is different due to the special constraints of the Ads Allocation problem. In addition, the Ads Allocation problem has hard constraints while [33] allowed penalized violations. We did not use the sampling method due to the size of the realistic problem instances and the all-or-nothing constraints that result in high sensitivity of the solution. However, some of the proposed algorithms used a specific sampling of the problem, i.e. the worst case, as part of the robust techniques.

3 Problem definition

In this section, we formally present the Ads Allocation problem, i.e. selecting a subset of possible ads that can be properly assigned to viewers in order to maximize profit during a planning horizon of several periods and considering uncertainty regarding the viewer's capacity. We describe the problem notation and main components in Sect. 3.1, introduce the formulation of the problem in Sect. 3.2 and present the uncertainty model in Sect. 3.3.

3.1 Problem notation

The problem contains two types of entities: ads and viewers. We consider that in this problem there is a set of n ads, i.e. Ads $= \{a_1, a_2, \ldots, a_n\}$, where each ad, a_i, has the following properties: a revenue it generates if allocated, p_i, a length, l_i, a required rating, r_i, and a required frequency,

f_i. Finally, each ad has a profile, $P(a_i)$, defining the target population for the ad. The ad frequency corresponds to the number of times the same viewer should view the ad in order to be considered assigned to that viewer. The ad rating indicates the required number of different viewers to whom the ad must be assigned in order to be considered allocated and be paid. The target population defines the set of viewers that are relevant for the ad.

The problem also considers a set of m viewers, i.e. Viewers $= \{v_1, v_2, \ldots, v_m\}$ where each viewer, v_j, has an estimated viewing capacity for a period of time, c_j. The actual viewing capacity of viewer v_j is uncertain but bounded, i.e. the viewing capacity value belongs to the range $[c_j(1 - u_j), c_j(1 + u_j)]$ for some uncertainty factor $0 \leq u_j \leq 1$ which assumed to be known, e.g. based on observations of viewers' past activity. Viewers also have a profile, $P(v_j)$, describing the categories to which the viewer v_j belongs.

The estimated viewers' viewing times can be based on viewing statistics. Such statistics are already available for some defined viewer groups. For example, according to BARB [1], in 2010 the average weekly viewing per person in Great Britain was 28:13 (h/min). According to Nielsen [23], in 2009–2010 the average American watched 35:34 of TV per week, kids aged 2–11 watched 25:48 of TV per week on average and adults over 65 watched 48:54 of TV per week. These statistics are averages of specific viewer groups, but similar statistics can be calculated for smaller viewer groups and even for the individual viewer based on observations of the viewer's past activity (taking into account privacy issues).

The profiles $P(a_i)$ and $P(v_j)$ contain values of the personalized categories. For example, we could consider categories of $\{age, gender, location\}$ where each of them could take one of a finite set of values, e.g. *age* can take values from $\{11-30, 31-50, +50\}$. In addition, we allow the ad profile vector to take the value *All* in every category (or a combination of values, e.g. 11–30 and 31–50). This allows a flexible and realistic definition of the ad's target population. An ad can only be allocated to a viewer if the viewer is part of the target population of the ad, that is if the viewer's profile belongs to the ad's profile. In this case we state that ad a_i is relevant to viewer v_j and use Relevant($a_{\rm i}$, v_j) to denote this. In other words,

$$\text{Relevant}(a_i, v_j) = \begin{cases} 1 & \text{if } P(v_j) \subseteq P(a_i) \\ 0 & \text{if } P(v_j) \nsubseteq P(a_i) \end{cases} \tag{1}$$

The goal of the Ads Allocation problem is to maximize the profit from a valid assignment of ads to viewers. A valid assignment which will lead to a payment should be personal, i.e. an allocation that suits the ad's target population and viewers' profiles, satisfies the ad rating and frequency requirements, and does not exceed the viewers' viewing capacities. Any violation of the ads constraints will result in no payment for the specific allocation of the ad to the viewer or even for all the allocations of the ad.

The ads selected by the problem should be viewed over a planning horizon of T periods. In each period (for instance a week) each viewer exhibits an actual viewing capacity, denoted by w_{jt} for viewer v_j in period t, that would deviate by up to u_j % from the estimated capacity. The allocation of ads to viewers can then be adjusted to better accommodate the viewing requirements that remain, by adding new ads if possible.

There are two decisions in the Ads Allocation problem: (a) whether to select ad a_i or not, denoted by the binary variable y_i, and (b) whether to assign ad a_i to viewer v_j, denoted by binary variable x_{ij}. We summarize the notations used in this problem in Table 1. An examples of a viewer instance and ad instance follows.

Example 1 Viewer instance: the profile $P(v_j)$ = [*age*: 31–50, *gender*: Male], and viewing capacity estimation of $c_j = 1{,}200$ s per period with the uncertainty factor $u_j = 10$ %.

Example 2 Ad instance: the target population profile $P(v_i)$ = [*age*: All, *gender*: Female], length of $l_i = 30$ s, rating of $r_i = 3{,}000$ viewers, payment of $p_i = \$10{,}000$ for all 3,000 viewers and frequency of $f_i = 10$ times per viewer.

Table 1 Notation summary

Ads	$\{a_1, a_2, \ldots, a_n\}$
$P(a_i)$	Profile of ad a_i target population
p_i	Revenue of ad a_i
l_i	Length of ad a_i
r_i	Rating of ad a_i
f_i	Frequency of ad a_i
Viewers	$\{v_1, v_2, \ldots, v_m\}$
$P(v_j)$	Profile of viewer v_j
c_j	Expected viewing capacity of viewer v_j during one period
u_j	The uncertainty factor regarding viewer's viewing capacity
Additional parameters	
T	Number of time periods
Relevant(a_i, v_j)	Equals 1 if relevant, i.e. $1 \Leftrightarrow P(v_j) \subseteq P(a_i)$
Decision variables	
y_i	Equals 1 iff ad a_i is selected for allocation, else 0
x_{ij}	Equals 1 iff ad a_i is allocated to viewer v_j, else 0

3.2 Problem formulation

If we assume that we know the actual viewing capacities of all viewers in every time period, say w_{jt} and the planning horizon is from T_0 until T, we can formulate the deterministic Ads Allocation problem as follows:

$$\begin{aligned}
\max \quad & \sum_{i=1}^{n} p_i \cdot y_i \\
\text{s.t.} \quad & \sum_{i=1}^{n} l_i \cdot f_i \cdot x_{ij} \leq \sum_{t=T_0}^{T} w_{jt} && \forall j = 1,\ldots,m \\
& \sum_{j=1}^{m} x_{ij} = r_i \cdot y_i && \forall i = 1,\ldots,n \\
& x_{ij} \leq \text{Relevant}(a_i, v_j) && \forall i = 1,\ldots,n; \quad j = 1,\ldots,m \\
& x_{ij} \in \{0,1\} && \forall i = 1,\ldots,n; \quad j = 1,\ldots,m \\
& y_i \in \{0,1\} && \forall i = 1,\ldots,n
\end{aligned} \tag{2}$$

This problem maximizes the reward obtained by the ads selected. The first constraint ensures that viewer v_j has enough capacity during the planning horizon to watch all ads assigned to it while taking into account their length and required frequency. The second constraint requires that ad a_i be assigned to r_i different viewers if the ad is selected. The third constraint only permits allocation of ads to relevant viewers. The last two constraints make our decision variables binary.

