Control Engineering Practice 75 (2018) 126–136

Contents lists available at ScienceDirect

Control Engineering Practice

journal homepage: www.elsevier.com/locate/conengprac

Feedback control of event rate in online advertising campaigns

Qian Sang [1], Niklas Karlsson *,[2], Jiaxing Guo [1]

Demand Platforms R&D, Oath Inc., 701 1st Ave., Bldg.#A, Sunnyvale, CA 94089, United States

ARTICLE INFO

Keywords:
Online advertising
Event rate
Beta actuation
PI control
Viewability
Video completion

ABSTRACT

An actuation and feedback control algorithm is proposed for the specific objective of scalable event rate control in online advertising systems. The actuator employs the beta actuation mechanism to address discontinuity in the plant under control with adjustable plant gain, while the feedback controller implements a PI mechanism to regulate the event rate to stay at or above a user-specified reference. Effectiveness of the proposed scheme is demonstrated via simulations and validated with the in-view rate control and the video completion rate control of real advertising campaigns on the AdLearn™ advertising optimization system, developed at AOL.

1. Introduction

Advertising, which is a US$600 billion industry (eMarketer, 2014, 2017), has in recent years come to rely heavily on feedback control for online applications. In fact, feedback control has become critically important for scalable optimization in such systems. Each advertiser wishes to spend an advertising budget in such a way that their specific branding and/or performance objective is optimized. Cooperation is not permitted and the advertisers compete over ad *impressions* (opportunities to show their advertisements to Internet users). In short, each advertiser wishes to deliver ads to those Internet users who can generate the highest ROI (return on investment) for the advertiser's advertising budget.

The allocation of ad impressions is handled in impression exchanges (Google, 2011). Any advertiser may submit bids for any opportunity to show an ad, but only the highest bidder is awarded the impression. The winner usually pays a second price as the cost for the impression awarded (Edelman, Ostrovsky, & Schwarz, 2007). The optimization problem turns into a problem of devising a bidding strategy that maximizes the overall returned value given a limited advertising budget. Given the extremely large number of Internet users browsing Internet every day and the large number of advertisers, it is an extraordinarily high-dimensional optimization problem. In addition to the scale, time-varying and stochastic traffic patterns and user behaviors add complexity to the optimization problem.

Feedback control has played a critical role in solving the above type of optimization problems for many years, see, e.g., Guo and Karlsson (2017), Karlsson and Zhang (2013) and Zhang, Rong, Wang, Zhu, and Wang (2016), for a high-level introduction to the control problem, and Karlsson (2014) and Wang, Zhang, and Yuan (2017), for an attempt at dealing with the unique challenges in this domain. The first deep dive into how the optimization problem is turned into a control problem and what some of the challenges are in order to solve the control problem was published in Karlsson (2016b).

However, the problem considered in Karlsson (2016b) is to maximize a value function given an advertising budget. A different problem is to control an average event rate (Karlsson & Sang, 2017), e.g., a campaign-level in-view rate or click-through rate, which is the focus of this paper. The event rate control problem is of particular interest in the online advertising context, due to the fact that the average event rate is often one of the KPI's (key performance indicator) to measure the success of an advertising campaign. A certain event rate usually is also specified for advertising agencies to meet in order to fulfill contracts with their clients, the advertisers. For example, advertisers may seek an average of 70% in-view rate for their display campaigns, per the recommendations by the Media Rating Council (MRC), see IAB (2014); otherwise, agencies lose money by serving additional impressions at no cost to advertisers until the in-view rate specification is met (make-good). For advertisers, event rate control is also beneficial as an additional lever to balance campaign performance over advertising cost: a higher event rate usually

* Corresponding author.
E-mail addresses: qian.sang@oath.com (Q. Sang), niklas.karlsson@oath.com (N. Karlsson), jiaxing.guo@oath.com (J. Guo).

[1] Senior research scientists with the Demand Platforms R&D group of Oath Inc., the parent company of AOL and Yahoo! and a fully owned subsidiary of Verizon Communications.

[2] Chief scientist and vice president of research with the Demand Platforms R&D group of Oath Inc.

https://doi.org/10.1016/j.conengprac.2018.03.010
Received 29 August 2017; Received in revised form 9 March 2018; Accepted 13 March 2018
Available online 4 May 2018

leads to better campaign performance; however, a higher event rate comes with an increased media cost.

A seemingly straightforward solution to the event rate control problem is threshold targeting; that is, bid on an incoming impression opportunity only if its event rate prediction is higher than the specified reference value. There are a few problems associated with such a solution. First, the resulting average event rate is often higher than specified, which may not be economically sound for all parties involved. For example, for agencies a higher event rate incurs additional cost, leading to reduced profit margin. Secondly, frequent manual adjustments to the threshold may be necessary to hit the rate reference, which in practice is error-prone and may actually lead to degraded performance. An automated algorithm that solves the event rate control problem is thus desirable.

Note, on the other hand, that some major players in the industry take a different approach. For example, Google and Facebook offer advertising products that guarantee a 100% viewability for their clients running branding campaigns, see FacebookBusiness (2015) and Moha (2015). However, it does not mean a 100% in-view rate is achieved; rather, clients are only charged by impressions that are considered viewable by Internet users. The cost associated with the non-viewable impressions is thus off-set by the higher viewability price. It is likely there is certain internal event rate control mechanism which can make this approach profitable, but no related information is generally available, possibly due to proprietary considerations.

To the best of the authors' knowledge, there has been no previous attempt at solving the event rate control problem in the context of online advertising and restricted to decentralized (scalable) feedback control. The proposed actuation and control scheme for event rate control provides more transparency in terms of cost to advertisers, and a lever at the advertiser's disposal to better cater their needs.

It is also worth mentioning that the applicability of the proposed event rate control scheme is not restricted to the type of the event of interest. The metrics for campaign effectiveness have been evolving as new measurement technologies mature. For instance, in-view rate and completion rate are two major KPI's of significant importance to advertisers running video campaigns. Recently, AVOC (audible and viewable on completion) emerges as a new metric that has attracted a lot of attention, see (AdExchanger, 2015). The proposed event rate control scheme can be applied to regulate such an AVOC rate with minimal control tuning and configuration adjustments.

The paper is organized as follows. Section 2 overviews the basics of programmatic advertising, to provide a background for the optimization and control problem considered in this paper. The control problem is defined in Section 3. By default the plant is discontinuous, but an actuation mechanism is proposed in Section 4 to effectively turn the input–output relationship of the plant continuous. Section 5 describes how to model and tune the plant. The information is used to establish a nominal plant model that is used in Section 6 to design a feedback controller. In Section 7 the control system is evaluated both in a simulated but realistic environment and on real advertising campaigns to assess the performance and the stability of the closed-loop control system. Finally, in Section 8 the paper is wrapped up with some concluding remarks.

2. Basics on programmatic advertising

Programmatic advertising is a game changing technology in the online advertising industry (Busch, 2016). It automates the ad request, purchase, and delivery process for highly-efficient online marketing. Programmatic advertising leverages latest progresses in the fields of artificial intelligence, machine learning, big data, etc., to deliver the right ads to the right audience at the right time. It benefits both Internet users and online marketers: Internet users receive useful information on products and services catering to their needs without a compromise in online experience, while marketers receive higher returned value on their advertising spending.

