Applications of Feedback Control in Online Advertising

Proceedings of the 2013 American Control Conference, Washington, DC, June 17 - 19, 2013

Niklas Karlsson[1] and Jianlong Zhang[2]

Abstract— **Internet advertising is a new and exciting area where feedback control has proven tremendously valuable and is expected to play an increasingly important role as the industry continues its rapid growth. The objective of algorithms in a display advertising network is to create value for advertisers, publishers, and the network owner, via optimal decisions on where, when, how, and to whom to show ads. Feedback controllers provide a system that learns from its mistakes and takes proper reactive and proactive actions to meet goals set by advertisers. This paper is a brief tutorial to the field with control engineers as the intended audience.**

I. INTRODUCTION

Online advertising is a form of promotion that uses the Internet and World Wide Web to deliver marketing messages to attract customers. Examples of online advertising include banner ads, contextual ads on search engine results pages, rich media ads, and so forth.

Online advertising may not sound like an industry with problem statements relevant to feedback control, but there are plenty of interesting and challenging feedback control problems, and the control systems used in online advertising are also often mission critical components that drive multi-billion dollar businesses. Moreover the industry is expected to continue growing rapidly for years.

Advertising in traditional media such as printed magazines is different in many ways from advertising online, but the fundamental problems facing the advertising professional are the same. According to John Wanamaker (1838-1922), the owner of the first department store in the United States: "*Half the money I spend on advertising is wasted; the trouble is I don't know which half*." Ultimately, all advertising is about showing the right ad to the right consumer at the right time. The difference between traditional and online advertising is that the data available to make fact-based decisions and the immediate publishing of information and content on Internet is not limited by geography or time. Carefully establishing objective functions, defining measurement and control signals, and designing a system with sufficiently high sampling rate and short delays results in tractable (yet challenging) feedback control problems.

Here we introduce control engineers to problem statements, challenges, and a primitive feedback control solution in the area of online advertising. Limited space allows us only to scratch the surface of this interesting new field.

[1]N. Karlsson is Vice President of R&D at Aol Networks, 395 Page Mill Road, Palo Alto, CA 94306, USA `niklas.karlsson at teamaol.com`

[2]J. Zhang is Technical Director at Aol Networks, 395 Page Mill Road, Palo Alto, CA 94306, USA `jianlong.zhang at teamaol.com`

In the context of a so-called display advertising network there are Internet users, publishers and advertisers plus an ad network optimization engine.

Publishers own web pages receiving traffic from Internet users. The traffic is monetized by selling advertisement space to advertisers or networks. The advertisement space is typically sold in the form of impressions, where an *impression* is one view of an ad. In addition to monetizing impression inventory, publishers want to protect their brand image and do not want to annoy site visitors. Hence, they are sensitive about what type of ads are shown on their web pages.

Although the objectives of advertisers vary, the common goal is to find the right internet users for developing brands, increasing sales, reducing ad expenditure, and spending ad budgets smoothly.

The network owner wants to maximize the network profit while meeting objectives for both advertisers and publishers. This is accomplished by purchasing impressions from the publisher at relatively low costs, and selling them to advertisers in a network of many publishers and Internet users. The optimization engines provides efficient audience targeting and ad allocation, and increases ROIs (return of investments) for both publishers and advertisers.

A contract between a network and a publisher, or between a network and an advertiser involves payment models, and the most common types of payment model are CPM, CPC, and CPA. CPM stands for cost per thousand impressions, CPC stands for cost per click, and CPA stands for cost per acquisition. CPC and CPA are also called performance-based pricing. CPA is different from CPC in that the actions are more closely associated with purchasing behaviors. If the contract between the network and a specific advertiser is a \$20 CPA, the advertiser pays nothing if an Internet user only views or clicks on the ad, but pays \$20 for each conversion that takes place.

Often publishers or advertisers define allowability constraints. One such constraint is when a publisher does not want a competitor's ad to be shown on their web pages; or when an advertiser only wants their ads to be shown to users of a specific gender or age, or from particular geographic locations. Another common advertiser constraint is called frequency cap, where an ad is allowed to be shown to the same user at most a certain number of times per day.

The paper is organized as follows. Section II provides a brief history of optimization in display advertising. Some common feedback control problems encountered in display advertising are introduced in Section III. Section IV discusses some of the important challenges that must be considered in order to solve the control problems. In Section V we provides a simplified problem statement, and a primitive solution to one of the control problems is derived in Section VI. Finally Section VII provides some experimental results.

