Allocation Algorithms for Interactive TV Advertisements *

Ron Adany
Department of Computer Science
Bar-Ilan University
Ramat-Gan 52900, Israel
adanyr@cs.biu.ac.il

ABSTRACT

In this research we consider the problem of allocating personalized advertisements (ads) to interactive TV viewers. We focus on the optimization problem of maximizing revenue while taking into account the special constraints and requirements of the TV ads industry. The research is part of studies towards a Ph.D. in Computer Science currently in the third year of the four year planned study. In this paper we define the research problem, present the achievements attained to date, detail the research plan and discuss the contribution of the work.

1. INTRODUCTION

Personalization is the next-generation in the world of advertisement. It is very attractive to all players. From the commercial companies' perspective, it offers the possibility to tailor their advertisements to specific audiences and to ensure that the target population receives the desired ads in the desired format. From the standpoint of the service suppliers, i.e. the media companies and the operators, whose major source of income is advertisement [13], it is a way to increase revenue [7, 13]. And, from the viewers' perspective, it allows them to view ads which best suit their profile, preferences and interests.

Ads' personalization is already extensively used in the Internet medium (e.g. Google AdsWords [9]), but not in the TV medium. There are several key points that distinguish TV ads, as we consider them in this research, from Internet ads, such as the environment, the method of exposure, the pricing method, allocation constraints, etc. Over the past few years we have witnessed progress in technology, infrastructure upgrade, increased use of alternative TV screens, e.g. cell-phones, and the penetration of interactive TV. This progress has given rise to real personalized services, and the assignment of advertisement to specific viewers, based on their interests and their relevance to the advertised content. The personalization problem becomes even more important for the mobile TV platform where there is no uncertainty with respect to who is watching [1, 2].

Many studies concern the problem of selecting personal advertisements most suitable to each individual viewer, e.g. [12, 15, 16], and many others focus on how to deliver them, e.g. [4, 16]. Our research supplements these studies, by using their results as input with the goal of optimizing the allocation of ads. The issue of optimizing the personal TV advertisement problem is still an open problem for which, to date, no adequate solution has been proposed. 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 iMedia system framework, which is designed for personal advertisement in the interactive TV environment [4]. Based on frameworks such as iMedia, the entire process of the personalized advertisement is as follows. Given advertisement requests, ads are allocated to viewers and playlists of ads are generated for the planned time periods in some centralized computing center. Then, advertisement contracts are signed with the advertisers according to the allocations, and the playlists are delivered and stored in the Set-Top-Box (STB) units with which each viewer is equipped, as is common today. During the planned time periods, viewers watch TV, and on commercial breaks each STB airs ads based on the viewer's playlist. At the end of each time period, each STB sends an ads' viewing report to the centralized computing center, detailing the actual airing of the ads from the viewer's playlist. At the end of all the planned time periods the billing process is activated according to the signed contracts.

2. PROBLEM DESCRIPTION

The Ads Allocation problem concerns the allocation of personal TV advertisements to viewers. We consider two versions of the Ads Allocation problem. The deterministic version, where the problem's data is known in advance, is presented in Section 2.1. The uncertain multi-period version, where there are multi allocation periods and uncertainty about the problem's data, is presented in Section 2.2.

2.1 Deterministic Version

The input for the Ads Allocation problem consists of a set of ads and a set of viewers. Each viewer is associated

*The research is part of the NEGEV Consortium [17] targeted at developing personalized content services, and directed by ***SintecMedia*** [20], which is a High-Tech company that designs and implements management systems for the TV broadcasting, cable and satellite industries.

EuroITV 2011 Lisbon, Portugal

with a viewing capacity, and a profile attributed to him. Each ad is associated with a transmission length, a required rating, a required airing frequency, a profit and a profile defining the target population. 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 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 target population defines the set of viewers that are relevant for the ad.