Fig. 1 provides a simplified, high-level overview of the several parties involved in the programmatic ad delivery process. The delivery of an online advertising campaign involves impressions, where an impression is one view of a certain advertisement. The process starts with an Internet user (*Audience*) trying to load a web page containing some advertising space by entering its URL to the web browser of a desktop/laptop computer or a mobile device such as a mobile phone or a tablet. Alternatively, it can be a user opening an app with advertising space inside on his/her mobile device. Immediately, the publisher of the web page (or mobile app), referred to as *Media* in Fig. 1, sends an impression request for an ad to an *Ad Exchange* (or via *SSP*'s, the supply-side platforms), along with relevant information about the audience. A sealed second-price standard auction (Krishna, 2002) is then held at the ad exchange, where the impression request is broadcast to all *DSP*'s (demand-side platforms) that interface with the exchange. On behalf of an *Advertiser* (or online marketer) interested in showing an ad to this Internet user, a DSP evaluates the impression request and submits a bid price b_i and a bid allocation $a_i \in [0,1]$ to the auction, where i is referred to as a user segment but may represent a specific user or a user partition. If the bid price b_i is higher than any competing price bids, then the impression is awarded to this advertiser with a probability equal to a_i. The winner is notified to send the ad over, to be placed in the adverting space on the media. This entire process is called *Real-Time Bidding* (RTB), and it typically takes about 50 milli-second.

Note that some other parties are left out in the discussion, e.g., data management platforms (DMP), advertising agencies, supply-side platforms (SSP), which are also important in the RTB process, but are out of the scope of discussion for this paper.

After the requested web page or mobile app fully loads with ads, an impression may turn into a value-bearing event with a probability of p_i. For instance, the user may click on the ad, in which case the impression turns into a click, or make a purchase on the advertiser's web site (directed from the media), in which case the impression turns into a conversion. The event rate p_i for the ith user segment is then the click-through rate and the conversion rate, respectively. For branding campaigns with which brand exposure to audience is of high priority, advertisers' focus is in-view rate or completion rate for video ads. An impression is considered viewable if 50% of the ad pixels are in view within the currently active browser for more than one second (more than two seconds for video ads), according to IAB (2014). A video ad is considered a complete view (completion), if the ad itself is played from beginning to end in its entirety within the browser.

The proposed actuation and control algorithm lies within a DSP in Fig. 1. Specifically, it is a solution to the advertisers' requirement that the campaign-level event rate p is no less than a prescribed value of $p^{ref} \in (0,1)$.

3. Problem formulation

As described in Section 2, the impression allocation for the ith user segment is governed by a sealed second price auction, where b_i is the bid price submitted to the auction and a_i is the bid allocation, or the sampled fraction of auctions the campaign chooses to participate in. It was shown in Karlsson (2016b) that the total marketing value given an advertising budget is maximized by submitting bid allocation values of $a_i = 1$, and a bid price b_i proportional to the event rate p_i, with a proportionality coefficient selected as the largest value for which neither the budget constraint nor the ROI constraint is violated.

The algorithm proposed in this paper enhances the solution to the above value maximization problem to deal with the additional event rate constraint, i.e., the average campaign-level event rate p is no less than a design specification $p^{ref} \in (0,1)$, provided by advertisers. The event rate p is the ratio of the total number of events n_{event} (defined by advertisers, e.g., viewable impressions, video completions, clicks) versus the total number of impressions n_{tot} for a certain campaign. The event rate p_i is defined as the probability that an impression from segment i turns into an event.

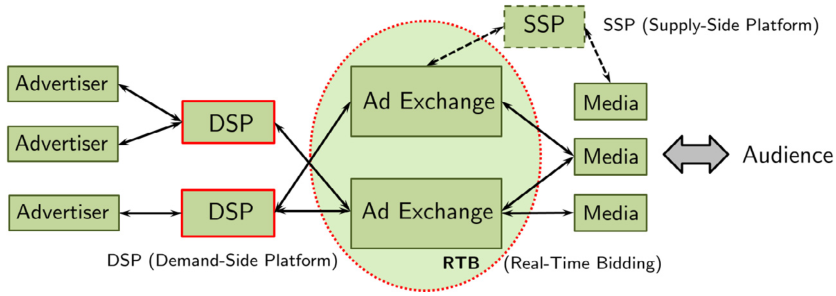

Fig. 1. A high-level overview of parties involved in programmatic advertising.

In this paper, the bid price b_i is assumed to be held constant (does not introduce additional dynamics), and individual bid allocation values a_i are to be adjusted in such a way that the campaign-level event rate p meets the design specification p^{ref}. The related problem of estimating p or p_i is addressed in Karlsson (2016a). The estimate of p (p_i, respectively) is denoted by $\hat{p}$ ($\hat{p}_i$, respectively).

Suppose the campaign is submitting competitive bid prices (i.e., it is the highest bidder) in segments labeled $i = 1, \ldots, m$, and suppose the total number of available impressions in segment i is $n^{rel}_{avail,i}$. Possible dynamic coupling between the computation of the bid price b_i and the bid allocation a_i is neglected. Site-level event rate estimates $\hat{p}_i \approx p_i$ are available and a computationally efficient (scalable) solution is required. The objective is to devise a feedback controller that adjusts a_i, $i = 1, \ldots, m$, such that the average observed event rate of the campaign is at or above a prescribed reference value p^{ref}.

Scalability is obtained by the decoupled solution shown in the block diagram in Fig. 2. *Actuator* is a static (memory-less) component processing the segment-level event rate estimates $\hat{p}_i$ and a campaign-level scalar control signal u. *Event Rate Controller* is a feedback based component consuming a campaign-level reference signal p^{ref} and a scalar feedback signal p representing an estimated campaign-level event rate. While the modularized solution provides scalability, it potentially leads to a discontinuous relationship between u and p. Indeed, if u is handled simply as a threshold value such that $a_i = \mathbb{I}_{\{p_i \geq u\}}$, where $\mathbb{I}_X$ is the indicator function satisfying $\mathbb{I}_X = 1$, if $X =$ true, and $\mathbb{I}_X = 0$, if $X =$ false; then the relationship between u and p is discontinuous.

Discontinuity of the plant brings challenges to the control and optimization problem. In the next section, the so-called Beta Actuation mechanism is proposed to render a smooth input–output relationship between u and p.

4. Beta actuation

The objective of the actuator is to map a campaign-level control signal u to adjustments of individual bid allocation values a_i in a manner that permits regulating the average campaign-level event rate p (see Fig. 2). At our disposal are segment-level event rate estimates $\hat{p}_i$, $i = 1, 2, \ldots, m$. Index i is referred to as segment, but may represent e.g. a site, an audience partition, or an individual user.

To make control possible, it is important that both the relationship from u to p, and the relationship from $\hat{p}_i$ to p are well-behaved. For example, small perturbations of u or $\hat{p}_i$ must result in only small perturbations of p. Furthermore, the relationship between u and p should be monotonic and continuous, and the range of values for u should map to the widest range possible for p, and ideally the range of u should be well-known, e.g., $[0, 1]$. Finally, to support scalability and to make dynamic analysis of the closed-loop system practically doable, it is preferable the actuator is static (memory-less) and computationally inexpensive to use.

The following requirements are imposed on the actuator mapping $a_i = g(\hat{p}_i, u)$, defined for $\hat{p}_i \in [0, 1]$ and $u \in [0, 1]$:

- g is a static (memory-less) mapping
- $0 \leq g(\hat{p}_i, u) \leq 1$ for all $\hat{p}_i \in [0, 1]$, $u \in [0, 1]$
- $g(\hat{p}_i, 0) = 1$ for all $\hat{p}_i \in [0, 1]$
- $g(\hat{p}_i, 1) = 0$ for $\hat{p}_i \in [0, 1)$
- $g(\hat{p}_i, u)$ is continuous in $\hat{p}_i$ and u
- $g(\hat{p}_i, u)$ is decreasing in u for $\hat{p}_i \in (0, 1)$
- $g(\hat{p}_i, u)$ is increasing in $\hat{p}_i$ for $u \in [0, 1)$
- $g(\hat{p}_i, u)$ is a computationally inexpensive mapping

It is assumed that $\hat{p}_i \approx p_i$, where p_i is the true event rate for the ith segment.