II. DISPLAY ADVERTISING 1998 – 2012

The first banner display ad on Internet is believed to have appeared in or around 1994, but it was not until several years later that solutions aimed at optimizing online advertising in an automated fashion came about. One of the successful early players in this space was TeknoSurf, which was founded in 1998. It was later renamed to Advertising.com, and eventually acquired by AOL.

TeknoSurf launched the so-called AdLearn system, which is an optimization engine configured to maximize some business objective function subject to various advertiser and publisher constraints and given historical information on impressions and user engagement.

A. The Beginning of Optimization

The first few versions of AdLearn were developed by solving a Linear Program problem [1], [2]. Indeed, the payment models, campaign objectives, and available measurement data back then allowed for this approach by utilizing the certainty equivalence principle (see Figure 1). The approach

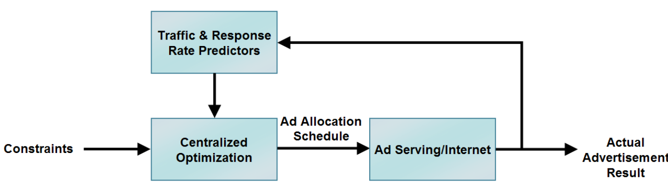

Fig. 1. Block diagram of the old centralized system for display advertising optimization.

was to linearize the optimization problem by implementing a stand-alone prediction system for all relevant stochastic quantities. The dimensionality of the optimization problem was thereafter reduced via various approximations. Next the resulting system was solved using a centralized Linear Program solver. The business was successful and grew rapidly; however, it was all possible because of a modest dimension of the problem (effectively, the number of publishers times the number of advertisers times the number of audience groups).

Over time, however, the initially adopted approach began to break down. First, as the network grew the dimension of the problem became a major challenge computationally. Moreover, the industry evolved towards more complex advertising constraints that did not fit well in the linear programming framework. Effectively, an increasing number of non-linear stochastic and difficult-to-estimate quantities entered into the problem statement, making it increasingly hard to solve. Before giving up on the Linear Program approach, a range of approximations and ad hoc mechanisms had to be implemented to handle the new challenges

The centralized optimization paradigm did not scale well enough and the various ad hoc approximations compromised the optimization results. Furthermore, imperfect predictions led to frequent violations of campaign delivery constraints, and the absence of a feedback mechanism meant that the optimization system did not learn from its mistakes.

B. The Second Generation of Optimization

A paradigm shift was necessary. The most obvious requirements to a new approach was that it must scale well and have the ability of learning from its mistakes and handle dynamic environments.

The approach chosen for a later version of AdLearn was a decentralized feedback control system with the control signals implemented as bids in an impression-based auction exchange [3], [4], [5]. The new system is illustrated in Figure 2. Each time an Internet user navigates to a web page

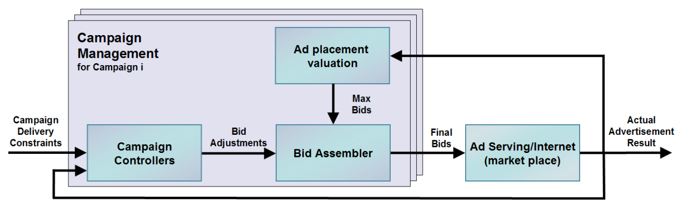

Fig. 2. Block diagram of the new decentralized system for display advertising optimization.

that belongs to a publisher in the network, an impression request is sent to the impression exchange of the network. Within a few milliseconds a market clearing takes place determining which ad to serve to the user. The bids used for the market clearing are calculated every few minutes by a separate campaign management system that contains two core algorithm modules and one bid assembler. The ad placement valuation recursively processes time series of impressions, clicks, and conversions for each ad to estimate the performance of the ad on various web pages and to different users or user segments. The campaign controller receives various campaign delivery constraints (see Section III) and relevant feedback data to assess at what degree the delivery constraints are met. The controller thereafter issues appropriate bid adjustments to ensure the campaign is delivered in the desirable manner

III. PROBLEM FORMULATION

Advertising network optimization involves e.g. impression valuation, market clearing, and data transmission, but we shall in this paper only consider what is referred to as campaign management handled by “Campaign Controllers” in Figure 2. This system adjusts the so-called Max Bids based on reference and feedback data. The Max Bid is calculated by the Ad Placement Valuation unit and represents the effective CPM value of each impression. Reference data captures any campaign delivery constraints and feedback data contains any information needed to assess how well the delivery constraints are met. The control signal may be a multiplicative or additive bid adjustment to the Max Bids.