An example of parameters of a viewer would be a viewing capacity of 20 hours a week with a profile of a male from London in the 35-40 age group. An example of an ad request would be a 30 second ad that needs to be allocated to 20,000 viewers, 10 times to each viewer, which will result in a profit of $10,000, and the target population is females from NYC in the 20-35 age group.

The goal of the Ads Allocation problem is to maximize the profit from a valid assignment of ads to viewers. A valid assignment that will result in payment should satisfies the ad rating and frequency requirements, does not exceed the viewers' viewing capacities and be personal, i.e. suits the ad's target population and viewers' profiles.

The deterministic version of the Ads Allocation problem is an extension of several well-studied optimization problems such as the General Assignment Problem (GAP) [14], the Multiple Knapsack Problem (MKP) [5], and the Multiple Knapsack problem with Assignment Restrictions (MKAR) [6, 19]. All of these problems are NP-hard and as an extension of them the Ads Allocation problem is also NP-Hard. Consequently, our proposed method for solving the problem is the heuristic approach (see Section 3.2) which is very common in solving instances of GAPs [14].

2.2 Multi-Period Uncertain Version

The multi-period uncertain version of the Ads Allocation problem is an extension of the deterministic case into a multi-period problem where the viewers' viewing capacities are uncertain. While data regarding the ads' requests, as well as data concerning the viewers' profile (e.g. by asking the viewers), are known in advance, the data on a viewer's viewing capacity is only a prediction of how much time a viewer will view TV within a certain period. Situations where viewers watch more or less time than expected are possible. The latter case, i.e. less viewing time than expected, is more problematic since there may be some ads that are not fulfilled which will cause a loss in revenue. However, the case of more viewing time than expected is also problematic, since knowing the actual viewing time in advance could result in allocation of more ads that in turn would increase revenue, which is our goal.

We assume that each estimated viewing time, c_j, of viewer v_j is given together with some uncertainty factor $0 \leq u_j \leq 1$, where the "real" viewing capacity is a value in the range of $[c_j * (1 - u_j), c_j * (1 + u_j)]$. This model allows more realistic representation of the TV viewers, since their viewing capacity is not stable but can be estimated with a bounded error. The estimated viewing times can be based on viewing statistics. Such statistics are already available for some defined viewer groups. For example, according to BARB [3], in 2010 the average weekly viewing per person in Great Britain was 28:13. According to Nielsen [18], in 2009-2010 the average American watched 35:34 (hours/minutes) of TV per week, kids ages 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).

In order to tackle uncertainty using an iterative approach of reallocation in future periods, the problem version is defined over multi-periods, where the ads allocations can be split over several periods instead of a single one, e.g. 4 weeks. This modification allows us to handle the uncertainty not only before it is revealed but also after.

3. RESULTS

We have been working on this research for almost three years. To date we have revealed several interesting results and published two papers. We provide a brief description of the theoretical results in Section 3.1 and in Section 3.2 we present a brief summary of the experimental results for the deterministic and multi-period uncertain versions.

3.1 Theoretical Results

We have presented the problem hardness and have developed several polynomial time approximation schemes (PTAS) for special instances of the deterministic version. A PTAS algorithm takes the problem instance and $\epsilon > 0$ as input and returns a solution with an approximation ratio of at most $(1 + \epsilon)$ from the optimal, i.e. a solution which is worse than the optimal by at most a factor of $(1 + \epsilon)$. The complexity of PTAS is polynomial in the instance size, but may be exponential in $1/\epsilon$.

We have shown that the problem is strongly NP-hard using a reduction from the 3-*partition* problem, which is strongly NP-hard [8]. This reduction holds even for instances where ad profits, ratings and frequency are equal to 1 and all ads can be allocated to all viewers, i.e. there is no personalization according to profiles. In addition, we have proven that personalization according to arbitrary profiles, i.e. arbitrary assignment restrictions, makes the problem APX-Hard, i.e. the problem has no PTAS unless $P = NP$. This proof was done using a reduction from the maximum 3-bounded 3-dimensional matching (3DM-3), which is APX-hard as shown in [11], and it holds even for instances where ad lengths, profits, ratings and frequency are equal to 1.