4.1. Beta distribution

The proposed actuator design makes use of the properties of the so-called beta distribution from mathematical statistics, see, e.g., Casella and Berger (2001). The beta distribution with parameters $\alpha > 0$ and $\beta > 0$ is a continuous probability distribution. If a random variable X follows the beta distribution, then $X \sim \mathrm{Beta}(\alpha, \beta)$. The probability density function of x is given by

$$f(x|\alpha, \beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)},$$

for $x \in [0, 1]$, where $B(\alpha, \beta)$ is the beta function (also called the Euler integral) defined by

$$B(\alpha, \beta) = \int_0^1 x^{\alpha-1}(1-x)^{\beta-1} dx.$$

Parameters $\alpha > 0$ and $\beta > 0$ are referred to as *shape* parameters. The expected value μ and the variance σ^2 of X are

$$\begin{aligned} \mu &:= \mathrm{E}(X) = \frac{\alpha}{\alpha+\beta}, \\ \sigma^2 &:= \mathrm{Var}(X) = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)}. \end{aligned}$$

The cumulative density function of x is given by

$$F(x|\alpha, \beta) = \frac{1}{B(\alpha, \beta)} \int_0^x t^{\alpha-1}(1-t)^{\beta-1} dt$$

and is more generally (beyond stochastic systems) called the regularized incomplete beta function.

It is easy to show that if $\sigma^2 > 0$, then

$$\begin{aligned} \alpha &= \frac{\mu^2(1-\mu)}{\sigma^2} - \mu \\ \beta &= (1-\mu)\left(\frac{\mu(1-\mu)}{\sigma^2} - 1\right) \end{aligned}$$

Leveraging on properties of the incomplete beta function, an actuator $a_i = g(\hat{p}_i, u)$ of the form $a_i = F\left(\hat{p}_i|\alpha, \beta\right)$ is proposed. If α and β are chosen wisely as functions of u, then the actuator satisfies the actuator requirements.

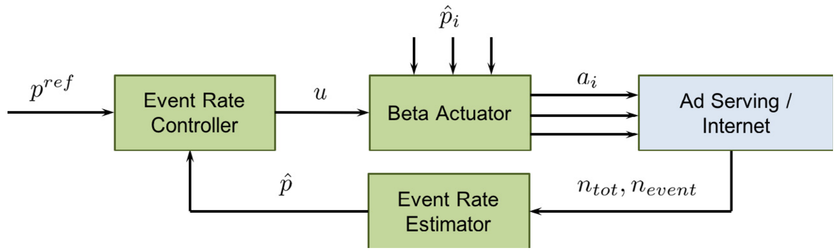

Fig. 2. Block diagram of the event rate control problem.

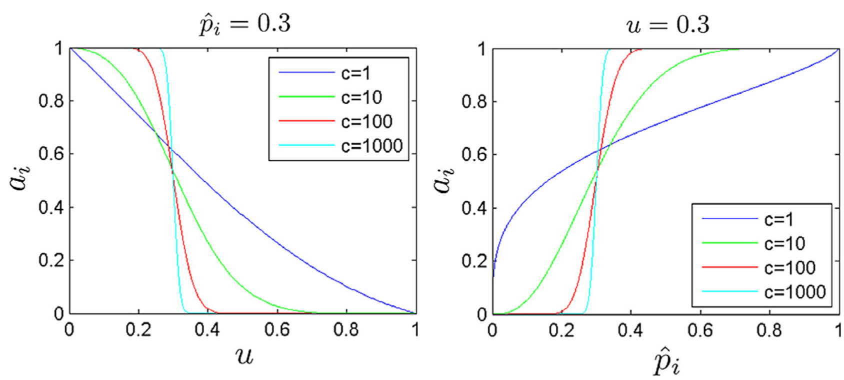

Fig. 3. The plots demonstrate how the bid allocation a_i for different values of c varies as a function of u for a fixed event rate $\hat{p}_i$ (left), and as a function of $\hat{p}_i$ for a fixed event rate u (right).

4.2. Beta actuator

Select $\alpha_c(u)$ and $\beta_c(u)$ parameterized by c such that the corresponding beta distribution with scale parameters $\alpha_c(u)$ and $\beta_c(u)$ has mean μ and variance σ^2 given by

$$\mu = u,$$
$$\sigma^2 = \frac{1}{c+1}u(1-u),$$

where $c > 0$ and $0 \le u \le 1$. Configuration parameter c is used to adjust the sensitivity of the actuator in response to variations in u and $\hat{p}_i$.

Using previously stated results for the beta distribution, it follows that

$$\alpha_c(u) = cu$$
$$\beta_c(u) = c(1-u)$$
$$a_i = F\left(\hat{p}_i|\alpha_c(u), \beta_c(u)\right)$$

if $0 < u < 1$; otherwise, $a_i = u$. The plots in Fig. 3 give an initial idea of how a_i depends on c, u, and $\hat{p}_i$. The left subplot shows that a_i goes from 1 to 0 as u goes from 0 to 1 at a rate that depends on the configuration parameter c, with most of the drop occurring when $u \approx \hat{p}_i$. The right subplot demonstrates the opposite behavior for a_i as a function of $\hat{p}_i$.

To underscore that the algorithm in *no* way is stochastic, and does *not* involve a cumulative density function in statistical sense, $B(\hat{p}_i|\alpha, \beta)$ is used to denote the regularized incomplete beta function. In particular, if $B(\alpha, \beta)$ denotes the beta function defined by

$$B(\alpha, \beta) = \int_0^1 t^{\alpha-1}(1-t)^{\beta-1}dt,$$

then

$$B(\hat{p}|\alpha, \beta) = \frac{1}{B(\alpha, \beta)}\int_0^{\hat{p}} t^{\alpha-1}(1-t)^{\beta-1}dt.$$

The actuator algorithm is summarized as in Algorithm 1.

The regularized incomplete beta function is a standard function in most math libraries, e.g., in Matlab it is called `'betainc'`.

To fully appreciate the properties of beta actuation, consider the following examples.

Algorithm 1 Beta actuation

Configuration parameters: c
Input signals: $\hat{p}_i$, u
Output signals: a_i

Computation:
  $\alpha = cu$
  $\beta = c(1-u)$
  for all i
    $a_i = B(\hat{p}_i|\alpha, \beta)$
  end

Example 4.1. Fig. 4 illustrates how the actuator responds gracefully to variations in the estimated event rate $\hat{p}_i$ for a select few values of u and for the specific value of $c = 50$. The graceful behavior is of importance since event rate estimates in online advertising typically are subject to significant noise, and the noise may otherwise introduce a destabilizing disturbance in the feedback loop. Note how $\hat{p}_i \to 0 \Rightarrow a_i \to 0$ and how $\hat{p}_i \to 1 \Rightarrow a_i \to 1$ regardless of the value of u. As shown, a_i is monotonically increasing as a function of $\hat{p}_i$, and a_i tends to increase most rapidly for values of $\hat{p}_i \approx u$. ■

Example 4.2. Fig. 5 demonstrates how a_i varies as a function of $\hat{p}_i$ for different values of u and c. Each subplot corresponds to one value of c ($c = 5, 50, 500, 5000$), and the curves in each subplot correspond to different values of u (from left to right they are $u = 0, 0.05, 0.1, \ldots, 1$). The bid allocation a_i changes less abruptly for small values of c and approaches the indicator function $\mathbb{I}_{\{\hat{p}_i \ge u\}}$ when $c \to \infty$, where $\mathbb{I}_X = 1$, if X = true, and $\mathbb{I}_X = 0$, if X = false. ■

Example 4.3. Fig. 6 shows an example of campaign-level relationship between control signal u and event rate p, depicted in the block diagram in Fig. 2. This relationship depends on the distribution of available impressions with different event rates. Suppose the number of available impressions n_i per segment-level event rate $\hat{p}_i$ is as displayed in the bar

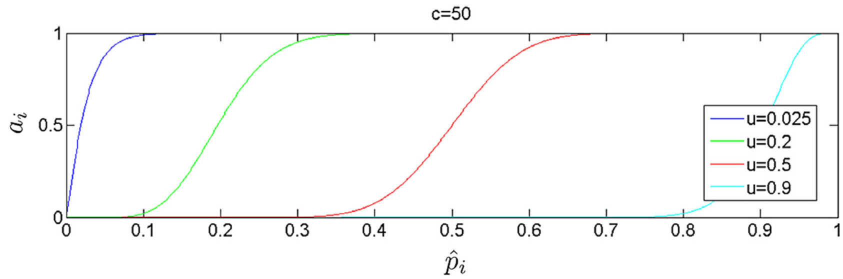

Fig. 4. The plot shows bid allocation a_i as a function of estimated event rate $\hat{p}_i$ for four different values of control signal u.