The viewing capacity, however, is uncertain in practice, which makes the first constraint difficult to enforce. If a viewer watches less TV than expected, i.e. the capacity is reduced, it may force some ads not to be shown the number of times required compromising the revenue with which it is associated. On the other hand, if the viewer has a larger capacity than expected, it is unclear how to adapt the solution to best utilize the extra viewing capacity. Luckily, the problem with uncertainty has natural re-optimization points: after every period we can observe the actual capacity each viewer had and which of the previous allocations were successful. Taking this current state of information into account, we can make a better decision as to how to adapt to uncertainty in the next time periods.

We can use the formulation in Problem (2) to represent solving the problem of optimizing the ads selection and assignment after t' periods have occurred. For this we need to adjust the data of the problem to include the information that has occurred up to this point, i.e. the ads which have been watched during the passed periods. In particular, we consider a new rating, r'_i, which is the rating of ad a_i minus the number of times ad a_i was fully assigned to a viewer.

Note, if after t' periods ad a_i has already been successfully assigned and viewed by r_i viewers, i.e. in periods $T_0,\ldots,t'$, then the adjusted rating value will be $r'_i = 0$ and the "memory" of the allocation will be represented by setting $y_i = 1$. In addition, we need to redefine ad frequency to f'_{ij}, which equals ad's a_i frequency f_i minus the number of times viewer v_j has seen ad a_i.

The modified formulation representing the updated problem for the remaining periods after t' periods, i.e. $T - t'$ periods, is as follows:

$$\begin{aligned}
\max \quad & \sum_{i=1}^{n} p_i \cdot y_i \\
\text{s.t.} \quad & \sum_{i=1}^{n} l_i \cdot f'_{ij} \cdot x_{ij} \leq \sum_{t=T_0+t'}^{T} w_{jt} && \forall j = 1,\ldots,m \\
& \sum_{j=1}^{m} x_{ij} = r'_i \cdot y_i && \forall i = 1,\ldots,n \\
& x_{ij} \leq \text{Relevant}(a_i, v_j) && \forall i = 1,\ldots,n; \quad j = 1,\ldots,m \\
& x_{ij} \in \{0,1\} && \forall i = 1,\ldots,n; \quad j = 1,\ldots,m \\
& y_i \in \{0,1\} && \forall i = 1,\ldots,n
\end{aligned} \tag{3}$$

The viewer capacity only appears in the first constraint and it is summed over all periods left to consider, i.e. all the remaining periods. If the capacity is deterministic, i.e. $u_j = 0$, then it is the same as a single-period problem with a total capacity of $\sum_{t=T_0+t'}^{T} w_{jt} = \sum_{t=T_0+t'}^{T} c_j = [T - (T_0 + t')] \cdot c_j$.

In Sect. 4, we describe the efficient heuristics we developed for both the deterministic, single-period version of the Ads Allocation problem and the stochastic multi-period version. In addition, since the last version of the stochastic multi-period problem is exactly the same problem faced by a single-period problem with uncertainty, the algorithms to solve a single-period problem with uncertainty are the same as the ones we developed for the multi-period problem.

3.3 The considered uncertainty

We consider a basic model of uncertainty for the viewing capacity while assuming that it belongs to a given bounded set without any further distribution assumptions, i.e. $w_{jt} \in [c_j(1-u_j), c_j(1+u_j)]$, $\forall\, j = 1,\ldots,m$ and $t = T_0,\ldots,T$.

There are a number of optimization under uncertainty approaches that can be used to solve Problem (3). For instance, if the distribution of w_{jt} is known, we could use a chance constrained model that would replace the first set of constraints with constraints that have to be satisfied with a high probability (see also Sect. 2.3). If a discrete distribution of the random capacities w_{jt} is known, possibly by sampling, this problem could be stated as a multistage stochastic programming problem. Both of these models require strong assumptions on the uncertainty in viewing capacities and lead to challenging computational models. Considering the size of the problem (selecting thousands of ads to assign to millions of viewers) and the presence of integer variables we have opted for a simpler approach.

The robust optimization approach is another common approach that does not require any special assumptions on the uncertainty other than knowing the bounds. We discuss the robust optimization approach, that optimizes the worst case, with heuristic solution methods for this problem as well as other heuristic algorithms in the next section (Sect. 4).

Example 3 demonstrates the considered model for periods $T_0, \ldots, T$ compared to the real viewing capacity in period T_0.

Example 3 Considered viewing capacity that is higher than the real viewing capacity

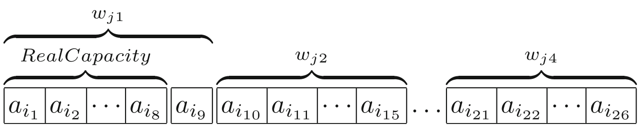

The considered viewing capacity in this example was higher than the real viewing capacity. Thus, only part of the planned ad playlist for the first week was watched by the viewer and reallocation is needed, e.g. a_{i_9} was not watched as planned.

The real viewing capacity can be much different and there are two approaches to handle this: the first considers the uncertainty in advance, i.e. while solving the allocation for periods $T_0, \ldots, T$, while the second adapts the future allocations after some of the uncertainty has been discovered, i.e. during reallocating for periods $T_0 + 1, \ldots, T$. We used both of these approaches in our heuristics as described in Sect. 4.2.

4 Heuristics

Heuristics are very common for solving instances of GAP [17] such as the Ads Allocation problem. In general terms, heuristics exploit special characteristics of the problem to obtain good solutions within a reasonable time. In this section, we describe the heuristics we proposed for the two versions of the problem, i.e. the deterministic single-period and the stochastic multi-period.