An ad campaign is a contract between an advertiser and the network for the network owner to deliver an ad budget according to an agreement. For example, the advertiser may agree to pay the network $5 per conversion (for each impression turning into a sale) over a four weeks time period to a maximum amount of $100,000. In general the advertiser also wants the pacing of the budget delivery to be even throughout the flight of the campaign so that most of the budget is not delivered at the beginning or end of the flight. The contract may also stipulate additional constraints: The advertiser may be interested only in showing ads to users in a specific geographic location like a city or state, may not want ads to be shown on weekends, or may not allow the same user to be shown the same ad more than a prescribed number of times per day.

Business requirements are often fuzzy and difficult to translate into mathematical equations, but in this section we outline the most common concrete campaign delivery control problems. Many times a campaign is configured with a combination of two or more delivery objectives.

A. Branding vs Performance

Broadly speaking an ad campaign is categorized as a branding and/or performance campaign. The objective of a branding campaign is to reach many users without necessarily selling a product in the near future. A branding campaign is typically set up via a CPM payment model since conversions are less important than simply showing ads.

The objective of a performance campaign is to show ads to users that are most likely to click or convert. A pure performance campaign is set up via a CPC or CPA payment model so that the advertiser does not pay for impressions that did not result in some form of engagement.

B. Smoothness Control

The smoothness objective is typically considered as the primary objective among all the possible objectives. Simply, if a campaign is assigned with the smoothness objective, then its daily delivery should match the pre-defined daily reference and bursts and dark hours should be avoided. Smoothness is applicable to both branding and performance campaigns, but it is more challenging for performance campaigns because of system properties that are discussed in Section IV.

C. Backend Performance Control

The backend performance goal is another important control objective, and is typically considered secondary in the sense that it may be violated if satisfying the primary objective requires that. A common set-up is for the primary objective to be CPM smoothness (the network is paid per impression), with a secondary performance model defined as maximum cost equivalence per click or conversion should be less than a pre-defined threshold. This ensures at least a certain ratio of the impressions should turn into clicks or conversions.

D. Partition Control

Partition control has gained interest among some advertisers in recent years. A campaign may consist of multiple ad creatives, where each creative is tailored to a certain user group. For example, different ads may be designed for male and female customers, or for different geographic regions. For a campaign with a partition objective its medias are divided into partition groups. The deliveries from these partitions should be maintained close to some pre-defined ratios, such as 50/50 between male and female, or 10% to people in California and 90% to people elsewhere.

IV. MODELING

Let us consider some properties of an ad network relevant for designing a feedback control system. No two ad campaigns are identical, so we opt to present the modeling qualitatively rather than quantitatively.

First, the relevant aspects of the network are those that influence the input-output relationship of the plant as perceived by an individual ad campaign. This relationship is defined by the mapping from 'bid adjustment' to 'actual advertisement result' in Figure 2.

The most obvious aspect of the plant is the time variability. The presence of people online is highly dependent on time. We expect a significant seasonal time-of-day pattern, but also a day-of-week pattern, and so forth [6]. These seasonal patterns together with traffic trends and spikes are reflected in a campaign's time series of impression, click, and conversion data. Figure 3 provides an example of impression volume

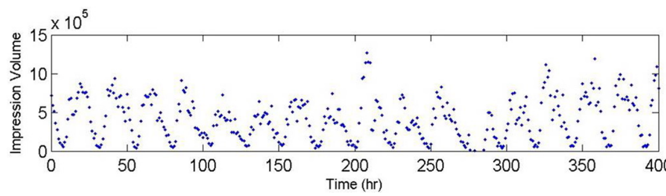

Fig. 3. Example of impression volume over time for a real campaign. Notice the distinct time-of-day pattern, but also the variability of volume variance.

over time for a real campaign. Notice the distinct time-of-day pattern, but also the variability of volume variance. This shows that the impression volume cannot be described by a linear time-invariant system driven by white noise. Instead. a better representation is likely a periodic log-Normal model [7], [8].