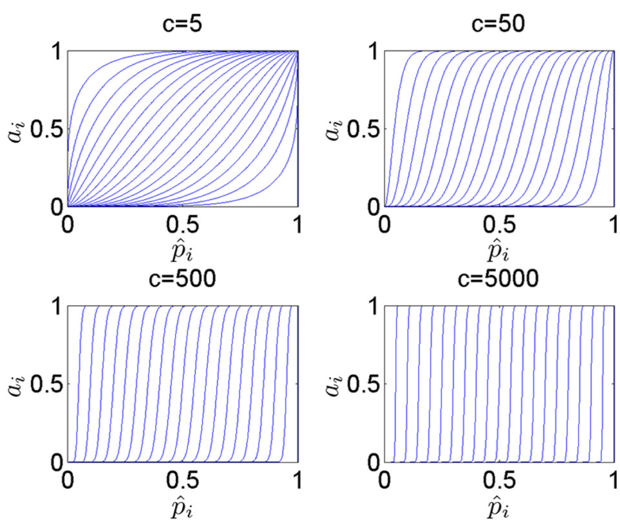

Fig. 5. The plots show bid allocation a_i as a function of estimated event rate $\hat{p}_i$ for $c = 5, 50, 500, 5000$, and for $u = 0, 0.05, 0.1, \dots, 1$ (left to right curve in each plot).

chart. All these impressions would have been awarded if $a_i = 1$ for all i. By adjusting u, which is the input to the beta actuator, a_i is regulated in such a way that the effective campaign-level event rate changes.

The four subplots on the right present the effective event rate p as a function of control signal u for $c = 5000, 500, 50, 5$. With $c = 5000$ the response curve is close to a discontinuous staircase function, while for a much smaller value of c, steps in the curve are virtually gone. In effect, the actuator makes the control problem less challenging. ■

5. Plant modeling and tuning

This section discusses plant modeling and tuning. The plant is defined by the mapping from the campaign-level control input u to the campaign-level output p as shown in Fig. 2. The input–output relationship $u \to p$ may be tuned using the beta actuation sensitivity parameter $c > 0$.

For simplicity of presentation and without loss of generality, in the sequel of this section *in-view rate control* for display advertising is considered. As mentioned in Section 2, an impression is considered viewable if 50% of the ad pixels are in view for more than one second (IAB, 2014). In the context of in-view rate control, an event is specifically an impression being measured as viewable by an Internet user.

The in-view rate is defined as a ratio of viewable impression volume to measured impression volume, where *measured impression volume* is the total number of served impressions that are measured by a certain viewability measurement technology (IAB, 2014).

The plant gain is first estimated based on data collected from a population of 200 eCPM[3] advertising campaigns. Fig. 7 shows the campaign-level in-view rate p versus control signal u (left) and the corresponding slopes dp/du vs u (right) in log scale, for four values of the beta actuator configuration parameter $c = 500, 500, 50, 5$. The slope value represents the effective plant gain and is of primary interest in what follows. Each curve in the plot is obtained by following the procedure as outlined in Example 4.3. In particular, the curve is generated by sweeping the control signal u from 0 to 1. For each value of u, a segment-level allocation signal a_i is calculated from Algorithm 1 for one specific configuration parameter c, according to the segment-level event rate estimate $\hat{p}_i$. The signal a_i is then used as a percentage to compute the viewable impression volume from the available measured impression volume. The campaign-level in-view rate (one point on the curve) for the specific u and c is obtained by aggregating the viewable and the measured impression volumes across all segments. Note that a smaller c value leads to smoother slope curves, and the choice of c is important in the tuning of the plant.

To obtain a generic model to use for control design when the same controller must work for any campaign, the percentile plots are further generated as shown in Fig. 8. Each point on the 95% curve in blue (as an example), is generated by sorting, from smallest to largest, the 200 data points for each specific u value, and selecting the 10th largest value. A larger c makes the control problem more challenging due to the large variations in the plant gain, while a smaller c may lead to a more conservative control design with sluggish control response. Here $c = 50$ is chosen, since it leads to a uniform plant gain over a large range of the control signal u, e.g., for u in between roughly 0.05 and 0.83.

A similar modeling and tuning procedure can be conducted for the rate control of other event types. For example, for a video campaign completion rate is defined as the total number of complete views versus the total number of impressions. The beta actuation sensitivity parameter $c = 50$ is selected for completion rate control, based on the percentile plots of the completion rate slope dp/du vs. control signal u, generated from data of a number of video eCPM campaigns. The profiling procedure is omitted here to avoid repetition. Experiment results on both in-view rate control and completion rate control will be shown in Section 7.

6. Control design

An event rate estimator is first presented that computes an estimate $\hat{p}$ of the campaign-level event rate p, as the feedback signal. A PI (proportional-integral) control scheme with windup protection is then employed for event rate control.

[3] An eCPM campaign is a campaign with an optimization objective of maximizing the total number of impressions for a given advertising budget. The "eCPM" stands for effective cost per thousand impressions.

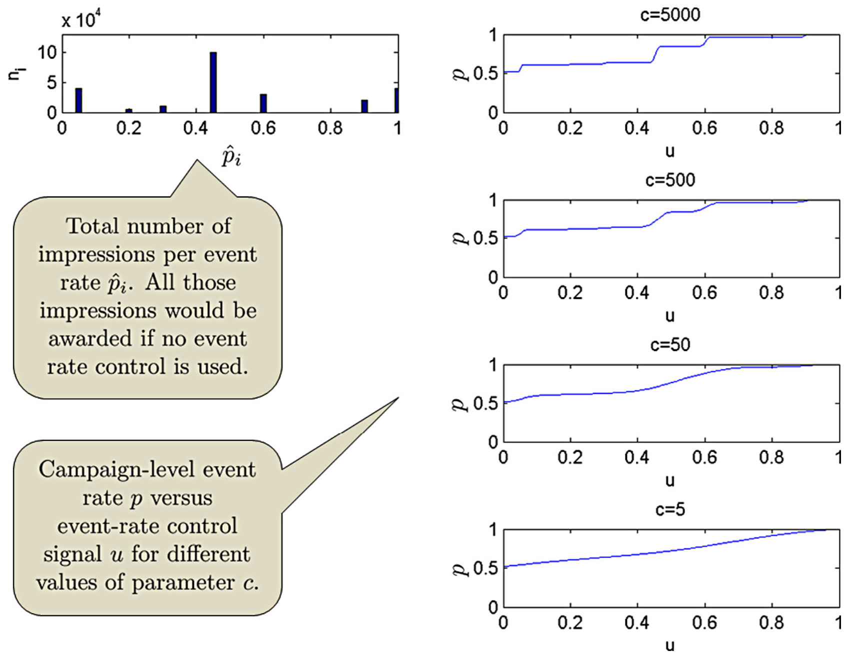

Fig. 6. Example of campaign-level relationship between control signal u and event rate p. The bar chart (left) shows the impression distribution n_i across event rates $\hat{p}_i$. The response curves on the right shows the effective campaign level event rate p as a function of u for different values of c.

