Advertising Revenue Optimization in Live Television Broadcasting

Pascale Crama[†], Dana G. Popescu[‡], Ajay S. Aravamudhan[†]

Abstract

In live broadcasting, the break lengths available for commercials may not always be fixed and known ex ante (e.g., strategic and injury time-outs are of variable duration in live sport transmissions). Because advertising represents a significant share of the broadcasters' revenue, broadcasters actively manage that revenue by jointly optimizing their advertising sales and scheduling policies. We characterize the optimal dynamic schedule in a simplified setting that incorporates stochastic break durations and advertisement lengths of 30 seconds and 15 seconds. The optimal policy is a greedy look-ahead rule that takes the remaining number of breaks into account. Under this setting, we find that there is no value to perfect information at the scheduling stage and knowing the duration of all the breaks will not change the schedule. When we incorporate diversity constraints (i.e., two ads from the same advertiser or for competing products cannot be shown during the same break), we characterize the optimal policy for a restricted set of stochastic break lengths. This policy combines the logic of the greedy look-ahead rule with the necessity to maintain an acceptable level of diversity in the ad portfolio. Finally, we also present heuristics that can be used to solve scheduling problems of greater complexity, and we recommend ways for broadcasters to balance their portfolio of booked ads. We run simulations, to test the performance of the heuristics under various scenarios. We find that two heuristic: myopic greedy and dynamic modified certainty equivalent (DMCE) perform close to optimal and outperform other heuristics currently used in practice.

Keywords: live broadcasting, advertising, scheduling, random capacity.

1 Introduction

Broadcasters generate a large part of their revenue through advertising. At CBS, the most watched US broadcast network, TV advertising accounted for two thirds of the total revenue (Bloomberg Businessweek 2010). Major sporting events—such as the Super Bowl, the Olympics, and the FIFA World Cup—strongly boost such revenues because advertisers are willing to pay a premium for their ads to air during the live broadcast of these events. In 2010, for instance, the cost of a 30-second spot during the Super Bowl was between $2.5 and $2.8 million, or 18 times higher than the corresponding prime-time advertising rates. Similarly, a 30-second spot during the Winter Olympics in the same year generated between $360,000 and $490,000, which was about 3 times the rate of an average prime-time spot (Bauer Insight 2010).

While live broadcasting of major sporting events can significantly boost revenues, selling and scheduling advertisements in that environment can be a challenging task, especially for sports events that involve unpredictable breaks during which ads can be shown. A case in point is cricket, a major sport in South Asia, whose matches have breaks of random duration

[†]Department of Operations Management, Singapore Management University
[‡]Department of Technology and Operations Management, INSEAD

in the action. The uncertainty about the duration of breaks creates an obvious problem for the broadcaster, namely how to schedule (live) the ads that have been sold while respecting the constraints on the schedule. The *diversity* constraints, i.e., two ads from the same advertiser or for competing products cannot be shown during the same break, which are commonly found in advertisement scheduling, are augmented by *capacity* constraints, i.e., the total duration of the ads scheduled during a break must not exceed the length of that break.

Suboptimal or infeasible schedules have many undesirable consequences for a broadcaster. If the schedule does not allow an ad to be shown in its entirety or if the schedule violates diversity constraints, no revenue will be earned and capacity will be wasted. A schedule that violates capacity constraints could lead to rescinding of the broadcast rights or other costly penalties, e.g., cricket broadcasting rights require the broadcaster to guarantee live coverage of every ball of every match. Moreover, showing an excessive number of ads at inopportune times will displease viewers and lead to lower future ratings. Thus, to generate the maximum possible revenue from live events, we look for the optimal dynamic scheduling policy under various scenarios of capacity and diversity constraints.

We model a television network that has a stochastic capacity of advertising airtime during a live event. This capacity consists of a number of commercial breaks of random duration. Breaks occur sequentially over a period of time and must be filled immediately upon arrival. Once a break occurs, its duration becomes known to the scheduler. We take as given the portfolio of booked ads which are to be aired during the live event. The ads have variable length and yields. For tractability, but also for practical relevance, we analyze a setting with ads of two lengths, 15 and 30 seconds.

In the base case, absent any diversity constraints, the optimal policy is a *greedy look-ahead* rule that takes the remaining number of breaks into account. Two surprising characteristics of the optimal policy are interesting to note. First, the optimal policy does not depend on the probability distribution of the break duration. This is counterintuitive, as one might expect the distribution of the remaining capacity to play a role. Second, and most importantly, perfect information is of no value; in other words, advance knowledge of the duration of all future commercial breaks does not change the network's revenue or schedule. Finally, we find that the optimal scheduling policy is not affected by service level penalties that are proportional to the ad yield or length.

Incorporating diversity constraints into the scheduling problem substantially complicates the scheduling algorithm. However, we are able to derive the optimal policy when the break durations mirror the ad lengths, and show that there is no value to perfect information and

---

In cricket, two batsmen attempt to score runs against the fielding team. The fielding team's bowlers throw six balls in succession, called an 'over', from opposite ends of the field. The fielding team can rearrange the players' positions in the field between every over, and ads can be shown during that time. As soon as the players have taken up their new positions, the game re-starts and the broadcaster resumes the live coverage of the game.

Contracts between broadcaster and advertiser typically specify that the advertiser will pay only if its ad is shown in full and not in a commercial break during which the same ad—or one for a competing product—is shown.

