OPTIMAL BUDGET ALLOCATION STRATEGIES FOR REAL TIME BIDDING IN DISPLAY ADVERTISING

Pavankumar Murali, Ying Li, Pietro Mazzoleni, Roman Vaculin

IBM T. J. Watson Research Center
1101 Kitchawan Road
Yorktown Heights, NY 10598, USA

ABSTRACT

We present a simulation model to determine optimal budget allocation strategies for real-time bidding (RTB) based display advertising. A common challenge across RTB exchanges is to optimize both budget spend and performance attainment. Our simulation model uses a stochastic dynamic programming approach based budget allocation to determine budget for each time instant. We report on results from a real-world pilot in which our approach delivered an average 18% performance gain.

1 INTRODUCTION

Real Time Bidding (RTB) is the process in which opportunities for ads placements (aka impressions) are auctioned to advertisers via programmatic auctions similar to financial markets. In RTB, an advertiser's ability to spend for an impression and, hence, their bidding price is strongly influenced by the available budget. Current bidding policies either use uniform pacing (equal allocation of budget to all time slots) or no pacing at all, resulting in low click or conversion rates (defined as the proportion of users who click on the ads or take another action such as purchase as a result of the won impression) [1]. Alternatively, we use a stochastic dynamic approach that allocates budget to time slots, based on historical win rates (defined as the ratio of impressions won to bids) and conversion rates. That is, for a given advertising campaign, we allocate higher budget to time period with a higher expected conversion rate. Here, we present a simulation methodology to test out this performance-based budget allocation algorithm and present results from a real-world pilot that uses this approach.

2 ESTIMATION AND SIMULATION MODELING FOR BUDGET ALLOCATION

We use a performance-based strategy to allocate more budget on time slots where an ad campaign has a higher likelihood to attain more clicks or conversions. For this, we measure the historical performance of the campaign during each time slot. Based on this, we build a discrete probability density function described by a list of conversion probabilities: $p_0, ..., p_T$ assuming T time slots per day, and $\sum_t^T p_t = 1$. At each time slot, we compute the budget allocated for the next time slot as follows:

$$b_{t+1}^p = \left(B - \sum_{m=1}^{t} s(m)\right) \frac{p_{t+1}.L(t+1)}{\sum_{m=t+1}^{T} p_m.L(m)} \qquad (1)$$

Where, B is the total budget available for the ad campaign, s(t) is the total money spent in time slots 0 to t, L(t) is the length of time slot t. To simulate the RTB process and estimate the bidding performance, we modeled the PDF of the following four variables, namely, bidding price, winning rate, clearing price and conversion rate, for a given time slot. Specifically, assuming that we want to allocate an ad campaign G's daily budget into n time buckets, then for each time bucket, we use historical bidding data of that particular time bucket to model the PDF of each variable by fitting the best statistical distribution. Figure 1 shows an example of the estimated PDFs of the four variables. Note that there are various ways of determining the n time buckets, where the buckets could be of same length or different lengths. The best time bucket definition is chosen based on simulation results.

978-1-4673-9743-8/15/$31.00 ©2015 IEEE

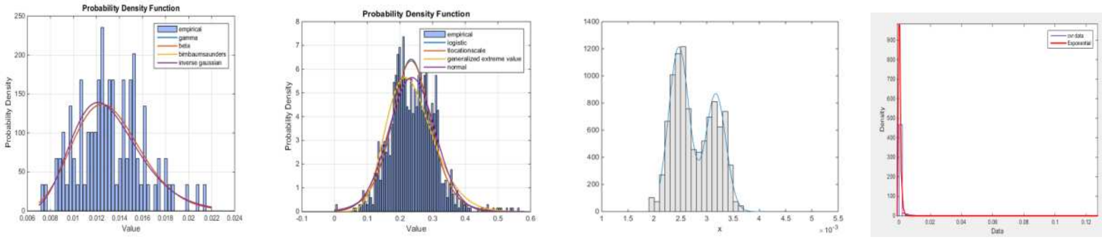

Figure 1: Examples of estimated PDFs of bidding price, winning rate, clearing price and conversion rate, respectively, where bar graph indicates the actual data, and superimposed curves are fitted distributions

A discrete-event simulator is built to estimate the number of bids made, the number of impressions won and the number of conversions for each time slot. For each time slot, the number of conversions reported is the average over 10,000 simulation runs. In each run, the bid amount, win rate and conversion rate for time slot t is sampled from the distributions estimated using the procedure explained above. The budget to be allocated to time slot t is computed using equation (1).

3 PILOT DESIGN AND RESULTS

Our budget allocation approach was piloted over 3 weeks at a Fortune 100 company for some of their Enterprise Cloud ad campaigns via a demand side platform (DSP). We used an A/B testing approach, where for each campaign we randomly split the audience into a control group and same sized test groups. For the test groups bidding parameters were adjusted based on the outputs of our simulation while the control group had parameters manually set by the marketers.

Table 1 tabulates the pilot results, where we report the RTB performance in terms of impressions bid, impressions won, total cost, total conversions, and cost per conversion, for each pair of test and control group. Note that out of the 3 campaigns selected for pilot, two of them are non-aggressive groups (historical spend < 50% of allocated budget); while the third campaign is aggressive (historically spends > 80% of the budget). We have also reported the gain or loss of each test group in terms of cost per conversion (CPC) as compared to the corresponding control group. From the table we see that the proposed time-based budget allocation approach has gained an average of 18% on CPC for both non-aggressive campaigns. Nevertheless, much worse performance was observed on the aggressive campaign, in particular, the test group has spent $44.71 more than the control group, yet with 16 fewer conversions. Such outcome likely indicates that our approach works better for non-aggressive campaigns. Another possibility could be that aggressive campaigns tend to have more dynamic bidding/spending behaviors, thus requiring models trained with more recent data. Another observation we have for the non-aggressive campaigns is that while we do have gains on CPC, the total conversions of test groups are much fewer than those of control groups.

Ad Group	Category	Performance					
		Impressions bid on	Impressions won	Total Cost	Total conversions	Cost per conversion	Gain/Loss
Non-aggressive Campaign 1	Test	1339292	178101	456.60	32	14.27	**20%**
	Control	4359578	343045	914.70	51	17.94	
Non-aggressive Campaign 2	Test	1087797	207249	883.77	52	17.00	**17%**
	Control	1971550	309188	1344.76	66	20.38	
Aggressive Campaign 1	Test	717899	366497	415.94	27	15.41	**-78%**
	Control	817549	311616	371.23	43	8.63	

Table 1: RTB performance outcome of the pilot

REFERENCES

[1] Lee, Kuang-Chih, Ali Jalali, and Ali Dasdan. "Real time bid optimization with smooth budget delivery in online advertising." Proceedings of the Seventh International Workshop on Data Mining for Online Advertising. ACM, 2013.