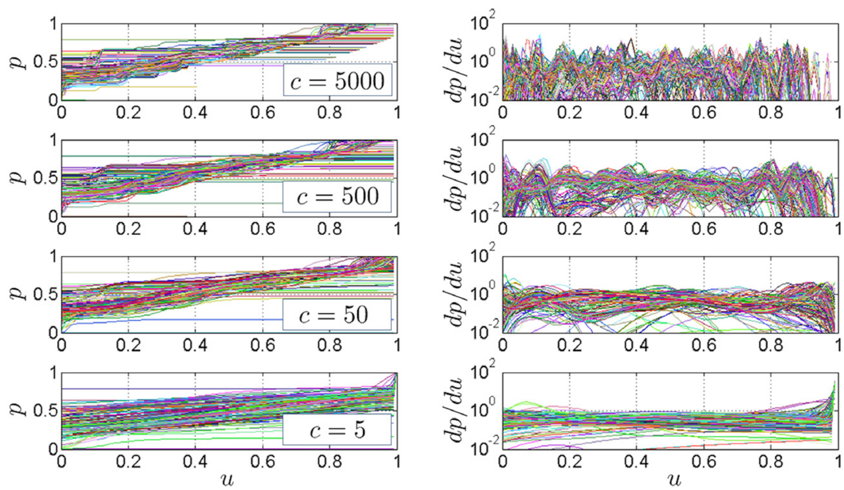

Fig. 7. In-view rate p (left) and in-view rate slope dp/du (right) vs. control signal u for select beta actuation configuration parameter c and for 200 representative ad campaigns.

6.1. Event rate estimator

Let $\{t_k\}, k = 0, 1, \ldots,$ denote the sampling time instants and h the sampling period; and let $n_{tot}(t_k)$ and $n_{event}(t_k)$ denote the total (across all segments) number of impressions and the total number of events, respectively, at time t_k. Let $\hat{p}(t_k)$ denote the campaign-level event rate estimate at time t_k. The estimate $\hat{p}(t_k)$ can be computed from the impression counts as follows (Karlsson, 2015, 2016a):

$$\alpha_p(t_k) = \lambda^h \alpha_p(t_{k-1}) + n_{event}(t_k), \quad \alpha_p(t_0) = \alpha_p^0$$
$$\beta_p(t_k) = \lambda^h \beta_p(t_{k-1}) + n_{tot}(t_k), \quad \beta_p(t_0) = \beta_p^0$$

where $\lambda \in (0, 1)$ is a design parameter, and

$$\hat{p}(t_k) = \frac{\alpha_p(t_k)}{\beta_p(t_k)}. \tag{1}$$

Note, if $n_{event}(t_k) \sim \text{Poisson}(n_{tot}(t_k)p)$ and our a priori belief of p satisfies $p \sim \text{Gamma}(\alpha_0, \beta_0)$, then the above estimator can be shown to be the optimal Bayesian estimator under a squared loss function, see Berger (1985) and Karlsson (2015, 2016a).

For in-view rate control, $n_{event}(t_k)$ is the total number of viewable impressions at time t_k, and $n_{tot}(t_k)$ is the total number of measured impressions at time t_k (or $n_{tot}(t_k)$ can simply be the total number of impressions for gross impression based in-view rate definition).

For completion rate control, $n_{event}(t_k)$ is the total number of complete views at time t_k, and $n_{tot}(t_k)$ is the total number of impressions at time t_k.

6.2. Event rate controller

The estimate $\hat{p}(t_k)$ is a measure of the system performance in terms of the average campaign-level event rate. The gap between this estimate and the user-specified event rate reference $p^{ref}(t_k) \in [0, 1]$ defines the error signal that drives the event rate controller.

A PI (proportional-integral) controller with windup protection (Åström & Hägglund, 2005) is employed to generate a control signal

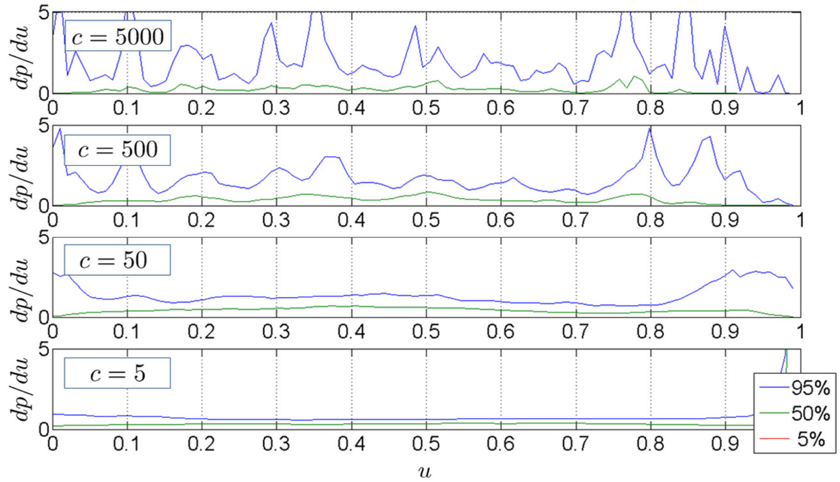

Fig. 8. Percentile plots of the in-view rate slope dp/du vs. control signal u for select Beta actuation sensitivity parameter c.

$u(t_k)$, to be used for the beta actuation. Let T_{int}^{norm} and T_{windup}^{norm} be design parameters that specify the time constants for the integrator and the correction as

$$T_{int} = T_{int}^{norm} h$$
$$T_{windup} = T_{windup}^{norm} h$$

The PI feedback control design is as follows (Åström & Hägglund, 2005):

$$e(t_k) = p^{ref}(t_k) - \hat{p}(t_k) \tag{2}$$
$$e_p(t_k) = bp^{ref}(t_k) - \hat{p}(t_k) \tag{3}$$
$$P(t_k) = K_p e_p(t_k)$$
$$I_{temp}(t_k) = I(t_{k-1}) + \frac{K_p h}{T_{int}} e(t_k), \quad I(t_0) = 0$$
$$u_{temp}(t_k) = P(t_k) + I_{temp}(t_k)$$

where b is the set-point weight, K_p is the proportional gain of the PI controller, and T_{int} is the integrator time constant. Let $\delta \in (0,1)$ be a parameter that specifies how much the control signal $u(t_k)$ is allowed to vary within a certain time unit, e.g., hour, and $u^{min}, u^{max} \in [0,1]$ with $u^{min} < u^{max}$ specify the hard limits $u(t_k)$ must be confined to (by default and in most practical situations $u^{min} = 0$ and $u^{max} = 1$). Note that u^{min} and u^{max} may change (infrequently) during a campaign flight. At each time instant t_k, $u_{low}(t_k)$ and $u_{high}(t_k)$ are defined as follows:

- if $u^{min} \geq u(t_{k-1}) + \delta h$ or $u^{max} \leq u(t_{k-1}) - \delta h$

$$u_{low}(t_k) = u^{min}, \quad u_{high}(t_k) = u^{max}$$

- else

$$u_{low}(t_k) = \max\left(u(t_{k-1}) - \delta h, u^{min}\right), \quad u(t_0) = u^{min}$$
$$u_{high}(t_k) = \min\left(u(t_{k-1}) + \delta h, u^{max}\right)$$

The control signal is then generated as

$$u(t_k) = \begin{cases} u_{low}(t_k), & \text{if } u_{temp}(t_k) < u_{low}(t_k) \\ u_{temp}(t_k), & \text{if } u_{low}(t_k) \leq u_{temp}(t_k) \leq u_{high}(t_k) \\ u_{high}(t_k), & \text{if } u_{temp}(t_k) > u_{high}(t_k) \end{cases}$$

Windup correction is added to the integrator term as

$$I(t_k) = I_{temp}(t_k) + \frac{h}{T_{windup}}\left(u(t_k) - u_{temp}(t_k)\right)$$

where T_{windup} is a design parameter.