4.1 The single-period deterministic problem

In the deterministic version of the problem we assume the viewing capacities are known in advance, i.e. for all the viewers the uncertainty factors are $u_j = 0$. This version can be relevant in situations where the viewing capacities are static and known in advance. For example, consider an array of static TV screens in public places where the "viewers" are the public TV screens and the viewing capacities are static and known in advance, e.g. 24 h a day. As mentioned before (see Sect. 3.3), since the deterministic multi-period version can be treated as a single-period problem, by multiplying the viewing capacities by the number of periods, we consider only the single-period version.

4.1.1 Basic heuristic

Our heuristics iteratively choose ads and pack them to viewers. In describing the heuristics, we use the following predicates which present a snapshot of the current allocation step:

- `Free_Viewer`(a_i, v_j): viewer v_j has at least $f_i \cdot l_i$ free viewing capacity to view ad a_i.
- `Available_Ad`(a_i): ad a_i has not yet been validly assigned to r_i viewers and it is available for allocation.

If for a given snapshot of the problem, these predicates are true as well as the predicates Relevant(a_i, v_j) (defined in Sect. 3.1) then, allocation of ad a_i to viewer v_j is possible.

The description of each of our proposed solutions follows the description of the basic ads allocation algorithm, Algorithm 1, which is the basis for all of them.

Algorithm 1 Basic Ads Allocation Algorithm

```
Input: $Ads$ and $Viewers$ sets
Rank $Ads$ [according to ads_rank_order]
for all $a_i$ in $Ads$ according to the rank order do
    $PossibleViewers_i \leftarrow \phi$
    for all $v_j$ in $Viewers$ do
        if Free_Viewer($a_i, v_j$) & Relevant($a_i, v_j$) then
            Add $v_j$ to $PossibleViewers_i$
        end if
    end for
    while Available_Ad($a_i$) & $PossibleViewers_i \neq \phi$ do
        Assign $a_i$ to the next viewer, $v_j$, in $PossibleViewers_i$
        Remove $v_j$ from $PossibleViewers_i$
    end while
end for
```

The algorithm iteratively chooses ads according to some ranking order. The ranking has significant impact on the performance and each heuristic uses different ranking orders. Each selected ad is allocated to viewers until it is no longer available, i.e. it is allocated to r_i viewers as required.

4.1.2 Greedy heuristic

This heuristic is greedy and prefers ads with the highest profit per second, i.e. highest $PaymentPerSec = payment/length$. This algorithm is different from the basic algorithm

(Algorithm 1) only in reference to the ad's ranking order, which is in a decreasing order of *PaymentPerSec* values. We denote this heuristic by *PaymentPerSec*.

4.1.3 Target populations heuristic

There are ads with very specific target populations and others with wider target populations. This algorithm gives preference to ads with the most specific target populations. This is done under the assumption that as the target population is more specific it is more rare and the number of relevant viewers is smaller, whereas the more general the target populations the larger the number of relevant viewers. Thus, in order to be able to assign the specific ads and be more profitable the algorithm first assigns the specific ads and then the general ones. We denote this heuristic *TargetPopulation*.

The preference for ads with specific target populations is done by counting the number of attributes in the target populations profile with the value *All*, which we denote *CountAllValue*. As the number of *CountAllValue* is higher, the ad's target population is wider and the preference will be lower. This algorithm is different from the basic algorithm (Algorithm 1) only in the ads ranking order which is in an increasing order of the *CountAllValue* values.

4.1.4 Backtrack heuristic

This algorithm comprises three steps. First, it runs the basic allocation algorithm (Algorithm 1) with a decreasing ads ranking order of *PaymentPerSec* values and then *CountAllValue* values. Thereafter it runs a backtrack phase, finds the partially assigned ads and eliminates their allocations. Finally, it runs the basic allocation algorithm (Algorithm 1) again with the same ranking order, i.e. first according to *PaymentPerSec* values and then *CountAllValue* values. The idea is to try to fully assign ads with the viewing capacity of the partially assigned ads, since there is no payment for partially assigned ads. We denote this heuristic algorithm *BacktrackHeuristic*.

4.2 The multi-period uncertain problem

In this version of the problem there are uncertain viewing capacities and thus multiple periods are considered. We use two approaches in order to handle the uncertainty: the data approach and the algorithm approach. The data approach represents a modification of the problem data while using the same algorithms used for the deterministic problem. In contrast, in the algorithm approach the problem data remains the same while the algorithm is modified to handle uncertainty. We used the algorithm approach to build a new basic algorithm and the data approach in order to improve this basic algorithm. In the rest of this section, we first provide a general description of the sequential solution procedure we propose and then present the developed heuristic solutions.

4.2.1 Sequential procedure model

Ad campaigns usually extend over several weeks and it is possible to modify the allocations of ads to viewers during these weeks. Therefore, as mentioned in Sect. 3.2, this problem naturally falls into the category of a multi-period problem where after each period, when some of the uncertainty has been revealed, the ads can be reallocated. In this section, we describe in detail a sequential solution procedure for the problem that is used for all the heuristic algorithms presented. The procedure performs the following three steps per period sequentially from the first to the last period:

1. The deterministic allocation problem is solved for all the remaining periods, i.e. some viewing capacity values are assumed and the allocation problem is solved. The solution is an allocation of ads to viewers for all the remaining periods.
2. For each viewer, we determine a schedule to order the airing of the ads allocated to him, i.e. build a playlist for each viewer from the ads which have been allocated to him. We perform this using a *scheduler* algorithm.
3. Following the allocation and schedule processes, the ads from the playlist per viewer are aired for one time period. At the end of the period, the actual viewing is revealed, i.e. the uncertainty of the period is realized, and the remaining demands of frequency and rating per ad are updated.

The scheduler algorithm should take into account some special constraints of the TV medium, e.g. avoid too many airings of the same ad within a period, enforce minimum timespan between airings of the same ad, etc. The scheduling itself is also an NP-Hard problem which has been addressed in many papers, e.g. [5, 11]. The scheduling problem is not in the scope of this paper, and since we do not consider any specific scheduler or any schedule constraints we assume a random scheduler. Notice that if we would have any a priori knowledge regarding the scheduler we could manipulate it and suit an allocation algorithm such that the uncertainty effect would decrease and the revenue obtained would increase.