For a performance campaign it is of significance that only one click or conversion per thousand or million of impressions occurs. This means that even for moderately large impression volumes, it is dangerous to make use of popular approximations based on the Central Limit Theorem in statistics [9], [10]. Indeed, the conversion volume would not demonstrate a symmetric and Gaussian-like pattern even if the impression volume is steady. A better representation is to assume conversions are the result of a Binomial experiment with the impression volume and the unknown conversion probability as input arguments. However, even this is an approximation since it assumes each impression is an independent Bernoulli experiment, which is not the case.

Next, *latency*, which is the time between when an impression is served to a user and the user converts (e.g. makes a purchase), varies dramatically from one ad campaign to another and makes campaign control a challenge. In particular, since both conversion probability and latency distribution are unknown properties of a campaign and must be estimated based on historical data, it is tricky to determine how to pace a performance campaign initially. If a constant number of impressions are shown every hour and we observe the resulting time series of conversions, we face the problem initially of not knowing whether the observed level of conversion volume is representative for the impression volume or if more conversions are to be expected later from the already served impressions. Figure 4 shows simulated data to explain the described dilemma. It is intuitive that

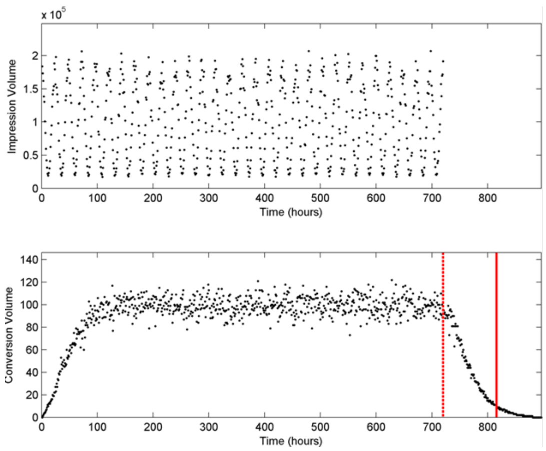

Fig. 4. Simulated data for observed impression and conversion data to illustrate the impact of latency.

some campaigns may have a long latency. For example, if a car maker is running a performance campaign we do not expect users to convert immediately. Clearly, most people considering to buy cars online would spend extensive time assessing their financial situations and purchase options.

Figure 5 gives two examples of real latency data. The top plot is a normalized histogram of latencies from a campaign where most conversions took place within one or a few hours after the corresponding impression. The bottom plot shows a normalized histogram of latencies for a long latency campaign. Super-imposed on the histograms are maximum-likelihood fitted Gamma probability density functions. The Gamma distribution is a versatile two-parametric function that is proven to describe the dominant shape of latency data well. But as can be seen for the long latency campaign, there is a 24 hour-periodicity that is not captured by the Gamma model. Capturing this aspect of the latency behavior would require additional parameters in the model.

By design, each campaign should be controlled by an *independent* control system. The purpose is to ensure a scalable system, but another reason is that an advertiser for various reasons may not want collaboration with others. Figure 6 shows one example of the so-called price-volume relationship. The curve shows how many impressions the campaign can win as the price increases, and is produced by determining the clearing price for all impressions the campaign is interested in. Clearly, the volume is a non-decreasing function of the price, but more interestingly, its slopes may differ dramatically at different price points. Small changes in bid price may result in significant increases in awarded impressions. The slope of the price-volume curve is referred to as the plant gain and influences the speed and robustness of the closed-loop system.

Another thing worthy of noting is that independent control does not mean the campaigns are de-coupled. Since all campaigns submit bids into just one impression market place competing for the same impressions, there is by necessity a strong coupling among campaigns. The competitive environment (price-volume curve) for a campaign changes when competing campaigns entering or leaving the network, or simply changing their bids. This means the campaign controllers must be robust and adaptive.

V. EXAMPLE PROBLEM

Due to space limit, we only discuss smoothness control of CPA campaigns with highly simplified models.