Although the problem has no PTAS, as described above, we developed PTASs for special instances of the problem:

- Instances of the Ads Allocation problem with no assignment restrictions and uniform ad lengths;
- Instances of the Ads Allocation problem with no assignment restrictions, ad rating values that are taken from a constant number of options, and ad length values that are a power of 2;
- Instances of the Ads Allocation problem with a bounded number of assignment restrictions, i.e. a constant number of profiles, and uniform ad lengths.

All of the PTASs we developed are based on generalizations of the PTAS for the multiple knapsack problem given in [5] and contain two steps: (1) selection of the ads to be assigned and (2) assigning the selected ads to viewers.

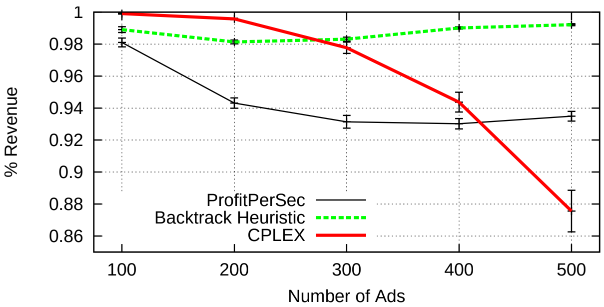

Figure 1: Algorithms' performances for the deterministic version over instances of 100-500 ads.

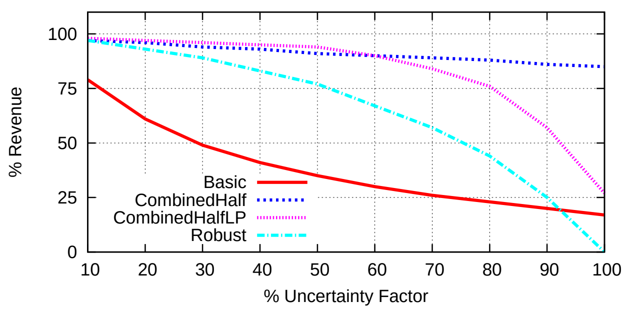

Figure 2: Algorithms' performances: average results for the multi-period uncertain version over instances of 1-8 periods.

3.2 Experimental Results

As a real world problem motivated by the industry, the focus of the research has been on developing algorithms that can be used and implemented (in contrast to pure theoretical research). Since the Ads Allocation problem is NP-Hard we developed heuristic algorithms which are common in solving instances of NP-Hard problems.

3.2.1 Deterministic Version Results

We developed several heuristic algorithms and evaluated them using simulations. Considering the size of the problem, i.e. thousands of ads and millions of viewers, we could not use an IP (Integer Programming) solver to solve the problem. In order to evaluate our heuristics we reduced the problem instances size and compared the results to those of a state-of-the-art IP solver, i.e. IBM ILOG CPLEX [10]. For the tested instances our heuristics returned results, within a few seconds, which were very close to the upper bound of the optimal value given by the solver (for some instances we could not find the optimal solution using CPLEX even without a runtime limit). Since realistic problem instances are much more, e.g. millions of viewers, with which CPLEX is unable to deal, the heuristic solutions, displaying very good performance, seem to be a good solution.

The instances we tested include different ratios of ads per viewers, and different kinds of ad profiles, e.g. general ads that are relevant to all viewers vs. specific ads relevant to a unique target population. The heuristics we propose can be split into three categories: payment oriented, target population oriented and backtrack oriented. One of the interesting heuristics we propose is the $BacktrackHeuristic$ algorithm which takes into account the personalization level and payment of the ads in addition to the backtrack process. Another interesting heuristic is the $ProfitPerSec$ algorithm, a greedy heuristic that prefers ads with a high profit per sec. Performance results of these algorithms are presented in figure 1. Our $BacktrackHeuristic$ outperforms the other heuristics and the IP solver, denoted as $CPLEX$, and on average attains 98% of the possible revenue. For the full and detailed results for this problem version see [1].