4.2.2 Basic uncertain heuristic

The basic uncertain allocation algorithm we used adapts the best allocation algorithm for the deterministic single-period version of the Ads Allocation problem to the multi-period setting, i.e. the *BacktrackHeuristic* algorithm (see results in Sect. 5.2). The adaptation takes into account that in the multi-period problem, the relevant airings of an ad depends on each viewer, as indicated by the f'_{ij} values in Problem (3). In addition, a new measurement for preference viewers was used, denoted *PaymentPerRemainSec*. This measurement is defined for each ad and viewer and represents the ad's payment for the total remaining seconds needed to be viewed by the viewer, i.e. $PaymentPerRemainSec = p_i/(l_i \cdot f'_{ij})$.

The *BasicU* algorithm runs the above algorithm (Algorithm 2) at the beginning of each period where the viewing capacity is assumed to be c_j per period. The *BasicU* algorithm was also used as a reference for solving the deterministic instance of the problem, i.e. the instance of the problem where the actual viewing capacities are known in advance (see Sect. 5.3).

Algorithm 2 Basic Uncertain Allocation Algorithm

```
Input: $Ads$ and $Viewers$
Sort $Ads$ in decreasing order of $PaymentPerSec$ values and then in decreasing order of $CountAll$ values.
for all $a_i$ in $Ads$ according to the order do
    POSSIBLE_VIEWERS ← all $v_j$ viewers which are Free_Viewer($a_i, v_j$) & Relevant($a_i, v_j$)
    Sort POSSIBLE_VIEWERS in decreasing order of $PaymentPerRemainSec$
    while Available($a_i$) & POSSIBLE_VIEWERS $\neq \phi$ do
        $v_j$ ← highest $PaymentPerRemainSec$ viewer
        Allocate $a_i$ to $v_j$
        Remove $v_j$ from POSSIBLE_VIEWERS
    end while
end for
```

4.2.3 Robust heuristics

We propose a series of algorithms all based on the robust optimization technique mentioned in Sect. 2.

The heuristic denoted *Robust* is an implementation of the robust optimization technique. It solves the worst case instance of the problem, where all the viewing capacities are equal to $c_j \cdot (1 - u_j)$. In other words, the input of the viewing capacity of the Robust algorithm is modified and the problem for all the periods is solved. Similar to the robust optimization technique, if the actual viewing capacity is longer than the lower bound considered in the robust planning, the extra time is not used since no allocations were prepared for such cases. The algorithm runs only once at the beginning of the first period and there is no need for reallocation at the beginning of each period as the original allocation will always be completed.

The second algorithm, denoted *AdaptiveRobust*, is similar to the robust algorithm with two major modifications. First, if the actual viewing capacity of a viewer is longer than expected, i.e. the worst case, the extra time is used and the allocations which were prepared for the next periods are aired to the viewer. In this manner there is no waste of viewing time but the allocations for the next periods may change. It is possible that all the original allocations of the viewer have already been viewed by him but there are more periods for viewing which can be used for extra allocations. Thus the second change is implementation of an adaptive approach. The Robust algorithm is called at the beginning of each period and not only once at the first period. That way the allocations are adapted to the actual viewing capacity as soon as they are revealed. The new allocation solution is used only if its revenue is higher than the original allocation's revenue, or else the original allocation solution is kept.

An interesting insight regarding this algorithm is that the revenue obtained will never be lower than the revenue of the original allocation. In other words, if at some time period the original allocation is replaced the revenue is guaranteed to increase. There are two reasons for this behavior. First, preparation for the worst case such that in each period all allocations can be fully allocated even in the worst case scenario. Second, future allocations will be adapted only if such adaptation will cause an increase in the revenue.

The third robust algorithm, denoted *LastRobust*, is implementation of the robust optimization technique only in the last period. It solves the remaining allocation problem at the beginning of each time period using the basic uncertain allocation algorithm, while the viewing capacity of each viewer is assumed to be c_j per period, though for the last period it is assumed to be $c_j \cdot (1 - u_j)$, i.e. the worst case. The motivation for this algorithm is based on the observation that in most of the cases the basic uncertain allocation algorithm is sufficient but its performance is dramatically damaged when the ads are not fully allocated due to uncertainty in the last period. In order to overcome this issue, the algorithm is prepared in advance for the worst case in the last period, such that the last period allocations will always be aired. Another motivation is that as the number of periods increases the average of the observed viewing capacity should converge to the expected capacity.

4.2.4 Modified ratings heuristic

The *ModifiedRatings* algorithm implements the data approach on the ads data. The ads rating requirements are modified, such that $r_i \cdot (1 + U_{a_i})$ is used instead of r_i, where U_{a_i} is the average uncertainty factor of all the relevant viewers of a_i. The viewing capacities remain the same, i.e. c_j.

The motivation of this algorithm is that in all cases, except the worst case, it is always possible that the prepared allocation will not be aired due to uncertainty. In order to neutralize the effect of such situations, i.e. cases where ads have not been fully allocated, this algorithm allocates the ads to more viewers than required. Thus, if for example there will be a dramatic change in the actual viewing capacity of a viewer and he will not view a specific ad there will be another viewer who has viewed the ad instead of the given viewer and the requirement for r_i full allocations will have been fulfilled.

4.2.5 *Combined uncertain heuristics*

The *CombinedHalf* algorithm combines both of the algorithms *ModifiedRatings* and *LastRobust* with some modifications. Combining the two algorithms as is would be considered an overkill, since the uncertainty is considered twice. Thus, only half of the uncertainty is considered in each of them. In other words, each rating value is changed to $r_i \cdot (1 + U_{a_i}/2)$, and each viewing capacity value of the last period is changed to $c_j \cdot (1 + u_j/2)$.

The *CombinedHalfLastPeriod* algorithm is another combination of modified ratings and the robust technique, shortly denoted as *CombinedHalfLP*. Throughout all the time periods this algorithm is similar to the *AdaptiveRobust* algorithm, excluding the last period where it is similar to the *CombinedHalf* algorithm. The motivation for this algorithm is based on the advantages of the *AdaptiveRobust* algorithm over all the periods excluding the last period. In the last period there are no future allocations that the *AdaptiveRobust* algorithm can use. Thus, the extra allocations performed by the *CombinedHalf* algorithm can help in the last period.

5 Experimental results

We tested the efficacy of our proposed algorithms through simulations. In this section, first we present the simulation environment, then the data used in the computational tests for each problem version and finally the results.