A. Key Notations

Suppose the system sampling time in one hour.

k	:	Current time
ℓ	:	Dummy variable for time
$u(\ell)$	:	Control signal effective during $[\ell, \ell+1)$
$n_I(\ell)$	:	Impression volume sourced to $[\ell, \ell+1)$
$n_A(\ell)$	:	Conversion volume sourced to $[\ell, \ell+1)$
$n_A(\ell, k)$	:	Conversion volume sourced to $[\ell, \ell+1)$ and reported before time k
$p(\ell)$	:	Conversion rate during $[\ell, \ell+1)$
$n_A^{daily,ref}$	:	Daily delivery goal for conversions

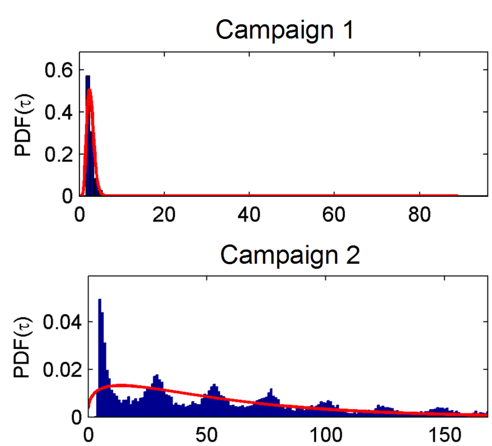

Fig. 5. Examples of real latency data together with maximum-likelihood fitted Gamma probability density functions. Top: short-latency campaign. Bottom: long-latency campaign.

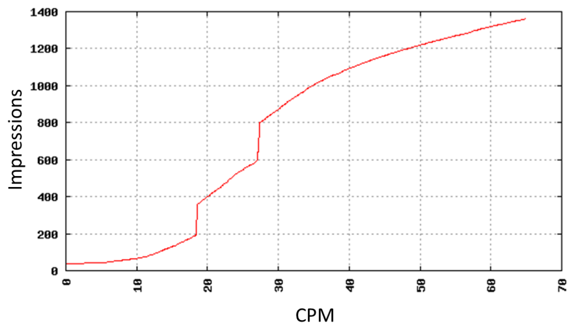

Fig. 6. Example of price-volume curve. It indicates how many impressions a campaign would have been awarded for different bid prices at a given time.

B. Simplified Plant Model

The principle relationship between $n_I(k)$ and $u(k)$ can be described by the following model:

$$n_I(k) = u(k)\left[\beta_0 + \beta_1 \sin\left(2\pi k/24 + \varphi_1\right)\right] e^{\epsilon(k)} \quad (1)$$

where $\epsilon(k)$ is a stationary, stable and zero-mean stochastic process, and β_0, β_1 and φ_1 are unknown model parameters which differ from campaign to campaign.

The relationship between $n_A(k)$ and $n_I(k)$ can be captured with a Binomial model:

$$n_A(k) \sim \text{Binomial}\left(n_I(k), p(k)\right) \quad (2)$$

Typically, $p(k)$ is in the range of $[10^{-4}, 10^{-6}]$, and some campaigns have strong time-of-day patterns in $p(k)$.

In this paper, we assume there is no delay for reporting impressions, and model the latency for reporting conversions, δ, by a Gamma distribution with parameters α and β:

$$\delta \sim \text{Gamma}\left(\alpha, \beta\right) \quad (3)$$

C. Control Objective

The objective is to design a feedback control signal $u(k)$ for a CPA campaign such that its daily delivery tracks the daily reference $n_A^{daily,ref}$; i.e.,

$$\lim_{k\to\infty} \sum_{\ell=k-23}^{\ell=k} n_A(\ell) = n_A^{daily,ref} \quad (4)$$

VI. EXAMPLE SOLUTION

The solution for the problem presented in Section V involves the estimation algorithms for latency, conversion rate and conversion volumes and the control algorithm for calculating feedback control signal.

A. Latency Estimation

Given a time series of latency observations $\delta_1, \delta_2, \ldots, \delta_t$, where δ_i denotes the time separation from an impression to the conversion for conversion i and t denotes the most recent conversion, we wish to estimate the latency probability distribution model (see (3) and Figure 5).