Table 1
Summary of design parameters.

K_p	T_{int}^{norm}	T_{windup}^{norm}	λ	c
0.17	3.33	3.17	0.9	50

6.3. Selection of design parameters

The choice of design parameters is of significant importance to the overall control system performance. For example, the controller gain K_p is critical and should be chosen appropriately, for while a large K_p leads to faster system response, it may cause instability of the closed-loop system. On the other hand, too small a K_p is undesirable due to sluggish system response. Enough gain margin (GM) is also required such that the controller can deal with system uncertainties not captured by the plant model. This section outlines the procedure to choose the design parameters for the in-view rate controller, currently deployed to AOL's AdLearn™ campaign optimization engine.

As can be seen from Fig. 8, the 95% curve with $c = 50$ provides an estimate of the plant gain (almost "worst case scenario"), and its maximum occurs at $u = 0.91$ with a plant gain of 2.93. According to the Nyquist stability criterion, the inverse of the plant gain estimate gives an upper bound on the controller gain K_p for closed-loop stability. To achieve a robust design, a 6 dB ($\approx 20\log_{10}2$) gain margin is selected, which is obtained with a proportional gain $K_p = 0.17$.

As a rule of thumb for the time constants of the integrator and the windup correction, $h/T_{int} \in [0.1, 0.3]$, and $T_{windup} < T_{int}$ (Åström & Hägglund, 2005). The two time constants are then chosen as $T_{int}^{norm} = 3.33$ and $T_{windup}^{norm} = T_{int}^{norm}/1.05$, such that $h/T_{int} = 0.3$ and $T_{windup} = T_{int}/1.05$. Furthermore, $\lambda = 0.9$ is chosen for the event rate estimator.

Table 1 summarizes the design parameter choices for the PI controller and the event rate estimator, to implement in-view rate control in AdLearn™.

7. Experiment results

In this section, the performance of the proposed event rate control scheme has been evaluated in a simulated environment for in-view rate control, and on real advertising campaigns for in-view rate control, as well as for video completion rate control, on the AdLearn™ advertising optimization engine by AOL.

7.1. Simulation result for in-view rate control

The proposed control system is first evaluated in a simulated environment for in-view rate control. The plant is defined as a campaign

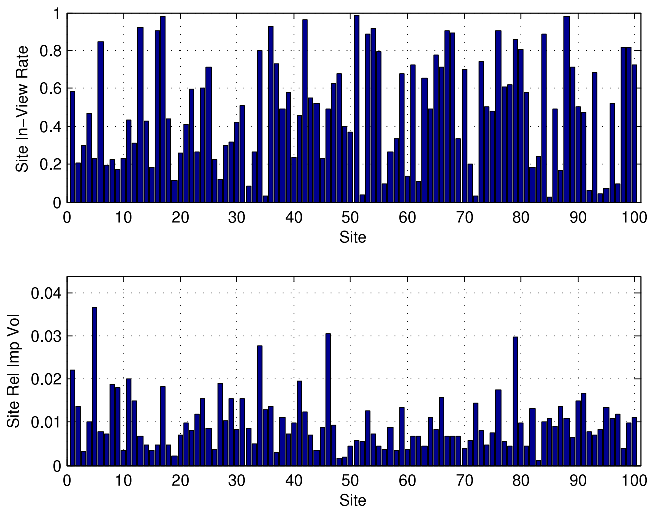

Fig. 9. Site-level in-view rates p_i and relative impression volume $n^{rel}_{avail,i}$. The bar chart at top shows the in-view rate for each of the 100 sites. The bar chart at the bottom shows the corresponding relative impression volume for each site.

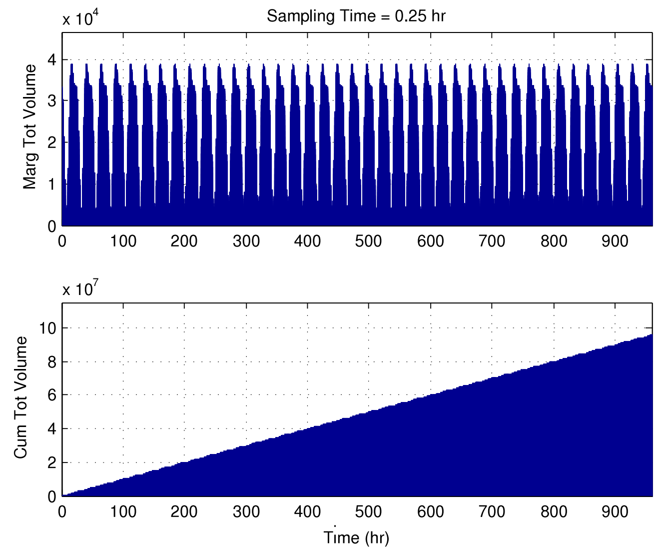

Fig. 10. Marginal (top) and cumulative (bottom) total impression volumes over all sites.

Table 2
Summary of simulation parameters.

b	δ	h	u^{min}	u^{max}	β_1	β_2	ϕ_1	ϕ_2
1	0.1	0.25	0	0.9	0.63	2.76	0.26	0.39

with a total of $n^{daily}_{avail} = 2.4 \times 10^6$ available measured impressions per day, randomly distributed over 100 sites (segments). The relative impression count per site is given by a (normalized) random number generated from a Gamma distribution with a relative standard deviation of 0.6. In particular, for each site a random number is drawn from $\mathrm{Gamma}(\alpha, \beta)$ with the shape parameter $\alpha = 1/\sigma^2$ and the scale parameter $\beta = \sigma^2$, where $\sigma = 0.6$. The site-level relative impression volume is given by the corresponding random number over the sum of all 100 random numbers. Site-level in-view rates are generated from a $\mathrm{Uniform}(0, 1)$ distribution. The resulting site-level in-view rates and relative available impression volume are shown in Fig. 9.

To capture a realistic time-of-day pattern in Internet traffic, the daily available impression counts are distributed throughout the day according to $n_{avail}(t_k) = \frac{n^{daily}_{avail}}{24}\left[1 + \beta_1 \sin\left(\frac{2\pi}{24}t_k + \phi_1\right) + \beta_2 \sin\left(\frac{2\pi}{12}t_k + \phi_2\right)\right]$, where the parameters have been summarized in Table 2, along with others (see also Table 1).

Fig. 10 shows the marginal (top) and cumulative (bottom) total impression volumes over all sites for a simulation duration of 960 h. (40 days). Note that the marginal impression volume displays a time-of-day (TOD) pattern.

The control performance is illustrated in Fig. 11 with the campaign-level average *in-view rate* (IVR) $\hat{p}$ (top) as computed in (1), the control signal u (middle), and the total awarded impression volume n_{meas} and viewable impression volume n_{view} (bottom). In particular, a case is simulated in which the advertiser changes the in-view rate reference signal p^{ref}, as shown with a red dashed line in Fig. 11 (top). By computing u to drive the beta actuator, the proposed event rate controller regulates $\hat{p}$ to p^{ref}.

Note when p^{ref} is set high, e.g., $p^{ref} = 0.95$, during the first 120 h, very few impressions from low IVR sites are awarded, which implies a low total awarded impression count. An under-delivery of the ad budget follows. This is due to insufficient impression inventories with relatively high IVR. In fact, since in the simulated scenario $u^{max} = 0.9$, the actuator is saturated. When p^{ref} is lowered to a less extreme level of 0.7 between hours 120 and 360, it can be tracked very well. However, if p^{ref} is set too low, the control signal u may be saturated to the low limit of $u^{min} = 0$ between hours 360 and 600 when $p^{ref} = 0.3$. The control scheme handles saturation well in either case and the system quickly recovers from saturation.