5.1 Simulation environment and data generation

We built a simulator to evaluate the efficacy of our proposed algorithms on synthetic data based on real distributions. The experiments were conducted as follows. First the problem instance was generated according to the required configurations. For the single-period deterministic version each of the allocation methods was run with the problem instance as input. For the multi-period with uncertainty version the sequential solution procedure was started whereby each iteration included (see Sect. 4.2.1): allocation of ads to viewers, random scheduling of the allocated ads, discovering the uncertainty of the current period and updating the problem instance. Note that we assume once a viewer begins to view an airing of an ad he will finish watching it. In other words, we assume that an airing of an ad cannot be cut in the middle. The algorithms were implemented in C# language using Visual Studio 2005 and .Net Framework version 2.0. All the experiments were run on Windows XP, on a 3.00 GHz with 4 GB RAM machine.

The generation of data for the different problem instances includes viewers and ads with values for all of their parameters. In addition, according to the viewers' uncertainty factors, the actual viewing capacity of each viewer in each period was randomly generated such that all the algorithms were tested on the same actual viewing capacities. The generation of data was done according to the following value distributions. The ad length values (in seconds), l_i, were {10, 20, 30, 40, 50, 60}, which were randomly selected according to real distributions given by SintecMedia [32]. The ad payment values, p_i, were {1, 5, 10, 15, 20}, which were randomly selected according to the distributions 10, 25, 30, 25, 10 %, respectively, e.g. the probability of the value 1 was 10 %. The ad frequency values, f_i, were {5, 10, 15} with equal distribution. The ad rating values, r_i, were {10, 30, 50} with equal distribution. The viewers' viewing capacity values per time period, c_j, were {300, 600, 900} seconds with equal distribution.

As for the profiles, we used three attributes each with a set of the following possible values: *Age*: {Adult, Child}, *Gender*: {Male, Female} and *Social Class*: {ABC1, ABC2}.

The viewers' profiles were distributed according to SintecMedia [32] and known social class distributions [12] and we did not consider variations of them. For the ad's target population profiles we defined three configurations:

- *General Case*: all ad's target population profiles are general, i.e. each ad can be allocated to each viewer. The values of all the profile attributes are *All*.
- *Specific Case*: each ad's target population is very specific, i.e. each ad is targeted to one combination of attribute values. The values of each profile attribute is randomly selected where all values have the same probability of being selected, i.e. 50 %.
- *Normal Case*: the more reasonable case where the target population profiles are a mix of the general and the specific cases:
  - *age*: {*Adult* = 70 %; *Child* = 10 %; *All* = 20 %}
  - *gender*: {*Male* = 30 %; *Female* = 30 %; *All* = 40 %}

- *social class*: $\{ABC1 = 30\ \%;\ ABC2 = 30\ \%;\ All = 40\ \%\}$.

For the multi-period with uncertainty version, the number of periods and viewers' uncertainty factors were also generated. The number of periods had the possible values $T = 1, 2, \ldots, 8$. For simplification, the uncertainty factors, u_j, of all the viewers were the same and the tested values were $0, 10, \ldots, 100\,\%$. The actual viewing capacity value per viewer and period was a uniform distributed random value from the range $[c_j \cdot (1 - u_j),\ c_j \cdot (1 + u_j)]$.

5.2 Deterministic problem results

In this section, we present the results of the deterministic version where there is no uncertainty and there is a single-time period.

We compare the results of the heuristics proposed with the results we obtained by solving the problem's IP model (see Sect. 3.2) using a generic IP Solver. We used ILOG CPLEX 11.1, with a runtime limit of 300 s, and use *CPLEX* to denote its solutions. Since some instances could not be optimally solved by *CPLEX* even without a runtime limit, we measured the performance of the algorithms as a percent of the upper bound of the optimal value given by *CPLEX*. As for the runtime, all of our heuristics ran within a few seconds, i.e. 1–3 s.

For each version of the problem we tested different configurations which are relevant to it. For the single-period deterministic version we tested all combinations of the ads' target population profiles with different combinations of the number of ads and the number of viewers. The number of ads and viewers that were tested included 1000 viewers and 100, 200, 300, 400 and 500 ads. For each combination we generated 100 random instances, and considered the average results with a confidence level of 95 % (the values on plots 1, 2, 3 are the average with the confidence intervals surrounding them). The most interesting results were achieved by the *ProfitPerSec*, *BacktrackHeuristic*, and *CPLEX* and thus we will focus only on those results.

In the normal case (Fig. 1), which is the most realistic, for a small number of ads, i.e. 100 and 200, *CPLEX* attained almost all the revenue possible, i.e. more than 99 %. Nonetheless as the number of the ads increased its performance dramatically dropped. The *PaymentPerSec* algorithm performs well for the smallest instances, i.e. 98 %, but badly for all the other instances, i.e. 93–94 %. The *BacktrackHeuristic* algorithm attained an average of 99 % and returned the best allocations for most of the instances (300–500 ads).

In the specific case (Fig. 2), as expected, all the algorithms' performance linearly decreased as the number of ads increased. The *CPLEX* algorithm returned the best allocations, more than 99 % on average. The *Backtrack-Heuristic* algorithm performed well with 96 % on average and the *PaymentPerSec* algorithm performed the worst with 93 % on average. The instances of this case seem easy to solve by *CPLEX*. The reason for this may be because each ad can be allocated only to a small subset of viewers, i.e. the number of options that need to be considered for each ad is small, thus CPLEX can quickly find a good allocation.

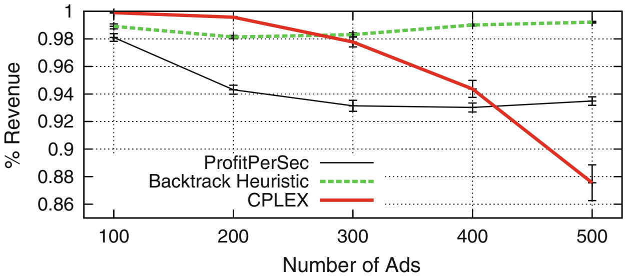

Fig. 1 Algorithms' performance—the normal case

In the general case (Fig. 3) the *BacktrackHeuristic* algorithm performed the best with more than 99 % on average. The performance of the *ProfitPerSec* algorithm was not far from this with 98 % on average. The interesting aspect is the performance of *CPLEX*. For the instances of 400 ads *CPLEX*'s performance dropped dramatically to 76 % on average, but then for the instances of 500 ads its performance increased to 93 % on average.

We also experimented with larger instances but the *CPLEX* algorithm was not able to solve them even after hours of run-time.