3.2.2 Multi-Period Uncertain Version Results

This version of the 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. Therefore, we present a sequential solution procedure for the problem and propose several heuristic algorithms for solving it. Through computational experiments on different scenarios of uncertainty and number of periods, the performances of the heuristics were normalized to the performance in the deterministic case where the actual viewing capacities are known in advance.

The heuristics we propose can be split into two categories: a robust heuristic which considers the worst case of viewing capacities and a modified rating heuristic which manipulates the ad ratings. Some of the interesting heuristics we propose adapt to the uncertainty and combine these two approaches. The $CombinedHalf$ heuristic considers the uncertainty of all periods in advance and allocates each ad to more viewers than needed, whereas the $CombinedHalfLP$ heuristic considers only the uncertainty of the last period. We use $Basic$ to denote the performance of the heuristic which only adapts to the uncertainty and does not consider it in advance.

As can be seen in figure 2, for higher uncertainty the performance of $CombinedHalfLP$ declines dramatically while the performance of the $CombinedHalf$ remains high, i.e. attaining above 85% of the revenue. Both algorithms obtain at least 87% of the revenue when the uncertainty is less than 60%, with some advantage to the $CombinedHalfLP$ heuristic. In general, it seems that the $CombinedHalfLP$ heuristic is preferred for lower uncertainty and the $CombinedHalf$ heuristic for higher uncertainty, where the break-even-point depends on the number of periods. For the full and detailed results of this problem version see [2].

4. RESEARCH PLAN

During the current year and the upcoming fourth year of research we plan to continue to investigate the problem, develop new algorithms and extend the evaluations. The detailed research plan, described below, includes completed, current and future work.

- Develop heuristics for the deterministic version of the Ads Allocation problem, and evaluate and compare them to the CPLEX IP solver. **Completed and published (see [1]).**

- Develop heuristics for the multi-period uncertain version of the Ads Allocation problem and evaluate them. **Completed and published (Best paper award winner, see [2]).**

- Develop approximation algorithms for the deterministic version of the Ads Allocation problem. Such results provide bounds on the attainable revenue resulting in the guaranteed quality of our heuristics. We already have several theoretical results including hardness results and polynomial time approximation schemes for special instances of the problem (see Section 3.1). **Current work, partially completed.**

- Extend the evaluation of the multi-period uncertain solutions for the Ads Allocation problem under different types of data environments, for instance, by altering the number of ads, the number of viewers, the uncertainty factors, the number of periods, etc. In addition, collect real data regarding viewers' viewing capacities and evaluate the solutions using this data. **Current work.**

- Address the Ads Allocation problem under relaxation of the all-or-nothing rating and frequency constraints. The all-or-nothing constraints, e.g. the request to allocate the ad to the exact number of required different viewers, seem to have a tremendous effect on the problem while in reality minor violations can be ignored. **Current work.**

- Consider other special constraints of interactive TV, such as past interactions with viewers, current watched content, user interactive limitations (e.g. problematic back-channel for participation TV or television commerce services), etc. **Future work.**

5. RESEARCH CONTRIBUTION

Since the aim of this research is to develop new algorithms to allow optimized ad personalization in the interactive TV environment and other enhanced TV mediums, its contribution to the interactive TV industry is consequential. This research will allow the industry to maximize revenues from advertising as well as deliver more relevant and interesting advertisements to the viewers. While other studies focus on selecting ads most suitable to the viewers, this research focuses on optimizing such allocations given the suitable ads for each viewer. As far as we know there is no other work underway which presents solutions to this problem while taking into account all the special constraints of the TV industry.

Along with its contribution to the industry, this work can also be relevant to other domains and industries faced with similar assignment problems, e.g. packing of containers with multi all-or-nothing constraints.