7.2. In-view rate control of a real advertising campaign

Fig. 12 illustrates the implementation of the proposed actuation and control mechanism as an enhancement to AOL's AdLearn™ optimization engine, to achieve event rate control. The bid price b_i and bid allocation a_i are computed by the AdLearn controller separately from the event rate control scheme, to fulfill the pacing and value maximization objective subject to advertising budget and/or ROI constraints, see Karlsson (2016b) for an overview of the optimization problem under consideration in AdLearn™. Our objective in this section is to demonstrate the control performance when integrating the proposed actuation and control mechanism to AdLearn™.

Note that with a slight abuse of notation, in Fig. 12 the bid allocation signal generated by the beta actuator is denoted by a^B_i. It is an adjustment to the bid allocation signal, a_i, generated by the AdLearn™ actuator. The final bid allocation is $a'_i = a_i \times a^B_i$, for the ith segment.

Fig. 13 shows the IVR control performance for a real advertising campaign. The control objective is to maximize the viewable impression volume, while delivering a given budget smoothly and in full, and keeping a campaign-level IVR at or above a specified reference level p^{ref}. The IVR control was activated on 10/08/2016 with a reference signal $p^{ref} = 0.5$ initially, which was then increased first to 0.6, then to 0.7, and finally to 0.8 (green line in the bottom left plot). From the control signal u (bottom right plot), the actuator was essentially saturated to the lower limit 0 until about 10/15/2016. This is because the targeted impression inventories of the campaign all have higher IVR than the specified reference of 0.5 and 0.6 during this time period. For the rest of the campaign flight, it is clear that $\hat{p}$ (red curve in the bottom left plot) tracks p^{ref} closely.

7.3. Completion rate control of a real advertising campaign

For simplicity of design and campaign management, the same design parameters as those used for in-view rate control have been employed, as summarized in Table 1. This section shows the control performance for a video campaign under completion rate control in AdLearn™.

The plots in Fig. 14 illustrate the performance of a video campaign with a 2-week campaign duration, starting from 05/17/2017 and ending

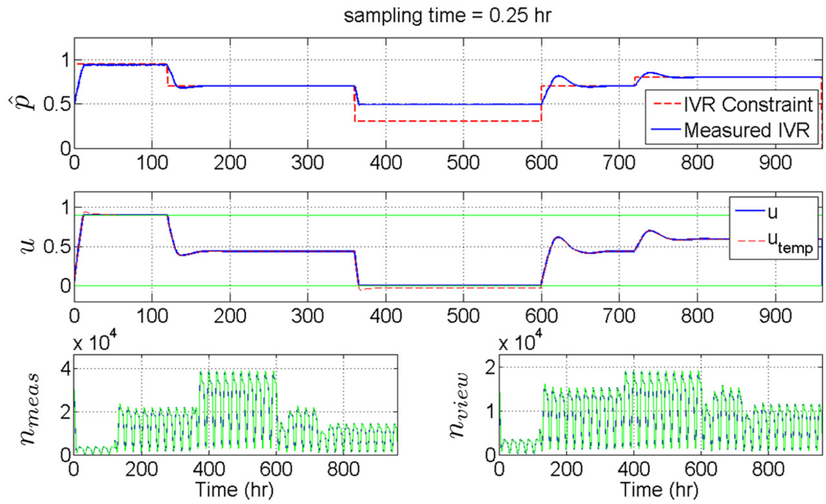

Fig. 11. Simulation results: campaign-level control system performance. The top plot shows the time history of estimated in-view rate (solid) as compared to the reference signal (dashed). The plot in the middle shows the control output (solid) of the in-view rate controller, as well as the signal before saturation. The plots on the bottom show the time history of the marginal total measured and viewable impression volumes.

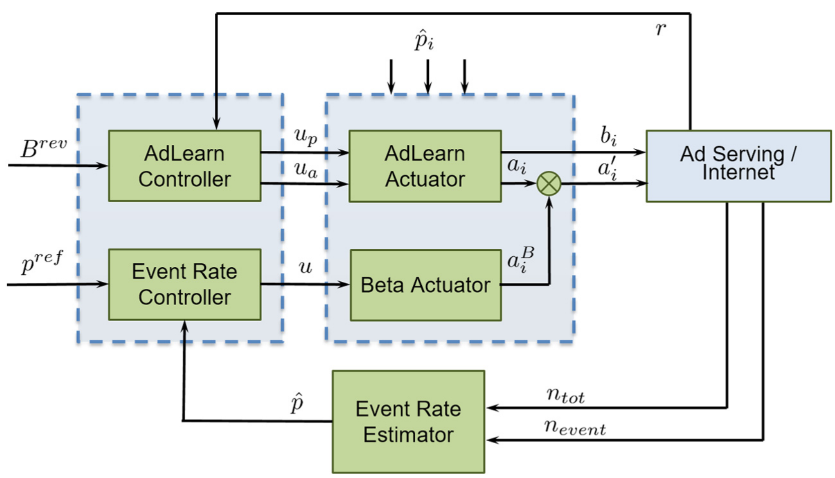

Fig. 12. Event rate control as an enhancement to AOL's AdLearn™ optimization engine.

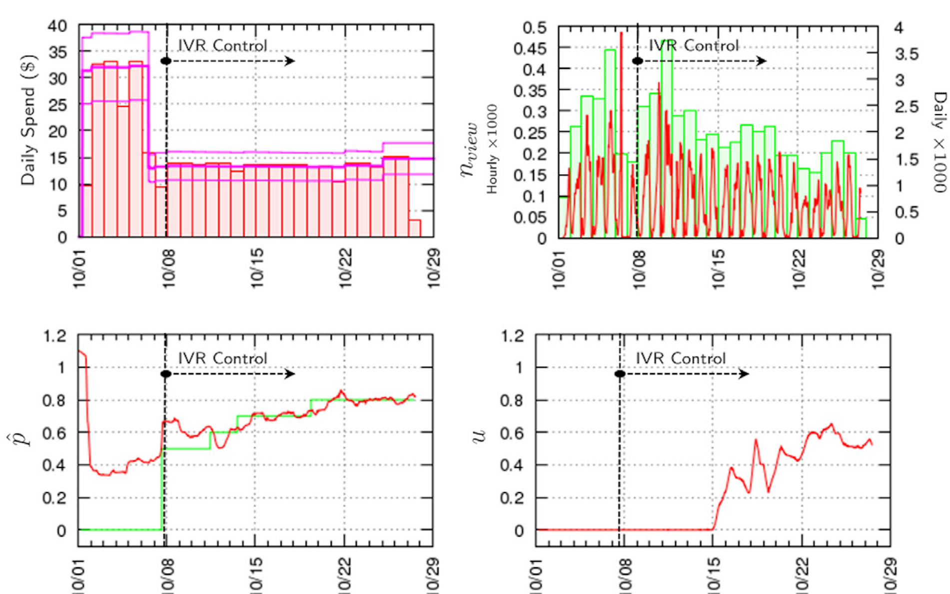

Fig. 13. Experiment results: campaign-level in-view rate control performance.