5.3 Uncertain problem results

In this section, we present the results of the uncertain version, i.e. single-period and multiple periods. We compared the revenue of the proposed solutions and all combinations of periods and uncertainty factors were examined with 100 ads, 300 viewers and ad's target population profiles according to the Normal Case (as described in Sect. 5.1).

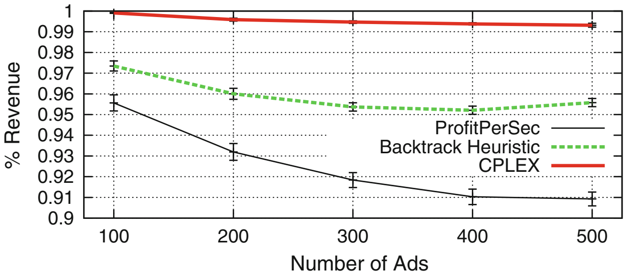

Fig. 2 Algorithms' performance—the specific case

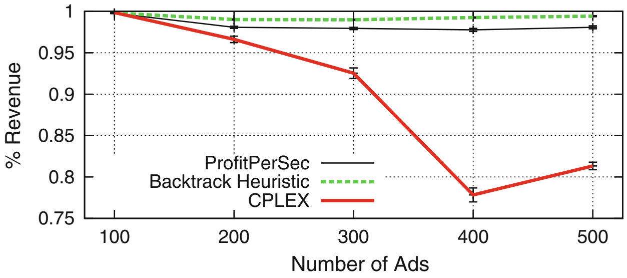

Fig. 3 Algorithms' performance—the general case

For each experiment combination we generated 60 random instances and considered the average result.

To compare the revenue obtained with each heuristic we normalized this value with respect to the revenue obtained by the *BasicU* algorithm for the deterministic instance in which the outcome of the viewing capacity was known. In other words, for each instance the obtained revenue values were normalized to the revenue of the *BasicU* algorithm for the same instance where the uncertain data was known in advance. Notice that although the *BasicU* algorithm solves the deterministic problem, it is only a heuristic algorithm and its solution is not guaranteed to be the best solution, i.e. the revenues obtained by some other heuristics are higher.

We will focus on the results which were achieved by the algorithms *BasicU*, *Robust*, *CombinedHalf* and *CombinedHalfLP*. The results of all the other proposed heuristics seem to follow the others but with lower achievements. The results of the *LastRobust* algorithm were always lower than those of *CombinedHalfLP* by 12.5 % on average; those of *ModifiedRatings* were always lower than those of *CombinedHalf* by 2 % on average; and those of *AdaptiveRobust* were always lower than those of *CombinedHalfLP* by 6.5 % on average. As for the runtime, all of our heuristics ran within a few seconds, i.e. 1–6 s.

5.3.1 Single-period

In the case of a single time period (Fig. 4) the results of *BasicU* were poor. Even in the cases where the uncertainty was low the results were poor and decreased to less than 10 % of the revenue already at an uncertainty factor of 30 %. The *Robust* algorithm performed the best with up to 40 % uncertainty, and for higher uncertainty both combined algorithms achieved the best results. When the uncertainty factor was 100 % the *Robust* algorithm achieved 0 % of the revenue since the viewing capacity considered was the worst case, namely 0 %, i.e. no viewing capacity at all. The combined algorithms achieved the same results for a single time period since the algorithms are the same for this instance.

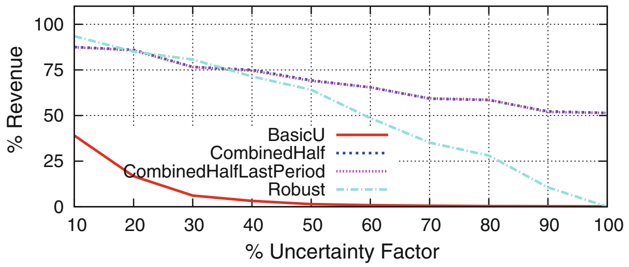

Fig. 4 Algorithms' performance—single period

5.3.2 Multi-periods

In the case of four time periods (Fig. 5) which is likely to be realistic, such as 4 weeks, the results were different from the results of uncertain with a single period (see Sect. 5.3.1). Up to 60 % uncertainty both combined algorithms performed very well achieving more than 90 % of the revenue, where the *CombinedHalfLP* performed the best. Between 60 and 100 % uncertainty the *CombinedHalf* algorithm performed the best where the increasing uncertainty had a minor effect on the results. Even for the case of 100 % uncertainty its results were higher than 85 % of the revenue. As before, *BasicU* algorithm's performance was very low and decreased to less than 50 % when the uncertainty factor reached less than 40 %.

In the case of eight time periods (Fig. 6) the problem seemed to be easy for cases of uncertainty lower than 60 %. As in the case of four periods, here also the *CombinedHalf* algorithm experienced a minor effect of the increasing uncertainty on the results which were higher than 99 % even for the case of 100 % uncertainty. The performances of the *CombinedHalfLP* algorithm were very high, up to 80 % uncertainty, and dramatically dropped with higher uncertainty values. The performances of the *Robust* algorithm were very high up to 50 % uncertainty and slowly dropped with higher uncertainty values down to 0 %.

In general, our algorithms obtained at least 87 % of the revenue with at least four periods compared to the results achieved with the common robust approach, i.e. the *Robust* algorithm. The *CombinedHalfLP* algorithm is preferred for

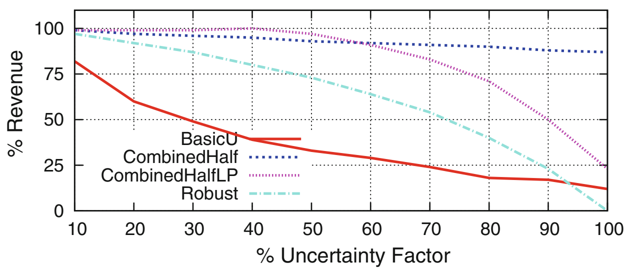

Fig. 5 Algorithms' performance—four periods

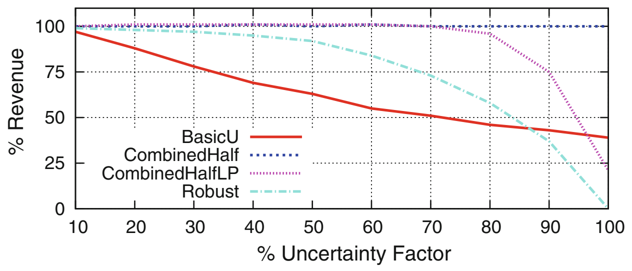

Fig. 6 Algorithms' performance—eight periods

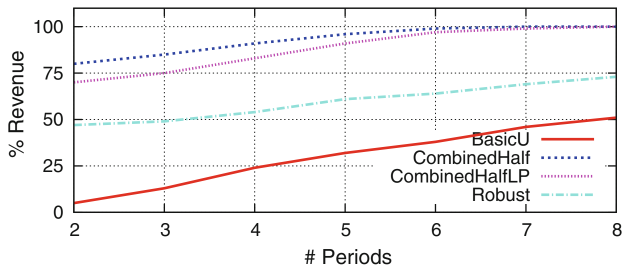

Fig. 8 Algorithms' performance—uncertainty factor of 70 %

lower uncertainty while the *CombinedHalf* algorithm is preferred for higher uncertainty, where the uncertainty break-even-point depends on the number of periods.