The likelihood function for t independent and identically-distributed Gamma random variables is

$$L(\alpha, \beta) = \prod_{i=1}^{t} f(\delta_i|\alpha, \beta)$$

where

$$f(\delta_i|\alpha, \beta) = \frac{1}{\beta^\alpha}\frac{1}{\Gamma(\alpha)}\delta_i^{\alpha-1}e^{-\delta_i/\beta},$$

and $\delta_i, \alpha, \beta > 0$. The maximum likelihood estimates for α and β can be found by setting the derivatives of the above likelihood function to zero. We obtain

$$\ln(\alpha) - \psi(\alpha) = \ln\left(\frac{1}{t}\sum_{i=1}^{t}\delta_i\right) - \frac{1}{t}\sum_{i=1}^{t}\ln(\delta_i) \quad (5)$$

$$\beta = \frac{1}{\alpha t}\sum_{i=1}^{t}\delta_i \quad (6)$$

where $\psi(\alpha) = \Gamma'(\alpha)/\Gamma(\alpha)$ is the digamma function. No closed-form solution α of (5) exists, but the function is well-behaved and can be easily calculated numerically.

As seen in (5) and (6) a minimum sufficient statistic for the Gamma parameters α and β is given by $\sum_{i=1}^{t}\delta_i$ and $\sum_{i=1}^{t}\log\delta_i$ [9]; i.e., we must not store all individual latency observations throughout the life of a campaign. It is sufficient to keep track of the most up-to-date values of these two sums. These sums represent the state variables of the latency estimator, and the solution of (5) and (6) is calculated in each iteration of the control system based on all available latency observations, producing $\hat{\alpha}$ and $\hat{\beta}$.

B. Conversion Rate Estimation

We assume that $p(k)$ in (2) is constant and that there is no latency for reporting conversions, then the following estimator provides the minimum variance estimate for $p(k)$.

$$\hat{p}(k) = \frac{\sum_{\ell=0}^{\ell=k-1} n_A(\ell)}{\sum_{\ell=0}^{\ell=k-1} n_I(\ell)} \quad (7)$$

In reality, at time k, we only have $n_A(\ell, k)$, which might be an immature measurement of $n_A(\ell)$. If we simply replace $n_A(\ell)$ with $n_A(\ell, k)$ in (7), $p(k)$ is under-estimated. To better leverage immature data, the conversion rate estimation should take into account of the latency estimates derived from (5) and (6). One candidate algorithm is

$$\hat{p}(k) = \frac{\sum_{\ell=0}^{\ell=k-1} n_A(\ell, k)}{\sum_{\ell=0}^{\ell=k-1} n_I(\ell)\int_{\ell}^{k} f(\tau|\hat{\alpha}, \hat{\beta})d\tau} \quad (8)$$

in which impression measurements are discounted based on the estimated latency model.

To deal with the fact that $p(k)$ may slowly change with time, the above algorithm can be further updated with forgetting factors or other robust modifications. Due to space limit, we do not discuss the algorithm for estimating time-of-day pattern in $p(k)$.

C. Conversion Volume Estimation

To smoothly pace a CPA campaign, we estimate the conversion volumes for the past runs. Below is one example solution for estimating $n_A(\ell)$ at time k:

$$\hat{n}_A(\ell) \quad = \quad n_A(\ell,k) + \hat{p}(k) n_I(\ell) \int_{k-\ell}^{\infty} f(\tau|\hat{\alpha},\hat{\beta}) d\tau \quad (9)$$

We have assumed the true conversion rate is either constant or slowly varying over time. Therefore, $\hat{p}(k)$ always represents the best knowledge for the conversion rate for any time before k.

D. Smoothness Controller

If we have perfect knowledge about the conversion rate, then the estimated delivery in conversions in (9) can be used as the feedback signal. The following feedback control law meets the control objective in the noise free scenario:

$$u(k) \quad = \quad u(k-1) + \gamma^I \left[n_A^{daily,ref} - \sum_{\ell=k-24}^{\ell=k-1} \left(n_A(\ell,k) + \hat{p}(k) n_I(\ell) \int_{k-\ell}^{\infty} f(\tau|\hat{\alpha},\hat{\beta}) d\tau \right) \right] \quad (10)$$

where γ^I is a design parameter. Control law (10) is a simple I-type controller with 24-hour-moving-average filter. It is a stabilizing controller as long as $\hat{p}(k)$ is stable and γ^I is small enough. For improved performance, one may look into linear-time-varying or non-linear control designs [11], [12]. The complete algorithm for updating feedback signal is described by (5), (6), (8) and (10).

VII. EXPERIMENTAL RESULTS

We conclude this tutorial on applications of estimation and control in online advertising with experiment on one representative CPA campaign. Figure 7 shows the pacing of this campaign for 3 weeks, managed by a controller similar to but more complicated than the algorithm presented in Section VI. In the top two plots we see that the estimated latency mean is between 14 and 18 hours and that the estimated latency standard deviation is between 9 and 11 hours. The variability may be the results of noisy measurement and/or a change in the underlying true but unknown latency distribution.