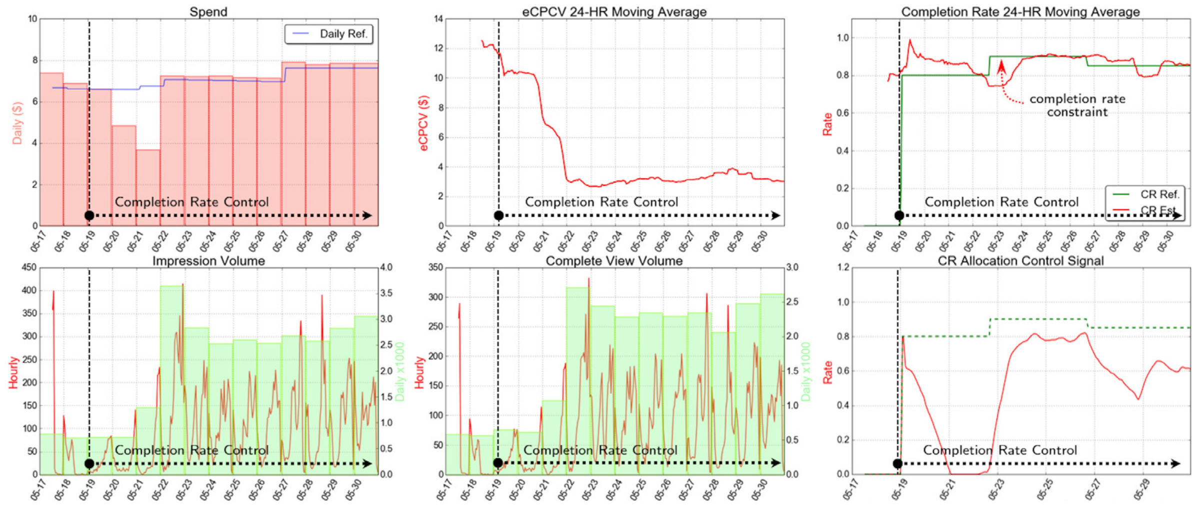

Fig. 14. Experiment results: campaign-level completion rate control performance. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

on 05/31/2017. Completion rate control was activated on 05/19/2017. It is clear from the "Spend" plot (top left) that pacing performance was satisfactory with daily budget delivery hitting daily reference almost every day after some initial transient period (05/20/2017–05/21/2017). Pacing was regulated by the AdLearn™ controller, separate from the completion rate controller. The average completion rate (the red line of the "Completion Rate 24-HR Moving Average" plot) met or beat the reference (the green line) most of the time, which was initially set at 0.8, increased to 0.9 on 05/22/2017, and reduced to 0.85 thereafter until the end of the campaign. There were times when the average completion rate was below the specification, e.g., during 05/22/2017–05/23/2017 and around 05/29/2017. This was due to sudden competitive bidding landscape changes. In spite of these adverse impacts, the completion rate control mechanism was able to bring the completion rate back up to track the reference within a day, by quickly ramping up the completion rate allocation control signal. Considering the extremely small budget of this campaign (about $6 daily spending), known to be difficult for performance optimization, the rate control performance is especially desirable.

8. Concluding remarks and future work

An approach to actuation and feedback control of the average event rate of online advertising campaigns has been proposed in this paper. In order to obtain a scalable solution, the proposed system consists of a static actuator module consuming segment-level information, and a feedback controller module consuming only campaign-level information. The actuator module effectively turns the input–output relationship of the controlled plant continuous, which reduces the challenges for feedback control. The feedback controller module employs a PI controller with windup protection to achieve reference tracking. The resulting control system has been evaluated with simulations, as well as on real advertising campaigns for in-view rate control and video completion rate control. Extensive simulation and experiment results demonstrate the excellent event rate control performance of the proposed scheme, which has been integrated to the production environment and offered as viewability and video completion optimization products in the AdLearn™ advertising optimization system, developed at AOL.

The proposed actuation and control scheme provides advertisers with a lever to balance campaign performance and advertising cost. Compared with threshold targeting, the proposed solution is more economically sound, and prevents potential errors and/or performance degradation by removing the need of any manual adjustments. Furthermore, the scheme is generic in that it can be readily configured to regulate the event rate of any campaign effectiveness metric of interest to advertisers.

In real advertising campaigns, we have found cases for which the event rate reference p^{ref} is set relative high, while the inventory is dominated by impressions with low event rate. The control signal u is adjusted higher until it saturates to the upper bound. As a result, the campaign may go dark, i.e., no impression is won, and it cannot recover on its own without human intervention. One remedy to this situation is to make the beta actuation sensitivity parameter c larger, to reduce the impact due to the dominance of low event rate impressions. As for future work, we are developing an adaptive scheme that can automatically adjust the sensitivity parameter c, depending on the event rate control performance.

References

AdExchanger, 2015. Viewability and the evolution of online video measurement. https://adexchanger.com/tv-and-video/viewability-and-the-evolution-of-online-video-measurement/. [Online] (accessed December 03, 2017).

Åström, K. J., & Hägglund, T. (2005). *Advanced PID control*. ISA, the International Society of Automation.

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis* (2nd ed.). Springer-Verlag.

Busch, O. (2016). *Programmatic Advertising: the Successful Transformation to Automated, Data-Driven Marketing in Real-Time* (1st ed.). Springer.

Casella, G., & Berger, R. (2001). *Statistical Inference* (2nd ed.). Duxbury.

Edelman, B., Ostrovsky, M., & Schwarz, M. (2007). *Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords*. Technical Report, National Bureau of Economic Research.

eMarketer, (2014). Advertisers will spend nearly $600 billion worldwide in 2015. http://www.emarketer.com/Article/Advertisers-Will-Spend-Nearly-600-Billion-Worldwide-2015/1011691.

eMarketer, (2017). US ad spending: the emarketer forecast for 2017. https://www.emarketer.com/Report/US-Ad-Spending-eMarketer-Forecast-2017/2001998.

Facebook, (2015). New: 100% in-view impression buying and ad analytics partnership with MOAT. https://www.facebook.com/business/news/100-in-view-impressions-and-moat-partnership.

Google, (2011). *The Arrival of Real-Time Bidding*. Technical Report. Google.

Guo, J., & Karlsson, N. (2017). Model reference adaptive control of advertising systems. In *Proceedings of the 2017 American Control Conference*, Seattle, June 2017, (pp. 5482–5487). http://dx.doi.org/10.23919/ACC.2017.7963807.

IAB, (2014). State of viewability transaction 2015. https://www.iab.com/guidelines/state-of-viewability-transaction-2015/.

Karlsson, N. (2014). Adaptive control using Heisenberg bidding. In *Proceedings of the 2014 American Control Conference* (pp. 1304–1309). http://dx.doi.org/10.1109/ACC.2014.6859107.

Karlsson, N. (2015). *Tutorial in Bayesian Decision Theory with Applications in Event Rate Estimation*. Technical Report, AOL Platforms R&D.

Karlsson, N. (2016a). Adaptive estimation of small event rates. In *Proceedings of the 55th IEEE Conference on Decision and Control* (pp. 3732–3737). http://dx.doi.org/10.1109/CDC.2016.7798831.

Karlsson, N. (2016b). Control problems in online advertising and benefits of randomized bidding strategies. *European Journal of Control*, *30*, 31–49.

Karlsson, N., & Sang, Q. (2017). Event rate control in online advertising. In *Proceedings of the 2017 American Control Conference*, (pp. 4215–4220). http://dx.doi.org/10.23919/ACC.2017.7963603.

Karlsson, N., & Zhang, J. (2013). Applications of feedback control in online advertising. In *Proceedings of the 2013 American Control Conference* (pp. 6008–6013). http://dx.doi.org/10.1109/ACC.2013.6580779.

Krishna, V. (2002). *Auction Theory*. Academic Press.

Moha, N. (2015). Toward viewability: You can't count what you haven't measured. https://www.thinkwithgoogle.com/marketing-resources/data-measurement/toward-viewability-advertising-measurement/. [Online] (accessed December 03, 2017).

Wang, J., Zhang, W., Yuan, S., (2016). Display advertising with real-time bidding (RTB) and behavioural targeting. ArXiv e-prints arXiv:1610.03013.

Zhang, W., Rong, Y., Wang, J., Zhu, T., & Wang, X. (2016). Feedback control of real-time display advertising. In *Proceedings of the 9th ACM International Conference on Web Search and Data Mining* (pp. 407–416). New York, NY, USA: ACM. http://dx.doi.org/10.1145/2835776.2835843.