Another type of analysis, where the uncertainty factor was fixed, was also performed. In the case of 30 % uncertainty (Fig. 7) the performances of the algorithms differed from each other up to three periods. The *CombinedHalfLP* algorithm seemed to outperform the *CombinedHalf* algorithms, which in turn outperformed the *Robust* algorithm with an average difference of 5 %. For five periods and more both combined algorithms seemed to perform similarly attaining more than 98 % of the revenue.

In the case of 70 % uncertainty (Fig. 8), as in the case of 30 % uncertainty, the performances of the algorithms differed from one another, but in this case with up to six periods. The *CombinedHalf* seemed to outperform the *CombinedHalfLP* with an average difference of 10 %. For six periods and more the combined algorithms seemed to perform similarly achieving more than 97 % of the revenue. The *Robust* performances were bad, attaining at most 75 % of the revenue for eight periods.

The extreme case of 100 % uncertainty (Fig. 9) is interesting where the *CombinedHalf* algorithm reached over 80 % of the revenue already within three periods. This case demonstrates the immunity of the *CombinedHalf* algorithm even in extreme uncertainty. The *Robust* algorithm reached 0 % of the revenue for all periods, since the considered viewing capacity in the worst case is zero, i.e. no viewing capacity at all. The performances of the *BasicU* and the *CombinedHalf* algorithms were very poor with less than 40 % of the revenue even when eight periods were considered. The reason for the unexpected behavior of the *CombinedHalf* algorithm, which obtained less revenue as the number of periods increased, is probably due to the fact that the algorithm considers the uncertainty only for the last period as the *CombinedHalf* algorithm does.

In general, for all uncertainty factors our proposed algorithms indicate a significant improvement over the *BasicU* and *Robust* algorithms. Whereas for instances with at least four periods our algorithms obtain at least 87 % of the revenue, compared to the achievements according to the common robust approach, i.e. by the *Robust* algorithm, which obtains much less revenue.

We also checked the standard deviation (STDEV) of our heuristics in order to identify the impact of the uncertainty on each heuristic. We focused on four periods and 20, 40, 60 % of uncertainty cases. For each case we generated 30 random instances as described in Sect. 5.1, and for each instance we repeatedly generated random viewing capacities 30 times according to the uncertainty factors. The STDEV was calculated over all the 30 repeats and the averages over the 30 instances are presented.

As depicted in Table 2, the STDEV values of *Robust* were the smallest, lower than 0.65 %. Since this algorithm considers the worst case, it was less affected by the generated random capacities. A deeper glance into the details reveals that *Robust* obtained the same revenue over all the repeats of the same instance. This fact should have led to a 0 value for STDEV, but the difference in the revenue obtained for the deterministic instance for each repeat caused it to be greater. The STDEV values for the *CombinedHalf* and the *CombinedHalfLP* were quite similar

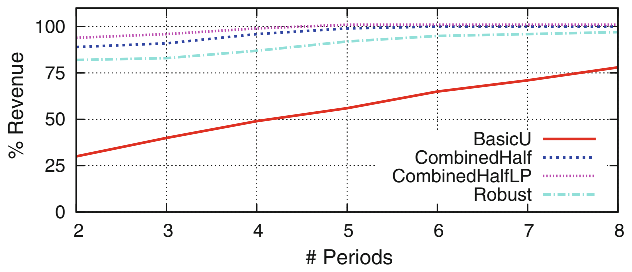

Fig. 7 Algorithms' performance—uncertainty factor of 30 %

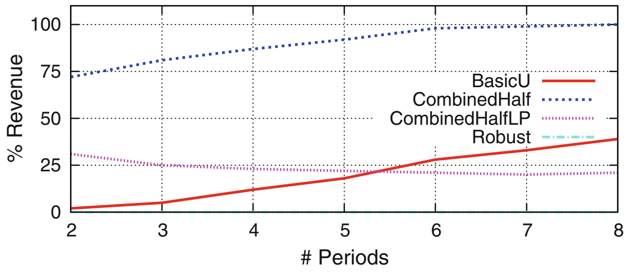

Fig. 9 Algorithms' performance—uncertainty factor of 100 %

Table 2 The average STDEV for instances with four periods and different uncertainty factors

Heuristics	Uncertainty factor (%)		
	20 %	40 %	60 %
BasicU	5.49	4.81	5.08
CombinedHalf	0.89	1.30	1.65
CombinedHalfLP	1.01	1.26	1.55
Robust	0.49	0.65	0.48

and low, i.e. lower than 1.65 %. These results support the claim that the combined heuristics are stable. As expected, the STDEV values of *BasicU* were significantly higher than the values of the other algorithms. These results can be explained easily since the algorithm does not take into account the uncertainty in advance.

6 Summary and conclusions

In this paper we considered the Uncertain Multi-Period Ads Allocation problem. We formally defined the problem and proposed a mathematical programming model. Since the problem is NP-hard, we presented a sequential solution procedure and proposed several heuristic algorithms to solve it. We have shown empirically that using natural heuristics for this multi-period problem it is possible to obtain a high percentage of the potential revenue, similar to the case where there is no uncertainty. We adopted two main proaches in order to handle the uncertainty: (a) adapt the allocations after some of the uncertainty is reveal and (b) consider the uncertainty in advance.

Our heuristics outperform the common robust approach for uncertainty problems and even for high uncertainty, e.g. 100 %, they obtained more than 75 % of the revenue. In general, we can conclude that instead of handling only the constraints which are uncertain, as is usually done, related constraints should also be considered even if they are not directly affected. For example, in the Ads Allocation problem, we considered the effects of the uncertain viewing capacities on the rating constraints. This approach, we believe, can be relevant also for other uncertain allocation problems.