For instance, in 2011 the Indian government issued a show-cause notice to the Ten Cricket channel for violating the country's advertising codes during its coverage of India's tour of South Africa, claiming that the broadcaster's ads had interfered with the program (ESPN Cricket Info, http://www.espncricinfo.com/).

This occurred during the 2008 Summer Olympics: the Australian network Seven's coverage was widely criticized on these grounds.

For a general description of the pricing and ad sales process in the US television advertising market see Bollapragada et al. (2002) and Phillips and Young (2010).

This assumption reflects the US market, in which more than 90% of the ads sold are in one of these two formats. For instance, data collected from three US cable networks and containing all plans sold by the networks on the upfront and scatter markets during 2006-2008, revealed that more than 94% of the ads were in either the 15-second or the 30-second format.

Nevertheless, such perfect knowledge would allow the broadcaster to improve its ad portfolio's composition in terms of the relative proportions of short and long advertisements.

the policy does not depend on the break length distribution. When the break durations are distributed over more than two values, the problem of optimal scheduling with diversity constraints becomes analytically intractable. Because ad scheduling transpires in real time during a live event, we therefore seek to derive simple and efficient heuristics that are fast and easy to implement. In Section 7 we propose several heuristics and compare their performance under scenarios characterized by various revenue ratios for long and short ads and overbooking levels.

A comparison of the expected revenue under the optimal policy (or perfect information, if optimal policy is intractable) and the greedy heuristic shows that the latter performs commendably well in many situations. Together with a clear understanding of the circumstances in which the greedy heuristic might fail, this result shows that the broadcaster does not lose much value by applying this simple algorithm. The greedy heuristic is adversely affected when short ads are selling at a *premium* (i.e., the yield of a short ad is, on average, higher than half the yield of a long ad): the revenue under the greedy heuristic might even decline as the premium on short ads increases and the total value of the portfolio increases. This results from the suboptimality in the scheduling which outweighs the benefit from the increased value of the short ads. In the presence of a diversity constraint, the performance of the greedy heuristic is further adversely affected when the ad portfolio displays a high *concentration* in the low-priced short ads, i.e., the low-priced short ads belong mostly to one advertiser or product category.

Finding the optimal solution to the scheduling problem described above, also allows to consider two more fundamental questions. First, the broadcaster has to decide how much airtime to sell. Random capacity and high prices push the broadcaster to sell in excess of airtime capacity: this lowers the service level (i.e., the ratio of ads aired to ads sold), which will lead to advertiser dissatisfaction and—in the case of contractual guarantees—to penalties. Selling less than the available airtime capacity, however, causes underutilization and a loss of revenue. In the presence of penalties, the broadcaster will have to choose his level of overbooking carefully to balance the trade-off between expected benefit and penalty payments. Second, the broadcaster must also consider the ad portfolio's diversity in terms of ad duration and number of advertiser or product categories. The portfolio composition plays a role when the broadcaster is scheduling ads for a live event because a judicious composition of the portfolio can help the scheduling policy perform better under capacity and diversity constraints. Taking ad prices as exogenous input to the model, we look for the ideal mix of short and long ads to sell depending on their respective revenue and conditional on implementing the optimal policy at the scheduling stage. We also investigate when high advertiser concentration (i.e., high percentage of ads sold to the same advertiser) becomes detrimental to revenue.

These insights on the portfolio composition are of paramount importance to the first stage of contract negotiation and ad sales. First, we find that the level of concentration can have a substantial effect on the total revenue. High advertiser concentration increases scheduling difficulties; with insufficient diversity, it may be impossible to schedule short ads in long breaks causing low service levels *for* and revenues *from* for such ads. This effect is particularly pronounced for high concentration of high-paying short ads.

Second, we consider the composition of the portfolio in terms of long and short ads. Because short ads enhance scheduling flexibility, we find that, in the absence of overbooking, a broadcaster should consider selling more short ads than the expected number of short breaks (or conversely, fewer long ads than the expected number of long breaks), even if short ads generate, on average, significantly less than half the revenue of long ads. The higher the variability in break duration, the larger the discount on short ads the broadcaster is willing to accept in order to retain the scheduling flexibility afforded by a higher number of short ads.

It is interesting to note the opposite effect that the ratio of long to short ad *revenues* has on the simplicity of the scheduling and portfolio composition problems. A high ratio ($\gg 2$) reduces the optimal scheduling policy to a simple *myopic greedy* algorithm, but it makes the portfolio composition problem more challenging as it is not clear how many of each ad types to sell; that

is, the optimal sales ratio is not self-evident: long ads are profitable but reduce the scheduling flexibility. For a low ratio ($\ll 2$), the optimal scheduling policy is non-trivial and the myopic greedy algorithm will perform poorly, in terms of total revenue, but the optimal sales ratio is evident: the broadcaster's goal is then to sell as many short ads as possible.

The rest of this paper is organized as follows. Section 2 reviews the literature. In Section 3, we set up the model (with two ad durations) and in Section 4 we derive and interpret the optimal scheduling policy. Section 5 presents extensions of the model that address break duration, contractual penalties and advertiser diversity. In Section 6, we address the portfolio composition problem and in Section 7 we describe several heuristics and test their performance relative to the case of perfect information. We conclude in Section 8. All (nontrivial) proofs are in the Appendix.