In addition to our heuristics algorithms we also present theoretical work which contributes to the theoretical investigation of the General Assignment Problem (GAP), the Multiple Knapsack Problem (MKP), and the Multiple Knapsack problem with Assignment Restrictions (MKAR).

6. REFERENCES

[1] R. Adany, S. Kraus, and F. Ordonez. Personal Advertisement Allocation for Mobile TV. In *International Conference on Advances in Mobile Computing & Multimedia*, 2009.

[2] R. Adany, S. Kraus, and F. Ordonez. Uncertain Personal Advertisement Allocation for Mobile TV. In *International Conference on Advances in Mobile Computing & Multimedia*, 2010.

[3] BARB. Barb reports: Monthly total viewing summary. `http://www.barb.co.uk`, 2010.

[4] T. Bozios, G. Lekakos, V. Skoularidou, and K. Chorianopoulos. Advanced Techniques for Personalized Advertising in a Digital TV Environment: The iMEDIA System. In *Proceedings of the eBusiness and eWork Conference*, pages 1025–1031, Venice, Italy, 2001.

[5] C. Chekuri and S. Khanna. A PTAS for the multiple knapsack problem. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, pages 213–222. Society for Industrial and Applied Mathematics Philadelphia, PA, USA, 2000.

[6] M. Dawande, J. Kalagnanam, P. Keskinocak, F. Salman, and R. Ravi. Approximation Algorithms for the Multiple Knapsack Problem with Assignment Restrictions. *Journal of Combinatorial Optimization*, 4(2):171–186, 2000.

[7] V. Dureau. Addressable advertising on digital television. In *Proceedings of the 2nd European conference on interactive television: enhancing the experience*, Brighton, UK, March–April 2004.

[8] M. R. Garey and D. S. Johnson. *Computers and In-tractability. A Guide to the Theory of NP-Completeness*. W.H. Freeman, New York, 1979.

[9] Google AdWards. `http://adwords.google.com`.

[10] IBM ILOG CPLEX Optimizer. `http://ibm.com`.

[11] V. Kann. Maximum bounded 3-dimensional matching is MAX SNP-complete. *Information Processing Letters*, 37(1):27–35, 1991.

[12] G. Kastidou and R. Cohen. An approach for delivering personalized ads in interactive TV customized to both users and advertisers. In *Proceedings of European conference on interactive television*, 2006.

[13] E. M. Kim and S. S. Wildman. A deeper look at the economics of advertiser support for television: the implications of consumption-differentiated viewers and ad addressability. *Journal of Media Economics*, 19:55–79, 2006.

[14] O. E. Kundakcioglu and S. Alizamir. Generalized assignment problem. In C. A. Floudas and P. M. Pardalos, editors, *Encyclopedia of Optimization*, pages 1153–1162. Springer, 2009.

[15] G. Lekakos and G. Giaglis. A Lifestyle-Based Approach for Delivering Personalized Advertisements in Digital Interactive Television. *Journal of Computer-Mediated Communication*, 9(2):00–00, 2004.

[16] M. López-Nores, J. Pazos-Arias, J. García-Duque, Y. Blanco-Fernández, M. Martín-Vicente, A. Fernández-Vilas, M. Ramos-Cabrer, and A. Gil-Solla. MiSPOT: dynamic product placement for digital TV through MPEG-4 processing and semantic reasoning. *Knowledge and Information Systems*, 22(1):101–128, 2010.

[17] Negev consortium. `http://www.negev-initiative.org`.

[18] Nielsen Media Research. Snapshot of television use in the u.s. `http://nielsen.com`, September 2010.

[19] Z. Nutov, I. Beniaminy, and R. Yuster. A (1–1/e)-approximation algorithm for the generalized assignment problem. *Operations Research Letters*, 34(3):283–288, 2006.

[20] SintecMedia. `http://sintecmedia.com`.