In future research we would like to address the Ads Allocation problem under relaxation of the all-or-nothing rating and frequency constraints. The all-or-nothing constraints seem to have a tremendous affect on the problem while in reality minor violations can be ignored. Another interesting direction would be to replace the random schedule process with a realistic one. Combining the allocation and schedule algorithms will probably result in even higher revenues.

The research is part of the NEGEV Consortium [21] targeted at developing personalized content services. It was initiated and directed by *SintecMedia* [32], a High-Tech company that designs and implements management systems for TV broadcasting industries.

References

1. BARB: Barb reports: Monthly total viewing summary. http://www.barb.co.uk (2010)
2. Ben-Tal, A., Nemirovski, A.: Robust solutions of uncertain linear programs. Operat. Res. Lett. **25**(1), 1–14 (1999)
3. Ben-Tal, A., Nemirovski, A.: Robust solutions of linear programming problems contaminated with uncertain data. Math. Programm. **88**(3), 411–424 (2000)
4. Birge, J., Louveaux, F.: Introduction to stochastic programming. Springer, Berlin (1997)
5. Bollapragada, S., Garbiras, M.: Scheduling commercials on broadcast television. Operat. Res. **52**(3), 337–345 (2004)
6. Bozios, T., Lekakos, G., Skoularidou, V., Chorianopoulos, K.: Advanced techniques for personalized advertising in a digital TV environment: the iMEDIA system. In: Proceedings of the eBusiness and eWork conference. Venice, pp. 1025–1031 (2001)
7. Chekuri, C., Khanna, S.: A PTAS for the multiple knapsack problem. In: Proceedings of the eleventh annual ACM-SIAM symposium on discrete algorithms, pp. 213–222. Society for Industrial and Applied Mathematics, Philadelphia (2000)
8. Chorianopoulos, K.: Personalized and mobile digital TV applications. Multimed. Tools Appl. **36**(1), 1–10 (2008)
9. Dawande, M., Kalagnanam, J., Keskinocak, P., Salman, F., Ravi, R.: Approximation algorithms for the multiple knapsack problem with assignment restrictions. J. Comb. Optim. **4**(2), 171–186 (2000)
10. Dureau, V.: Addressable advertising on digital television. In: Proceedings of the 2nd European conference on interactive television: enhancing the experience, Brighton, UK (2004)
11. Fleming, P.J., Pashkevich, M.A.: Optimal advertising campaign generation for multiple brands using MOGA. IEEE Trans. Syst. Man Cybernet. **37**(6), 1190–1201 (2007)
12. Gilbert, D.: The American class structure in an age of growing inequality. Wadsworth Publishing, Belmont (2002)
13. Google AdWards: http://adwords.google.com (2011)
14. Juniper Research: Annual mobile adspend to pass $1 billion in 2008, with mobile TV to drive revenues to almost $7.6 billion by 2013. http://juniperresearch.com (2008)
15. Kastidou, G., Cohen, R.: An approach for delivering personalized ads in interactive TV customized to both users and advertisers. In: Proceedings of European conference on interactive television (2006)
16. Kim, E.M., Wildman, S.S.: A deeper look at the economics of advertiser support for television: the implications of consumption-differentiated viewers and ad addressability. J. Media Econ. **19**, 55–79 (2006)
17. Kundakcioglu, O.E., Alizamir, S.: Generalized assignment problem. In: Floudas, C.A., Pardalos, P.M. (eds) Encyclopedia of Optimization, pp. 1153–1162. Springer, Berlin (2009)
18. Lekakos, G., Giaglis, G.: A lifestyle-based approach for delivering personalized advertisements in digital interactive television. J. Comput. Mediat. Commun. **9**(2) (2004)
19. Lin, X., Janak, S., Floudas, C.: A new robust optimization approach for scheduling under uncertainty: I. bounded uncertainty. Comput. Chem. Eng. **28**(6–7), 1069–1085 (2004)

20. López-Nores, M., Pazos-Arias, J., García-Duque, J., Blanco-Fernández, Y., Martín-Vicente, M., Fernández-Vilas, A., Ramos-Cabrer, M., Gil-Solla, A.: MiSPOT: dynamic product placement for digital TV through MPEG-4 processing and semantic reasoning. Knowl. Inform. Syst. **22**(1), 101–128 (2010)
21. Negev consortium: http://www.negev-initiative.org (2011)
22. Nielsen Media Research: Advertising fact sheet. http://blog.nielsen.com (2010)
23. Nielsen Media Research: Snapshot of television use in the U.S. http://blog.nielsen.com (2010)
24. Nielsen Media Research: Three Screen Report, Q1 2010, vol. 8, http://blog.nielsen.com (2010)
25. NSR: Latest NSR Report Finds Mobile TV and Video Growth to be Driven by Broadcasters and Consumers. http://nsr.com (2008)
26. Nutov, Z., Beniaminy, I., Yuster, R.: A (1-1/e)-approximation algorithm for the generalized assignment problem. Operat. Res. Lett. **34**(3), 283–288 (2006)
27. Oksman, V., Ollikainen, V., Noppari, E., Herrero, C., Tammela, A.: Podracing: experimenting with mobile TV content consumption and delivery methods. Multimed. Syst. **14**(2), 105–114 (2008)
28. Park, J., Lim, B., Lee, Y.: A Lagrangian dual-based branch-and-bound algorithm for the generalized multi-assignment problem. Manag. Sci. **44**, 271–282 (1998)
29. Paulson, L.: TV comes to the mobile phone. IEEE Comput. **39**(4), 13–16 (2006)
30. Pramataris, K., Papakyriakopoulos, D., Lekakos, G., Mulonopoulos, N.: Personalized Interactive TV Advertising: The iMEDIA Business Model. Electron. Mark. **11**, 1–9 (2001)
31. Shmoys, D., Tardos, É.: An approximation algorithm for the generalized assignment problem. Math. Program. **62**(1), 461–474 (1993)
32. SintecMedia: http://sintecmedia.com (2011)
33. Toktas, B., Yen, J., Zabinsky, Z.: Addressing capacity uncertainty in resource-constrained assignment problems. Comput. Operat. Res. **33**, 724–745 (2006)
34. Velusamy, S., Gopal, L., Bhatnagar, S., Varadarajan, S.: An efficient ad recommendation system for TV programs. Multimed. Syst. **14**(2), 73–87 (2008)