The estimated conversion probability is displayed in the middle left plot. Notice how the conversion probability varies between 0.1 and 0.2 % during most of the time interval. The drop in observed sourced conversion probability on September 23/24 is the result of the latency which means that most conversions sourced to impressions during those days have not yet occurred. The estimated conversion probability is doing a better job by accounting for future conversions.

The successive increase in control signal throughout the first two weeks of September, shown in the middle right plot, is an indication of an increasingly competitive environment. Most likely there were other campaigns entering the market place bidding for the same impressions. The bottom left plot shows the hourly and daily number of impressions

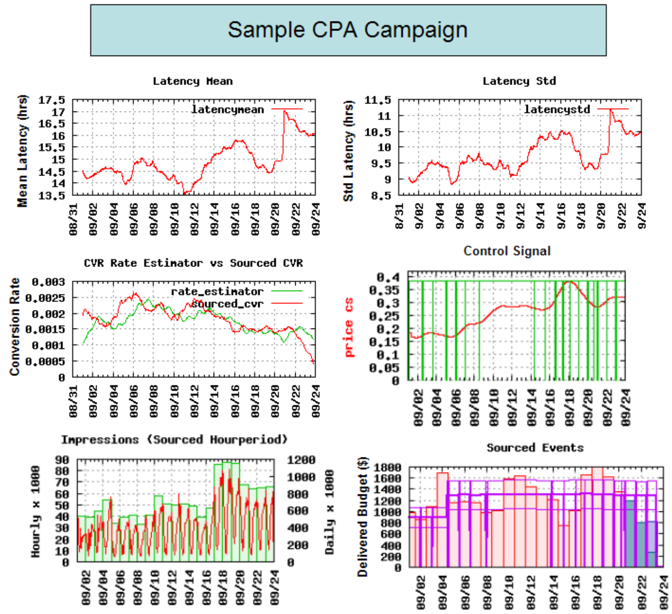

Fig. 7. Example of CPA campaign managed by the controller developed in Section VI. Top: estimated latency mean and standard deviation. Middle left: observed sourced and estimated conversion probability. Middle right: Control signal. Bottom left: hourly and daily awarded impression volume. Bottom right: daily budget delivery and reference signal.

awarded to the campaign and the bottom right plot shows the daily delivery of the advertising budget. The distinct day of week pattern follows the usual user activity online and advertisers in general do prefer that their campaigns follow the same pattern (similar to how the campaign also show a distinct time of day pattern). This last plot also shows a (mostly) horizontal thick reference line and two thinner lines indicating a $\pm 20\%$ delivery acceptance band. The blue bars on September 21-24 indicate days with an immature conversion report, due to long latency.

REFERENCES

[1] Dimitri P. Bertsekas. *Nonlinear Programming*. Athena Scientific, second edition, 1999.

[2] David G. Luenberger. *Linear and Nonlinear Programming*. Springer, second edition, 1972.

[3] John Gittins, Kevin Glazebrook, and Richard Weber. *Multi-armed Bandit Allocation Indices*. Wiley, 2 edition, 2011.

[4] Vijay Krishna. *Auction Theory*. Academic Press, 2002.

[5] Niklas Karlsson, Jianlong Zhang, and Yonggang Xu. *Methods for controlling an advertising campaign, United States Patent 7,835,937*. USPTO, 2010.

[6] Niklas Karlsson, Jianlong Zhang, and Jaime Batiz-Mendoza. *System identification, estimation, and prediction of advertising-related data, United States Patent 7,827,128*. USPTO, 2010.

[7] P. McCullaph and J.A. Nelder. *Generalized Linear Models*. Chapman & Hall, second edition, 1989.

[8] G.A. Seber and C.J. Wild. *Nonlinear Regression*. Wiley, first edition, 2003.

[9] George Casella and Roger L. Berger. *Statistical Inference*. Duxbury, 1990.

[10] Patrick Billingsley. *Probability and Measure, 3rd Edition*. Wiley-Interscience, 1995.

[11] Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, third edition, 2001.

[12] Jean-Jacques Slotine and Weiping Li. *Applied Nonlinear Control*. Prentice Hall, 1991.